

---

# Greedy Optimization Provably Wins the Lottery: Logarithmic Number of Winning Tickets is Enough

---

**Mao Ye** <sup>\*</sup>  
UT Austin  
my21@cs.utexas.edu

**Lemeng Wu** <sup>\*</sup>  
UT Austin  
lmwu@cs.utexas.edu

**Qiang Liu**  
UT Austin  
lqiang@cs.utexas.edu

## Abstract

Despite the great success of deep learning, recent works show that large deep neural networks are often highly redundant and can be significantly reduced in size. However, the theoretical question of how much we can prune a neural network given a specified tolerance of accuracy drop is still open. This paper provides one answer to this question by proposing a greedy optimization based pruning method. The proposed method has the guarantee that the discrepancy between the pruned network and the original network decays with exponentially fast rate w.r.t. the size of the pruned network, under weak assumptions that apply for most practical settings. Empirically, our method improves prior arts on pruning various network architectures including ResNet, MobilenetV2/V3 on ImageNet.

## 1 Introduction

Large-scale deep neural networks have achieved remarkable success on complex cognitive tasks, including image classification, e.g., He et al. (2016), speech recognition, e.g., Amodei et al. (2016) and machine translation, e.g., Wu et al. (2016). However, a drawback of the modern large-scale DNNs is their low inference speed and high energy cost, which makes it less appealing to deploy those models on edge devices such as mobile phones and Internet of Things (Cai et al., 2019).

It has been shown that network pruning (Han et al., 2015) is an effective technique to reduce the size of the DNNs without a significant drop of accuracy. However, most existing works on network pruning are based on heuristics, leaving the theoretical questions largely open on what kind of network can be effectively pruned, how much we can prune a DNN given a specified tolerance of accuracy drop and how to achieve it with a practical and computationally efficient procedure.

Recently, a line of works on network pruning with theoretical guarantees have emerged, including sensitivity-based methods (Baykal et al., 2019b; Liebenwein et al., 2020), coresnet-based methods (Baykal et al., 2019a; Mussay et al., 2020), greedy forward selection (Ye et al., 2020). Both the sensitivity-based and coresnet-based methods prune the network by sampling and bound the error caused pruning via concentration inequalities. They show that the error introduced by pruning decays with an  $\mathcal{O}(n^{-1})$  rate w.r.t. the size  $n$  of pruned network. This is comparable to the asymptotic error obtained by directly training a neural network of size  $n$  with gradient descent descent, which is also  $\mathcal{O}(n^{-1})$  following the mean field analysis of Mei et al. (2018); Araújo et al. (2019); Sirignano & Spiliopoulos (2019). More recently, Ye et al. (2020) proposed the first pruning method that achieves a faster  $\mathcal{O}(n^{-2})$  error rate and is hence provably better than direct training with gradient descent. See Table 1 for a summary on those works.

However, the analysis of Ye et al. (2020) only applies to two-layer networks and requires the original network to be sufficiently over-parameterized. In this paper, we proposed a new greedy optimization based pruning method, which learns sub-networks of size  $n$  with a significantly smaller  $\mathcal{O}(\exp(-cn))$

---

<sup>\*</sup>Equal Contribution

	Rate	No Over-param	Deep Net
Baykal et al. (2019b); Liebenwein et al. (2020)	$\mathcal{O}(n^{-1})$	✓	✓
Baykal et al. (2019a); Mussay et al. (2020)	$\mathcal{O}(n^{-1})$	✓	✓
Ye et al. (2020)	$\mathcal{O}(n^{-2})$	✗	✗
This paper	$\mathcal{O}(\exp(-cn))$	✓	✓

Table 1: Overview on theoretical guaranteed pruning methods. Rate above gives how the error due to pruning decays as the size of the pruned network ( $n$ ) increases. Column ‘No Over-param’ denotes whether the method applies to an original network that is not over-parameterized in order to obtain the rate. Column ‘Deep net’ denotes whether the analysis applies to deep networks.

error rate, improving the rate from polynomial to exponential. In addition, our theoretical rate only requires weak assumptions that hold for most networks in practice, without requiring the the original networks to be overparameterized as Ye et al. (2020). Different from the Lottery Ticket Hypothesis (Frankle & Carbin, 2018), which selects the winning tickets that give good performance when trained in isolation from initialization, our approach finds the tickets (that already won) from a fully converged network.

Practically, our algorithm is simple and easy to implement. In addition, we introduce practical speedup techniques to further improve the time efficiency. Empirically, our method improves the prior arts on network pruning under various network structures including ResNet-34 (He et al., 2016), MobileNetV2 (Sandler et al., 2018) and MobileNetV3 (Howard et al., 2019) on ImageNet (Deng et al., 2009) as well as DGCNN (Wang et al., 2019) on ModelNet40 (Wu et al., 2015) on point cloud classification.

**Notation** We use notation  $[N] := 1, \dots, N$  for the set of the first  $N$  positive integers. All the vector norms  $\|\cdot\|$  are assumed to be  $\ell_2$  norm. We denote the vector  $\ell_0$  norm by  $\|\cdot\|_0$ .  $\|\cdot\|_{\text{Lip}}$  denotes the Lipschitz norm for functions.  $\mathbb{I}\{\cdot\}$  indicates the indicator function.

## 2 Background and Method

**Problem Setup** Given a pre-trained deep neural network with  $L$  layers:  $F(\mathbf{x}) = F_L \circ F_{L-1} \circ \dots \circ F_2 \circ F_1(\mathbf{x})$ , where the  $\ell$ -th layer  $F_\ell$  consisting of  $N$  neurons forms a mapping of form

$$F_\ell(\mathbf{z}) = \frac{1}{N} \sum_{i=1}^N \sigma(\boldsymbol{\theta}_i^\ell, \mathbf{z}),$$

with  $\mathbf{z}$  as a proper input of the  $\ell$ -th layer, which is the output of the previous  $\ell - 1$  layers. Here  $\sigma(\boldsymbol{\theta}, \cdot)$  is a general nonlinear map parameterized by  $\boldsymbol{\theta}$  that represents a neuron or other module in the network. For example, in a fully connected layer, we have  $\sigma(\boldsymbol{\theta}, \mathbf{z}) = w_1 \sigma_+(\mathbf{w}_2^\top \mathbf{z})$  with  $\boldsymbol{\theta} = [w_1, w_2]$  and  $\sigma_+$  an activation function such as ReLU or Tanh. In a convolution layer, we have  $\sigma(\boldsymbol{\theta}, \mathbf{z}) = w_1 \sigma_+(\mathbf{w}_2 * \mathbf{z})$ , where  $*$  denotes the convolution operator. In this paper we may call  $\sigma(\boldsymbol{\theta}_i^\ell, \cdot)$  the neuron  $i$  or the  $i$ -th neuron for simplicity. Without loss of generality, we assume each layer in the given deep network has the same number of neurons using the same activation function.

The goal of network pruning is to construct a thinner network by replacing each layer with a subset of  $n < N$  neurons. For simplicity of presentation, we focus on pruning a single layer  $F_\ell$  for now and we discuss how to apply our algorithm in a layer-wise fashion to prune the whole network in section 2.4.

To prune the  $\ell$ -th layer, the goal is to replace  $F_\ell$  with a thinner layer  $f_{\ell, A}$  with  $n < N$  neurons:

$$f_{\ell, A}(\mathbf{z}) = \sum_{i=1}^N a_i \sigma(\boldsymbol{\theta}_i^\ell, \mathbf{z}), \quad \mathbf{A} = [a_1, \dots, a_N] \in \Omega_N, \|\mathbf{A}\|_0 \leq n,$$

where  $\Omega_{[N]}$  is the probability simplex on the  $N$  neurons, that is,

$$\Omega_N = \left\{ \mathbf{v} : \mathbf{v} = [v_1, \dots, v_N] \in \mathbb{R}^N, \quad v_i \geq 0, \quad \forall i \in [N] \quad \text{and} \quad \sum_{i=1}^N v_i = 1 \right\}.$$

By enforcing that  $\mathbf{A} \in \Omega_N$ , we prune the layer by finding the best convex combination of a subset of neurons. The constraint that  $\sum_i a_i = 1$  ensures that the overall magnitude of the output of the layer after pruning matches that of the original network even when a lot neurons are moved. We denote the network with the  $\ell$ -th layer replaced by  $f_{\ell,A}$  as  $f_A$ , i.e.,

$$f_A = F_L \circ \dots \circ F_{\ell+1} \circ f_{\ell,A} \circ F_{\ell-1} \circ \dots \circ F_1.$$

Given an observed dataset  $\mathcal{D}_m := (\mathbf{x}^{(i)}, y^{(i)})_{i=1}^m$  with  $m$  data points. We want to choose  $\mathbf{A}$  such that the pruned network  $f_A$  is close to the original  $F$  as much as possible, measured by the regression discrepancy loss,

$$\mathbb{D}[f_A, F] = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_m} \left[ (f_A(\mathbf{x}) - F(\mathbf{x}))^2 \right].$$

Our algorithm and theoretical analysis can be extended to other discrepancy losses such as the cross-entropy. The problem of pruning the  $\ell$ -th layer can be formulated by the following constraint problem

$$\min_{\mathbf{A}} \mathbb{D}[f_A, F], \quad \text{s.t.} \quad \mathbf{A} \in \Omega_N, \quad \|\mathbf{A}\|_0 \leq n. \quad (1)$$

This yields a challenging sparse optimization problem, which we address using greedy optimization, yielding algorithms that are both theoretically guaranteed and practically efficient.

## 2.1 Pruning with Greedy Local Imitation

We first introduce a simply greedy algorithm via local imitation for searching a good solution of problem (1). The pruned network can be viewed as

$$H \circ f_{\ell,A}(\mathbf{z}) = H \circ \left( \sum_{i=1}^N a_i \sigma(\theta_i^\ell, \mathbf{z}) \right).$$

Here  $\mathbf{z}$  is the output of the  $\ell - 1$ -th layers and  $H = F_L \circ \dots \circ F_{\ell+1}$  is the mapping of the later layers. Denote  $\mathbf{z}^{(i)} = F_{\ell-1} \circ \dots \circ F_1(\mathbf{x}^{(i)})$ . The set  $\mathcal{D}_m^\ell := (\mathbf{z}^{(i)})_{i=1}^m$  denotes the distribution of training data pushed through the first  $\ell - 1$  layers. Suppose  $H$  is Lipschitz continuous, which typically holds for neural networks, we are about to upper bound  $\mathbb{D}[f_A, F]$  by  $\mathbb{D}[f_A, F] \leq \|H\|_{\text{Lip}}^2 \bar{\mathbb{D}}[f_{\ell,A}, F_\ell]$ , where  $\bar{\mathbb{D}}$  is the local discrepancy loss measuring the discrepancy on the output of the  $\ell$ -th layer between the pruned and original network

$$\bar{\mathbb{D}}[f_{\ell,A}, F_\ell] = \mathbb{E}_{\mathbf{z} \sim \mathcal{D}_m^\ell} \|f_{\ell,A}(\mathbf{z}) - F_\ell(\mathbf{z})\|^2.$$

In local imitation, we construct  $f_{\ell,A}$  such that its output well imitates the output of  $F_\ell$ , i.e.,

$$\min_{\mathbf{A}} \bar{\mathbb{D}}[f_{\ell,A}, F_\ell], \quad \text{s.t.} \quad \mathbf{A} \in \Omega_N, \quad \|\mathbf{A}\|_0 \leq n. \quad (2)$$

Importantly, different from the original loss  $\mathbb{D}[\cdot]$  in (1), the layer-wise local discrepancy loss  $\bar{\mathbb{D}}[\cdot]$  is convex w.r.t.  $\mathbf{A}$  and enjoys good geometric property for enabling fast exponential error rate via greedy optimization, as we show in sequel. On the other hand, as the final discrepancy  $\mathbb{D}$  is controlled by the local discrepancy  $\bar{\mathbb{D}}$ , i.e., minimizing  $\bar{\mathbb{D}}$  effectively minimizes  $\mathbb{D}$ .

The local imitation is a bi-directional greedy optimization for solving (2). It starts with an empty layer, and sequentially adds, removes or adjusts neurons that yield the largest decrease of the loss. Specifically, denote by  $f_{\ell,A(k)}$  the layer we obtained at the  $k$ -th iteration with  $\mathbf{A}(k) = [a_1(k), \dots, a_N(k)]$ . We start with selecting the single best neuron that minimizes the loss:

$$f_{\ell,A(0)} = \sigma(\theta_{i_0^*}^\ell, \cdot), \quad \text{with} \quad i_0^* = \arg \min_{i \in [N]} \bar{\mathbb{D}}[\sigma(\theta_i^\ell, \cdot), F_\ell(\cdot)] \quad (3)$$

where  $i_0^*$  is the index of the selected neuron; correspondingly, we have  $a_i(0) = \mathbb{I}\{i = i_0^*\}$ .

