

---

# RL-PRUNER: STRUCTURED PRUNING USING REINFORCEMENT LEARNING FOR CNN COMPRESSION AND ACCELERATION

---

A PREPRINT

**Boyao Wang** \* **Volodymyr Kindratenko**  
University of Illinois Urbana-Champaign  
boyaow2@illinois.edu, kindrtnk@illinois.edu

## ABSTRACT

Convolutional Neural Networks have demonstrated exceptional performance in recent years. Compressing these models not only reduces storage requirements, making deployment to edge devices feasible, but also accelerates inference, thereby reducing latency and computational costs. Structured pruning, which removes filters at the layer level, directly modifies the model architecture. This approach achieves a more compact architecture while maintaining target accuracy, ensuring that the compressed model retains good compatibility and hardware efficiency. Our method is based on a key observation: filters in different layers of a neural network have varying importance to the model’s performance. When the number of filters to prune is fixed, the optimal pruning distribution across different layers is uneven to minimize performance loss. Layers that are more sensitive to pruning should account for a smaller proportion of the pruning distribution. To leverage this insight, we propose RL-Pruner, which uses reinforcement learning to learn the optimal pruning distribution. RL-Pruner can automatically extract dependencies between filters in the input model and perform pruning, without requiring model-specific pruning implementations. We conducted experiments on models such as GoogleNet, ResNet, and MobileNet, comparing our approach to other structured pruning methods to validate its effectiveness. Our code is available on <https://github.com/Beryex/RLPruner-CNN>.

## 1 Introduction

Convolutional neural networks (CNNs) have demonstrated outstanding performance across a range of computer vision tasks, including image classification, detection, and segmentation [Krichen, 2023, Bharadiya, 2023, O’Shea and Nash, 2015, Gu et al., 2018, Li et al., 2022]. As these architectures become wider and deeper, they gain an enhanced ability to extract complex features from input data. However, this increase in parameters leads to substantial computational costs, making inference both resource-intensive and slow. Moreover, there is a growing need to deploy CNNs on edge devices, which are often limited in computational power and memory [Rashid et al., 2022, Stahl et al., 2021]. Consequently, effective methods for compressing CNNs are increasingly in demand, aiming to reduce parameter counts and accelerate inference while maintaining the original model’s performance. Such compression techniques are crucial for creating efficient and deployable CNNs that can meet the challenges of real-world applications.

Several techniques have been proposed to compress CNNs, most of which fall into one of four categories: structured and unstructured pruning, quantization, low-rank factorization, and knowledge distillation [Cheng et al., 2020, Deng et al., 2020]. Structured pruning [Fang et al., 2023, He and Xiao, 2024, Anwar et al., 2017] removes entire filters from neural networks and directly modifies the model’s architecture, enabling both compression and realistic acceleration on standard hardware. Unstructured pruning [Zhang et al., 2018, Ma et al., 2022], also known as weight pruning, removes specific unimportant elements from the weight matrix, which requires hardware or libraries that support sparse computation to achieve practical speedup. Low-rank factorization [Swaminathan et al., 2020, Idelbayev and Carreira-Perpinan, 2020] approximates the model’s weight matrix by decomposing it into a product of low-rank matrices. Quantization [Wu et al., 2016, Gong et al., 2014, Zhou et al., 2017] reduces the bitwidth of the weight data in a model, achieving significant compression, but also requires hardware support to realize theoretical speedups for low-bit quantization, such as binary or ternary quantization. Knowledge distillation [Gou et al., 2021, Cho and Hariharan, 2019] transfers knowledge from a larger, more advanced teacher model to a smaller student model, which

---

\*Corresponding author

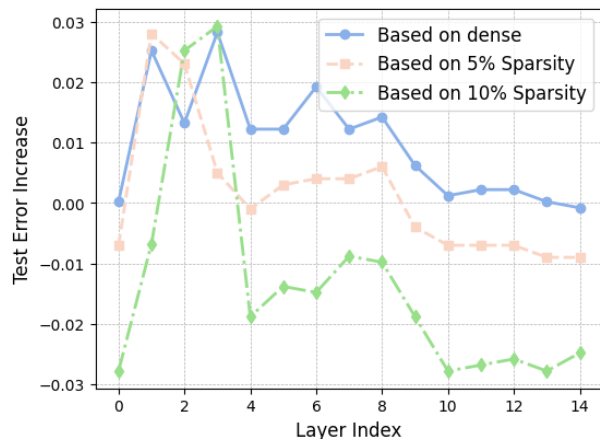


Figure 1: Increase in test error for a compressed VGG-16 model after pruning 10% of filters from each individual layer. The pruning is conducted based on three different baselines: the dense model, a model with 5% of filters pruned from all layers, and a model with 10% of filters pruned from all layers. The results highlight that different layers have varying sensitivity to pruning, and this sensitivity changes dynamically throughout the pruning process. A negative increase indicates that performance improves.

may not generate a new compressed architecture but helps recover the performance loss due to compression. In this paper, we primarily focus on structured pruning, though we hope the insights from our methods will also inform future work in other categories.

In structured pruning, while it is crucial to identify the least important filters to prune from weight matrices with minimal performance degradation, it is equally important to assess the importance of different layers to the model’s performance and determine which layers should be pruned more or less. This involves learning the sparsity distribution across layers. As shown in Figure 1, the importance of different convolutional and linear layers to the overall model performance varies significantly. Moreover, after a certain extent of pruning, the relative importance among different layers also changes. To leverage this insight, we can assign higher filter sparsity to less important layers and lower sparsity to more important layers, adjusting the sparsity distribution dynamically throughout the pruning process to minimize the performance drop.

In this paper, we introduce a novel approach called RL-Pruner. Unlike other structured pruning methods that learn the sparsity distribution among layers during sparse training [Liu et al., 2017a, Ding et al., 2021, Fang et al., 2023], RL-Pruner is a post-training structured pruning method that utilizes reinforcement learning with sampling to learn the optimal sparsity distribution across layers for several pruning steps. At each step, the current model architecture is defined as the state, while the pruning sparsity distribution acts as the policy. RL-Pruner adds Gaussian noise to the policy distribution to generate the real pruning sparsity distribution as an action, producing the corresponding

compressed architecture as the next step’s state. Each generation is treated as a sampling process. For each step, RL-Pruner maintains a replay buffer that stores the pruning sparsity distribution actions and their corresponding Q values, computed by the reward function. The reward function can be defined flexibly, such as based on the compressed model’s test error, a combination of test error and parameter reduction ratio, or other criteria depending on the specific use case. After each pruning step, RL-Pruner updates the sparsity distribution policy to learn the optimal sparsity distribution. When computational resources allow, post-training stages are periodically applied after several pruning steps to recover any performance loss caused by pruning. We employ knowledge distillation in these post-training stages, with the original model acting as the teacher and the compressed model as the student.

