

Practical Network Acceleration with Tiny Sets: Hypothesis, Theory, and Algorithm

Guo-Hua Wang, and Jianxin Wu, *Member, IEEE*

Abstract—Due to data privacy issues, accelerating networks with tiny training sets has become a critical need in practice. Previous methods achieved promising results empirically by filter-level pruning. In this paper, we both study this problem theoretically and propose an effective algorithm aligning well with our theoretical results. First, we propose the *finetune convexity hypothesis* to explain why recent few-shot compression algorithms do not suffer from overfitting problems. Based on it, a theory is further established to explain these methods for the first time. Compared to naively finetuning a pruned network, feature mimicking is proved to achieve a lower variance of parameters and hence enjoys easier optimization. With our theoretical conclusions, we claim dropping blocks is a fundamentally superior few-shot compression scheme in terms of more convex optimization and a higher acceleration ratio. To choose which blocks to drop, we propose a new metric, *recoverability*, to effectively measure the difficulty of recovering the compressed network. Finally, we propose an algorithm named PRACTISE to accelerate networks using only tiny training sets. PRACTISE outperforms previous methods by a significant margin. For 22% latency reduction, it surpasses previous methods by on average 7 percentage points on ImageNet-1k. It also works well under data-free or out-of-domain data settings. Our code is at <https://github.com/DoctorKey/Practise>.

Index Terms—Deep Neural Network, Network Acceleration, Few-Shot Compression, Knowledge Distillation, Image Classification.

1 INTRODUCTION

DEEP networks have achieved remarkable success, but suffer from high computational costs. To accelerate them, many methods have been proposed, such as network pruning [1], [2], [3], [4], decoupling [5], [6] and quantization [7], [8]. However, most previous methods rely on *the original training set* (i.e., all the training data) to recover the compressed model’s accuracy. But, to preserve data privacy and/or to achieve fast deployment, only scarce training data may be available. For example, a customer often asks an algorithmic provider to accelerate their networks, but the whole training data *cannot* be made available due to privacy concerns. Only the uncompressed model and a few training examples are provided. In some extreme cases, not even a single data point is accessible. Algorithmic engineers need to synthesize images or collect some out-of-domain training images by themselves. Hence, to accelerate a deep network with only very few samples is emerging as a critical problem.

Previous work [9], [10], [11] have achieved promising results. However, most of them address few-shot compression by empirically designing algorithms but leave many fundamental questions unanswered, such as

- Why is it possible to finetune a pruned model with very few samples without suffering from overfitting?
- Why must knowledge distillation (KD) especially feature-based mimicking be adopted in finetuning? What benefits does KD or feature mimicking enjoy?
- Which compression scheme is best in few-shot compression? What properties should one compression

scheme possess to achieve better performance?

This paper aims at answering all these key questions from both theoretical and empirical perspectives.

Training a deep network from scratch is highly non-convex and leads to a complex optimization landscape [12], [13], [14]. But we find that finetuning a pruned model is surprisingly convex. Based on lots of experiments, we propose a *finetune convexity hypothesis*, which surmises that the loss landscape of finetuning a pruned model is convex w.r.t. the parameters. Furthermore, the degree of convexity decides the amount of training data required during finetuning. The convexity property allows few-shot finetuning of the pruned model without serious overfitting. To better understand why this convexity occurs, we reformulate the model compression process to insert linear adaptor modules around pruned structures and *only* finetune these modules with few examples. Experiments show that its optimization landscape and accuracy are almost the same as those of end-to-end finetuning. That is, *the difference between the original model and the pruned model (before finetuning) is convex or even close to linear in the latent representation/feature space*.

Although the convexity property reduces few-shot finetuning’s difficulty, there are still large accuracy gaps between different few-shot compression methods. Compared to vanilla finetuning, existing work find that knowledge distillation is effective in practice, especially feature mimicking (let the student mimic the penultimate feature of the teacher [11]). However, this crucial empirical finding lacks a theoretical foundation. On top of our finetune convexity hypothesis, we establish a theory for feature mimicking, which studies the variance of learned parameters. We prove that compared to vanilla finetuning, parameters will achieve lower variance via feature mimicking. Hence, feature mimicking makes the final model more stable. To the best of our knowledge, it

• G.-H. Wang and J. Wu are with the State Key Laboratory for Novel Software Technology, Nanjing University, Nanjing 210023, China. J. Wu is the corresponding author. This paper was partly supported by the National Natural Science Foundation of China under Grant 62276123 and Grant 61921006.
E-mail: wangguohua@lamda.nju.edu.cn, wujx2001@gmail.com




Figure 1. Comparison of different structure compression schemes with 500 training images. Our scheme (‘Block’) outperforms previous schemes dominantly for the latency-accuracy tradeoff. The ResNet-34 model was compressed on ImageNet-1k and tested on an NVIDIA TITAN Xp GPU.

Table 1

Top-1 validation accuracy (%) on ImageNet-1k for different compression schemes. ResNet-34 was accelerated by reducing 16% latency with 50 training images. Previous methods prune filters with the ‘normal’ style. Our block-level scheme, ‘BP (blocks)’, simply removes the first few blocks and finetune the pruned network by back propagation.

KD [15]	FSKD [9]	CD [10]	MiR [11]	BP (blocks)
44.5	45.3	56.2	64.1	66.5

is the *first theory to successfully show why feature mimicking outperforms vanilla finetuning*.

The next crucial problem is how to prune the network structure. Almost all previous few-shot compression work [9], [10], [11] adopted filter-level pruning. In this paper, we are the first to reveal that block-level pruning is fundamentally superior in few-shot compression. First, pruning blocks enjoys a *more convex optimization landscape*, which requires fewer data compared to pruning filters. Second, it *achieves much faster speed on real-world computing devices* than filter pruning when the FLOPs are the same. Fig. 1 shows that dropping blocks dominantly outperforms previous compression schemes for the latency vs. accuracy tradeoff. Table 1 further reports that *an embarrassingly simple dropping block baseline (i.e., finetune without any other processing) already surpasses existing methods which use complicated techniques*. The baseline, ‘BP (blocks)’, simply removes the first few blocks and finetunes the pruned network with only the cross-entropy loss.

To achieve better block pruning, we study the strategy to choose which blocks to drop with only scarce training samples. Several criteria [16], [17], [18] have been proposed for pruning blocks, but some [16], [17] require a large amount of data and others [18] only evaluate the output difference before/after block removal. We notice that *although dropping some blocks significantly changes the feature maps, they are easily recovered by a finetuning even with a tiny training set*, because of the finetune convexity hypothesis. Simply measuring the difference between pruned/original networks is at best suboptimal. We need to not only consider the loss landscape’s properties to obtain higher accuracy, but also real accelerating ratios. To take into account factors of the loss

landscape, a new concept *recoverability* is proposed in this paper, which is computed efficiently with only a few training images. Our recoverability metric is surprisingly consistent with the accuracy of the finetuned network. Combined with the acceleration ratio, our final pruning score outperforms previous ones significantly.

Finally, we propose PRACTISE, namely Practical network acceleration with tiny sets of images, to effectively accelerate a network with scarce data. PRACTISE significantly outperforms previous few-shot pruning methods. For 22.1% latency reduction, PRACTISE surpasses the previous state-of-the-art (SOTA) method on average by 7.0% (percentage points, not relative improvement) Top-1 accuracy on ImageNet-1k. It enjoys high generalization ability, which can be used on synthesized/out-of-domain images. Our contributions are:

- **Hypothesis.** We propose the finetune convexity hypothesis which unveils the mystery why *few-shot* finetuning a pruned model does not suffer from the overfitting problem. We point out that convexity is tightly correlated with the effectiveness of finetuning.
- **Theory.** We establish a theory for feature mimicking, proving that learnable parameters achieve lower variance via feature mimicking compared to vanilla finetuning. It aligns very well with empirical results.
- **Algorithm.** Although dropping blocks is previously known, we are *the first to reveal its great potential in few-shot compression*, which has higher convexity than filter pruning. By considering properties of the optimization landscape (such as the convexity and initial accuracy), a new concept *recoverability* is proposed to measure the difficulty of recovering each block, and in determining the priority to drop blocks. Then, we propose PRACTISE, an algorithm for accelerating networks with tiny sets of images. With 22.1% latency reduction, PRACTISE outperforms previous SOTAs by 7.0% Top-1 accuracy on ImageNet.

A preliminary version of the PRACTISE algorithm has appeared as a conference publication [19].