At iteration  $k$ , we search for the best neuron  $i_k^*$  and step size  $\gamma_k^*$  that minimizes the loss most, i.e.,

$$[i_k^*, \gamma_k^*] = \arg \min_{i \in [N], \gamma \in U_i} \bar{\mathbb{D}}[(1 - \gamma)f_{\ell,A(k)} + \gamma\sigma(\theta_i^\ell, \cdot), F_\ell]. \quad (4)$$

Then we update  $f_{\ell, \mathbf{A}(k)}$  to  $f_{\ell, \mathbf{A}(k+1)} = (1 - \gamma_k^*) f_{\ell, \mathbf{A}(k)} + \gamma_k^* \sigma(\boldsymbol{\theta}_i^\ell, \cdot)$ . Here  $U_i$  in (4) is the search interval of the step size  $\gamma$ . We set  $U_i = [0, 1]$  if the  $i$ -th neuron has not been selected yet (i.e.,  $a_i(k) = 0$ ) and  $U_i = [-a_i(k)/(1 - a_i(k)), 1]$  if the neuron has already been added (i.e.,  $a_i(k) > 0$ ). Therefore, this update can correspond to adding or removing a neuron, or simply adjusting the weight of existing neurons: the  $i_k^*$ -th neuron is added into the pruned network  $f_{\ell, \mathbf{A}}$  if we have  $a_{i_k^*}(k) = 0$ , and it is removed from  $f_{\ell, \mathbf{A}}$  if we have  $\gamma_k^* = -a_{i_k^*}(k)/(1 - a_{i_k^*}(k))$ ; no new neuron is added or removed if otherwise.

We stop the iteration when a convergence criterion, i.e.,  $\bar{\mathbb{D}} \leq \epsilon$ , is met.

**Solving Greedy Optimization in (4)** A naive way to solve problem (4) is by enumerating each neuron and solving the corresponding inner minimization on  $\gamma$ . This is computational costly as it requires computing the forward pass in neural network many times. However, given  $i$ , the local discrepancy loss is a quadratic function w.r.t.  $\gamma$ . Combined with some special property of the local imitation algorithms, we are able to solve (4) with only computing the forward pass in network once. We refer readers to Appendix 5.1 for details.

### 2.1.1 Greedy Local Imitation Decays Error Exponentially Fast

Now we proceed to give the convergence rate for the proposed local imitation algorithm. We introduce the following assumption.

**Assumption 1** Assume that for any  $i \in [N]$ ,  $\mathbf{z}^{(j)} \in \mathcal{D}_m^\ell$ , we have  $\|\sigma(\boldsymbol{\theta}_i^\ell, \mathbf{z}^{(j)})\| \leq c_1$  and  $\|H\|_{Lip} \leq c_1$  for some  $c_1 < \infty$ .

Assumption 1 holds when network parameters and data are bounded and the activation is Lipschitz continuous, which is very mild and holds for most network in practice. The following theorem characterizes the convergence of local imitation showing that the error caused by pruning decays exponentially fast when the number of neurons in the pruned model increases.

**Theorem 1 (Convergence Rate)** Under assumption 1, at each step  $k$  of the greedy optimization in (2), we obtain a layer with no more than  $k$  neurons (i.e.,  $\|\mathbf{A}(k)\|_0 \leq k$ ), whose loss satisfies  $\mathbb{D}[f_{\mathbf{A}(k)}, F] \leq \|H\|_{Lip}^2 \bar{\mathbb{D}}[f_{\ell, \mathbf{A}(k)}, F_\ell] = \mathcal{O}(\exp(-\lambda_\ell k))$ , where  $\lambda_\ell > 0$  is a strictly positive constant. That is, the loss decays exponentially with the number of neurons in  $f_{\ell, \mathbf{A}(k)}$ .

**Remark** Minimizing  $\gamma$  over  $U_i$  can be viewed as line searching the optimal step size for adjusting neuron  $i$ . We may also consider to choose a proper fixed step size, e.g.,  $\gamma = 1/(k+1)$  instead of line searching, in which case the optimization in (2) is simplified into

$$\min_{i \in [N]} \bar{\mathbb{D}} \left[ (k f_{\ell, \mathbf{A}(k)} + \sigma(\boldsymbol{\theta}_i^\ell, \cdot)) / (k+1), F_\ell \right], \quad (5)$$

which can be shown to give an  $\mathcal{O}(k^{-2})$  error at the  $k$ -th step under the same assumption as Theorem 1. See Appendix 5.2 for more details.

## 2.2 Pruning with Greedy Global Imitation

The local imitation method uses a surrogate local discrepancy loss which is convex w.r.t.  $\mathbf{A}$  to prune the networks. Despite the good property of local imitation, the use of surrogate loss can be ineffective for some layers. For example, during the iteration of local imitation, the best neuron that minimizes the surrogate loss is not necessarily the best one that minimizes the actual discrepancy loss.

We propose a second pruning method, which directly minimizing the original discrepancy loss. This method follows the similar greedy fashion as the local imitation. We initialize the network by  $f_{\mathbf{A}(0)} = H_2 \circ f_{\ell, \mathbf{A}(0)} \circ H_1 = H_2 \circ (\sum_{i=1}^N a_i(0) \sigma(\boldsymbol{\theta}_i^\ell, \cdot)) \circ H_1$ , where

$$a_i(0) = \mathbb{I}\{i = i_0^*\}, \quad i_0^* = \arg \min_{i \in [N]} \mathbb{D}[H_2 \circ \sigma(\boldsymbol{\theta}_i^\ell, \cdot) \circ H_1, F], \quad (6)$$

and  $H_1 = F_{\ell-1} \circ \dots \circ F_1$  and  $H_2 := F_L \circ \dots \circ F_{\ell+1}$ . Similarly, at each iteration, We adjust the network in a greedy way by solving the following problem:

$$\min_{i \in [N]} \min_{\gamma \in U_i} \mathbb{D} \left[ H_2 \circ ((1 - \gamma) f_{\ell, \mathbf{A}(k)} + \gamma \sigma(\boldsymbol{\theta}_i^\ell, \cdot)) \circ H_1, F \right]. \quad (7)$$

However, solving problem (7) is computationally costly as the loss is non-convex w.r.t.  $\mathbf{A}$  and thus solving the inner minimization on  $\gamma$  requires exhaustive search. To reduce the computational cost, in iteration  $k$ , we instead consider the following problem

$$\min_{i \in [N]} \mathbb{D} [H_2 \circ ((1 - \gamma_k) f_{\ell, \mathbf{A}(k)} + \gamma_k \sigma(\boldsymbol{\theta}_i^\ell, \cdot)) \circ H_1, F], \quad \gamma_k = (1 + k)^{-1}. \quad (8)$$

Suppose that  $i_k^*$  gives that solution of problem (8), we update the network by setting

$$a_i(k+1) = (1 - \gamma_k) a_i(k) + \gamma_k \mathbb{I}\{i = i_k^*\}.$$

And we end the iteration when convergence criterion is met. Notice that different from local imitation, the algorithm adjusts  $A$  based on the final output of the network instead of the ‘local’ output of the pruned layer and thus we name it global imitation.

Different from the local imitation, due to the nonlinear of  $H_2$ , besides Assumption 1, obtaining a convergence rate for the global imitation requires several additional assumptions characterizing the linearity of  $H_2$  as well as the geometric property of the pruned layer.

**Theorem 2** *Under Assumption 1 and some additional assumptions, specified in Appendix 5.3, on the linearity of  $H_2$  and initialization, we have  $\mathbb{D}[f_{\mathbf{A}(k)}, F] = \mathcal{O}(k^{-2})$  and  $\|\mathbf{A}(k)\|_0 \leq k$ .*

### 2.2.1 Accelerating Global Imitation via Taylor Approximation

A native way to solve problem (8) is by enumerating all the neurons and calculating  $\mathbb{D} [H \circ ((1 - \gamma_k) f_{\ell, \mathbf{A}(k)} + \gamma_k \sigma(\boldsymbol{\theta}_i^\ell, \cdot)), F]$ , which has at least  $\mathcal{O}(Nn)$  time complexity for pruning a layer with  $N$  neurons to  $n$  neurons. Here we propose a technique to reduce the computational cost via Taylor approximation. At iteration  $k$ , for any neuron  $i \in [N]$ , we have

$$\mathbb{D} [H \circ [(1 - \gamma_k) f_{\ell, \mathbf{A}(k)} + \gamma_k \sigma(\boldsymbol{\theta}_i^\ell, \cdot)], F] = \frac{1}{k+1} gr_{\mathbf{A}(k),i} + \mathcal{O}((k+1)^{-2}),$$

where we define

$$gr_{\mathbf{A}(k),i} = \frac{\partial}{\partial \gamma} \mathbb{D} [H \circ [(1 - \gamma) f_{\ell, \mathbf{A}(k)} + \gamma \sigma(\boldsymbol{\theta}_i, \cdot)], F] \Big|_{\gamma=0}.$$

Thus, when  $k$  is large enough (which we find 25 is sufficient in practice), this approximation allows us to find the (near) optimal solution with small error of problem (8) by finding the neuron with the largest  $gr_{\mathbf{A},i}$ . Simple algebra shows that

$$gr_{\mathbf{A},i} = 2 \sum_{j=1}^n (\mathbb{I}\{j = i\} - a_j) r_{\mathbf{A},i}, \quad \text{where}$$

$$r_{\mathbf{A},i} := \mathbb{E}_{\mathbf{z} \sim \mathcal{D}_m^\ell} [(H \circ f_{\ell, \mathbf{A}}(\mathbf{z}) - H \circ F_\ell(\mathbf{z})) H'(f_{\ell, \mathbf{A}}(\mathbf{z})) \sigma(\boldsymbol{\theta}_i^\ell, \mathbf{z})].$$

Therefore, we can easily calculate  $gr_{\mathbf{A}(k),i}$  for all  $i \in [N]$  once we obtain  $r_{\mathbf{A}(k),i}$  for all  $i \in [N]$ . In appendix, we show that  $r_{\mathbf{A}(k),i}$  can be easily computed with automatic differentiation function in common deep learning libraries by introducing some ancillary parameters into the model. See Appendix 5.4 for details. If we choose to use this approximation when  $k > \tilde{k}$  for some  $\tilde{k} > 0$ , we reduce the complexity from  $\mathcal{O}(Nn)$  to  $\mathcal{O}(n)$ .

### 2.3 Pruning vs GD: Numerical Verification of the Rate

Our result implies that the subnetwork  $f_{\mathbf{A}}$  obtained by pruning gives  $\mathbb{D}[f_{\mathbf{A}}, F] = \mathcal{O}(\exp(-\lambda n))$  where  $n$  is the number of neurons remained in the pruned layer. In comparison, the mean field analysis (Araújo et al., 2019; Mei et al., 2018) suggests that directly train a network with same size as the pruned model gives  $\mathcal{O}(n^{-1})$  discrepancy loss. This suggests that pruning is provably better than training. We conduct a numerical experiment to verify the theoretical result. Given some simulated dataset, we firstly train a two hidden layer neural network with 100 neurons for each layer. And we prune the layer close to the input to different number of neurons using the local and global imitation. We also train the network with different number of neurons for the pruned layer and 100 neurons for the other one. Figure 1 plots the discrepancy loss and the number of neurons of the pruned layer. The empirical result matches our theoretical findings. We refer readers to Appendix 5.5 for more details.

---

**Algorithm 1** The Greedy Local/Global Imitation

---

- 1: **Input:** A pretrained network  $F$  with  $L$  layers. The targeted layer index  $\ell$  for pruning. Method =  $\in \{\text{local, global}\}$
- 2: Initialize  $f_{\ell, A}$  using (3) for local imitation else (6) for global imitation.
- 3: **while** convergence criterion is not met **do**
- 4:     Randomly sample a mini-batch data  $\hat{D}$ .
- 5:     Update  $f_{\ell, A}$  by solving (4) for local imitation or (8) for global imitation, using data  $\hat{D}$ .
- 6: **end while**
- 7: **Return:** The pruned layer  $f_{\ell, A}$ .