RL-Pruner can automatically extract dependencies between filters in different layers of the input model by tracking tensor computations. Currently, our approach supports several popular CNN architectures, including residual connections, concatenation connections and skip connections. As a result, our method does not require model-specific pruning implementations and can perform pruning autonomously, enhancing the method’s generality.

To validate the effectiveness of RL-Pruner, we apply the proposed method to several popular CNNs used for image classification, including VGGNet [Simonyan and Zisserman, 2014], ResNet [He et al., 2016], GoogLeNet [Szegedy et al., 2015], and MobileNet [Howard, 2017], using the CIFAR-10 and CIFAR-100 datasets [Krizhevsky et al., 2009]. According to the experimental results, RL-Pruner achieves 60% channel sparsity for VGG-19 on CIFAR-100 and 40% channel sparsity for GoogLeNet and MobileNetV3-Large on CIFAR-100, all with a performance drop of less than 1%.

## 2 Related Work

**Pruning** Pruning is one of the mainstream neural network compression methods and has made impressive progress [Fang et al., 2023, He and Xiao, 2024, Anwar et al., 2017, Zhang et al., 2018, Ma et al., 2022]. Formally, given a neural network containing  $L_c$  convolutional layers and  $L_f$  fully connected layers, that is,  $\mathcal{N} = \{W_l \in R^{N_l^{out} \times N_l^{in} \times K_l \times K_l}, 1 \leq l \leq L_c\} \cup \{W_l \in R^{N_{l+1} \times N_l}, 1 \leq l \leq L_f\}$ , where  $N_l^{out}$ ,  $N_l^{in}$  denotes the output and input channel number of the  $l_{th}$  layer respectively. Each layer contains  $N_{l+1}$  filters, with each filter having a shape of  $N_l \times K_l \times K_l$  for convolutional layers and a shape of  $N_l$  for fully connected layers. Unstructured pruning removes individual weights, resulting in sparsified filter matrices, which do not lead to inference speedups unless specialized hardware is available. In contrast, structured pruning operates at a larger granularity, removing entire filters and achieving both compression and realistic acceleration on standard hardware. Several pruning criteria have been proposed, including norm-based methods that

prune weights or filters with the least norms [Li et al., 2017, He et al., 2018a] and activation-based methods that use the activation map to identify unimportant weights or filters [He et al., 2017, Lin et al., 2020a, Luo et al., 2017], and regularization-based methods that learn structured sparse networks by adding different sparsity regularizers [Hoefler et al., 2021, Liu et al., 2017b]. In this paper, we adapt the Taylor criterion, which combines the gradient of a weight, calculated from a calibration dataset, with its norm to determine the importance of weight filters.

**Neural Architecture Search** Structured pruning removes filters and generates compressed subnetworks, which can be viewed as a form of neural architecture search (NAS). Since different layers in a model have varying importance to its overall performance, several works have focused on automatically assigning layer-wise sparsity. Among these, [He et al., 2018b, Yu et al., 2021, 2022] deploy reinforcement learning agents to automatically select the appropriate layer-wise pruning ratio by searching over the action space, i.e., the pruning ratio, without requiring manual sensitivity analysis. Gradient-based methods [Guo et al., 2020, Ning et al., 2020] modify the gradient update rule to make the optimization problem with sparsity constraints differentiable with respect to weights. Evolutionary methods [Liu et al., 2019, Lin et al., 2020b] use evolutionary algorithms to explore and search for sparse subnetworks. Our method primarily uses reinforcement learning to determine the optimal sparsity distribution across layers.

**Knowledge Distillation** Pruning typically reduces a network’s performance, so post-training is often employed to recover this loss. Knowledge distillation [Hinton, 2015] transfers knowledge from a larger teacher model to a smaller student model. Response-based knowledge distillation allows the student model to mimic the predictions from the teacher model’s final output layer by adding a distillation loss between their predictions to the optimization function. Feature-based knowledge distillation, on the other hand, enables the student model to learn intermediate layer representations from the teacher model. In our method, we adopt response-based knowledge distillation during the recovery post-training stage, as pruning alters the dimensions of intermediate layers, making knowledge transfer between these layers less straightforward.

### 3 Method

RL-Pruner first builds a dependency graph between layers in the model, then performs pruning in several steps. In each step: 1) A new pruning sparsity distribution is generated as an action based on the base distribution, which serves as the policy; 2) Each layer is pruned using the Taylor criterion according to the corresponding sparsity; 3) The compressed model is evaluated to obtain a reward, and the action and reward are stored in a replay buffer. After each step, the base distribution is updated according to the replay buffer, and if computational resources are sufficient,

a post-training stage is applied to the compressed model using knowledge distillation, where the original model acts as the teacher. Figure 2 illustrates our method.

#### 3.1 Build Dependency Graph

**Definition** Given a convolutional neural network  $\mathcal{N} = \{W_l \in R^{N_l^{out} \times N_l^{in} \times K_l \times K_l}, 1 \leq l \leq L_c\} \cup \{W_l \in R^{N_l^{out} \times N_l^{in}}, 1 \leq l \leq L_f\}$ , we define convolutional and linear layers as prunable layers  $\{\mathcal{L}_i^P, 1 \leq i \leq L_c + L_f\}$ , while activation layers, and pooling layers are considered non-prunable layers. Batch normalization layers are also considered non-prunable, as they are only pruned if their preceding convolutional or linear layers are pruned. Formally, we establish the dependency between any two layers  $\mathcal{L}_i, \mathcal{L}_j, i \leq j$  by,

$$DG_{(i,j)}(k_{out}) = [k_{in}], \forall 1 \leq k_{out} \leq N_i^{out} \quad (1)$$

if a tensor flows from  $\mathcal{L}_i$  to  $\mathcal{L}_j$  with no intermediate layer between them. This means that if we prune a specific output channel  $k_{out}$  in  $\mathcal{L}_i$ , the corresponding input channel  $k_{in}$  in  $\mathcal{L}_j$  must also be pruned. If there is no dependency between  $\mathcal{L}_i$  and  $\mathcal{L}_j$ , we have:

$$DG_{(i,j)}(k_{out}) = [-1], \forall 1 \leq k_{out} \leq N_i^{out} \quad (2)$$

Our goal is to build a dependency graph  $DG$  among all layers, and record the relationships among their channel index mappings.