2 RELATED WORK

Filter-level pruning accelerates networks by removing filters in convolutional layers. Different criteria for choosing filters have been proposed [1], [2], [3], [18], [20], [21], along with different training strategies [22], [23], [24], [25]. However, most of these methods rely on the whole training data to train the network. When facing data privacy issues, these filter-level pruning methods suffer from poor latency-accuracy performance with tiny training sets [9], [10]. In this paper, we argue that the drawbacks of filter pruning are its flat optimization landscape and low acceleration ratio. It requires pruning lots of parameters and FLOPs to reduce latency. That results in a large capacity gap between the pruned and the original networks, and it is challenging to recover the pruned network’s accuracy on only a tiny training set. Instead, we advocate block-level pruning for few-shot compression. Dropping blocks enjoys both a more convex optimization and higher acceleration ratio.

Block-level pruning removes several complete blocks (e.g., residual blocks) in a network. Some work have been




Figure 2. The pipeline for traditional few-shot compression.

proposed for the whole training data case. DBP [16] proposes using linear probing to evaluate the accuracy of each block’s features, and dropping blocks with low accuracy. ϵ -ResNet adds a sparsity-promoting function to discard the block if all responses of this block are less than a threshold ϵ . Both DBP and ϵ -ResNet require a large dataset for training and testing. CURL [18] uses a proxy dataset to evaluate the KL-divergence change before/after block removal. However, it neglects the finetuning process. In fact, the accuracy of the pruned network *after finetuning* is vital, but there are no existing criteria to approximate it well. In this paper, we propose a new *recoverability* to evaluate if the network pruned by dropping blocks can recover the accuracy well.

Few-shot or data-free knowledge distillation aims at training a student network by a pretrained teacher with limited or even no original training data. Few-shot KD (FSKD) [9] inserts 1×1 conv. after the pruned conv. layer and trains each layer by making the pruned network’s feature maps mimic the original network’s. Layer-wise training can obtain more supervision from the teacher, but results in error accumulation. CD [10] proposes cross distillation to reduce the accumulated errors. MiR [11] proposes a mimicking then replacing framework to optimize the pruned network holistically. When not even a single original training sample is available, data free knowledge distillation (DFKD) [26] synthesizes alternatives of the original training data. Subsequent DFKD methods accelerate the synthesis process [27] and enhance the performance by multi-teacher [28]. Pruning and quantization [29], [30], [31] are two main applications of DFKD. However, most FSKD and DFKD methods adopt filter-level pruning and result in poor latency-accuracy performance. Our PRACTISE works robustly on synthesized/out-of-domain images and improves the accuracy in the data-free scenario by a large margin.

3 FINETUNE CONVEXITY HYPOTHESIS

For few-shot compression, several effective algorithms [9], [10], [11] have been proposed. However, they are designed empirically and lack theoretical analysis. The key question is: why is it possible to finetune the network with only a few training samples without serious overfitting? In this section, we study the loss landscapes w.r.t. the parameters. Surprisingly, we find that finetuning with few samples appears to behave like a convex problem, which alleviates the overfitting issue. We observed that for different compression schemes, the convexity of their optimization landscapes is different. Pruning blocks enjoys a higher convexity, hence it requires less training data than filter pruning.

We formulate traditional few-shot compression into three steps as illustrated in Fig. 2. First, some parameters are pruned by setting them to zero. For dropping a block, we prune all parameters inside the residual connection. For pruning filters, the pruned parameters are set to zero. Second, the network architecture is pruned guided by these zero parameters, while the network functionality is not changed. Third, the compressed network is finetuned on data D .

Following [12], [13], [14], we measure the convexity via interpolating two models. Considering the first step $(\mathcal{A}; \theta) \rightarrow (\mathcal{A}; \theta')$ in Fig. 2, we simply set some parameters to zero while keeping the architecture unchanged. Hence, it is able to merge the raw/pruned models by linear interpolation. Given $\lambda \in [0, 1]$, the interpolated model is

$$f_\lambda = \lambda f' + (1 - \lambda)f \quad (1)$$

$$\iff (\mathcal{A}; \theta_\lambda) = (\mathcal{A}; \lambda\theta' + (1 - \lambda)\theta). \quad (2)$$

We take ResNet-34 as the raw model (f), and prune it by dropping filters ($f^{\setminus F}$) and blocks ($f^{\setminus B}$), respectively. Fig. 3a shows the test loss of interpolated models. The loss curve roughly describes the property of the loss landscape. Surprisingly, these curves are smooth and monotonic. More importantly, they are roughly convex. Pruning blocks achieves a more convex curve compared to pruning filters. Fig. 3c presents this observation qualitatively. Compare to the path $f \rightarrow f^{\setminus F}$, the path $f \rightarrow f^{\setminus B}$ is more convex. With these observations, we propose the following conjecture.

Conjecture 1. By linearly interpolating between the raw and the pruned model, the convexity of the loss landscape indicates the difficulty of finetuning, which reflects the amount of training data required, the speed of convergence, etc.

To check this conjecture, we finetune both $f^{\setminus B}$ and $f^{\setminus F}$ on a training set D , and obtain the finetuned models $f_D^{\setminus B}$ and $f_D^{\setminus F}$, respectively. Then we evaluate interpolations between pruned and finetuned models. Fig. 3b shows results. First, with a more convex loss landscape, $f_D^{\setminus B}$ achieves a lower loss than $f_D^{\setminus F}$ in all cases. Intuitively, a more convex loss landscape means it is easier for SGD to optimize models rapidly. Second, for $\lambda \in [0.9, 1]$, the dash lines overlap on the solid lines, that is because the interpolation equation $\theta_\lambda = \lambda\theta^{\setminus B} + (1 - \lambda)\theta$ can be considered as the SGD updating formula, where θ is the optimal gradient direction and $1 - \lambda$ is the small learning rate. This further demonstrates the validity of Conj. 1 based on the close relationship between linear interpolation and SGD finetuning.

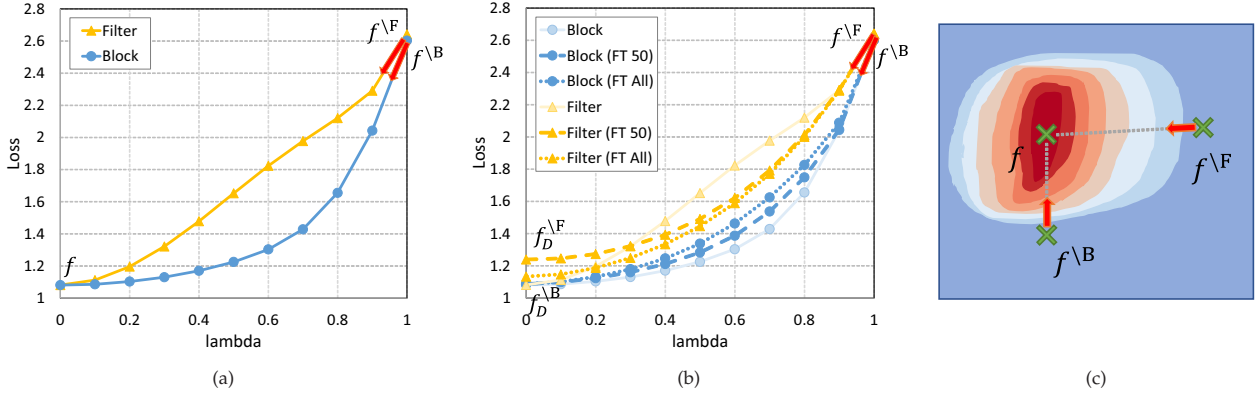


Figure 3. (a) presents the loss of interpolated models between the raw model f ($\lambda = 0$) and the pruned model f' ($\lambda = 1$). Then we finetune the pruned model f' on datasets D containing 50 (“FT 50”) or all (“FT All”) training images, and obtain the finetuned models f'_D . (b) presents the loss of interpolated models between these two models. (c) illustrates the loss landscape qualitatively. Finetuning a pruned model enjoys a convex loss landscape w.r.t. model parameters. Compared to pruning filters, the loss landscape of pruning blocks is more convex.