---

**Algorithm 2** Layer-wise Prune

---

- 1: **Input:** pretrained network  $F$  with  $L$  layers.
- 2: **for**  $\ell = 1, 2, \dots, L$  **do**
- 3:     Obtain the pruned layer  $f_{\ell, A_{\text{local}}}$  by local imitation on  $F$  with target layer  $\ell$ .
- 4:     Obtain the pruned layer  $f_{\ell, A_{\text{global}}}$  by global imitation on  $F$  with target layer  $\ell$ .
- 5:     Replace the  $\ell$ -th layer  $F_\ell$  of  $F$  with  $f_{\ell, A_{\text{local}}}$  if local imitation is better, else  $f_{\ell, A_{\text{global}}}$ .
- 6: **end for**
- 7: **Return:** The pruned network  $F$ .

---

## 2.4 Practical Algorithm: Pruning All Layers

In section 2.1 and 2.2 we introduce how to use the greedy local/global imitation to prune a certain layer in a network. In order to prune the whole network, we apply the greedy optimization scheme in a layer-wise fashion. Starting from the full network  $F$ , we apply the pruning method to prune the first layer (the one that is closest to the input)  $F_1$  to  $f_1$ , which returns a pruned network  $F^{\text{prune},1} = F_L \circ F_{L-1} \circ \dots \circ F_2 \circ f_1$ . We then apply the pruning method to prune the second layer  $F_2$  in  $F^{\text{prune},1}$  and continue until all the layers are pruned. By applying the pruning algorithm in this manner we prune the whole network.

The local imitation and global imitation perform differently when pruning different layers. To combine their advantages, when pruning each layer, both methods are applied individually with same convergence criterion and the one gives better performance is picked up. In this paper, we stop pruning when the discrepancy loss  $\mathbb{D}$  of the pruned model is smaller than a user specified threshold. The method prunes more neurons at convergence is selected. If both methods prune the same number of neurons, then the one with smaller discrepancy loss is chosen. Algorithm 1 summarizes the procedure of local and global imitation and Algorithm 2 gives the layer-wise scheme on pruning the whole deep network.

The exponential decay rate can be obtained by iteratively applying our theory on each layer. Suppose the pruned model  $F^{\text{pruned}}$  has  $n$  neurons at each layer, we have  $\mathbb{D}[F^{\text{pruned}}, F] = \mathcal{O}(\exp(-\lambda n))$ . We refer reader to Appendix 5.6 for details. Also notice that the exponential decay rate also holds for Algorithm 1 as it chooses the method with smaller loss.

## 3 Experiment

### 3.1 Comparing the Local and Global Imitation

Our first experiment aims to analyze the performance of local and global imitation for pruning deep neural network for image classification. We first apply both methods to a pretrained VGG-11 on CIFAR-10 dataset. We prune all the 8 convolution layers individual (when pruning one layer, the

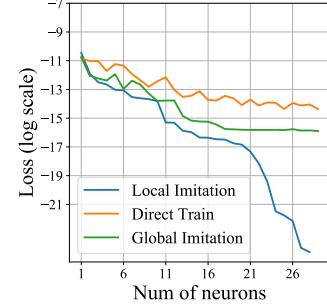


Figure 1: Discrepancy loss of the pruned model and train-from-scratch network with different sizes. The loss is in logarithm scale.

6

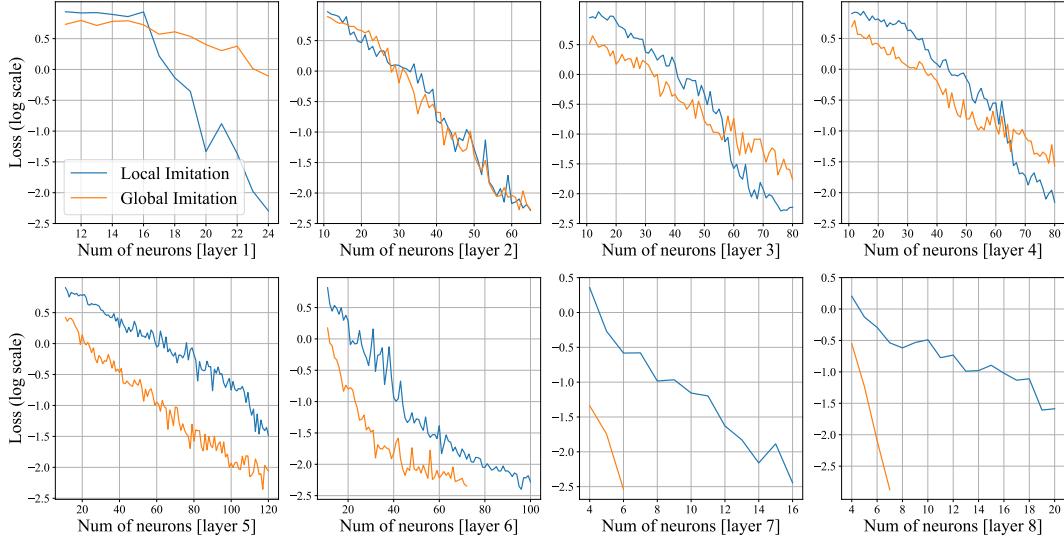


Figure 2: Pruning convolution layers on VGG11 using local and global imitation. From left to right and from top to bottom corresponds to the first (the one close to input) to the last convolution layers.

other layers remain unpruned) using both local and global imitation in order to compare these two methods side by side. Code for reproducing can be found at <https://github.com/lushleaf/Network-Pruning-Greedy-Forward-Selection>.

**Settings** The full network is trained with SGD optimizer with momentum 0.9. We use 128 batch size with initial learning rate 0.1 and train the model for 160 epochs. We decay the learning rate by 0.1 at the 80-th and the 120-th epochs. During pruning we use 128 batch size. We do not apply the Taylor approximation tricks to global imitation for this experiment. we use cross entropy between the pruned model and original model as discrepancy loss.

**Result** Figure 2 summarizes the result. Overall, we find the local and global imitation performs differently on different layers. The local imitation tends to decreases the loss faster on layer that is more close to input and with less neurons. While global imitation tends to performs much better than local imitation on layer that are close to output.

**Combining Local and Global Imitation Outperforms Both** In practice, we find purely pruning with local imitation tends gives worse result than pruning only with global imitation. However, combining the local imitation with global imitation performs better than pruning only with global imitation. To show this, we apply local and global imitation with the same setting as in section 2.4 on pruning ResNet34 and MobileNetV2 on ImageNet. For comparison, pruning with only global imitation is also applied. The local+global setting achieves 73.5 top1 accuracy on the pruned ResNet34 with 2.2G FLOPs and 72.2 top1 accuracy on the pruned MobileNetV2 with 245M FLOPs. While pruning with only global imitation only achieve 73.2 top1 accuracy on ResNet34 and 72.1 top1 accuracy on MobileNetV2 with same FLOPs. The experimental settings are in Section 3.2.

### 3.2 Imagenet Experiment

We use ILSVRC2012, a subset of ImageNet (Deng et al., 2009) which consists of about 1.28 million training images and 50000 validation images with 1000 different classes.

**Setting** We apply our method on pruning ResNet34 (traditional large architecture) (He et al., 2016), MobileNetV2 (efficient architecture) (Sandler et al., 2018) and MobileNetV3-small (an very small efficient architecture) (Howard et al., 2019) on ImageNet.

We use batch size 64 for both local and global imitation. We set the algorithm to converge when the gap between the cross entropy training loss before pruning and after pruning is smaller than  $\epsilon$ .

Model	Method	Top-1 Acc	Size (M)	FLOPs
ResNet34	Full Model (He et al., 2016)	73.4	21.8	3.68G
	$L_1$ norm (Li et al., 2017)	72.1	-	2.79G
	Neural Imp (Molchanov et al., 2019)	72.8	-	2.83G
	Rethink (Liu et al., 2018)	72.9	-	2.79G
	More is Less (Dong et al., 2017)	73.0	-	2.75G
	GFS (Ye et al., 2020)	73.5	17.2	2.64G
	Ours	<b>73.5</b>	<b>14.9</b>	<b>2.20G</b>
	SPF (He et al., 2018a)	71.8	-	2.17G
	FPGM (He et al., 2019)	72.5	-	2.16G
MobileNetV2	GFS (Ye et al., 2020)	72.9	14.7	2.07G
	Ours	<b>73.3</b>	<b>13.5</b>	<b>1.90G</b>
	Full Model (Sandler et al., 2018)	72.2	3.5	314M
	GFS (Ye et al., 2020)	71.9	3.2	258M
	Ours	<b>72.2</b>	<b>3.2</b>	<b>245M</b>
MobileNetV3-Small	Uniform (Sandler et al., 2018)	70.4	2.9	220M
	AMC (He et al., 2018b)	70.8	2.9	220M
	Meta Pruning (Liu et al., 2019)	71.2	-	217M
	LeGR (Chin et al., 2019)	71.4	-	224M
	GFS (Ye et al., 2020)	71.6	2.9	220M
	Ours	<b>71.7</b>	<b>2.9</b>	<b>218M</b>
	ThiNet (Luo et al., 2017)	68.6	-	175M
	DPL (Zhuang et al., 2018)	68.9	-	175M
MobileNetV3-Small	GFS (Ye et al., 2020)	70.4	2.3	170M
	Ours	<b>70.5</b>	<b>2.3</b>	<b>170M</b>
	Full Model (Howard et al., 2019)	67.5	2.5	64M
	Uniform (Howard et al., 2019)	65.4	2.0	47M
MobileNetV3-Small	GFS (Ye et al., 2020)	65.8	2.0	49M
	Ours	<b>66.4</b>	<b>2.0</b>	<b>48M</b>

Table 2: Result on pruning deep neural networks on ImageNet.

We vary  $\epsilon$  to get pruned model with different sizes. When conducting global imitation, we use the Taylor approximation trick introduced in section 2.2.1 to accelerate global imitation. We start the approximation when the number of neurons is larger than 25 and we evaluate the top 5 neurons with largest  $gr_{A,i}$  and pick up the best one to adjust. We find that this setting is able to produce the same pruning result as the exact version while substantially reduces the computation cost.

We finetune the pruned models with standard SGD optimizer with momentum 0.9 and weight decay  $5 \times 10^{-5}$ . All the pruned models are finetuned for 150 epochs with batch size 256 using cosine learning rate decay (Loshchilov & Hutter, 2016). We use initial learning 0.001 for ResNet34 and 0.01 for MobileNetV2 and MobileNetV3. We resize images to  $224 \times 224$  resolution and adopt the standard data augmentation scheme (mirroring and shifting).

**Result** Table 2 reports the top1 accuracy, FLOPs and model size of the pruned network. Our algorithm consistently improves prior arts on network pruning.

### 3.3 DGCNN Experiment

We conduct experiment on the point cloud classification tasks on ModelNet40. Since the network structure used to extract the global information in point cloud usually requires to aggregate features from neighbor points, the high feature dimension heavily influence the forward time. We deploy our method on DGCNN. We compare with several baselines, including PointNet (Qi et al., 2017a), PointNet++ (Qi et al., 2017b), DGCNN with different width multipliers, and signed splitting steepest descent (Wu et al., 2020), which obtains a compact DGCNN by growing a extremely thin model. Table 3 shows that our method produces networks with comparable accuracy while with much less inference time. We refer readers to Appendix 5.7 for details of experiment settings.

Model	Acc.	Forward time (ms)	# Param (M)
PointNet (Qi et al., 2017a)	89.2	32.19	2.85
PointNet++ (Qi et al., 2017b)	90.7	331.4	0.86
DGCNN (1.0x)	92.6	60.12	1.81
DGCNN (0.75x)	92.4	48.06	1.64
S3D (Wu et al., 2020)	<b>92.9</b>	42.06	1.51
Ours	<b>92.9</b>	<b>37.43</b>	1.49
DGCNN (0.5x)	92.3	38.90	1.52
Ours	<b>92.7</b>	<b>28.06</b>	1.31
DGCNN (0.25x)	91.8	30.90	1.42
Ours	<b>92.5</b>	<b>24.06</b>	1.24

Table 3: Results on the ModelNet40 classification task.

## 4 Related Work

**Greedy Method** Our method is highly related to Ye et al. (2020), which is also a greedy method with  $\mathcal{O}(n^{-2})$  error rate for pruning over-parameterized two layer network. In comparison, we obtain exponential decay rate for pruning deep neural network with no requirement on the over-parameterization of full model. Our local imitation method is also related to Frank Wolfe algorithm (Frank & Wolfe, 1956). Compared with it, our local imitation is a bi-level greedy joint optimization method while Frank Wolfe first searches for best direction and then conduct descent greedily.