**Dependency Graph Construction** To construct the dependency graph, we fed the model an example input image and extract all layers  $\{\mathcal{L}_i, 1 \leq i \leq L\}$  along with their input tensor  $\{T_l^{in}, 1 \leq l \leq L\}$  and output tensor  $\{T_l^{out}, 1 \leq l \leq L\}$ , where  $L$  is the number of all layers. We then build the dependency graph by matching the input and output tensors of each layer. A basic dependency occurs when the output tensor of one layer  $\mathcal{L}_i$  is directly fed into layer  $\mathcal{L}_j$ , which implies for basic dependency

$$DG_{(i,j)}(k_{out}) = [k_{out}], \forall 1 \leq k_{out} \leq N_i^{out} \quad (3)$$

We first iterate over  $\{T^{in}\}$  and  $\{T^{out}\}$  to check whether any two tensors match, in order to detect all the basic dependencies. Next, we check for any unused input tensor, e.g.  $T_a^{in}$ , which may be caused by tensor modifications (e.g., flattening) or special connections between layers (e.g., residual or concatenation connections). Flattening is commonly used between convolutional and linear layers. To detect the use of tensor flattening, we also flatten all output tensors and compare them to the unused input tensors to build the dependency. Let us denote the detected output tensor as  $T_a^{out}$ , which represents the tensor before flattening. To determine the mapping relationship after flattening, we need to check the output area  $A_{out}$  of the preceding layer that produces  $T_a^{out}$ . For flattening, we have

$$DG_{(i,j)}(k_{out}) = [k_{out} * A_{out}, (k_{out} + 1) * A_{out}] \quad (4)$$

which indicates the output channel  $k_{out}$  of  $\mathcal{L}_i$  and the input channel from  $k_{out} * A_{out}$  to  $(k_{out} + 1) * A_{out}$  must be

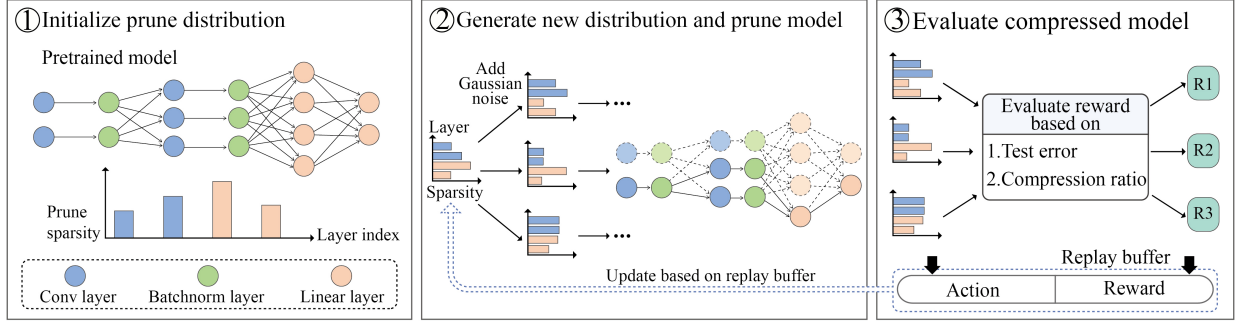


Figure 2: Illustration of RL-Pruner: We start by initializing a uniform pruning distribution based on the output channel density of each prunable layer. Next, we iteratively generate a new distribution, evaluate the corresponding compressed model to obtain a reward, and update the pruning distribution using the replay buffer.

pruned simultaneously. To detect concatenation connections, we iterate over  $\{T^{out}\}$  to check whether there exists  $T_b^{out}$  such that

$$T_a^{in} = \text{concat}(T_b^{out}, T_j^{out}), 1 \leq j \leq L, j \neq b \quad (5)$$

This can be detected by slicing  $T_b^{out}$  over  $T_a^{in}$  such that  $T_a^{in}[i_b, i_b + N_b^{out}] = T_b^{out}$ . We can then build the dependency as

$$DG_{(b,a)}(k_{out}) = [i_b + k_{out}] \quad (6)$$

To detect residual connections, we iterate over  $\{T^{out}\}$  to check whether there exists  $T_b^{out}, T_c^{out}$  such that

$$T_a^{in} = T_b^{out} + T_c^{out}, a \neq b, a \neq c, b \neq c \quad (7)$$

We can then build the dependency as

$$DG_{(b,a)}(k_{out}) = [k_{out}] \quad (8)$$

$$DG_{(c,a)}(k_{out}) = [k_{out}] \quad (9)$$

This also applies to the Squeeze-and-Excitation Module, where

$$T_a^{in} = T_b^{out} * T_c^{out}, a \neq b, a \neq c, b \neq c \quad (10)$$

For residual connections and the Squeeze-and-Excitation Module, tensors from layers  $\mathcal{L}_b$  and  $\mathcal{L}_c$  are either added or multiplied and then passed to  $\mathcal{L}_a$ , meaning that pruning output channel  $k_b^{out}$  of  $\mathcal{L}_b$  also requires pruning output channel  $k_c^{out}$  of  $\mathcal{L}_c$  and vice versa. As a result, layers  $\mathcal{L}_b$  and  $\mathcal{L}_c$  must be pruned simultaneously. After constructing the dependency graph, we divide the layers that need to be pruned simultaneously into disjoint sets, ensuring that layers within the same set are pruned together.

### 3.2 Assigning Layer-wise Sparsity

After automatically extracting the layer dependencies and corresponding channel index mappings, we still need to determine how to distribute sparsity among layers for each pruning steps to minimize the performance drop caused by pruning while achieving the desired overall model sparsity. Let the pruning distribution  $PD = \{\mathbf{p} \in [0, 1]^L \mid \sum_{i=1}^{L_c+L_f} p_i = 1\}$  represent the sparsity distribution among

prunable layers, where we assign sparsity  $p_i * S$  to the  $i_{th}$  prunable layer, with  $S$  being the target overall model sparsity. Since we can only prune output channels in integer quantities, the optimization problem of finding the optimal  $PD$  is not differentiable.