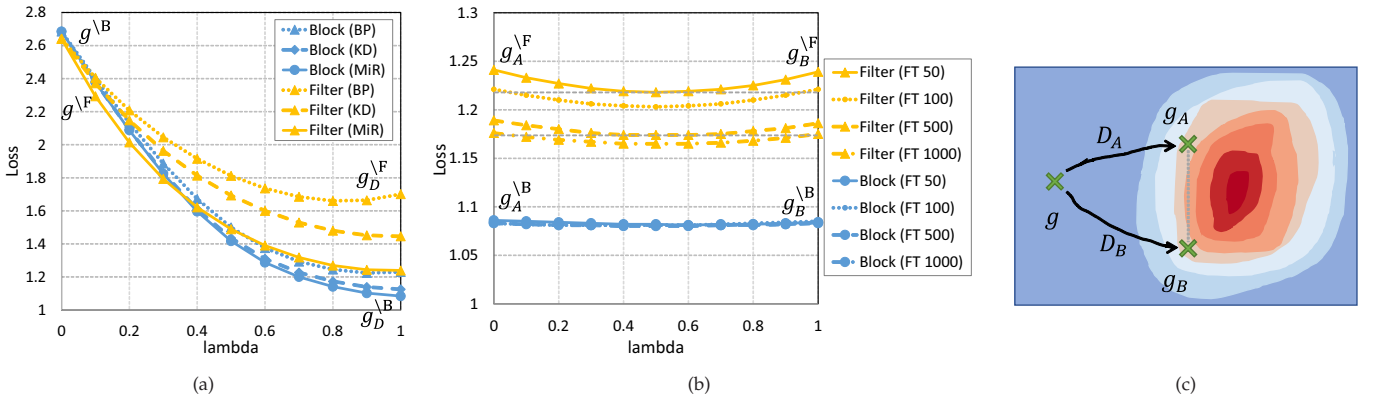


Figure 4. (a) presents the loss of interpolated models with different finetuning methods. All models are finetuned by 50 samples. For (b), we finetune models on two datasets A and B with the same amount, respectively. Then the loss of interpolations is evaluated. (c) illustrates the loss landscape of finetuning pruned models.

Then we consider the remaining steps in few-shot compression, where the architecture will be pruned. Assume f^B and f^F become g^B and g^F by pruning architecture, respectively. We finetune these pruned models on a tiny data D , and obtain g_D^B and g_D^F . Fig. 4a shows the results of linear interpolation. All curves of different finetuning methods are convex, and pruning blocks outperforms pruning filters for lower loss significantly. Based on both Fig. 3b and 4a, we propose the following conjecture.

Conjecture 2. [Finetune convexity hypothesis (weak)]: The loss landscape of finetuning a pruned model is convex w.r.t. the parameters. That means

$$\mathcal{L}(g'_{\Delta D}) \leq \lambda \mathcal{L}(g) + (1 - \lambda) \mathcal{L}(g'_D), \quad (3)$$

where $g'_{\Delta D}$ is an interpolated model:

$$g'_{\Delta D} = (\mathcal{A}'; \theta' + (1 - \lambda)\Delta\theta_D), \quad (4)$$

$$g = (\mathcal{A}'; \theta'), \quad (5)$$

$$g'_D = (\mathcal{A}'; \theta' + \Delta\theta_D). \quad (6)$$

This conjecture answers the question of why it is able to finetune the whole network with only a few training samples. With the convexity property, optimization hardly suffers from overfitting even if only scarce training data is available. Note




Figure 5. (a) shows the testing loss curve of pruned model g^B on the whole ImageNet-1k training set. (b) presents the loss of interpolated models between the 5th epoch model ($\lambda = 0$) and the last epoch model ($\lambda = 1$), i.e., those indicated by green arrows in (a).

that this conclusion does not hold for training from scratch with the whole training data. As introduced in previous work [12], [13], [14], a loss barrier exists on linear interpolation between the raw and finetuned models. We also evaluate it by training a pruned network from scratch. Fig. 5 shows results. Training from scratch suffers from a complex optimization landscape. Linearly interpolating cannot achieve a convex loss curve.

Conj. 2 only considers the connection between a model g and its finetuned version g' . We further finetune the pruned model g on different data sets A and B to study

the relationship between different finetuned versions. Fig. 4b presents results of linear interpolations between g'_A and g'_B . With different tiny sets, the finetuned models are in the same basin, and there is no loss barrier on interpolation. Fig. 4c presents this observation qualitatively. Hence, we have the following strong conjecture:

Conjecture 3. [Finetune convexity hypothesis (strong)]:

With few-shot training datasets A and B , the finetuning solutions lie in the same optimization basin and the functional \mathcal{L} is convex w.r.t. the parameters θ . That means

$$\mathcal{L}(g'_{A \leftrightarrow B}) \leq \lambda \mathcal{L}(g'_A) + (1 - \lambda) \mathcal{L}(g'_B), \quad (7)$$

where

$$g'_{A \leftrightarrow B} = (\mathcal{A}; \theta' + \lambda \Delta \theta_A + (1 - \lambda) \Delta \theta_B), \quad (8)$$

$$g'_A = (\mathcal{A}; \theta' + \Delta \theta_A), \quad (9)$$

$$g'_B = (\mathcal{A}; \theta' + \Delta \theta_B). \quad (10)$$

We define the loss leakage as

$$\mathcal{L}_{leak}(g'_{A \leftrightarrow B}) = \max_{\lambda} \lambda \mathcal{L}(g'_A) + (1 - \lambda) \mathcal{L}(g'_B) - \mathcal{L}(g'_{\lambda}). \quad (11)$$

If the training data is increased from A to $A \cup B$, the loss reduction is proportional to the loss leakage. That is

$$\mathcal{L}(g'_A) - \mathcal{L}(g'_{A \cup B}) \propto \mathcal{L}_{leak}(g'_{A \leftrightarrow B}). \quad (12)$$

Our loss leakage indicates the pruned network's thirst for training data. With a large loss leakage, the pruned network requires more data to finetune. As illustrated in Fig. 4b, compared to pruning blocks, pruning filters needs more data to reduce the loss. In particular, for A and B containing 50 images each, the linear interpolation model ($\lambda = 0.5$) achieves a similar loss with the model finetuned on 100 images. In other words, averaging models finetuned on each set separately is on-par with finetuning the model on the union set. This observation supports our Conj. 3 and further supports our Conj. 2. Note that Conj. 3 can be considered as an extension of Conj. 2 by setting an empty dataset A .

In summary, we have unraveled the mysteries proposed at the beginning of this section. Thanks to the convexity of the optimization landscape, finetuning with few training samples is possible without serious overfitting. The loss landscape of pruning blocks is more convex compared to pruning filters. Hence, it requires less training data and enjoys a faster convergence. Next, we will propose a theory for different finetuning methods based on our convexity hypothesis.

4 THEORY ON FINETUNING

As presented in Fig. 4a, different finetuning methods lead to different final accuracies in practice. The naive method ("BP") simply uses the cross-entropy loss, which results in inferior accuracy as previous work pointed out [9], [10]. Knowledge distillation [15] (KD), especially feature-level distillation [32], [33], can improve accuracy. MiR [11] proposes to use the features before the global average pooling for the pruned model to mimic and achieved state-of-the-art performance with few training samples. Note that all these algorithms were designed empirically. There is no theory to explain




Figure 6. (a) and (b) present the vanilla discriminative model and feature mimicking model, respectively.

why the feature mimicking method MiR outperforms other methods. In this section, we propose one that explains the behaviors of existing few-shot distillation approaches, and shows that the loss function of MiR is optimal among them. Furthermore, we show that feature mimicking achieves a more stable model compared to optimizing cross-entropy.

We formulate traditional classification as the discriminative model in Fig. 6a. Given the input x , this model estimates the conditional probability distribution $p(y|x)$. With KD, $p(y|x)$ can be viewed as the teacher's distribution, while the student distribution is $q_{\theta}(y|x)$, where θ denotes parameters. KL-divergence is adopted to measure the distance between these two distributions:

$$KL [p(y|x) || q_{\theta}(y|x)]. \quad (13)$$

For simplicity, we only consider the single variate case, but it is natural to extend the analyses to multivariate.

More information about intermediate features is introduced in feature mimicking, shown as Fig. 6b. We know the feature distribution $p(f|x)$ in addition to the teacher's prediction $p(y|x, f)$. Now, the target is the joint distribution of both f and y :

$$KL [p(y, f|x) || q_{\theta}(y, f|x)], \quad (14)$$

where p and q denote the teacher and student, respectively. Intuitively, once we successfully optimize this target, we can assume $q_{\theta}(y, f|x)$ approaches $p(y, f|x)$, denoted by $q_{\theta}(y, f|x) \rightarrow p(y, f|x)$. Therefore,

$$q_{\theta}(y|x) = \int q_{\theta}(y, f|x) df \rightarrow \int p(y, f|x) df = p(y|x), \quad (15)$$

which means we also optimize Eq. 13 successfully and obtain a discriminative model for classification. The following claim establishes a strict connection between Eq. 13 and Eq. 14. Its proof is in Appendix A.