**Theory on Lottery Ticket Hypothesis** (Malach et al., 2020) aims to prove the existence of subnetwork inside a random network that well approximates an unknown target network which has finite width and depth. It shows that a sufficiently large random network (with a specific structure) contains such a subnetwork with width of higher order compared with the targeted network. Later Pensia et al. (2020); Orseau et al. (2020) improve the result by reducing the size of the original random network. Elesedy et al. (2020) also gives analysis of Lottery Ticket Hypothesis in linear model using the tool from compressive sensing. Compared with our method, their theoretical results require strong structure assumptions on the full model and pruned model. Besides, they fail to give an efficient algorithm to search for the subnetwork for deep learning model in practice. Notice that Ye et al. (2020) also gives analysis on Lottery Ticket Hypothesis and it is straightforward to combine their framework and our analysis to give faster rate.

**Structured Pruning** Existing methods on structured pruning includes the sparsity regularization based training methods, e.g., Molchanov et al. (2017a); Liu et al. (2017); Ye et al. (2018); Huang & Wang (2018); criterion based methods, e.g., Molchanov et al. (2017b); Li et al. (2017); Molchanov et al. (2019), reconstruction error based method, e.g., He et al. (2017); Luo et al. (2017); Zhuang et al. (2018); Yu et al. (2018) and direct search method, e.g., He et al. (2018b); Liu et al. (2019). Our local imitation falls into the class of reconstruction error based method. Compared with those existing works, the proposed greedy optimization method enjoys good convergence property under weak assumptions and achieve better practical performance. Zhou et al. (2020) proposes a layer-wise imitation based training method for training deep and thin network, which is able to reduce the optimization error caused by the depth of the network. Their work is orthogonal to our work as we focus on reducing the error caused by small width.

## 5 Conclusion

This paper proposes a greedy optimization based pruning method, which is guaranteed to find a set of winning tickets (neurons) that approximates the fully trained unpruned network with exponential decay error rate w.r.t the number of selected tickets. The proposed pruning method is efficient with small time and space complexity and can be generally applied to various modern deep learning models.

**Broader Impact Statement** This work proposes a greedy optimization based pruning method, which has strong theoretical guarantee and good empirical performance. It gives positive improvement to the community of network efficiency. Our work do not have any negative societal impacts that we can foresee in the future.

**Acknowledgement** This paper is supported in part by NSF CAREER 1846421, SenSE 2037267 and EAGER 2041327.

## References

Amodei, Dario, Ananthanarayanan, Sundaram, Anubhai, Rishita, Bai, Jingliang, Battenberg, Eric, Case, Carl, Casper, Jared, Catanzaro, Bryan, Cheng, Qiang, Chen, Guoliang, et al. Deep speech 2: End-to-end speech recognition in english and mandarin. In *International conference on machine learning*, pp. 173–182, 2016.

Araújo, Dyego, Oliveira, Roberto I, and Yukimura, Daniel. A mean-field limit for certain deep neural networks. *arXiv preprint arXiv:1906.00193*, 2019.

Baykal, Cenk, Liebenwein, Lucas, Gilitschenski, Igor, Feldman, Dan, and Rus, Daniela. Data-dependent coresets for compressing neural networks with applications to generalization bounds. In *International Conference on Learning Representations*, 2019a. URL <https://openreview.net/forum?id=HJfwJ2A5KX>.

Baykal, Cenk, Liebenwein, Lucas, Gilitschenski, Igor, Feldman, Dan, and Rus, Daniela. Sipping neural networks: Sensitivity-informed provable pruning of neural networks. *arXiv preprint arXiv:1910.05422*, 2019b.

Cai, Han, Zhu, Ligeng, and Han, Song. ProxylessNAS: Direct neural architecture search on target task and hardware. In *International Conference on Learning Representations*, 2019. URL <https://arxiv.org/pdf/1812.00332.pdf>.

Chin, Ting-Wu, Ding, Ruizhou, Zhang, Cha, and Marculescu, Diana. Legr: Filter pruning via learned global ranking. *arXiv preprint arXiv:1904.12368*, 2019.

Deng, Jia, Dong, Wei, Socher, Richard, Li, Li-Jia, Li, Kai, and Fei-Fei, Li. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.

Dong, Xuanyi, Huang, Junshi, Yang, Yi, and Yan, Shuicheng. More is less: A more complicated network with less inference complexity. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5840–5848, 2017.

Elesedy, Bryn, Kanade, Varun, and Teh, Yee Whye. Lottery tickets in linear models: An analysis of iterative magnitude pruning. *arXiv preprint arXiv:2007.08243*, 2020.

Frank, Marguerite and Wolfe, Philip. An algorithm for quadratic programming. *Naval research logistics quarterly*, 3(1-2):95–110, 1956.

Frankle, Jonathan and Carbin, Michael. The lottery ticket hypothesis: Finding sparse, trainable neural networks. *arXiv preprint arXiv:1803.03635*, 2018.

Han, Song, Pool, Jeff, Tran, John, and Dally, William. Learning both weights and connections for efficient neural network. In *Advances in neural information processing systems*, pp. 1135–1143, 2015.

He, Kaiming, Zhang, Xiangyu, Ren, Shaoqing, and Sun, Jian. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

He, Yang, Kang, Guoliang, Dong, Xuanyi, Fu, Yanwei, and Yang, Yi. Soft filter pruning for accelerating deep convolutional neural networks. *arXiv preprint arXiv:1808.06866*, 2018a.

He, Yang, Liu, Ping, Wang, Ziwei, Hu, Zhilan, and Yang, Yi. Filter pruning via geometric median for deep convolutional neural networks acceleration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4340–4349, 2019.

He, Yihui, Zhang, Xiangyu, and Sun, Jian. Channel pruning for accelerating very deep neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1389–1397, 2017.

He, Yihui, Lin, Ji, Liu, Zhijian, Wang, Hanrui, Li, Li-Jia, and Han, Song. Amc: Automl for model compression and acceleration on mobile devices. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 784–800, 2018b.

Howard, Andrew, Sandler, Mark, Chu, Grace, Chen, Liang-Chieh, Chen, Bo, Tan, Mingxing, Wang, Weijun, Zhu, Yukun, Pang, Ruoming, Vasudevan, Vijay, et al. Searching for mobilenetv3. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1314–1324, 2019.

Huang, Zehao and Wang, Naiyan. Data-driven sparse structure selection for deep neural networks. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 304–320, 2018.

Li, Hao, Kadav, Asim, Durdanovic, Igor, Samet, Hanan, and Graf, Hans Peter. Pruning filters for efficient convnets. *The International Conference on Learning Representations*, 2017.

Liebenwein, Lucas, Baykal, Cenk, Lang, Harry, Feldman, Dan, and Rus, Daniela. Provable filter pruning for efficient neural networks. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=BJxk01SYDH>.

Liu, Zechun, Mu, Haoyuan, Zhang, Xiangyu, Guo, Zichao, Yang, Xin, Cheng, Kwang-Ting, and Sun, Jian. Metapruning: Meta learning for automatic neural network channel pruning. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 3296–3305, 2019.

Liu, Zhuang, Li, Jianguo, Shen, Zhiqiang, Huang, Gao, Yan, Shoumeng, and Zhang, Changshui. Learning efficient convolutional networks through network slimming. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2736–2744, 2017.

Liu, Zhuang, Sun, Mingjie, Zhou, Tinghui, Huang, Gao, and Darrell, Trevor. Rethinking the value of network pruning. *arXiv preprint arXiv:1810.05270*, 2018.

Loshchilov, Ilya and Hutter, Frank. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016.

Luo, Jian-Hao, Wu, Jianxin, and Lin, Weiyao. Thinet: A filter level pruning method for deep neural network compression. In *Proceedings of the IEEE international conference on computer vision*, pp. 5058–5066, 2017.

Malach, Eran, Yehudai, Gilad, Shalev-Shwartz, Shai, and Shamir, Ohad. Proving the lottery ticket hypothesis: Pruning is all you need. *arXiv preprint arXiv:2002.00585*, 2020.

Mei, Song, Montanari, Andrea, and Nguyen, Phan-Minh. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33):E7665–E7671, 2018.

Molchanov, Dmitry, Ashukha, Arsenii, and Vetrov, Dmitry. Variational dropout sparsifies deep neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 2498–2507. JMLR. org, 2017a.

Molchanov, Pavlo, Tyree, Stephen, Karras, Tero, Aila, Timo, and Kautz, Jan. Pruning convolutional neural networks for resource efficient inference. *International Conference on Learning Representations*, 2017b.

Molchanov, Pavlo, Mallya, Arun, Tyree, Stephen, Frosio, Iuri, and Kautz, Jan. Importance estimation for neural network pruning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 11264–11272, 2019.

Mussay, Ben, Osadchy, Margarita, Braverman, Vladimir, Zhou, Samson, and Feldman, Dan. Data-independent neural pruning via coresets. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=H1gmHaEKwB>.

Orseau, Laurent, Hutter, Marcus, and Rivasplata, Omar. Logarithmic pruning is all you need. *arXiv preprint arXiv:2006.12156*, 2020.

Pensia, Ankit, Rajput, Shashank, Nagle, Alliot, Vishwakarma, Harit, and Papailiopoulos, Dimitris. Optimal lottery tickets via subsetsum: Logarithmic over-parameterization is sufficient. *arXiv preprint arXiv:2006.07990*, 2020.

Qi, Charles R., Su, Hao, Mo, Kaichun, and Guibas, Leonidas J. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017a.

Qi, Charles Ruizhongtai, Yi, Li, Su, Hao, and Guibas, Leonidas J. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30*, pp. 5099–5108. 2017b.

Sandler, Mark, Howard, Andrew, Zhu, Menglong, Zhmoginov, Andrey, and Chen, Liang-Chieh. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4510–4520, 2018.

Sirignano, Justin and Spiliopoulos, Konstantinos. Mean field analysis of deep neural networks. *arXiv preprint arXiv:1903.04440*, 2019.

Wang, Yue, Sun, Yongbin, Liu, Ziwei, Sarma, Sanjay E, Bronstein, Michael M, and Solomon, Justin M. Dynamic graph cnn for learning on point clouds. *ACM Transactions on Graphics (TOG)*, 38(5):1–12, 2019.

Wu, Lemeng, Ye, Mao, Lei, Qi, Lee, Jason D, and Liu, Qiang. Steepest descent neural architecture optimization: Escaping local optimum with signed neural splitting. *arXiv preprint arXiv:2003.10392*, 2020.

Wu, Yonghui, Schuster, Mike, Chen, Zhifeng, Le, Quoc V, Norouzi, Mohammad, Macherey, Wolfgang, Krikun, Maxim, Cao, Yuan, Gao, Qin, Macherey, Klaus, et al. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*, 2016.

Wu, Zhirong, Song, Shuran, Khosla, Aditya, Yu, Fisher, Zhang, Linguang, Tang, Xiaouou, and Xiao, Jianxiong. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1912–1920, 2015.

Ye, Jianbo, Lu, Xin, Lin, Zhe, and Wang, James Z. Rethinking the smaller-norm-less-informative assumption in channel pruning of convolution layers. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=HJ94fqApW>.

Ye, Mao, Gong, Chengyue, Nie, Lizhen, Zhou, Denny, Klivans, Adam, and Liu, Qiang. Good subnetworks provably exist: Pruning via greedy forward selection. *arXiv preprint arXiv:2003.01794*, 2020.

Yu, Ruichi, Li, Ang, Chen, Chun-Fu, Lai, Jui-Hsin, Morariu, Vlad I, Han, Xintong, Gao, Mingfei, Lin, Ching-Yung, and Davis, Larry S. Nisp: Pruning networks using neuron importance score propagation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 9194–9203, 2018.

Zhou, Denny, Ye, Mao, Chen, Chen, Meng, Tianjian, Tan, Mingxing, Song, Xiaodan, Le, Quoc, Liu, Qiang, and Schuurmans, Dale. Go wide, then narrow: Efficient training of deep thin networks. *arXiv preprint arXiv:2007.00811*, 2020.

Zhuang, Zhuangwei, Tan, Mingkui, Zhuang, Bohan, Liu, Jing, Guo, Yong, Wu, Qingyao, Huang, Junzhou, and Zhu, Jinhui. Discrimination-aware channel pruning for deep neural networks. In *Advances in Neural Information Processing Systems*, pp. 875–886, 2018.