To address this, we adopt a Monte Carlo sampling strategy [Rubinstein and Kroese, 2016] combined with the Q-learning algorithm [Watkins and Dayan, 1992] from reinforcement learning. Specifically, Monte Carlo sampling is used to explore different pruning distributions, while Q-learning updates the policy, represented by the base pruning distribution  $PD$ . Thus, each pruning step consists of multiple sampling stages followed by a pruning stage.

**Sampling** We define each compressed model’s architecture as the state and  $PD$  as the policy. In each sampling, we add a Gaussian noise vector  $\mathbf{z}$  to  $PD$  to generate the real pruning action  $a$ :

$$\mathbf{z} = [z_1, z_2, \dots, z_L], z_i \sim \mathcal{N}(0, v), \forall i = 1, 2, \dots, L \quad (11)$$

$$a = PD + \mathbf{z} \quad (12)$$

where  $v$  is the variance of Gaussian noise that controls the exploration volume around  $PD$ . A larger  $v$  indicates greater exploration, allowing the pruning policy to explore a wider range of potential architectures. Let  $s$  denote the current model architecture, then, inspired by Bellman equation [Bellman, 1966], we evaluate the Q value of each  $a_i$  and the corresponding compressed model’s architecture  $s'_i$ :

$$Q(s, a_i) = \mathbb{E} \left[ r + \gamma \max_{a'_j} Q(s'_i, a'_j) \mid s, a_i \right], 1 \leq i, j \leq N_S \quad (13)$$

$$r = \mathcal{R}(s'_i) \quad (14)$$

where  $\gamma$  is the discount factor and  $N_S$  is the sampling number at each time step. To achieve the sampling efficiency, we sample  $T$  time steps in each sampling stage to

approximate the expectation in equation 13:

$$Q(s, a_i) \approx \frac{1}{T} \sum_{t=1}^T \left[ r_t + \gamma \max_{a'_{(t,j)}} Q(s'_{(t,i)}, a'_{(t,j)}) \right], 1 \leq i, j \leq N_S \quad (15)$$

$$r_t = \mathcal{R}(s'_{(t,i)}) \quad (16)$$

The reward function is determined by the compressed model’s architecture, modeled by

$$\mathcal{R}(s) = T_e(s) + \alpha * C_F(s) + \beta * C_P(s) \quad (17)$$

where  $T_e$ ,  $C_F$ ,  $C_P$  represent the test accuracy, FLOPs compression ratio, and parameter count compression ratio of the input model, respectively. The hyperparameters  $\alpha, \beta$  can be set to prioritize either accelerating inference or reducing parameter numbers, depending on the specific objective.

**Update Policy** After each sampling stage, we store each action  $a_i$  along with its Q value  $Q(s, a_i)$  in a replay buffer  $RB$ , replacing the entry with the lowest Q value in  $RB$ . If the lowest Q value exceeds  $Q(s, a_i)$ , the corresponding sampling data is discarded. We then select an action  $a^*$  from the replay buffer to update the policy  $PD$ . To balance exploration and exploitation, we use an  $\epsilon$ -greedy strategy: with probability  $\epsilon$ , we select  $a^*$  randomly from  $RB$ . Otherwise, we select the action with the highest Q value:

$$a^* = \arg \max_{a_i} RB(s, a_i) \quad (18)$$

As the model architecture is initially more redundant and becomes more compact through pruning, we use a high exploration  $\epsilon$  value during the initial pruning steps, gradually reducing  $\epsilon$  in the later steps. Several  $\epsilon$  decay strategies are considered, including constant, linear, and cosine decay. We then adjust the policy distribution toward the selected action distribution by a step size  $\lambda$ . To ensure stability in the update process, we adopt the idea of proximal policy optimization (PPO) [Schulman et al., 2017], which constrains the change in the policy distribution:

$$PD_{new} = PD + \lambda * a^* \quad (19)$$

$$PD = clip\left(\frac{PD_{new}}{PD}, 1 - \delta, 1 + \delta\right) * PD \quad (20)$$

where each element’s change ratio is limited within  $[1 - \delta, 1 + \delta]$ . As a result, after several sampling stages, the policy  $PD$  will converge toward an optimal pruning distribution across the layers.

### 3.3 Layer Pruning

Now that we have the dependency graph and the pruning sparsity for each layer, the next question is how to select which output channel within each layer to prune in order to minimize performance drop. For each convolutional layer  $W_l \in R^{N_l^{out} \times N_l^{in} \times K_l \times K_l}$  and linear layer  $R^{N_l^{out} \times N_l^{in}}$  that containing a total of  $N_l^{out}$  output channels, each corresponds to a weight matrix, we use a calibration dataset

extracted as a subset from the training dataset to help determine which weight matrix is less important. According to [LeCun et al., 1989, Ma et al., 2023], the importance of the  $i_{th}$  weight matrix  $W_i$  in layer  $\mathcal{L}$  is given by:

$$I_{W_i} = |\Delta \mathcal{L}(\mathcal{D})| \quad (21)$$

$$= |\mathcal{L}_{W_i}(\mathcal{D}) - \mathcal{L}_{W_i=0}(\mathcal{D})| \quad (22)$$

$$= \left| \frac{\partial \mathcal{L}^\top(\mathcal{D})}{\partial W_i} W_i - \frac{1}{2} W_i^\top H W_i + \mathcal{O}(\|W_i\|^3) \right| \quad (23)$$

where  $H = \frac{\partial^2 \mathcal{L}(\mathcal{D})}{\partial W_i^2}$  is the Hessian matrix. Since computing the Hessian matrix requires  $\mathcal{O}(N^2)$  computational resources, we omit it to accelerate the pruning process. We also disregard  $\mathcal{O}(\|W_i\|^3)$ , as its value is typically small compared to the first term. Thus, the estimated weight structure importance becomes:

$$\hat{I}_{W_i} = \left| \frac{\partial \mathcal{L}^\top(\mathcal{D})}{\partial W_i} W_i \right| \quad (24)$$

For layers that are pruned simultaneously, we calculate the average importance across the grouped layers to assess the importance of each output channel. Once each output channel’s importance is evaluated, we sort them and prune the least important channels according to the specified sparsity level.