Claim 1. Eq. 13 is upper bounded by Eq. 14, that is,

$$KL [p(y|x) || q_{\theta}(y|x)] \leq KL [p(y, f|x) || q_{\theta}(y, f|x)]. \quad (16)$$

Claim 1 tells us optimizing Eq. 14 is reasonable for obtaining a classifier. In the rest of this section, we will show that optimizing Eq. 14 is easier than Eq. 13, and it results in a more stable model. First, we introduce how to optimize Eq. 14 by the following claim, whose proof is in Appendix B.




Figure 7. Different ways to optimize the classifier (“FC”) in KD. Blue modules with solid boundaries denote teacher modules, while orange modules with dash boundaries denote student modules.

Claim 2. Eq. 14 can be decoupled into a classification loss and a feature mimicking loss, in the form

$$KL [p(y, f|x) || q_\theta(y, f|x)] = C + \mathbb{E}_{f \sim p(f|x)} \left\{ \underbrace{-\ln q_\theta(f|x)}_{\text{Feature Mimicking}} + \underbrace{KL [p(y|f, x) || q_\theta(y|f, x)]}_{\text{Classification}} \right\}, \quad (17)$$

where C is a constant. Assume $q_\theta(f|x)$ is a Gaussian with mean $\mu_\theta(x)$ and standard deviation $\frac{1}{\sqrt{2\beta}}$, then

$$\begin{aligned} & \mathbb{E}_{f \sim p(f|x)} [-\ln q_\theta(f|x)] \\ &= \mathbb{E}_{f \sim p(f|x)} \left[\beta (f - \mu_\theta(x))^2 \right] + C'. \end{aligned} \quad (18)$$

That is, using the L2 loss for the feature mimicking component is optimal under certain reasonable assumptions. In practice, this loss is often adopted as the feature-based distillation loss. β is then a hyperparameter to balance feature mimicking and classification, and $\mu(x)$ is our student backbone network which generates features given input x .

For classification loss $\mathbb{E}_{f \sim p(f|x)} KL [p(y|f, x) || q_\theta(y|f, x)]$, we notice f , which is generated from the teacher $p(f|x)$, appears in $q_\theta(y|f, x)$. Hence, the student can utilize the teacher’s features during training. We summarize three ways to implement this loss, as illustrated in Fig. 7.

- 1) To ignore the teacher’s f , and directly model y according to the input x (as in Fig. 7a). This is the traditional manner widely used in previous work [9], [15], [32], [33]. Without the help of the teacher’s feature, the model achieves inferior accuracy.
- 2) Second, cross distillation (as in Fig. 7b). The student classifier utilizes the teacher features to reduce the classification loss. Cross distillation outperforms vanilla distillation. This approach was used in CD [10] and grafting [34] methods.
- 3) Directly use the teacher’s classifier (as in Fig. 7c). Because we already know $p(y|f)$ (which is the teacher’s classifier or “FC”), we can directly use it for $q_\theta(y|f)$. This is the optimal solution, making $\mathbb{E}_{f \sim p(f|x)} KL [p(y|f, x) || q_\theta(y|f, x)] = 0$. This approach was used in MiR [11], and also used in this paper as the finetuning method.

4.1 Variance on model parameters

Next, we analyze why Eq. 14 is better than Eq. 13. Our main result is: feature mimicking (Eq. 14) achieves a lower variance on the model’s parameters. In other words, it requires less training data to obtain a better estimate of parameters. We start from the following claim (cf. Appendix C for the proof).

Claim 3. Given n training samples x_1, x_2, \dots, x_n , assume the loss function for finetuning follows the canonical form

$$\mathcal{L} = -\ln q_\theta(f|x). \quad (19)$$

Then we approximately have

$$\hat{\theta} \sim \mathcal{N}(\tilde{\theta}, \frac{1}{n\mathcal{I}(\tilde{\theta})}), \quad (20)$$

where $\tilde{\theta}$ denotes the teacher parameters, and

$$\mathcal{I}(\theta) = \mathbb{E}_{x \sim p(x)} \int \left(\frac{\partial q_\theta(f|x)}{\partial \theta} \right)^2 \frac{1}{q_\theta(f|x)} df. \quad (21)$$

Claim 3 establishes an estimate for the variance. According to Eq. 20, the variance of $\hat{\theta}$ is calculated by $\frac{1}{n\mathcal{I}(\tilde{\theta})}$. In the following, we compute this term for feature mimicking and classification in Claims 4 and 5, respectively. The proofs can be found in Appendix D and E, respectively.

Claim 4. If a network is finetuned by feature mimicking with the loss in Eq. 18, and assume the last layer before the output features is linear (i.e., the convolution layer):

$$\mu_{w,b}(x) = w \cdot x + b, \quad (22)$$

then by Claim 3, we have

$$\hat{w} \sim \mathcal{N}(\tilde{w}, \frac{1}{2n\beta\mathbb{E}[x^2]}), \quad (23)$$

$$\hat{b} \sim \mathcal{N}(\tilde{b}, \frac{1}{2n\beta}). \quad (24)$$

Claim 5. If a network is finetuned by optimizing the loss in Eq. 13, and assume the last layer is

$$q_{w,b}(\mathbf{y}|f) = \sigma(\mathbf{w} \cdot f + \mathbf{b}), \quad (25)$$

where σ is the softmax operation

$$\sigma(\mathbf{x})_i = \frac{\exp(x_i)}{\sum_{j=1}^C \exp(x_j)}, \quad (26)$$

then by Claim 3, for the i -th entry of \mathbf{w} and \mathbf{b} , we have

$$\hat{w}_i \sim \mathcal{N}(\tilde{w}_i, \frac{1}{n\epsilon_i\mathbb{E}[f^2]}), \quad (27)$$

$$\hat{b}_i \sim \mathcal{N}(\tilde{b}_i, \frac{1}{n\epsilon_i}), \quad (28)$$

where ϵ_i is the i -th entry of ϵ and

$$\epsilon = q_{w,b}(\mathbf{y}|f)[1 - q_{w,b}(\mathbf{y}|f)]. \quad (29)$$

According to Claims 4 and 5, the variances of the weight parameters for feature mimicking and classification are $\frac{1}{2n\beta\mathbb{E}[x^2]}$ and $\frac{1}{n\epsilon\mathbb{E}[f^2]}$, respectively. Because normalization layers exist in networks, we can assume both $\mathbb{E}[x^2]$ and $\mathbb{E}[f^2]$ are bounded by a constant. If a large-scale training set is available, both variances are small due to a large n .

But, if n is small, the variance for classification (Claim 5) is large, because $\epsilon \rightarrow 0$ empirically (about 10^{-3}). For the feature mimicking, the typical value of β is $\frac{1}{2}$, and hence the variance (Claim 4) is still small. That means in the few-shot scenario, *feature mimicking is more stable than classification*.

Finally, our theory also explains traditional KD [15], which applies a temperature T within the softmax operation:

$$q_{w,b}(\mathbf{y}|f) = \sigma\left(\frac{\mathbf{w} \cdot f + b}{T}\right). \quad (30)$$

If $T > 1$, it will soften the prediction and make ϵ larger. This explains why KD is very helpful in few-shot compression.

In short, our theory establishes that feature mimicking achieves lower variance on model parameters compared to optimizing the cross-entropy loss. This explains why feature mimicking is a superior finetuning method.

5 THE PRACTISE ALGORITHM

Now we move on to propose a new algorithm for few-shot compression aligned with our hypothesis and theory. First, we analyze the benefits of dropping blocks (e.g., residual blocks in ResNet). Compared with previous few-shot compression schemes, dropping blocks enjoys a more convex optimization landscape and a high acceleration ratio that achieves superior latency-accuracy performance with tiny training sets. Second, we propose a new recoverability criterion to choose which blocks to drop. Different from previous criteria, recoverability measures the hardness of finetuning a pruned network, which is not only closely and more directly related to the model’s accuracy, but also is efficient to compute and suits the few-shot scenario well. Based on the recoverability, we propose PRACTISE, an algorithm for practical network acceleration with tiny training sets. PRACTISE does not need image labels, and it is even able to work in the data-free setting.