## Appendix

We introduce the following extra notations which are used in several parts of the Appendix. Suppose that  $\theta_i^\ell, i \in [N]$  is the weight of the  $N$  neurons of the  $\ell$ -th layer in the original network  $F$ . We simplify the notation by denoting  $\theta_i^\ell = \theta_i, i \in [N]$ . Suppose  $\sigma(\theta, z) \in \mathbb{R}^d$ , we define the data-dependent feature by

$$\phi_{\text{Matrix}}(\theta) = \left[ \sigma(\theta, z^{(1)}), \dots, \sigma(\theta, z^{(m)}) \right]^\top \in \mathbb{R}^{m \times d},$$

and its vectorization

$$\phi(\theta) = \left[ \text{vec}^\top \left( \sigma(\theta, z^{(1)}) \right), \dots, \text{vec}^\top \left( \sigma(\theta, z^{(m)}) \right) \right]^\top \in \mathbb{R}^{md}.$$

We also define  $\mathbf{h}_i = \phi(\theta_i), i \in [N]$ ,  $\bar{\mathbf{h}} = \frac{1}{N} \sum_{i=1}^N \mathbf{h}_i$  and  $\mathbf{h}_A = \sum_{i=1}^N a_i \mathbf{h}_i$  with  $A = [a_1, a_2, \dots, a_N]$ . Define  $\mathcal{M} = \text{conv}(\{\phi(\theta_i) \mid i \in [N]\})$  as the convex hull generated by the set  $\{\phi(\theta_i) \mid i \in [N]\}$ . Given some set  $M \subseteq \mathbb{R}^d$ , we denote the relative interior of  $M$  by  $\text{ri}M$ , the closure of  $M$  by  $\text{cl}M$  and the affine hull of  $M$  by  $\text{Aff}M$ . We define  $\mathcal{B}(\mathbf{x}_0, r)$  as the ball centered at  $\mathbf{x}_0$  with radius  $r$ .

### 5.1 Details on Solving (4) for Local Imitation

Now we describe the approach to solve problem (4) with only compute one forward pass. Define

$$V_{i,A}(\gamma) = \bar{\mathbb{D}}[(1 - \gamma)f_{\ell,A} + \gamma\sigma(\theta_i^\ell, \cdot), F_\ell] = \gamma^2 g_{i,A} - 2\gamma q_{i,A} + \bar{\mathbb{D}}[f_{\ell,A}, F_\ell],$$

$$\text{where } q_{i,A} = \mathbb{E}_{z \sim \mathcal{D}_m^\ell} [(F_\ell - f_{\ell,A}) (\sigma(\theta_i^\ell, \cdot) - f_{\ell,A})], \quad g_{i,A} = \mathbb{E}_{z \sim \mathcal{D}_m^\ell} [\left( \sigma(\theta_i^\ell, \cdot) - f_{\ell,A} \right)^2].$$

Notice that it is easy to obtain  $\sigma(\theta_i^\ell, z^{(j)})$  for all  $i \in [N]$  and  $j \in [m]$  by feeding the dataset into the neural network once, as it is the output of neuron  $i$  in layer  $j$ . And thus  $q_{i,A(k)}$  and  $g_{i,A(k)}$  can be calculated cheaply given  $\sigma(\theta_i^\ell, z^{(j)}), i \in [N]$  and  $j \in [m]$ . Define  $\tilde{\gamma}_{i,A(k)} = q_{i,A(k)} / g_{i,A(k)}$ , which is the optimum of  $V_{i,A(k)}(\gamma)$  w.r.t.  $\gamma$  given  $i$  (here the optimization of  $\gamma$  is unconstrained). The following theorem shows some properties of the greedy local imitation method.

**Theorem 3** *Under assumption 1, if  $\bar{\mathbb{D}}[f_{\ell,A(k)}, F_\ell] > 0$ , then we have  $\tilde{\gamma}_{\ell,i_{\ell,k}^*} < 1$ .*

Now we proceed to show how to obtain  $i_k^*$  and  $\gamma_k^*$  efficiently. Suppose at iteration  $k$ ,  $\bar{\mathbb{D}}[f_{\ell,A(k)}, F_\ell] > 0$  (otherwise the algorithm has converged). Given some neuron  $i$  with  $a_i(k) = 0$ , we first calculate  $\tilde{\gamma}_{i,A(k)}$ . If  $\tilde{\gamma}_{i,A(k)} \in (0, 1)$ , then we define the score of this neuron by the decrease of loss with this neuron selected, i.e.,

$$\text{score}_{A(k)}(i) := \bar{\mathbb{D}}[f_{\ell,A}, F_\ell] - \min_{\gamma \in U_i} V_{i,A(k)}(\gamma) = -q_{i,A(k)}^2 / g_{i,A(k)}.$$

If  $\tilde{\gamma}_{i,A(k)} \geq 1$ , then from Theorem 3, we know that  $i \neq i_k^*$ . If  $\tilde{\gamma}_{i,A(k)} \leq 0$  and if  $i = i_k^*$ , we have  $\gamma_k^* = 0$ , which implies that  $\bar{\mathbb{D}}[f_{\ell,A(k+1)}, F_\ell] = \bar{\mathbb{D}}[f_{\ell,A(k)}, F_\ell]$ . It makes contradiction to Theorem 3, which implies that  $i \neq i_k^*$ . In this two cases, since neuron  $i$  is not the optimal neuron to select, we can safely set  $\text{score}_{A(k)}(i) = 0$ . For neuron  $i$  with  $a_i(k) > 0$ . Similarly, if  $\tilde{\gamma}_{i,A(k)} \geq 1$ , then  $i \neq i_k^*$  and thus we set  $\text{score}_{A(k)}(i) = 0$ . If  $\tilde{\gamma}_{i,A(k)} \in U_i$ , then similarly  $\text{score}_{A(k)}(i) = -q_{i,A(k)}^2 / g_{i,A(k)}$ . If  $\tilde{\gamma}_{i,A(k)} < -a_i(k)/(1 - a_i(k))$ , then

$$\text{score}_{A(k)}(i) = V_{i,A(k)}(-a_i(k)/(1 - a_i(k))).$$

And thus we have  $i_k^* = \arg \max_{i \in [N]} \text{score}_{A(k)}(i)$ . Notice that the score of most neuron can be calculated cheaply using  $q_{i,A(k)}$  and  $g_{i,A(k)}$ . The only exception are neuron with  $a_i(k) > 0$  and  $\tilde{\gamma}_{i,A(k)} < -a_i(k)/(1 - a_i(k))$ . However, its score can be calculated using  $\sigma(\theta_i^\ell, \cdot)$  and thus no extra forward pass is required.

## 5.2 Local Imitation with Fixed Step Sizes

In this section we give detailed discussion on local imitation with a fixed step size scheme shown in Section 2.2.1. Different from the greedy optimization (4), in this scheme, as the step size is fixed, the solution returned in each iteration is no better than the one of (4). As a consequence, it gives slower convergence rate.

**Theorem 4** *Under Assumption 1, at each step  $k$  of the greedy optimization in 5, we have  $\mathbb{D}[f_{\mathbf{A}(k)}, F] \leq \|H\|_{Lip}^2 \bar{\mathbb{D}}[\ell_{\mathbf{A}(k)}, F_\ell] = \mathcal{O}((k+1)^{-2})$  and  $\|\mathbf{A}(k)\|_0 \leq k+1$ .*

## 5.3 Theory on Greedy Global Imitation

Now we give the theoretical result on greedy global imitation. Denote  $\kappa_1 = \sup_{\mathbf{A} \in \Omega_N} \frac{\|H \circ \bar{\mathbf{h}} - H \circ \mathbf{h}_A\|}{\|\bar{\mathbf{h}} - \mathbf{h}_A\|}$ ,  $\kappa_2 = \sup_{\mathbf{A} \in \Omega_N} \frac{\|\bar{\mathbf{h}} - \mathbf{h}_A\|}{\|H \circ \bar{\mathbf{h}} - H \circ \mathbf{h}_A\|}$  and  $D$  as the diameter of  $\mathcal{M}$ , which is defined in Lemma 5. Notice that  $\kappa_1 \kappa_2 \geq 1$ . Using Lemma 3, we know that  $\bar{\mathbf{h}} \in \text{ri}\mathcal{M}$ , which indicate that there exists some  $\lambda > 0$  such that

$$\mathcal{B}(\bar{\mathbf{h}}, \lambda) \cap \text{Aff}\mathcal{M} \subseteq \mathcal{M},$$

where  $\mathcal{B}(\bar{\mathbf{h}}, \lambda)$  denotes the ball with radius  $\lambda$  centered at  $\bar{\mathbf{h}}$ .

**Theorem 5 (Complete Version of Theorem 2)** *Suppose Assumption 1 holds. Further suppose that 1.  $D^2 \geq \kappa_1^2 \kappa_2^2 (D^2 - \lambda^2)$ ; 2. at initialization  $\|\bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(0)}\| \leq R$ ; 3.  $\kappa_1 D \leq R$ , where we define  $R = \frac{\kappa_1^2 \kappa_2 \lambda + \kappa_1 \sqrt{\kappa_1^2 \kappa_2^2 (\lambda^2 - D^2) + D^2}}{(\kappa_1^2 \kappa_2^2 - 1)}$  ( $R = +\infty$  if  $\kappa_1 \kappa_2 = 1$ ). Then we have  $\mathbb{D}[f_{\mathbf{A}(k)}, F] = \mathcal{O}((k+1)^{-2})$ , and  $\|\mathbf{A}(k)\|_0 \leq k+1$ .*

**Remark** Here the descending property of global imitation is influenced by the non-linear mapping  $H$ . As a consequence, the algorithm gives good convergence property when the whole dynamics is guaranteed to stay in a proper convergence region ( $R$ ). The first extra assumption assumes the existence of this convergence region; The second extra assumption assumes a good initialization to ensure the dynamics stays in the convergence region at initialization; The third assumption can be roughly interpreted as assuming the dynamics will not jump out of the convergence region during descending. Notice that the extra assumptions holds when  $\kappa_1$  and  $\kappa_2$  is sufficiently close to 1.

## 5.4 Details on Taylor Approximation Tricks

In this section we give details on the computation of Taylor approximation tricks. Notice that

$$\begin{aligned} gr_{\mathbf{A},i} &= \frac{\partial}{\partial \gamma} \mathbb{D}[H \circ [(1-\gamma)f_{\ell,\mathbf{A}} + \gamma\sigma(\boldsymbol{\theta}_i, \cdot)], F] \Big|_{\gamma=0} \\ &= \frac{\partial}{\partial \gamma} \mathbb{E}_{\mathbf{z} \sim \mathcal{D}_m^\ell} (H \circ [(1-\gamma)f_{\ell,\mathbf{A}}(\mathbf{z}) + \gamma\sigma(\boldsymbol{\theta}_i, \mathbf{z})] - H \circ F_\ell(\mathbf{z}))^2 \Big|_{\gamma=0} \\ &= 2\mathbb{E}_{\mathbf{z} \sim \mathcal{D}_m^\ell} (H \circ [(1-\gamma)f_{\ell,\mathbf{A}}(\mathbf{z}) + \gamma\sigma(\boldsymbol{\theta}_i, \mathbf{z})] - H \circ F_\ell(\mathbf{z})) \frac{\partial}{\partial \gamma} (H \circ [(1-\gamma)f_{\ell,\mathbf{A}}(\mathbf{z}) + \gamma\sigma(\boldsymbol{\theta}_i, \mathbf{z})]) \Big|_{\gamma=0} \\ &= 2\mathbb{E}_{\mathbf{z} \sim \mathcal{D}_m^\ell} (H \circ f_{\ell,\mathbf{A}}(\mathbf{z}) - H \circ F_\ell(\mathbf{z})) H'(f_{\ell,\mathbf{A}}(\mathbf{z})) (\sigma(\boldsymbol{\theta}_i, \mathbf{z}) - f_{\ell,\mathbf{A}}(\mathbf{z})) \\ &= 2\mathbb{E}_{\mathbf{z} \sim \mathcal{D}_m^\ell} (H \circ f_{\ell,\mathbf{A}}(\mathbf{z}) - H \circ F_\ell(\mathbf{z})) H'(f_{\ell,\mathbf{A}}(\mathbf{z})) \left( \sigma(\boldsymbol{\theta}_i, \mathbf{z}) - \sum_{j=1}^N a_j \sigma(\boldsymbol{\theta}_j, \mathbf{z}) \right) \\ &= 2 \sum_{j=1}^N (\mathbb{I}\{j = i\} - a_j) r_{\mathbf{A},j}. \end{aligned}$$

Thus the key quantities we want to obtain is

$$r_{\mathbf{A},i} = \mathbb{E}_{\mathbf{z} \sim \mathcal{D}_m^\ell} [(H \circ f_{\ell,\mathbf{A}}(\mathbf{z}) - H \circ F_\ell(\mathbf{z})) H'(f_{\ell,\mathbf{A}}(\mathbf{z})) \sigma(\boldsymbol{\theta}_i^\ell, \mathbf{z})].$$