### 3.4 Post-train with Knowledge Distillation

Given that the model is pruned multiple times to achieve the desired overall sparsity, and that pruning inevitably reduces performance, it is beneficial to periodically apply post-training after each pruning step to recover lost performance. Knowledge distillation [Hinton, 2015] facilitates the transfer of knowledge from a more advanced teacher model to a smaller or compressed student model. To preserve the performance of the compressed models, we use the original model as the teacher and the corresponding compressed model as the student. Specifically, during post-training, we introduce a penalty term to the loss function, which measures the distance between the probability distributions of the teacher model,  $p_t$ , and the student model,  $p_s$ , based on the input image  $x$ . The modified loss function is:

$$Loss = \tau * Loss(p_t(x), p_s(x)) \quad (25)$$

$$+ (1 - \tau) * Loss(label, p_s(x)) \quad (26)$$

where  $\tau$  is a hyperparameter that controls the extent to which the student model learns from the teacher. As a result, the compressed model benefits from the knowledge of the original model, helping to maintain more of its performance.

## 4 Experiment

### 4.1 Settings

**Setup** We build RL-Pruner from the scratch using PyTorch [Paszke et al., 2019]. Our experiments utilize an

Architecture (Accuracy, FLOPs, Para. Num.)	Reward Strategy	Pruned Accuracy / FLOPs Ratio / Para. Num. Ratio at Different Sparsity Levels			
		20%	40%	60%	80%
VGG-19 (72.89, 418.63M, 39.33M)	Accuracy-Based	73.09 / 22.62 / 35.30	73.11 / 35.11 / 62.18	73.00 / 43.55 / 79.46	71.25 / 55.71 / 89.34
	FLOPs-Based	73.05 / 26.09 / 36.58	72.65 / 39.40 / 63.12	72.50 / 47.33 / 81.04	71.51 / 53.96 / 90.38
	Parameter-Based	73.00 / 24.54 / 36.02	72.08 / 38.94 / 63.70	71.98 / 47.43 / 81.98	71.67 / 56.84 / 91.44
ResNet-56 (72.25, 127.93M, 0.86M)	Accuracy-Based	70.79 / 51.01 / 37.67	68.33 / 73.36 / 66.19	59.65 / 87.02 / 89.32	37.74 / 95.51 / 97.77
	FLOPs-Based	71.21 / 49.92 / 34.53	68.36 / 75.69 / 68.81	61.49 / 84.50 / 88.13	40.22 / 94.06 / 97.28
	Parameter-Based	71.41 / 51.05 / 37.28	68.46 / 75.86 / 70.80	60.68 / 87.19 / 88.32	40.02 / 95.65 / 97.69
DenseNet-121 (79.83, 907.93M, 7.05M)	Accuracy-Based	78.60 / 38.59 / 36.22	77.08 / 67.26 / 64.60	73.33 / 85.72 / 82.60	69.10 / 95.26 / 94.50
	FLOPs-Based	79.08 / 44.15 / 37.91	76.43 / 72.77 / 67.41	71.67 / 87.51 / 83.84	67.71 / 95.44 / 94.67
	Parameter-Based	78.44 / 39.76 / 39.10	76.21 / 68.32 / 67.25	73.85 / 85.98 / 84.14	69.73 / 94.95 / 95.01
GoogLeNet (77.62, 535.66M, 6.40M)	Accuracy-Based	77.61 / 31.78 / 34.09	76.78 / 53.88 / 60.55	74.17 / 72.65 / 81.05	69.96 / 89.85 / 95.23
	FLOPs-Based	77.62 / 35.67 / 34.01	77.00 / 58.45 / 61.06	74.16 / 75.66 / 82.17	69.37 / 86.96 / 95.39
	Parameter-Based	77.64 / 32.21 / 34.49	76.47 / 55.57 / 60.75	73.85 / 71.10 / 81.18	69.47 / 88.73 / 94.88
MobileNetV3-Large (56.18, 7.47M, 4.03M)	Accuracy-Based	56.26 / 33.01 / 36.22	55.32 / 61.63 / 63.95	52.66 / 79.98 / 83.81	42.41 / 92.78 / 95.85
	FLOPs-Based	56.26 / 33.01 / 36.22	55.32 / 61.63 / 63.95	53.28 / 80.17 / 83.74	41.65 / 93.14 / 95.91
	Parameter-Based	56.76 / 33.12 / 36.51	55.74 / 61.46 / 65.23	53.08 / 79.54 / 84.53	44.10 / 92.57 / 96.08

Table 1: Pruned accuracy, FLOPs compression ratio, and parameter reduction ratio on CIFAR-100 for various architectures under different reward strategies and sparsity configurations. The absolute FLOPs count and parameter number of the original architectures are listed in millions, while all other values are presented as percentages.

NVIDIA 4090D GPU with 24GB of memory and Intel(R) Xeon(R) Platinum 8481C CPU.

For implementing our methods, we conduct experiments using default hyperparameter values: noise variance  $z = 0.04$ , step size for policy update  $\lambda = 0.1$ , discount factor  $\gamma = 0.9$ , sample steps  $T = 1$  and sample number  $N_S = 10$  for each sampling stage, with each pruning step containing 10 sampling stages. Proximal policy optimization is employed with  $\delta = 0.2$ . We utilize three reward strategies in our experiments: accuracy-based, FLOPs-based, and parameter-based, corresponding to  $(\alpha = 0, \beta = 0)$ ,  $(\alpha = 0.25, \beta = 0)$  and  $(\alpha = 0, \beta = 0.25)$  in equation 17 respectively. For pruning each layer using the Taylor method, we extract 100 samples from the training dataset as the calibration dataset. A quick post-training is performed with a coefficient  $\tau = 0.75$  for the loss between the teacher model and the student model. Furthermore, we switch the teacher model if a compressed model with higher accuracy is found during pruning. For the exploration strategy, we initialize the exploration parameter  $\epsilon = 0.4$ , which decreases in a cosine fashion over the first 10% of pruning steps. Detailed hyperparameter analysis is provided 4.3.