5.1 The motivation to drop blocks

For accelerating neural networks in real-world applications, latency and accuracy are the two most important metrics. Lower latency (instead of FLOPs) means the model runs faster *in reality*. To accelerate the model with tiny training sets, many compression schemes have been proposed. Fig. 8 compares the acceleration ratios of them. FSKD [9] prunes filters within residual blocks according to L_1 norm (‘normal’). CD [10] proposes pruning conv. layers only in shallow layers, and keeps deeper layers unchanged (‘CD-style’). MiR [11] trims the residual connection (‘residual’). All these pruning schemes suffer from inefficient acceleration ratios. As shown in Fig. 8, with about 30% FLOPs reduction, compressed models achieve only 16.1% latency reduction (41.7→35 ms). Simply resizing the input image’s resolution (namely ‘Resolution’) achieves better acceleration. At last, dropping blocks outperforms all these methods. To achieve 35 ms latency, it only needs to reduce 12.7% FLOPs, significantly less than the 30% in pruning filters.

We also care about the accuracy of the finetuned model. As discussed in Sec. 3, pruning blocks enjoys a more convex optimization landscape compared to pruning filters, hence requires less training samples (cf. Fig. 4).

Obviously, dropping blocks is more effective for model acceleration than pruning filters, but it is neglected in few-shot compression. One possible reason is that pruning filters has achieved extraordinary performance with the *whole* training dataset. However, when only a tiny training set is available, we should consider pruning efficiency more




Figure 8. The relationship between latency reduction and FLOPs reduction. Dropping blocks enjoys the highest acceleration ratio compared with other compression schemes.




Figure 9. Illustration of the recoverability of each block and different criteria for dropping blocks. ResNet-34 has 12 blocks that can be dropped (‘layer1.1’ to ‘layer4.2’). We drop one block each time and evaluate the Top-1 error with or without finetuning. Scores of different criteria are computed for each block and presented in this figure as different curves, too. Note that the Top-1 error (%) is evaluated on the ImageNet-1k validation set (50000 images), while finetuning and evaluating criteria take only 500 training images. Our PRACTISE predicts the finetuned network’s error almost perfectly well.

carefully. Existing work [9], [11] prune filters with large FLOPs reductions but suffer from a less convex optimization landscape, hence require more training data. On the other hand, dropping blocks enjoys a high acceleration ratio and keeps more parameters and capacity from the original model. With a more convex loss landscape, it requires less data for finetuning and achieves a superior latency-accuracy tradeoff with the *tiny* training set (cf. Fig. 1). We have demonstrated that even a naive dropping block method has already outperformed most existing methods.

5.2 The recoverability of the pruned model

To further improve dropping blocks, we study how to choose the blocks to drop. Existing criteria perform poorly in the few-shot scenario, as they either require a large training set [16], [17] or fail to predict the finetuned accuracy [18]. We argue that an effective metric should be consistent with the finetuned accuracy. But, as discussed in Sec. 3, it has a complex relationship to the convexity of loss landscape, initial accuracy, etc. In addition, traditional pruning steps in Fig. 2 will change the model’s optimization landscape. Hence,

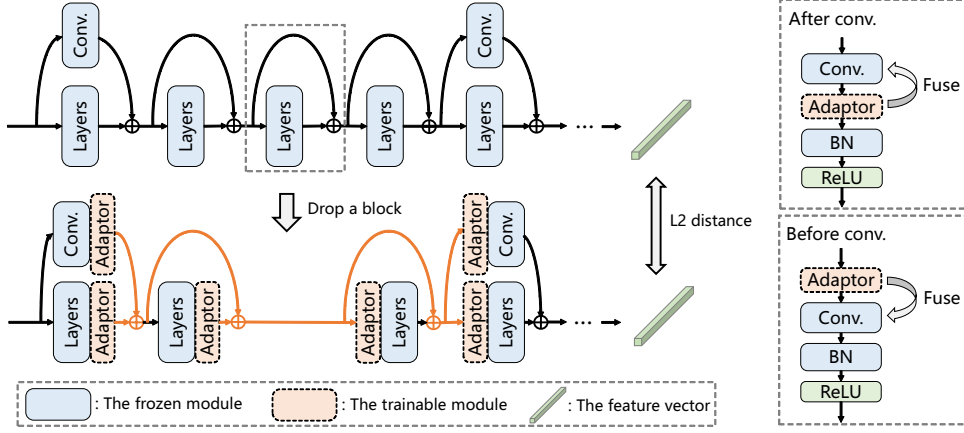


Figure 10. Illustration of our method for determining which block to drop. One block is dropped, and adaptors are inserted around the dropped position to recover from dropping this block. Then, the L2 distance of pruned/original networks’ features are computed as the *recoverability* of this dropped block. Finally, we drop several blocks that are easily recovered to obtain the pruned network.

it is inferior to simply evaluate the property of $f \leftrightarrow f'$ to indicate the finetuning of $g \leftrightarrow g'$. Instead, we will propose an algorithm, PRACTISE, to efficiently evaluate the difficulty of finetuning. In short, PRACTISE inserts adaptor layers around the pruning structure and optimizes the adaptor *with other parts fixed*. Surprisingly, for few-shot finetuning, training only adaptors achieves almost the same optimization properties as training all parameters end-to-end.

First, we propose a new concept called *recoverability*, which measures the ability of a pruned model to recover accuracy. As shown in Fig. 9, we will take the ‘layer1.1’ block as an example. Simply dropping this block results in a 54.3% Top-1 error. Both KL-divergence (‘CURL’) and ‘L2 distance’ variations before/after block removal are large, showing the model’s outputs indeed change dramatically. But, the accuracy is recovered very effectively by finetuning with even a tiny training set. Finetuning is ignored in existing criteria, but it is finetuned accuracy that directly determines the compression quality. To overcome this difficulty, we propose a method to evaluate the recoverability of each dropped block. As illustrated in Fig. 9, ours (‘PRACTISE’) enjoys a high consistency with the Top-1 error of the finetuned model.

Fig. 10 presents our method. Given the original model \mathcal{M}_O , the pruned model $\mathcal{M}_{P(\mathcal{B}_i)}$ is obtained by dropping a block \mathcal{B}_i . We insert adaptors in the positions connected to this block. These adaptors are conv. layers with kernel size 1×1 and are placed before/after the original conv. layers according to different positions. For blocks in front of the dropped block, adaptors are inserted after conv. layers. On the contrary, adaptors are inserted before conv. layers. Because convolutions are linear, all adaptors can be fused in neighboring conv. layers while keeping the network functionality unchanged. That is, these adaptors will not become overhead to the pruned model. We use the tiny training set to optimize these adaptors to minimize the distance of original/pruned models’ features (i.e., feature mimicking). The recoverability is calculated as

$$\mathcal{R}(\mathcal{B}_i) = \min_{\alpha} \mathbb{E}_{x \sim p(x)} \|\mathcal{M}_O(x; \theta) - \mathcal{M}_{P(\mathcal{B}_i)}(x; \theta \setminus b_i, \alpha)\|_F^2, \quad (31)$$

where θ means parameters in the original model, and $\setminus b_i$ means excluding the parameters in the dropped block \mathcal{B}_i , and α denotes parameters in adaptors. This metric, as Fig. 9 shows, estimates the recoverability of each block very well.

Its success aligns well with our hypothesis and theory, as shown in Fig. 11. Comparing Fig. 11a and 11b, we conclude that pruning architecture (the second step in Fig. 2) changes the loss landscape. On the contrary, optimizing only adaptors ($\mathcal{M}_{P(\mathcal{B})+A'} \leftrightarrow \mathcal{M}_{P(\mathcal{B})+A}$) performs almost the same as finetuning the entire model ($\mathcal{M}'_{P(\mathcal{B})} \leftrightarrow \mathcal{M}_{P(\mathcal{B})}$). As demonstrated in Fig. 11d, the finetuned models $\mathcal{M}'_{P(\mathcal{B})}$ and $\mathcal{M}_{P(\mathcal{B})+A'}$ even lie in the same optimization basin without obvious loss barriers. In other words, *optimizing only adaptors is almost the same as the end-to-end optimization*. They give strong empirical evidences that PRACTISE indeed finds recoverability precisely. In Fig. 11c, dropping ‘‘layer1.1’’ results in the most convex curve, which means it is easy to recover the accuracy by few samples. In contrast, ‘‘layer4.2’’ suffers from a flat curve. With a low initial accuracy, it struggles with poor performance.

Fig. 12 further examines the outcome when we drop more blocks. It demonstrates that these adaptors recover most accuracy loss and achieve accuracy on-par with finetuning all parameters in the few-shot scenario.