And once we obtain  $r_{\mathbf{A},i}$ , we are able to calculate  $gr_{\mathbf{A},i} = 2 \sum_{j=1}^N (\mathbb{I}\{j=i\} - a_j) r_{\mathbf{A},j}$ . Now we introduce how to calculate  $r_{\mathbf{A},i}$  efficiently by introducing an ancillary variable. Suppose when pruning layer  $\ell$ , we have

$$f_{\ell,\mathbf{A}}(\mathbf{z}) = \sum_{i=1}^N a_i \sigma(\boldsymbol{\theta}_i^\ell, \mathbf{z}) = \sum_{i=1}^N (a_i + b_i) \sigma(\boldsymbol{\theta}_i^\ell, \mathbf{z}), \text{ where } b_i = 0 \forall i \in [N].$$

Here  $b_i$  is the introduced ancillary variable, which always takes 0 value. We have

$$\begin{aligned} & \frac{\partial}{\partial b_i} \mathbb{E}_{\mathbf{z} \sim \mathcal{D}_m^\ell} \left( H \circ \left( \sum_{i=1}^N (a_i + b_i) \sigma(\boldsymbol{\theta}_i^\ell, \mathbf{z}) \right) - H \circ F_\ell(\mathbf{z}) \right)^2 \Big|_{b_i=0} \\ &= \mathbb{E}_{\mathbf{z} \sim \mathcal{D}_m^\ell} (H \circ f_{\ell,\mathbf{A}}(\mathbf{z}) - H \circ F_\ell(\mathbf{z})) \frac{\partial}{\partial b_i} (H \circ f_{\ell,\mathbf{A}}) \Big|_{b_i=0} \\ &= \mathbb{E}_{\mathbf{z} \sim \mathcal{D}_m^\ell} (H \circ f_{\ell,\mathbf{A}}(\mathbf{z}) - H \circ F_\ell(\mathbf{z})) H'(f_{\ell,\mathbf{A}}) \frac{\partial}{\partial b_i} f_{\ell,\mathbf{A}} \Big|_{b_i=0} \\ &= \mathbb{E}_{\mathbf{z} \sim \mathcal{D}_m^\ell} (H \circ f_{\ell,\mathbf{A}}(\mathbf{z}) - H \circ F_\ell(\mathbf{z})) H'(f_{\ell,\mathbf{A}}) \frac{\partial}{\partial b_i} \left( \sum_{i=1}^N (a_i + b_i) \sigma(\boldsymbol{\theta}_i^\ell, \mathbf{z}) \right) \Big|_{b_i=0} \\ &= \mathbb{E}_{\mathbf{z} \sim \mathcal{D}_m^\ell} (H \circ f_{\ell,\mathbf{A}}(\mathbf{z}) - H \circ F_\ell(\mathbf{z})) H'(f_{\ell,\mathbf{A}}) \sigma(\boldsymbol{\theta}_i^\ell, \mathbf{z}) \\ &= r_{\mathbf{A},i}. \end{aligned}$$

This for implementation in practice, we can introduce  $b_i$  with its value fixed 0 and calculate its gradient using, which is  $r_{\mathbf{A},i}$  using the auto differentiate operator in common deep learning libraries.

## 5.5 Details on Numeric Verification of Rate

In this section, we give details on the toy experiment on verifying numeric rate. We first introduce the problem setup for the comparison between pruning and direct gradient training in obtaining small network. We use the two-hidden-layer deep mean field network formulated by Araújo et al. (2019). Suppose that the second hidden layer (the one close to output) has 50 neurons; the first hidden layer (the one close to input) has  $n \leq 50$  neurons with 50 dimensional feature map; and the input has 100 dimension. That is, we consider the following deep mean field network

$$F^n(\mathbf{x}) = F_2 \circ F_1^n(\mathbf{x}),$$

where

$$F_1^n(\mathbf{x}) = \frac{1}{n} \sum_{i=1}^n a_{1,i} \text{ReLU}(\mathbf{b}_{1,i}^\top \mathbf{x})$$

with  $\mathbf{x} \in \mathbb{R}^{100}$ ,  $\mathbf{b}_{1,i} \in \mathbb{R}^{100 \times 50}$ ,  $a_{1,i} \in \mathbb{R}$ . And

$$F_2(\mathbf{z}) = \frac{1}{50} \sum_{i=1}^{50} a_{2,i} \text{ReLU}(\mathbf{b}_{2,i}^\top \mathbf{z}),$$

with  $\mathbf{z} \in \mathbb{R}^{50}$ ,  $\mathbf{b}_{2,i} \in \mathbb{R}^{50}$ ,  $a_{2,i} \in \mathbb{R}$ . Suppose that we train the original network  $F^N$  with  $N = 50$  neuron at the first hidden layer using gradient descent defined in Araújo et al. (2019) for  $T$  time ( $T < \infty$ ) with random initialization. To obtain a small network with  $n$  neurons at the first hidden layer, we consider two approaches. In the first approach, we prune the first hidden layer of the trained  $F^N$  using local imitation to obtain  $F_{\text{local}}^n$  where  $n$  indicates the number of neurons remained in the first hidden layer. In the second approach, we direct train the network  $F_{\text{direct train}}^n$  with  $n$  neurons in the first layer using the same gradient descent dynamics, initialization and training time as that in training  $F^N$ . By the analysis in Araújo et al. (2019), we have  $\mathbb{D}[F_{\text{direct train}}^n, F^N] = \mathcal{O}(n^{-1})$ . And by Theorem 1, we have  $\mathbb{D}[F_{\text{local}}^n, F^N] = \mathcal{O}(\exp(-\lambda n))$  for some  $\lambda > 0$ . This implies that pruning is provably much better than directly training in obtaining compact neural network.

Now we introduce the experiment settings. To simulate the data, we first generate a random network

$$F_{\text{gen}}(\mathbf{x}) = (\exp(\mathbf{w}_2/10) - 0.5)^\top \text{Tanh}(\sin(2\pi \mathbf{w}_1)^\top \mathbf{x}/5)/1000,$$

where  $\mathbf{w}_1 \in \mathbb{R}^{1000 \times 100}$  and  $\mathbf{w}_2 \in \mathbb{R}^{1000}$  is generated by randomly sampling from uniform distribution  $\text{Unif}[0, 1]$  (each element is sampled independently). And then we generate the training data by sampling feature  $\mathbf{x}$  from  $\text{Unif}[0, 1]$  (each coordinate is sampled independently) and then generate label  $y = F_{\text{gen}}(\mathbf{x})$ . The simulated training dataset consists of 200 data points. We initialize the parameters of  $F^N$  and  $F_{\text{direct train}}^n$  from standard Gaussian distribution with variance 1,  $\mathcal{N}(0, 1)$  (each element are initialized independently) and both  $F^N$  and  $F_{\text{direct train}}^n$  are trained using the same and sufficiently long time to ensure convergence. We also include the pruned model using global imitation, which is denoted as  $F_{\text{global}}^n$ . The pruned models are not finetuned. We vary different  $n$  and summarize the discrepancy.

## 5.6 Theory on Pruning All Layers

In the main text, we mainly discuss the convergence rate of pruning one layer. In this section, we discuss how to apply our convergence rate for single layer pruning to obtain an overall convergence rate. Following the layer-wise procedure introduced in Section 2.4, suppose that the algorithms prunes  $F_\ell$  to  $f_{\ell, \mathbf{A}_\ell}$ ,  $\ell \in [L]$ . And thus, during the layer-wise pruning, the algorithm generates a sequence of pruned networks

$$\begin{aligned} f_{[0]} &= F_L \circ F_{L-1} \circ \dots \circ F_3 \circ F_2 \circ F_1 \\ f_{[1]} &= F_L \circ F_{L-1} \circ \dots \circ F_3 \circ F_2 \circ f_{1, \mathbf{A}_1} \\ f_{[2]} &= F_L \circ F_{L-1} \circ \dots \circ F_3 \circ f_{2, \mathbf{A}_2} \circ f_{1, \mathbf{A}_1} \\ &\vdots \\ f_{[L-1]} &= F_L \circ f_{L-1, \mathbf{A}_{L-1}} \circ \dots \circ f_{3, \mathbf{A}_3} \circ f_{2, \mathbf{A}_2} \circ f_{1, \mathbf{A}_1} \\ f_{[L]} &= f_{L, \mathbf{A}_L} \circ f_{L-1, \mathbf{A}_{L-1}} \circ \dots \circ f_{3, \mathbf{A}_3} \circ f_{2, \mathbf{A}_2} \circ f_{1, \mathbf{A}_1} \end{aligned}$$

Thus here  $f_{[\ell]}$  is the network with the first  $\ell$  layers pruned,  $f_{[L]}$  is the final pruned network with all layers pruned and  $f_{[0]}$  is the original network. Notice that  $f_{[\ell]}$  is obtained by pruning the  $\ell$ -th layer of  $f_{[\ell-1]}$ . In this step, we suppose that we try both greedy local and global imitation and obtain  $f_{\ell, \mathbf{A}_\ell^{\text{local}}}$  and  $f_{\ell, \mathbf{A}_\ell^{\text{global}}}$  with  $\|\mathbf{A}_\ell^{\text{local}}\|_0 = \|\mathbf{A}_\ell^{\text{global}}\|_0$ . And if  $\mathbb{D}[F_L \circ \dots \circ F_{\ell+1} \circ f_{\ell, \mathbf{A}_\ell^{\text{local}}} \circ \dots \circ f_{1, \mathbf{A}_1}, f_{[\ell-1]}] \leq \mathbb{D}[F_L \circ \dots \circ F_{\ell+1} \circ f_{\ell, \mathbf{A}_\ell^{\text{global}}} \circ \dots \circ f_{1, \mathbf{A}_1}, f_{[\ell-1]}]$ , we set  $\mathbf{A}_\ell = \mathbf{A}_\ell^{\text{local}}$ , else we set  $\mathbf{A}_\ell = \mathbf{A}_\ell^{\text{global}}$ . Define  $H_{[\ell]} = F_L \circ F_{L-1} \circ \dots \circ F_{\ell+1}$ ,  $\ell \in [L-1]$  (here  $H_{[L-1]} = F_L$ ) and  $\mathbf{z}_{[\ell]}^{(i)} = f_{\ell-1, \mathbf{A}_{\ell-1}} \circ \dots \circ f_{1, \mathbf{A}_1}(\mathbf{x}^{(i)})$ ,  $\ell \in [L-1]$  (here we define  $\mathbf{z}_{[1]}^{(i)} = \mathbf{x}^{(i)}$ ). The set  $\mathcal{D}_m^{[\ell]} := \left( \mathbf{z}_{[\ell]}^{(i)} \right)_{i=1}^m$  denotes the distribution of training data pushed through the first  $\ell - 1$  layers.

We introduce the following assumption on the boundedness.

**Assumption 2** Assume that for any  $i \in [N]$ ,  $\ell \in [L-1]$ ,  $\mathbf{z}_{[\ell]}^{(j)} \in \mathcal{D}_m^{[\ell]}$ , we have  $\|\sigma(\theta_i^\ell, \mathbf{z}_{[\ell]}^{(j)})\| \leq c_2$  and  $\|H_{[\ell]}\|_{\text{Lip}} \leq c_2$  for some  $c_2 < \infty$ .

**Theorem 6 (Overall Convergence)** Under assumption 2, we have  $\sqrt{\mathbb{D}[f_{[L]}, F]} = \mathcal{O}\left(\sum_{\ell=1}^L \exp\left(-\frac{\lambda_\ell}{2} \|\mathbf{A}_\ell\|_0\right)\right)$ , with  $\lambda_\ell > 0$  for all  $\ell \in [L]$  depending on  $f_{[\ell-1]}$ .

## 5.7 DGCNN Experiment

We deploy our method on DGCNN (Wang et al., 2019). DGCNN contains 4 EdgeConv layers that use K-Nearest-Neighbor(KNN) to aggregate the information from the output of convolution operation. Pruning the convolution operation in EdgeConv can significantly speed up the KNN operation and therefor make the whole model more computational efficient.

**Settings** The full network is trained with SGD optimizer with momentum 0.9 and weight decay  $1 \times 10^{-4}$ . We train the model using 64 batch size with an initial learning rate 0.1 for 250 epochs. We apply cosine learning rate scheduler during the training and decrease the learning rate to 0.001 at the final epoch. During the pruning, we use 32 batch sizes and the other settings keep the same as our ImageNet experiment in Section 3.2.

## Technical Lemmas

We introduce several technical Lemmas that are useful for proving the main theorems.