**Models and Datasets** In this study, we focus on classification tasks, although our proposed methods are readily extendable to other tasks, such as segmentation. To demonstrate the general applicability of RL-Pruner for CNNs, we evaluate it on several widely used classification models, including VGG-19 [Simonyan and Zisserman, 2014], ResNet-56 [He et al., 2016], GoogLeNet [Szegedy et al., 2015], DenseNet121 [Huang et al., 2017], and MobileNetV3-Large [Howard, 2017]. These models encompass diverse architectural features such as residual connections, concatenation connections, and Squeeze-and-Excitation modules. For datasets, we utilize CIFAR-10 and CIFAR-100 [Krizhevsky et al., 2009] to assess performance. Floating point operations (FLOPs) determine

computational complexity and inference speed, while the parameter count reflects memory usage during inference. In our experiments, we primarily use three criteria to evaluate the compressed model: test Top-1 accuracy, FLOPs compression ratio, and parameter count compression ratio. These ratios are computed as  $1 - \frac{\text{FLOPs}(\text{pruned model})}{\text{FLOPs}(\text{base model})}$  and  $1 - \frac{\text{Para. Num.}(\text{pruned model})}{\text{Para. Num.}(\text{base model})}$  respectively.

**Baseline Setup** To validate the effectiveness of our method, we compare it against other structured pruning techniques, including DepGraph [Fang et al., 2023] and GReg [Wang et al., 2020]. Additionally, we benchmark our method against other neural architecture search approaches, such as GNN-RL [Yu et al., 2022]. We evaluate all methods at three channel sparsity ratios—25%, 50%, and 75%—to assess their performance under various levels of sparsity.

## 4.2 Results

**Performance** We begin by evaluating our methods on popular CNN architectures using CIFAR-100. As shown in Table 1, we can compress VGG-19 to 60% channel sparsity, and GoogleNet and MobileNetV3-Large to 40% channel sparsity, with a performance drop of less than 1%. This indicates that there is significant redundancy in these architectures that can be pruned without substantial performance loss, given the considerable reduction in FLOPs and parameter counts at these sparsity levels. We also note that ResNet-56’s performance declines rapidly as pruning progresses. This may be due to ResNet-56’s many layers with relatively few parameters, resulting in a low number of channels per layer. Consequently, the optimal sparsity distribution learned by our algorithm is easily rounded to zero during pruning, which limits its ability to represent each layer’s relative importance and leads to a performance drop.

Architecture (Accuracy, FLOPs, Para. Num.)	Method	Pruned Accuracy / FLOPs Ratio / Para. Num. Ratio at Different Sparsity Levels		
		25%	50%	75%
VGG-19 (72.89, 418.63M, 39.33M)	GReg	72.09 / 29.17 / 44.34	71.55 / 52.72 / 71.31	53.41 / 82.35 / 89.24
	DepGraph	72.49 / 36.64 / 45.69	<b>72.69 / 43.97 / 71.64</b>	72.05 / 75.38 / 91.32
	GNN-RL	66.42 / 17.29 / 33.39	59.63 / 74.96 / 57.81	N/A
	RL-Pruner	<b>†73.13 / 26.50 / 42.83</b>	<b>†72.65 / 39.29 / 72.05</b>	<b>73.14 / 50.42 / 87.08</b>
ResNet-56 (94.20, 127.93M, 0.86M)	GReg	93.36 / 52.56 / 40.17	89.31 / 79.28 / 74.17	84.35 / 95.37 / 93.72
	DepGraph	93.69 / 48.47 / 45.83	92.44 / 74.01 / 77.90	83.85 / 92.86 / 94.36
	GNN-RL	93.67 / 25.07 / 30.50	<b>92.45 / 84.40 / 53.97</b>	<b>87.04 / 95.30 / 74.63</b>
	RL-Pruner	<b>†93.71 / 53.44 / 49.36</b>	91.76 / 80.09 / 84.25	86.63 / 91.39 / 95.97

Table 2: Pruned accuracy, FLOPs compression ratio, and parameter reduction ratio on for VGG-19 on CIFAR-100 and ResNet-56 on CIFAR-10 under different reward strategies and channel sparsity configurations. The absolute FLOPs count and parameter numbers for the original architectures are reported in millions, while all other values are expressed as percentages. The GNN-RL method prunes only the convolutional layers, while all other methods prune both convolutional and linear layers. RL-Pruner methods use an accuracy-based reward strategy. Each method undergoes a quick 100-epoch post-training phase without knowledge distillation. The highest accuracy at each sparsity level is highlighted in bold. † indicates that post-training resulted in degraded performance, in which case we present the results without post-training.

**Structured Pruning Methods Comparison** Following previous work [Fang et al., 2023], we compare our method to other structured pruning and neural architecture search methods using two popular architectures: VGG-19 on CIFAR-100 and ResNet-56 on CIFAR-10. As shown in Table 2, our method achieves higher performance for VGG-19 across all three tested levels of channel sparsity. For the VGG-19 architecture, our method identifies the linear layers as less important and allocates more channel sparsity to them, resulting in a significantly lower FLOPs compression ratio at 75% sparsity. Note that the GNN-RL method does not support pruning VGG-19 at 75% sparsity in our reproduction. For ResNet-56, our method achieves higher performance at 25% sparsity, along with greater FLOPs and parameter compression ratios. However, we observe suboptimal results at 50% and 75% sparsity, likely due to the low number of channels per layer in ResNet-56.

### 4.3 More Analysis

In this section, we explore the effects of various methods employed during the pruning process, such as reward strategies and exploration ( $\epsilon$ ). To minimize errors introduced by complex architectures, we conduct our experiments on the simpler VGG-19 architecture to evaluate these methods.

**Effects of Reward Strategy** We first evaluate how different reward strategies influence pruning results. We track the sparsity changes of specific layers in VGG-19 over 75 pruning steps. As shown in Figure 3, the parameter-based reward strategy tends to prune more filters in the convolutional layers, while the accuracy-based reward strategy favors pruning more neurons in the linear layers. It is challenging to establish a definitive principle linking the reward strategy to which layers are pruned more. Instead, our method identifies the optimal architecture based on the given reward strategy.