Our method is efficient, too. Thanks to the limited computations and parameters in adaptors, calculating the recoverability by Eq. 31 requires only a few training samples and very short training time.

The other factor for dropping blocks we must take care of is the latencies of different blocks. With the same recoverability, the block with a higher latency should have a higher priority to drop. To this end, we calculate the acceleration ratio of the block \mathcal{B}_i by

$$\tau(\mathcal{B}_i) = \frac{\text{lat}_{\mathcal{M}_O} - \text{lat}_{\mathcal{M}_{P(\mathcal{B}_i)}}}{\text{lat}_{\mathcal{M}_O}}, \quad (32)$$

where *lat* denotes the latency. Finally, we define the pruning score for each block \mathcal{B}_i as

$$s(\mathcal{B}_i) = \frac{\mathcal{R}(\mathcal{B}_i)}{\tau(\mathcal{B}_i)}. \quad (33)$$

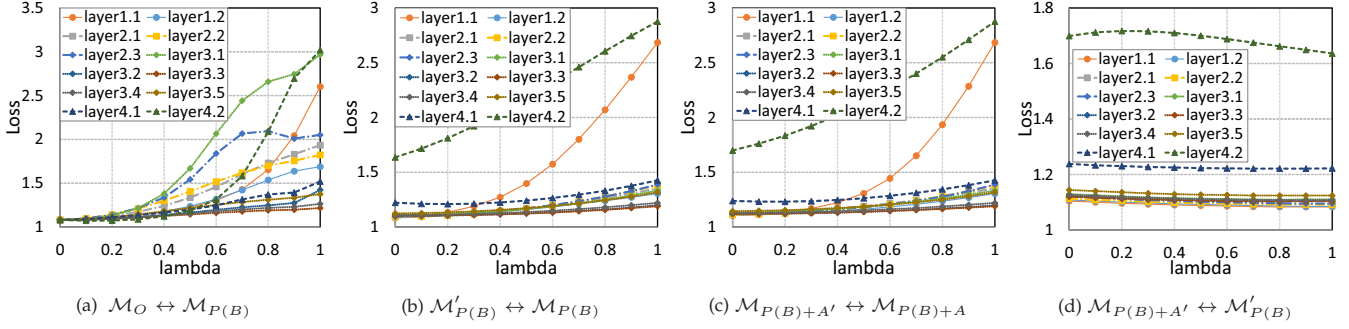


Figure 11. The loss of interpolated models between different pruned models. Each curve denotes pruning one block. Given the original model \mathcal{M}_O , the pruned model $\mathcal{M}_{P(B)}$ is obtained by dropping the block B , then finetuned on a tiny set to $\mathcal{M}'_{P(B)}$. (a) and (b) compare these models' loss landscapes. Due to the change of architecture, the loss landscapes also change significantly, as (a) and (b) are significantly different. Note that $\mathcal{M}'_{P(B)}$ is desired target. By inserting adaptors to get model $\mathcal{M}_{P(B)+A'}$, our PRACTISE only trains these adaptors to achieve $\mathcal{M}_{P(B)+A'}$. *optimizing adaptors almost perfectly achieves the desired target curve*—(b) and (c) aligns precisely. In particular, (d) reveals that the finetuned models from these two training methods lie in the same optimization basin without obvious loss barriers.




Figure 12. Comparison of training all and only adaptors. ResNet-34 was pruned by dropping different numbers of blocks with only 500 training samples. Top-1 error (%) was evaluated on the ImageNet-1k validation set. Note that the Top-1 error of the original model is 26.98%.

Our pruning score considers both recoverability and latency of each block. A lower score means the block has a higher priority to drop. We compute the pruning score for each block and drop the top k blocks with smallest scores.

5.3 Recover the accuracy of the pruned model

Once the pruned model's structure is determined, the last problem is how to recover the accuracy. The advantages of MiR has been proved in Sec. 4. Hence, following MiR, we finetune the pruned model by minimizing

$$\mathcal{L} = \|\mathcal{M}_O(x; \theta_O) - \mathcal{M}_P(x; \theta_P)\|_F^2, \quad (34)$$

where θ_O denotes the frozen parameters of the original model and θ_P denotes the trainable parameters of the pruned model.

Overall, the whole algorithm of PRACTISE is presented in Alg. 1. Our PRACTISE enjoys *zero* extra hyperparameters. With feature mimicking, we accelerate models *without* using training labels.

PRACTISE can even work in data-free scenarios. One choice is to adopt the synthesized images from DFKD [26] as training images. Most existing DFKD work [26], [28] adopt filter pruning, whereas PRACTISE improves the latency-accuracy performance by a significant margin. That helps a lot in data-free scenarios. Another choice is to collect out-of-domain data. On the other hand, with a large amount of

Algorithm 1: PRACTISE

Input: The original model \mathcal{M}_O , the number of dropped blocks k , the tiny training data \mathcal{D}_T

Test the latency of \mathcal{M}_O ;

for each block \mathcal{B}_i do

Drop \mathcal{B}_i to obtain the pruned model $\mathcal{M}_{P(\mathcal{B}_i)}$;
 Test latency of $\mathcal{M}_{P(\mathcal{B}_i)}$ and find $\tau(\mathcal{B}_i)$ (Eq. 32);
 Insert adaptors;
 Compute $\mathcal{R}(\mathcal{B}_i)$ with \mathcal{D}_T (Eq. 31);
 Compute the score $s(\mathcal{B}_i)$ (Eq. 33);
 Add \mathcal{B}_i back and remove all adaptors;

Choose the top k blocks with the smallest scores;

Drop these k blocks to obtain \mathcal{M}_P ;

Finetune \mathcal{M}_P with \mathcal{D}_T by minimizing \mathcal{L} (Eq. 34);

return The pruned model \mathcal{M}_P

out-of-domain images, the accuracy of the pruned network is even close to that of using original training images. That demonstrates the high generalization ability of PRACTISE.

6 EXPERIMENTAL RESULTS

In this section, we present PRACTISE results on both few-shot and data-free compression. Following previous work [9], [10], [11], ResNet-34 [35] and MobileNetV2 [36] will be pruned on tiny training sets of ImageNet-1k [37]. Then, to test the generalization ability of PRACTISE, we will prune ResNet-50 with synthesized images and out-of-domain data, respectively. Finally, ablation studies are conducted .

Implementation details. As presented in Alg. 1, PRACTISE requires computing the latency and recoverability of each pruned model $\mathcal{M}_{P(\mathcal{B}_i)}$, then finetuning the pruned model \mathcal{M}_P . In this paper, we only consider dropping the block with the same input and output dimensionality. For the latency, we tested the model with $64 \times 3 \times 224 \times 224$ input 500 times and record the mean latency. All latency numbers were tested on the same computer with an NVIDIA TITAN Xp GPU. To compute recoverability, we used SGD with batch size 64 to optimize adaptors by 1000 iterations. The initial learning rate was 0.02 and decreased by a factor of 10 per 40% iterations. When the size of training data \mathcal{D}_T is less than

Table 2

Top-1/Top-5 validation accuracy (%) on ImageNet-1k for pruning ResNet-34. All models reduced 16% latency with 50, 100, 500, and 1000 training samples. Previous methods pruned filters within the residual block. Top-1/Top-5 accuracy of the original ResNet-34 are 73.31%/91.42%.

Method	Latency (ms)	50	100	500	1000
BP (filter)	35.1 (15.8%↓)	39.0 \pm 1.41/68.9 \pm 1.17	41.0 \pm 0.33/70.5 \pm 0.66	51.8 \pm 0.30/78.1 \pm 0.38	57.8 \pm 0.30/81.5 \pm 0.18
BP (block)	34.9 (16.3%↓)	66.5\pm0.81/78.4\pm0.44	66.8\pm0.23/87.7\pm0.23	68.6\pm0.18/88.8\pm0.09	69.8\pm0.12/89.3\pm0.07
KD [15]	35.1 (15.8%↓)	44.5 \pm 1.20/72.3 \pm 0.87	46.4 \pm 0.34/74.0 \pm 0.58	54.7 \pm 0.26/79.7 \pm 0.19	57.9 \pm 0.21/81.6 \pm 0.12
FSKD [9]	35.1 (15.8%↓)	45.3 \pm 0.77/71.5 \pm 0.62	51.2 \pm 0.30/76.8 \pm 0.23	57.6 \pm 0.21/81.6 \pm 0.15	59.4 \pm 0.13/82.7 \pm 0.06
CD [10]	35.1 (15.8%↓)	56.2 \pm 0.37/80.8 \pm 0.31	59.1 \pm 0.22/82.8 \pm 0.11	63.7 \pm 0.18/86.0 \pm 0.05	64.4 \pm 0.03/86.3 \pm 0.07
MiR [11]	35.1 (15.8%↓)	64.1 \pm 0.10/86.3 \pm 0.11	65.1 \pm 0.19/87.0 \pm 0.11	67.0 \pm 0.09/88.1 \pm 0.07	67.8 \pm 0.06/88.5 \pm 0.02
PRACTISE	34.9 (16.3%↓)	70.3\pm0.16/89.6\pm0.06	71.5\pm0.74/90.3\pm0.37	72.5\pm0.04/90.9\pm0.03	72.5\pm0.05/91.0\pm0.02

Table 3

Top-1/Top-5 validation accuracy (%) on ImageNet-1k for pruning ResNet-34. Our model reduced 22.1% latency. Previous methods pruned filters both inside and outside the residual connection (i.e., 'residual' in Fig. 8). Top-1/Top-5 accuracy of the original ResNet-34 are 73.31%/91.42%.

Method	Latency (ms)	50	100	500	1000
BP (filter)	33.8 (18.9% ↓)	24.2 \pm 0.92/52.7 \pm 1.36	27.6 \pm 0.41/56.7 \pm 0.62	42.9 \pm 0.28/70.5 \pm 0.27	51.2 \pm 0.32/76.5 \pm 0.16
BP (block)	32.5 (22.1% ↓)	60.6\pm0.62/83.5\pm0.42	61.6\pm0.31/84.3\pm0.36	65.0\pm0.19/86.5\pm0.20	66.8\pm0.18/87.5\pm0.13
KD [15]	33.8 (18.9% ↓)	30.1 \pm 0.69/57.7 \pm 1.10	33.1 \pm 0.43/61.0 \pm 0.53	45.7 \pm 0.26/72.2 \pm 0.25	50.5 \pm 0.29/75.9 \pm 0.23
FSKD [9]	33.8 (18.9% ↓)	31.1 \pm 0.90/56.5 \pm 1.10	36.6 \pm 0.44/63.1 \pm 0.46	42.8 \pm 0.49/69.1 \pm 0.58	44.9 \pm 0.20/70.5 \pm 0.29
MiR [11]	33.8 (18.9% ↓)	59.9 \pm 0.30/83.2 \pm 0.31	62.1 \pm 0.22/84.8 \pm 0.18	65.4 \pm 0.07/87.0 \pm 0.03	66.6 \pm 0.05/87.7 \pm 0.04
PRACTISE	32.5 (22.1% ↓)	68.0\pm1.36/88.2\pm0.77	70.4\pm0.42/89.7\pm0.23	71.8\pm0.07/90.5\pm0.02	71.9\pm0.05/90.6\pm0.04

Table 4

Top-1 validation accuracy (%) on ImageNet-1k for pruning ResNet-34. The model was pruned at three different compression ratios with 500 training samples. Previous methods prune filters by CD-style, while we drop blocks. The Top-1 accuracy and the latency of the original ResNet-34 are 73.31% and 41.7 ms, respectively.

Latency (ms)	BP (filter)	KD [15]	FSKD [9]	CD [10]	MiR [11]	PRACTISE
32.5 (22.1%↓)	47.54 \pm 0.41	49.34 \pm 0.25	33.19 \pm 0.60	59.65 \pm 0.12	68.14 \pm 0.04	71.75\pm0.07
34.9 (16.3%↓)	58.94 \pm 0.36	61.01 \pm 0.24	62.56 \pm 0.13	68.17 \pm 0.07	70.53 \pm 0.10	72.52\pm0.04
38.3 (8.2%↓)	65.02 \pm 0.30	67.22 \pm 0.18	69.59 \pm 0.09	71.12 \pm 0.06	71.95 \pm 0.07	73.04\pm0.06

64, batch size was the data size. For the final finetuning, all parameters in the pruned network were updated by SGD to minimize Eq. 34 (i.e., feature mimicking) with 2000 iterations by default. Batch size and learning rate schedules were the same as those for adaptors. We used the data argumentation strategy supplied by PyTorch official examples, which is the same as that of previous work [9], [10], [11]. All experiments were conducted with PyTorch [38].

To compare with other few-shot pruning methods, we are mainly concerned about the latency-accuracy tradeoff of the pruned model. We tested the latency of networks pruned by previous methods and directly cite their accuracy numbers. The results of PRACTISE were run by five times with different sampled tiny training sets. We report the mean accuracy along with standard deviation.

6.1 Different amounts of training data

We compare PRACTISE with state-of-the-art few-shot pruning methods. Following previous work, we prune ResNet-34 with different amounts of training images on ImageNet-1k. Table 2 summarizes results. All previous methods pruned filters within residual blocks (cf. 'normal' in Fig. 8), and both our PRACTISE and dropping block scheme achieve significantly better results. First, we removed the first few

blocks to reach the latency goal and then simply finetuned the network with the cross-entropy loss. This simple baseline for dropping blocks, 'BP (block)', has already outperformed previous methods. Note that this is our contribution, too, which reveals the advantage of dropping blocks in the few-shot scenario. PRACTISE further improves results. With similar latency reduction, it dramatically outperforms previous SOTA by an average of 5.7% Top-1 accuracy on ImageNet-1k. Table 3 compares these methods with a larger latency reduction. PRACTISE surpasses MiR by a significant margin again, on average 7.0% Top-1 accuracy. Note that our model is faster than previous ones by 1.3 ms, too.

6.2 Different acceleration ratios

Table 4 compares PRACTISE with previous methods for different acceleration ratios. For 8.2% latency reduction, PRACTISE outperforms MiR by 1.1% Top-1 accuracy. In further reducing latency by 22.1%, PRACTISE surpasses MiR by 3.6% Top-1 accuracy. That indicates PRACTISE enjoys higher accuracy than others when the acceleration ratio becomes larger. Fig. 1 presents curves for the latency-accuracy tradeoffs of different methods. Our PRACTISE outperforms previous methods dominantly.




Figure 13. The data-latency-accuracy tradeoff. ResNet-34 was pruned by PRACTISE on ImageNet-1k. Numbers in the legend mean different amounts of training images.

Table 5

Comparison of PRACTISE and few-shot pruning methods on ImageNet-1k. MobileNetV2 was pruned with 500 training samples.

Method	Latency (ms)	Top-1/Top-5
Original	37.6	71.9/90.3
BP (filter)	31.5 (16.2% ↓)	45.0 \pm 0.34/71.8 \pm 0.38
KD [15]	31.5 (16.2% ↓)	48.4 \pm 0.34/73.9 \pm 0.32
MiR [11]	31.5 (16.2% ↓)	67.6 \pm 0.05/87.9 \pm 0.04
PRACTISE	30.4 (19.1% ↓)	69.3\pm0.05/88.9\pm0.05
BP (filter)	34.1 (9.3% ↓)	55.5 \pm 0.16/80.3 \pm 0.26
KD [15]	34.1 (9.3% ↓)	59.1 \pm 0.17/82.5 \pm 0.15
MiR [11]	34.1 (9.3% ↓)	69.7 \pm 0.04/89.2 \pm 0.03
PRACTISE	31.9 (15.2% ↓)	70.3\pm0.03/89.5\pm0.03

6.3 The data-latency-accuracy tradeoff

Next, for a better understanding of the data-latency-accuracy tradeoff in the few-shot compression scenario, we pruned ResNet-34 by PRACTISE with different latency reductions and various numbers of training images. Fig. 13 presents the results. With less latency reduction, the accuracies for different amounts of training data are comparable. But for a large latency reduction, it is challenging to recover the accuracy by only a tiny training set. And the accuracy gap becomes large w.r.t. different amounts of training data.

6.4 Results on MobileNetV2

MobileNetV2 [36] is a lightweight model and is popularly applied on mobile devices. It has 10 blocks that can be dropped by PRACTISE. Table 5 summarizes MobileNetV2 results. PRACTISE outperforms previous methods by a significant margin. Compared with MiR, PRACTISE obtains pruned models with both lower latency and higher accuracy.

6.5 Train with synthesized/out-of-domain images

In some extreme scenarios, not even a single original training sample is to be provided. Zero-shot pruning is required. Because PRACTISE does not need ground-truth labels, it is able to accelerate networks with images synthesized by data-free knowledge distillation methods. Table 6 summarizes results. Note that most zero-shot pruning methods adopt

Table 6

Comparison of PRACTISE and data-free methods on ImageNet-1k. Latency (ms) and Top-1 validation accuracy (%) on ImageNet-1k are reported. * denotes the method reimplemented by ourselves.

Network	Method	Pruning	Latency	Top-1
ResNet-50	Original		83.8	76.1
	DI [26]	filter	-	72.0
	MixMix [28]	filter	-	69.8
	ADI [26]	filter	-	73.3
	ADI* [26]	filter	79.9 (4.7%↓)	73.5
	PRACTISE	block	66.2 (21.0%↓)	74.8
MobileNetV2	Original		37.6	71.9
	DI [26]	filter	-	15.3
	MixMix [28]	filter	-	42.5
	ADI* [26]	filter	30.8 (18.1%↓)	62.8
	PRACTISE	block	30.4 (19.1%↓)	68.0

Table 7

Top-1 validation accuracy (%) on ImageNet-1k for different out-of-domain training datasets. ResNet-50 was reduced 21% latency by PRACTISE.

Dateset	50	500	1000	5000	All
ImageNet [37]	74.22	74.58	74.58	75.14	75.24
ADI [26]	69.85	72.68	73.01	74.40	74.79
CUB [39]	72.49	73.71	73.94	74.86	74.92
Place365 [40]	72.80	74.10	74.18	75.05	75.21

filter pruning, which are inferior to pruning blocks as we have shown. We adopt synthesized images produced by ADI [26] as the training set. For both ResNet-50 and MobileNetV2, our PRACTISE achieves higher Top-1 accuracy with more latency reductions. We advocate that we need to adopt dropping blocks for DFKD in the future.

Another choice for zero-shot pruning is to collect out-of-domain training images. PRACTISE is also robust to work with these data. Table 7 presents results. The original ResNet-50 was trained on ImageNet-1k, and we pruned it on other datasets by PRACTISE. ADI consists of images synthesized by DeepInversion [26]. CUB [39] contains images of birds with 200 categories. Place365 [40] consists of scene pictures. Our PRACTISE enjoys a high generalization ability to work with all these datasets. Another benefit of out-of-domain data is the unlimited number of images. We notice the accuracy is boosted by using more training samples and even close to that of using the entire original ImageNet-1k training set.

6.6 Different criteria for dropping blocks

Finally, we compare PRACTISE with other criteria for dropping blocks. Fig. 9 already shows results for removing only one block. Obviously, PRACTISE is better than others, and is very consistent with the finetuned models' accuracies. Most existing methods mainly measure the gap between the original network and the pruned network without finetuning. They neglect the recoverability of each dropped block, hence resulting in an inferior strategy for dropping blocks. Our PRACTISE pays attention to the recoverability and the acceleration ratio of each block. Therefore the pruned network enjoys higher accuracy and lower latency.

Table 8 compares different criteria for dropping more blocks. As the number of dropped blocks increases, our

Table 8

Top-1 validation accuracy (%) on ImageNet-1k for pruning ResNet-34 with 500 images. We compare different criteria for dropping different numbers of blocks.

#Dropped blocks	1	2	3	4	5
CURL [18]	72.33	71.08	69.14	65.48	64.97
ϵ -ResNet [17]	72.51	71.20	69.00	67.90	64.75
L2 distance	72.51	71.20	69.00	68.61	64.93
PRACTISE	73.02	72.52	71.92	70.86	69.35

PRACTISE outperforms others by even larger margins. To sum up, PRACTISE finds the best set of blocks to drop compared with other methods.

7 CONCLUSIONS AND FUTURE WORK

This paper aims at accelerating networks with tiny training sets. We revealed a crucial observation in the few-shot compression: finetuning a pruned network is almost a convex problem. This conjecture successfully explains why related algorithms do not suffer from overfitting problems. We believe this finding makes significant progress in few-shot model compression. We further established a theory for feature mimicking. Compared to the vanilla method, it achieves lower parameter variances. Based on our theoretical results, for the first time, we revealed that dropping blocks is fundamentally better than previous filter-level pruning in this scenario, in terms of higher convexity and acceleration ratio. To determine which blocks to drop, we proposed a new concept, *recoverability*, to measure the difficulty of recovering the pruned network’s accuracy with few samples. Combining recoverability and feature mimicking, PRACTISE was proposed. It enjoys high latency-accuracy performance and is robust to deal with synthesized/out-of-domain images. Extensive experiments demonstrated that PRACTISE outperforms previous methods by a significant margin (on average 7% Top-1 accuracy on ImageNet-1k for 22% latency reduction).

Our conjecture has limitations such as it is confined to finetuning a pruned model. It needs future explorations for other scenarios like transfer learning, domain adaptation, etc. Our theory is proposed mainly for the discriminative model. How to distill the generative model is still an open problem. For network acceleration, how to compute the recoverability of other compression schemes is also an interesting problem. Recently, finetuning and accelerating a large pretrained network on downstream tasks are emerging as critical needs. Applying PRACTISE for tuning a pretrained model on downstream tasks is worth exploring, too.

REFERENCES

- [1] H. Li, A. Kadav, I. Durdanovic, H. Samet, and H. P. Graf, “Pruning filters for efficient convnets,” *arXiv preprint arXiv:1608.08710*, 2016.
- [2] J.-H. Luo, J. Wu, and W. Lin, “ThiNet: A filter level pruning method for deep neural network compression,” in *Int. Conf. Comput. Vis.*, 2017, pp. 5058–5066.
- [3] Z. Liu, J. Li, Z. Shen, G. Huang, S. Yan, and C. Zhang, “Learning efficient convolutional networks through network slimming,” in *Int. Conf. Comput. Vis.*, 2017, pp. 2736–2744.
- [4] Z. Liu, X. Zhang, Z. Shen, Y. Wei, K.-T. Cheng, and J. Sun, “Joint multi-dimension pruning via numerical gradient update,” *IEEE Trans. Image Process.*, vol. 30, pp. 8034–8045, 2021.
- [5] J. Guo, Y. Li, W. Lin, Y. Chen, and J. Li, “Network decoupling: From regular to depthwise separable convolutions,” in *Brit. Mach. Vis. Conf.*, 2018, p. 248.
- [6] S. Lin, R. Ji, C. Chen, D. Tao, and J. Luo, “Holistic CNN compression via low-rank decomposition with knowledge transfer,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 41, no. 12, pp. 2889–2905, 2018.
- [7] S. Han, H. Mao, and W. J. Dally, “Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding,” in *Int. Conf. Learn. Represent.*, 2016, pp. 1–14.
- [8] R. Banner, I. Hubara, E. Hoffer, and D. Soudry, “Scalable methods for 8-bit training of neural networks,” in *Adv. Neural Inform. Process. Syst.* 31, 2018, pp. 5151–5159.
- [9] T. Li, J. Li, Z. Liu, and C. Zhang, “Few sample knowledge distillation for efficient network compression,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020, pp. 14 639–14 647.
- [10] H. Bai, J. Wu, I. King, and M. Lyu, “Few shot network compression via cross distillation,” in *AAAI Conf. Arti. Intell.*, vol. 04, 2020, pp. 3203–3210.
- [11] H. Wang, J. Liu, X. Ma, Y. Yong, Z. Chai, and J. Wu, “Compressing models with few samples: Mimicking then replacing,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2022, pp. 701–710.
- [12] F. Draxler, K. Veschgini, M. Salmhofer, and F. Hamprecht, “Essentially no barriers in neural network energy landscape,” in *Int. Conf. Mach. Learn.*, 2018, pp. 1309–1318.
- [13] J. Frankle, G. K. Dziugaite, D. Roy, and M. Carbin, “Linear mode connectivity and the lottery ticket hypothesis,” in *Int. Conf. Mach. Learn.*, 2020, pp. 3259–3269.
- [14] S. K. Ainsworth, J. Hayase, and S. Srinivasa, “Git Re-Basin: Merging models modulo permutation symmetries,” in *Int. Conf. Learn. Represent.*, 2023, pp. 1–29.
- [15] G. E. Hinton, O. Vinyals, and J. Dean, “Distilling the knowledge in a neural network,” *arXiv preprint arXiv:1503.02531*, 2015.
- [16] W. Wang, S. Zhao, M. Chen, J. Hu, D. Cai, and H. Liu, “DBP: Discrimination based block-level pruning for deep model acceleration,” *arXiv preprint arXiv:1912.10178*, 2019.
- [17] X. