

Supplementary Material

1. Structure of the policy network

We use a VGG-16 style architecture as our policy network. Different from the vanilla VGG-16, which is designed for image classification, we use 1D convolution instead. The detailed architecture of the policy network is presented in Table 1.

Table 1: The network architecture of the policy network. N is the number of filters in the CNN to be pruned.

Index	Layer	Type	Feature map	Kernel size	Stride	Output size	Activation
0	Input	Feature	1	-	-	$N \times 7$	-
1	conv1_1	1D conv	64	3×3	1	$N \times 64$	ReLu
2	conv1_2	1D conv	64	3×3	1	$N \times 64$	ReLu
3	pool1	Pool	-	2×2	2	$(N/2) \times 64$	Max
4	conv2_1	1D conv	128	3×3	1	$(N/2) \times 128$	ReLu
5	conv2_2	1D conv	128	3×3	1	$(N/2) \times 128$	ReLu
6	pool2	Pool	-	2×2	2	$(N/4) \times 128$	Max
7	conv3_1	1D conv	256	3×3	1	$(N/4) \times 256$	ReLu
8	conv3_2	1D conv	256	3×3	1	$(N/4) \times 256$	ReLu
8	conv3_3	1D conv	256	3×3	1	$(N/4) \times 256$	ReLu
10	pool3	Pool	-	2×2	2	$(N/8) \times 256$	Max
11	conv4_1	1D conv	512	3×3	1	$(N/8) \times 512$	ReLu
12	conv4_2	1D conv	512	3×3	1	$(N/8) \times 512$	ReLu
13	conv4_3	1D conv	512	3×3	1	$(N/8) \times 512$	ReLu
14	pool4	Pool	-	2×2	2	$(N/16) \times 512$	Max
15	conv5_1	1D conv	512	3×3	1	$(N/16) \times 512$	ReLu
16	conv5_2	1D conv	512	3×3	1	$(N/16) \times 512$	ReLu
17	conv5_3	1D conv	512	3×3	1	$(N/16) \times 512$	ReLu
18	pool5	Pool	-	2×2	2	$(N/32) \times 512$	Max
19	fc6_1	fc (π)	-	-	-	N	Softmax
19	fc6_2	fc (v)	-	-	-	1	Tanh

2. Pseudocodes of key components in our approach.

Algorithm 1 Get the improved policy π after MCTS search: `getPolicyPi(s_i)`

Input: Current configuration of the network to be pruned s_i , number of MCTS simulations per action n_{mcts} , total number of filters in the network to be pruned n_f , temperature τ .

Output: π_i

```
1: for  $i$  in range( $n_{mcts}$ ) do
2:   MCTS( $s_i$ )
3: Get  $N(s_i, a)$  after MCTS simulations.
4: if  $\tau = 0$  then
5:   bestAction =  $\operatorname{argmax}_a N(s_i, a)$ 
6:    $\pi[\text{bestAction}] = 1$ 
7: else
8:    $\pi = \frac{N(s_i, a)^{(1/\tau)}}{\sum_b N(s_i, b)^{(1/\tau)}}$ 
return  $\pi$ 
```

Algorithm 2 Get training samples from a single iteration: `getTrainSamples(s_0)`

Input: The raw network to be pruned s_0 , pruning ratio γ , trainingAccBaseline b ,

Output: trainSamples (s_i, π_i, v)

```
1:  $t = 0, s_t = s_0$ 
2: trainSamples = []
3: while  $\text{FLOPs}(s_t)/\text{FLOPs}(s_0) > \gamma$  do
4:    $\pi_t = \text{getPolicyPi}(s_t)$ 
5:   trainSamples.append( $[s_t, \pi_t]$ )
6:   nextAction = randomChoice( $\pi_t$ )
7:    $s_{t+1} = \text{pruneFilter}(s_t, \text{nextAction})$ 
8:    $t = t + 1$ 
9: if  $\text{trainAcc}(s_t) > b$  then
10:   $v = 1$ 
11: else
12:   $v = -1$ 
13: trainSamples =  $[(x[0], x[1], v)$  for  $x$  in trainSamples]
14: return trainSamples
```

Algorithm 3 Learn to get the slimmed CNN with RL and MCTS

Input: The raw network to be pruned s_0 , neural network for pruning action selection f_θ , number of self-play simulations n_{sim} , maximum training queue length L .

Output: The optimal slimmed CNN s_p

```
1: totalTrainingQueue = []
2: while stopCounter <  $n_s$  do
3:   for  $i$  in range( $n_{sim}$ ) do
4:     Initialize MCTS
5:     trainingSamples = getTrainSamples( $s_0$ )
6:     if len(totalTrainingQueue) >  $L$  then
7:       totalTrainingQueue.pop()
8:     totalTrainingQueue += trainingSamples
9:      $f_\theta = \text{RLTrain}(\text{totalTrainingQueue}, f_\theta)$ 
10:    Get the slimmed network  $s'$  by pruning  $s_0$  with  $f_\theta$ 
11:    if trainAcc( $s'$ ) >  $b$  then
12:       $b = \text{trainAcc}(s')$ 
13:       $s_p = s'$ 
14:      stopCounter = 0
15:    else
16:      stopCounter += 1
17: return  $s_p$ 
```