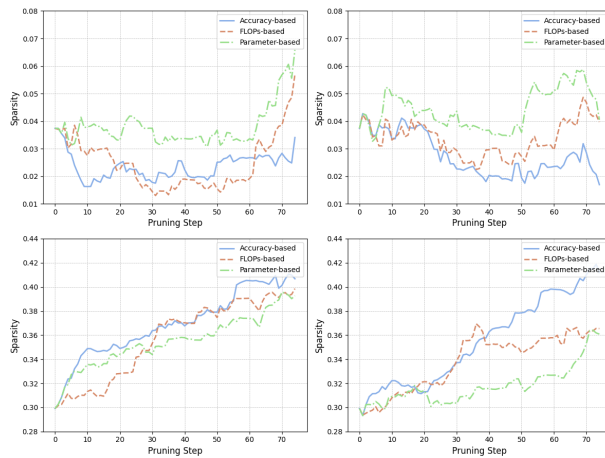


Figure 3: Sparsity changes across different layers during the pruning of VGG-19 on CIFAR-100 over 75 pruning steps with various reward strategies. The first row (from left to right) shows the sparsity changes in the 14<sup>th</sup> and 16<sup>th</sup> convolutional layers, while the second row (from left to right) depicts the sparsity changes in the 1<sup>st</sup> and 2<sup>nd</sup> linear layers of the architecture.

**Exploration  $\epsilon$  Choice** We evaluate different initial values of exploration  $\epsilon$  and reduce it to 0 over the first 10% of pruning steps using various decay strategies for VGG-19 on CIFAR-100. As shown in Table 3, adopting an exploration strategy can lead to higher performance by exploring a broader range of architectures. However, it may also occasionally converge to a suboptimal architecture, resulting in reduced performance.

**Compared to Scratch Training** We then validate the distillation effects of our method. Specifically, we train a compressed version of the VGG-19 model, where the architecture is initialized to match the compressed model, from scratch, and compare its performance to that of the

Exploration $\epsilon$	Decay Strategy	Acc. at Different Sparsity Levels		
		25%	50%	75%
0	constant	72.94	72.33	72.00
0.2	constant	72.71	72.24	<b>72.53</b>
	linear	† 73.10	† 72.32	† 72.03
	cosine	† 73.10	† 72.32	† 72.03
0.4	constant	72.98	72.51	72.19
	linear	73.10	72.32	72.03
	cosine	<b>73.46</b>	<b>72.84</b>	71.60

Table 3: Comparison of different exploration  $\epsilon$  choices. The highest accuracy at each sparsity level is highlighted in bold. †: the results are the same caused by randomness.

Model / Dataset	Sparsity	Compressed	Trained	$\Delta$ Acc.
VGG19	40%	73.11	72.60	+0.51
CIFAR100	60%	73.00	72.10	+0.90
	80%	71.25	69.50	+1.75

Table 4: Comparison to models training from the scratch

compressed model. As shown in Table 4, training the model initialized with the compressed architecture from scratch does not achieve the same performance as the compressed model. Moreover, the performance gap increases as channel sparsity grows. This suggests that pruning combined with post-training not only accelerates inference and reduces parameter count but also guides the model to learn features more efficiently, leading to higher knowledge density in the neural network—something that cannot be achieved through training alone. This phenomenon becomes more pronounced as sparsity increases.

**Computational Cost** Since our method constructs dependency graphs by tracking and comparing tensors, building the dependency graph for complex architectures (e.g., ResNet and DenseNet) takes several minutes, while for simpler architectures (e.g., VGG), it takes less than a minute. Given that our approach involves dozens of pruning steps, with each step comprising multiple sampling stages, the entire pruning process takes several hours under our experimental settings. There is potential to accelerate the sampling process by parallelizing it, as each sampling operation within a sampling stage is independent.

## 5 Conclusion

In this paper, we present RL-Pruner, a structured pruning method that learns the optimal sparsity distribution across layers and supports general pruning without model-specific modifications. We hope our approach, which recognizes that each layer has a different relative importance to the model’s performance, will influence future work in neural network compression, including unstructured pruning and quantization.

## References

- Moez Krichen. Convolutional neural networks: A survey. *Computers*, 12(8), 2023. ISSN 2073-431X. doi:10.3390/computers12080151. URL <https://www.mdpi.com/2073-431X/12/8/151>.
- J Bharadiya. Convolutional neural networks for image classification. *International Journal of Innovative Science and Research Technology*, 8(5):673–677, 2023.
- Keiron O’Shea and Ryan Nash. An introduction to convolutional neural networks, 2015. URL <https://arxiv.org/abs/1511.08458>.
- Jiuxiang Gu, Zhenhua Wang, Jason Kuen, Lianyang Ma, Amir Shahroudy, Bing Shuai, Ting Liu, Xingxing Wang, Gang Wang, Jianfei Cai, and Tsuhan Chen. Recent advances in convolutional neural networks. *Pattern Recognition*, 77:354–377, 2018. ISSN 0031-3203. doi:<https://doi.org/10.1016/j.patcog.2017.10.013>. URL <https://www.sciencedirect.com/science/article/pii/S0031320317304120>.
- Zewen Li, Fan Liu, Wenjie Yang, Shouheng Peng, and Jun Zhou. A survey of convolutional neural networks: Analysis, applications, and prospects. *IEEE Transactions on Neural Networks and Learning Systems*, 33(12):6999–7019, 2022. doi:10.1109/TNNLS.2021.3084827.
- Nafiu Rashid, Berken Utku Demirel, and Mohammad Abdullah Al Faruque. Ahar: Adaptive cnn for energy-efficient human activity recognition in low-power edge devices. *IEEE Internet of Things Journal*, 9(15):13041–13051, 2022. doi:10.1109/JIOT.2022.3140465.
- Rafael Stahl, Alexander Hoffman, Daniel Mueller-Gritschneider, Andreas Gerstlauer, and Ulf Schlichtmann. Deepertings: Fully distributed cnn inference on resource-constrained edge devices. *International Journal of Parallel Programming*, 49:600–624, 2021.
- Yu Cheng, Duo Wang, Pan Zhou, and Tao Zhang. A survey of model compression and acceleration for deep neural networks, 2020. URL <https://arxiv.org/abs/1710.09282>.
- Lei Deng, Guoqi Li, Song Han, Luping Shi, and Yuan Xie. Model compression and hardware acceleration for neural networks: A comprehensive survey. *Proceedings of the IEEE*, 108(4):485–532, 2020. doi:10.1109/JPROC.2020.2976475.
- Gongfan Fang, Xinyin Ma, Mingli Song, Michael Bi Mi, and Xinchao Wang. Depgraph: Towards any structural pruning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16091–16101, June 2023.
- Yang He and Lingao Xiao. Structured pruning for deep convolutional neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(5):2900–2919, 2024. doi:10.1109/TPAMI.2023.3334614.
- Sajid Anwar, Kyuyeon Hwang, and Wonyong Sung. Structured pruning of deep convolutional neural networks. *J.*