**Lemma 1** *Given some convex set  $M \subset \mathbb{R}^d$ , for any  $\mathbf{q}_1 \in riM$  and  $\mathbf{q}_2 \in clM$ . Then all the points from the half-segment  $[\mathbf{q}_1, \mathbf{q}_2]$  belongs to the relative interior of  $M$ , i.e.,*

$$[\mathbf{q}_1, \mathbf{q}_2] \hat{=} \{(1 - \lambda)\mathbf{q}_1 + \lambda\mathbf{q}_2 \mid 0 \leq \lambda < 1\} \subseteq riM.$$

**Lemma 2** *Let  $M$  be a convex set in  $\mathbb{R}^d$ , then if  $M$  is nonempty, then the relative interior of  $M$  is nonempty.*

Lemma 1 and 2 are classic results from convex optimization.

**Lemma 3** *Define*

$$M = \text{conv} \{ \mathbf{q} \mid \mathbf{q} \in S \},$$

where  $S = \{\mathbf{q}_1, \dots, \mathbf{q}_n\} \subseteq \mathbb{R}^d$  with  $1 \leq n < \infty$ . Define  $\bar{\mathbf{q}} = \frac{1}{n} \sum_{i=1}^n \mathbf{q}_i$ , then  $\bar{\mathbf{q}} \in riM$ .

**Lemma 4** *Suppose that for some  $\lambda > 0$  such that  $(\mathcal{B}(\bar{\mathbf{h}}, \lambda) \cap \text{Aff}\mathcal{M}) \subseteq \mathcal{M}$ , then  $\max_{\mathbf{s} \in \mathcal{M}} \langle \bar{\mathbf{h}} - \mathbf{h}_A, \bar{\mathbf{h}} - \mathbf{s} \rangle \geq \lambda \|\bar{\mathbf{h}} - \mathbf{h}_A\|$ .*

**Lemma 5** *Under Assumption 1, for any  $\mathbf{h}, \mathbf{h}' \in \mathcal{M}$ ,  $\|\mathbf{h} - \mathbf{h}'\| \leq D$  for some  $D \leq 2\sqrt{mc_1}$ . Here  $D$  can be viewed as the diameter of  $\mathcal{M}$ .*

**Lemma 6** *Under assumption 1, suppose  $\tilde{\mathbf{s}}_k^* = \arg \min_{\mathbf{s} \in \mathcal{M}} \langle \bar{\mathbf{h}} - \mathbf{h}_{A(k)}, \mathbf{s} - \mathbf{h}_{A(k)} \rangle$  and  $\tilde{\gamma}_k^* = \arg \min_{\gamma \in [0, 1]} \|\mathbf{h}_{A(k)} + \gamma (\tilde{\mathbf{s}}_k^* - \mathbf{h}_{A(k)}) - \bar{\mathbf{h}}\|^2$ , then*

$$\|\mathbf{h}_{A(k)} + \tilde{\gamma}_k^* (\tilde{\mathbf{s}}_k^* - \mathbf{h}_{A(k)}) - \bar{\mathbf{h}}\|^2 \leq \rho \|\mathbf{h}_{A(k)} - \bar{\mathbf{h}}\|^2,$$

for some  $\rho \in (0, 1)$ .

**Lemma 7** *Consider the following number sequence  $x_{k+1}^2 \leq ax_k^2 - bx_k + c$ . Suppose that this number sequence satisfies the following conditions: (1)  $a > 1$ ,  $b \geq 0$ ,  $c \geq 0$ ; (2)  $x_k \geq 0$  for any  $k$ ; (3)  $(a-1)x^2 - bx^2 + c$  has two real roots  $z_1 \leq z_2$ ; (4)  $\sqrt{c} \leq z_2$ ; (5)  $x_0 \leq z_2$ . Then  $\sup_k x_k \leq z_2$ .*

## Proof of Main Theorems

### 5.7.1 Proof of Theorem 1

Using Lemma 3, we know that  $\bar{\mathbf{h}} \in ri\mathcal{M}$ , which indicate that there exists some  $\lambda > 0$  such that

$$\mathcal{B}(\bar{\mathbf{h}}, \lambda) \cap \text{Aff}\mathcal{M} \subseteq \mathcal{M},$$

where  $\mathcal{B}(\bar{\mathbf{h}}, \lambda)$  denotes the ball with radius  $\lambda$  centered at  $\bar{\mathbf{h}}$ . Define  $\text{Extre}(\mathcal{M})$  as the set of extreme points of  $\mathcal{M}$ , we know that  $\text{Extre}(\mathcal{M}) \subseteq \{\mathbf{h}_1, \dots, \mathbf{h}_N\}$ . Consider the following problem

$$\min_{\mathbf{s} \in \mathcal{M}} \langle \bar{\mathbf{h}} - \mathbf{h}_{A(k)}, \mathbf{s} - \mathbf{h}_{A(k)} \rangle.$$

As the objective  $\langle \bar{\mathbf{h}} - \mathbf{h}_{A(k)}, \mathbf{s} - \mathbf{h}_{A(k)} \rangle$  is linear w.r.t.  $\mathbf{s}$ , we know that  $\mathbf{s} \in \text{Extre}(\mathcal{M}) \subseteq \{\mathbf{h}_1, \dots, \mathbf{h}_N\}$ . Also, for any  $i \in [N]$ , we have  $[0, 1] \subseteq U_i$ . This gives that

$$\begin{aligned} \min_{i \in [N]} \min_{\gamma \in U_i} \bar{\mathbb{D}}[(1 - \gamma)f_{\ell, A(k)} + \gamma\sigma(\boldsymbol{\theta}_i, \cdot), F_\ell] &\leq \min_{\gamma \in [0, 1]} \|\mathbf{h}_{A(k)} + \gamma (\tilde{\mathbf{s}}_k^* - \mathbf{h}_{A(k)}) - \bar{\mathbf{h}}\|^2 \\ &\leq (1 - \lambda^2/D^2) \|\mathbf{h}_{A(k)} - \bar{\mathbf{h}}\|^2, \end{aligned}$$

where  $\tilde{\mathbf{s}}_k^* = \arg \min_{\mathbf{s} \in \mathcal{M}} \langle \bar{\mathbf{h}} - \mathbf{h}_{A(k)}, \mathbf{s} - \mathbf{h}_{A(k)} \rangle$ . Here the last inequality is by Lemma 6. This gives that

$$\|\mathbf{h}_{A(k+1)} - \bar{\mathbf{h}}\|^2 \leq (1 - \lambda^2/D^2) \|\mathbf{h}_{A(k)} - \bar{\mathbf{h}}\|^2.$$

And thus we have

$$\|\mathbf{h}_{\mathbf{A}(k)} - \bar{\mathbf{h}}\|^2 \leq (1 - \lambda^2/D^2)^k \|\mathbf{h}_{\mathbf{A}(0)} - \bar{\mathbf{h}}\|^2.$$

And thus we have

$$\mathbb{D}[f_{\mathbf{A}(k)}, F] \leq c_1^2 (1 - \lambda^2/D^2)^k \|\mathbf{h}_{\mathbf{A}(0)} - \bar{\mathbf{h}}\|^2 \leq c_1^2 (1 - \lambda^2/D^2)^{\|\mathbf{A}(k)\|_0} \|\mathbf{h}_{\mathbf{A}(0)} - \bar{\mathbf{h}}\|^2,$$

where the last inequality is by  $\|\mathbf{A}(k)\|_0 \leq k$ .

### Proof of Theorem 3

Notice that if we have  $\frac{\langle \bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k)}, \mathbf{h}_{i_k^*} - \mathbf{h}_{\mathbf{A}(k)} \rangle}{\|\mathbf{h}_{i_k^*} - \mathbf{h}_{\mathbf{A}(k)}\|^2} \geq 1$ , then  $\gamma_k^* = 1$  and in this case,

$$\|\mathbf{h}_{\mathbf{A}(k+1)} - \bar{\mathbf{h}}\| = \bar{\mathbb{D}}[\sigma(\boldsymbol{\theta}_{i_k^*}, \cdot), F(\cdot)] \geq \bar{\mathbb{D}}[\sigma(\boldsymbol{\theta}_{i_0^*}, \cdot), F(\cdot)].$$

On the other hand, since  $0 < \|\bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k)}\|$ , by the argument in proving Theorem 1, we have

$$\|\bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k+1)}\| \leq \sqrt{1 - \lambda^2/D^2} \|\bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k)}\| < \|\bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k)}\|.$$

This gives that

$$\|\bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k+1)}\| < \|\bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k)}\| \leq \bar{\mathbb{D}}[\sigma(\boldsymbol{\theta}_{i_0^*}, \cdot), F(\cdot)],$$

which makes contradiction.

### Proof of Theorem 4

Using Lemma 3, we know that  $\bar{\mathbf{h}} \in \text{ri}\mathcal{M}$ , which indicate that there exists some  $\lambda > 0$  such that

$$\mathcal{B}(\bar{\mathbf{h}}, \lambda) \cap \text{Aff}\mathcal{M} \subseteq \mathcal{M},$$

where  $\mathcal{B}(\bar{\mathbf{h}}, \lambda)$  denotes the ball with radius  $\lambda$  centered at  $\bar{\mathbf{h}}$ . Following the same argument of Ye et al. (2020) in proving theorem 2, we have  $\|\bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k)}\|^2 = \mathcal{O}((k+1)^{-2})$ . The result that  $\|\mathbf{A}(k)\|_0 \leq k+1$  is obvious as in each iteration, the number of nonzero elements in  $\mathbf{A}$  at most increases by 1.

### Proof of Theorem 5

Suppose that at iteration  $k$ , we have  $\mathbf{h}_{\mathbf{A}(k)}$ . And the global imitation algorithm returns  $\mathbf{h}_{\mathbf{A}(k+1)}$  with  $f_{\ell, \mathbf{A}(k+1)} = H_2 \circ [(1 - \gamma_k) f_{\ell, \mathbf{A}(k)} + \sigma(\boldsymbol{\theta}_{i_k^*}, \cdot)] \circ H_1$ . We also define  $i'_k$  as the solution of local imitation. And we let  $f_{\mathbf{A}'(k+1)} = H_2 \circ [(1 - \gamma_k) f_{\ell, \mathbf{A}(k)} + \sigma(\boldsymbol{\theta}_{i'_k}, \cdot)] \circ H_1$ . Define  $\mathbf{w}_{k+1} = (k+1)(\bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k)})$ ,  $\mathbf{w}'_{k+1} = (k+1)(\bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}'(k)})$ ,  $\mathbf{W}_{k+1} = (k+1)(H \circ \bar{\mathbf{h}} - H \circ \mathbf{h}_{\mathbf{A}(k)})$  and  $\mathbf{W}'_{k+1} = (k+1)(H \circ \bar{\mathbf{h}} - H \circ \mathbf{h}_{\mathbf{A}'(k)})$ . We have

$$\begin{aligned} \|\mathbf{W}_{k+1}\|^2 &\leq \|\mathbf{W}'_{k+1}\|^2 \\ &= (k+1)^2 \|H \circ \bar{\mathbf{h}} - H \circ \mathbf{h}_{\mathbf{A}'(k)}\|^2 \\ &\stackrel{(1)}{\leq} \kappa_1^2 (k+1)^2 \|\bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}'(k)}\|^2 \\ &= \kappa_1^2 \|\mathbf{w}'_{k+1}\|^2 \\ &\stackrel{(2)}{\leq} \kappa_1^2 (\|\mathbf{w}_k\|^2 - 2\lambda \|\mathbf{w}_k\| + D^2) \\ &\stackrel{(3)}{\leq} \kappa_1^2 (\kappa_2^2 \|\mathbf{W}_k\|^2 - \kappa_2 2\lambda \|\mathbf{W}_k\| + D^2) \\ &= \kappa_1^2 \kappa_2^2 \|\mathbf{W}_k\|^2 - 2\kappa_1^2 \kappa_2 \lambda \|\mathbf{W}_k\| + \kappa_1^2 D^2. \end{aligned}$$

Here  $D$  is the quantities defined in Lemma 5, (1) and (3) use the definition of  $\kappa_1$  and  $\kappa_2$  and (2) is by the argument of Ye et al. (2020) in proving Theorem 2 (notice that their argument also applies to the case that  $\bar{\mathbf{h}}$  is in the relative interior of  $\mathcal{M}$ , which is proved by Lemma 3, instead

of that  $\bar{\mathbf{h}}$  is in the interior of  $\mathcal{M}$ ). By the assumption that  $D^2 \geq \kappa_1^2 \kappa_2^2 (D^2 - \lambda^2)$ , the formula  $\kappa_1^2 \kappa_2^2 x^2 - 2\kappa_1^2 \kappa_2 \lambda x + \kappa_1^2 D^2 = x^2$  has two real root, denoted by  $z_1 \leq z_2$ , where

$$z_1 = \frac{\kappa_1^2 \kappa_2 \lambda - \kappa_1 \sqrt{\kappa_1^2 \kappa_2^2 (\lambda^2 - D^2) + D^2}}{(\kappa_1^2 \kappa_2^2 - 1)}$$

$$z_2 = \frac{\kappa_1^2 \kappa_2 \lambda + \kappa_1 \sqrt{\kappa_1^2 \kappa_2^2 (\lambda^2 - D^2) + D^2}}{(\kappa_1^2 \kappa_2^2 - 1)}.$$

We define  $q_1 = \kappa_1^2 \kappa_2^2$  and  $q_2 = \kappa_1^2 \kappa_2$ , and we have

$$\|\mathbf{W}_{k+1}\|^2 \leq q_1 \|\mathbf{W}_k\|^2 - 2q_2 \lambda \|\mathbf{W}_k\| + \kappa_1^2 D^2.$$

If  $q_1 = 1$ , then the rate holds by directly applying the argument of [Ye et al. \(2020\)](#) in proving Theorem 2. If  $q_1 > 1$ , we know that  $2q_2 \lambda \geq 0$  and  $\kappa_1^2 D^2 \geq 0$ ;  $\|\mathbf{W}_k\| \geq 0$  for any  $k$  by its definition; the formula  $q_1 x^2 - 2q_2 \lambda x + \kappa_1^2 D^2 = x^2$  has two real roots  $z_1 \leq z_2$ ;  $z_2 \geq \kappa_1 D$  by the assumption; and  $\|\mathbf{W}_{k+1}\| \leq z_2$  by the assumption. Using Lemma 7, we have, for any  $k$ ,

$$\|\mathbf{W}_k\| \leq z_2,$$

which implies that

$$\|H \circ \bar{\mathbf{h}} - H \circ \mathbf{h}_{\mathbf{A}(k)}\|^2 = \mathcal{O}((k+1)^{-2}),$$

and thus  $\mathbb{D}[f_{\mathbf{A}(k)}, F] = \mathcal{O}(k^{-2})$ . The result that  $\|\mathbf{A}(k)\|_0 \leq k+1$  is obvious as in each iteration,  $\|\mathbf{A}(k)\|_0$  at most increase 1.

## Proof of Theorem 6

When pruning the  $\ell$ -th layer, if this layer is pruned by local imitation, by applying Theorem 1 on the  $\ell$ -th layer of  $f_{[\ell-1]}$ , we have

$$\sqrt{\mathbb{D}[f_{[\ell]}, f_{[\ell-1]}]} = \mathcal{O}\left(\exp\left(-\frac{\lambda_\ell}{2} \|\mathbf{A}_\ell\|_0\right)\right),$$

for some  $\lambda_\ell > 0$ . Else if this layer is pruned by global imitation, we have

$$\begin{aligned} \sqrt{\mathbb{D}[f_{[\ell]}, f_{[\ell-1]}]} &\leq \sqrt{\mathbb{D}[F_L \circ \dots \circ F_{\ell+1} \circ f_{\ell, \mathbf{A}_\ell^{\text{local}}} \circ \dots \circ f_{1, \mathbf{A}_1}, f_{[\ell-1]}]} \\ &= \mathcal{O}\left(\exp\left(-\frac{\lambda_\ell}{2} \|\mathbf{A}_\ell^{\text{local}}\|_0\right)\right) = \mathcal{O}\left(\exp\left(-\frac{\lambda_\ell}{2} \|\mathbf{A}_\ell\|_0\right)\right). \end{aligned}$$

Using triangle inequality, we know that

$$\sqrt{\mathbb{D}[f_{[L]}, F]} \leq \sum_{\ell=1}^L \sqrt{\mathbb{D}[f_{[\ell]}, f_{[\ell-1]}]} = \mathcal{O}\left(\sum_{\ell=1}^L \exp\left(-\frac{\lambda_\ell}{2} \|\mathbf{A}_\ell\|_0\right)\right),$$

with  $\lambda_\ell > 0$  for all  $\ell \in [L]$ .

## Proof of Technical Lemmas

### Proof of Lemma 3

The case that  $n = 1$  is trivial and we consider the case that  $n \geq 2$ . By the definition, we know that  $M$  is a non-empty and closed convex set. And thus by Lemma 2,  $\text{ri}M$  is not empty. Define

$$\tilde{\mathbf{q}} \in \text{ri}M, \tilde{\mathbf{q}} = \sum_{i=1}^n \alpha_i \mathbf{q}_i, \quad \sum_{i=1}^n \alpha_i = 1 \text{ and } \alpha_i \geq 0 \ \forall i \in [n].$$

We define  $\alpha_{\max} = \max_{i \in [n]} \alpha_i$ . Notice that  $\alpha_{\max} \geq 1/n$ , otherwise, if  $\alpha_{\max} < 1/n$ , we have  $\sum_{i=1}^n \alpha_i \leq n\alpha_{\max} < 1$ , which makes contradiction. If  $\alpha_{\max} = 1/n$ , then  $\alpha_i = 1/n$  for all  $i \in [n]$ , otherwise,  $\sum_{i=1}^n \alpha_i < 1$ , which makes contradiction. In the case that  $\alpha_{\max} = 1/n$ , we have already obtained the desired result.

Now we assume  $\alpha_{\max} > \frac{1}{n}$ . Define  $\lambda = 1 - \frac{1}{n\alpha_{\max}} \in [0, 1)$  and  $\beta_i = \frac{\alpha_{\max} - \alpha_i}{n\alpha_{\max} - 1}$ . Notice this gives that

$$\sum_{i=1}^n \beta_i = \sum_{i=1}^n \frac{\alpha_{\max} - \alpha_i}{n\alpha_{\max} - 1} = \frac{n\alpha_{\max} - \sum_{i=1}^n \alpha_i}{n\alpha_{\max} - 1} = 1 \quad \text{and } \beta_i \geq 0 \quad \forall i \in [n].$$

We define  $\mathbf{q}' = \sum_{i=1}^n \beta_i \mathbf{q}_i$  and by the property of  $\beta_i$  and the definition of  $M$ , we have  $\mathbf{q}' \in M = \text{cl}M$ . Notice that

$$\bar{\mathbf{q}} = \frac{1}{n} \sum_{i=1}^n \mathbf{q}_i = (1 - \lambda) \sum_{i=1}^n \alpha_i \mathbf{q}_i + \lambda \sum_{i=1}^n \beta_i.$$

Using Lemma 1, we know that  $\bar{\mathbf{q}} \in \text{ri}M$ .

## 5.8 Proof of Lemma 4

Notice that by choosing  $s' = \bar{\mathbf{h}} - \lambda \frac{\bar{\mathbf{h}} - \mathbf{h}_A}{\|\bar{\mathbf{h}} - \mathbf{h}_A\|} \in \mathcal{M}$ , we have

$$\max_{s \in \mathcal{M}} \langle \bar{\mathbf{h}} - \mathbf{h}_A, \bar{\mathbf{h}} - s \rangle \geq \langle \bar{\mathbf{h}} - \mathbf{h}_A, \bar{\mathbf{h}} - s' \rangle = \lambda \|\bar{\mathbf{h}} - \mathbf{h}_A\|.$$

## 5.9 Proof of Lemma 5

Notice that for any  $i \in [N]$ ,

$$\|\mathbf{h}_i\| = \sqrt{\sum_{j=1}^m \sigma^2(\boldsymbol{\theta}_i, \mathbf{z}^{(j)})} \leq \sqrt{mc_1}.$$

And for any  $\mathbf{h} \in \mathcal{M}$ , we have  $\mathbf{h} = \sum_{i=1}^N \beta_i \mathbf{h}_i$ , for some  $\beta_i \geq 0$  and  $\sum_{i=1}^N \beta_i = 1$ , which gives that

$$\|\mathbf{h}\| = \left\| \sum_{i=1}^N \beta_i \mathbf{h}_i \right\| \leq \sum_{i=1}^N \beta_i \|\mathbf{h}_i\| \leq \sqrt{mc_1}.$$

## Proof of Lemma 6

Proof of this Lemma follows standard argument in analyzing Frank Wolfe algorithm. We include it for the completeness. Notice that

$$\tilde{\mathbf{s}}_k^* = \arg \min_{s \in \mathcal{M}} \langle \bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k)}, s - \mathbf{h}_{\mathbf{A}(k)} \rangle = \arg \min_{s \in \mathcal{M}} \langle \bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k)}, s - \bar{\mathbf{h}} \rangle = -\arg \max_{s \in \mathcal{M}} \langle \bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k)}, \bar{\mathbf{h}} - s \rangle.$$

Using Lemma 4, we know that  $\langle \bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k)}, \bar{\mathbf{h}} - \tilde{\mathbf{s}}_k^* \rangle \leq -\lambda \|\bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k)}\|$ . Notice that

$$\begin{aligned} & \langle \bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k)}, \tilde{\mathbf{s}}_k^* - \mathbf{h}_{\mathbf{A}(k)} \rangle \\ &= \langle \bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k)}, \bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k)} + \tilde{\mathbf{s}}_k^* - \bar{\mathbf{h}} \rangle \\ &= -\langle \bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k)}, \bar{\mathbf{h}} - \tilde{\mathbf{s}}_k^* \rangle + \|\bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k)}\|^2 \\ &\leq -2 \langle \bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k)}, \bar{\mathbf{h}} - \tilde{\mathbf{s}}_k^* \rangle + \|\bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k)}\|^2 + \|\bar{\mathbf{h}} - \tilde{\mathbf{s}}_k^*\|^2 \\ &= \|\bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k)} - (\bar{\mathbf{h}} - \tilde{\mathbf{s}}_k^*)\|^2 = \|\mathbf{h}_{\mathbf{A}(k)} - \tilde{\mathbf{s}}_k^*\|^2, \end{aligned}$$

where the last inequality uses the fact that  $\langle \bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k)}, \bar{\mathbf{h}} - \tilde{\mathbf{s}}_k^* \rangle \leq 0$ . This gives that  $0 \leq \langle \bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k)}, \tilde{\mathbf{s}}_k^* - \mathbf{h}_{\mathbf{A}(k)} \rangle \leq \|\mathbf{h}_{\mathbf{A}(k)} - \tilde{\mathbf{s}}_k^*\|^2$ . And thus we have

$$\begin{aligned} & \min_{\gamma \in [0, 1]} \|\mathbf{h}_{\mathbf{A}(k)} - \bar{\mathbf{h}}\|^2 - 2\gamma \langle \bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k)}, \tilde{\mathbf{s}}_k^* - \mathbf{h}_{\mathbf{A}(k)} \rangle + \gamma^2 \|\tilde{\mathbf{s}}_k^* - \mathbf{h}_{\mathbf{A}(k)}\|^2 \\ &= \|\mathbf{h}_{\mathbf{A}(k)} - \bar{\mathbf{h}}\|^2 - \frac{\langle \bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k)}, \tilde{\mathbf{s}}_k^* - \mathbf{h}_{\mathbf{A}(k)} \rangle^2}{\|\mathbf{h}_{\mathbf{A}(k)} - \tilde{\mathbf{s}}_k^*\|^2} \\ &\leq \|\mathbf{h}_{\mathbf{A}(k)} - \bar{\mathbf{h}}\|^2 - \lambda^2 \frac{\|\bar{\mathbf{h}} - \mathbf{h}_{\mathbf{A}(k)}\|^2}{\|\mathbf{h}_{\mathbf{A}(k)} - \tilde{\mathbf{s}}_k^*\|^2} \\ &\leq (1 - \lambda^2/D^2) \|\mathbf{h}_{\mathbf{A}(k)} - \bar{\mathbf{h}}\|^2, \end{aligned}$$

where the last inequality is by Lemma 5.

### Proof of Lemma 7

Define  $f(x) = ax^2 - bx + c$ . By assumption (1) and assumption (3), for any  $z \in [z_1, z_2]$ ,  $f(z) - z^2 \leq 0$ . We proof the desired result by induction. Suppose that  $x_k \in [0, z_2]$ . Case 1:  $x_k \in [z_1, z_2]$  and in this case,

$$x_{k+1}^2 \leq f(x_k) \leq x_k^2 \leq z_2^2.$$

Case 2:  $x_k \in [0, z_1]$  and in this case

$$x_{k+1}^2 \leq f(x_1) \leq \max_{z \in [0, z_1]} f(z) \leq \max(f(0), f(z_1)) = \max(c, z_1^2).$$

This gives that  $x_{k+1} \leq \max(\sqrt{c}, z_1) \leq \max(z_2, z_1) = z_2$ . The desired result follows by induction.