- Emerg. Technol. Comput. Syst.*, 13(3), February 2017. ISSN 1550-4832. doi:10.1145/3005348. URL <https://doi.org/10.1145/3005348>.
- Tianyun Zhang, Shaokai Ye, Kaiqi Zhang, Jian Tang, Wujie Wen, Makan Fardad, and Yanzhi Wang. A systematic dnn weight pruning framework using alternating direction method of multipliers. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018.
- Xiaolong Ma, Sheng Lin, Shaokai Ye, Zhezhi He, Linfeng Zhang, Geng Yuan, Sia Huat Tan, Zhengang Li, Deliang Fan, Xuehai Qian, Xue Lin, Kaisheng Ma, and Yanzhi Wang. Non-structured dnn weight pruning—is it beneficial in any platform? *IEEE Transactions on Neural Networks and Learning Systems*, 33(9):4930–4944, 2022. doi:10.1109/TNNLS.2021.3063265.
- Sridhar Swaminathan, Deepak Garg, Rajkumar Kannan, and Frederic Andres. Sparse low rank factorization for deep neural network compression. *Neurocomputing*, 398:185–196, 2020. ISSN 0925-2312. doi:<https://doi.org/10.1016/j.neucom.2020.02.035>. URL <https://www.sciencedirect.com/science/article/pii/S0925231220302253>.
- Yerlan Idelbayev and Miguel A. Carreira-Perpinan. Low-rank compression of neural nets: Learning the rank of each layer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- 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 (CVPR)*, June 2016.
- Yunchao Gong, Liu Liu, Ming Yang, and Lubomir Bourdev. Compressing deep convolutional networks using vector quantization, 2014. URL <https://arxiv.org/abs/1412.6115>.
- Aojun Zhou, Anbang Yao, Yiwen Guo, Lin Xu, and Yurong Chen. Incremental network quantization: Towards lossless cnns with low-precision weights, 2017. URL <https://arxiv.org/abs/1702.03044>.
- Jianping Gou, Baosheng Yu, Stephen J Maybank, and Dacheng Tao. Knowledge distillation: A survey. *International Journal of Computer Vision*, 129(6):1789–1819, 2021.
- Jang Hyun Cho and Bharath Hariharan. On the efficacy of knowledge distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- 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 (ICCV)*, Oct 2017a.
- Xiaohan Ding, Tianxiang Hao, Jianchao Tan, Ji Liu, Jun-gong Han, Yuchen Guo, and Guiguang Ding. Resrep: Lossless cnn pruning via decoupling remembering and forgetting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4510–4520, October 2021.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
- Andrew G Howard. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets, 2017. URL <https://arxiv.org/abs/1608.08710>.
- Yang He, Guoliang Kang, Xuanyi Dong, Yanwei Fu, and Yi Yang. Soft filter pruning for accelerating deep convolutional neural networks. *arXiv preprint arXiv:1808.06866*, 2018a.
- Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. In *Proceedings of the IEEE international conference on computer vision*, pages 1389–1397, 2017.
- 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*, pages 1529–1538, 2020a.
- Jian-Hao Luo, Jianxin Wu, and Weiyao Lin. Thinet: A filter level pruning method for deep neural network compression. In *Proceedings of the IEEE international conference on computer vision*, pages 5058–5066, 2017.
- Torsten Hoefler, Dan Alistarh, Tal Ben-Nun, Nikoli Dryden, and Alexandra Peste. Sparsity in deep learning: Pruning and growth for efficient inference and training in neural networks. *Journal of Machine Learning Research*, 22(241):1–124, 2021.
- 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, 2017b.
- 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 *Proceedings of the European conference on computer vision (ECCV)*, pages 784–800, 2018b.
- Sixing Yu, Arya Mazaheri, and Ali Jannesari. Auto graph encoder-decoder for neural network pruning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6362–6372, 2021.
- Sixing Yu, Arya Mazaheri, and Ali Jannesari. Topology-aware network pruning using multi-stage graph embedding and reinforcement learning. In *International conference on machine learning*, pages 25656–25667. PMLR, 2022.
- Shaopeng Guo, Yujie Wang, Quanquan Li, and Junjie Yan. Dmcp: Differentiable markov channel pruning for neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1539–1547, 2020.
- Xuefei Ning, Tianchen Zhao, Wenshuo Li, Peng Lei, Yu Wang, and Huazhong Yang. Dsa: More efficient budgeted pruning via differentiable sparsity allocation. In *European Conference on Computer Vision*, pages 592–607. Springer, 2020.
- Zechun Liu, Haoyuan Mu, Xiangyu Zhang, Zichao Guo, Xin Yang, Kwang-Ting Cheng, and Jian Sun. Metapruning: Meta learning for automatic neural network channel pruning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3296–3305, 2019.
- Mingbao Lin, Rongrong Ji, Yuxin Zhang, Baochang Zhang, Yongjian Wu, and Yonghong Tian. Channel pruning via automatic structure search. *arXiv preprint arXiv:2001.08565*, 2020b.
- Geoffrey Hinton. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- Reuven Y Rubinstein and Dirk P Kroese. *Simulation and the Monte Carlo method*. John Wiley & Sons, 2016.
- Christopher JCH Watkins and Peter Dayan. Q-learning. *Machine learning*, 8:279–292, 1992.
- Richard Bellman. Dynamic programming. *science*, 153 (3731):34–37, 1966.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms, 2017. URL <https://arxiv.org/abs/1707.06347>.
- Yann LeCun, John Denker, and Sara Solla. Optimal brain damage. *Advances in neural information processing systems*, 2, 1989.
- Xinyin Ma, Gongfan Fang, and Xinchao Wang. Llm-pruner: On the structural pruning of large language models, 2023. URL <https://arxiv.org/abs/2305.11627>.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- Huan Wang, Can Qin, Yulun Zhang, and Yun Fu. Neural pruning via growing regularization. *arXiv preprint arXiv:2012.09243*, 2020.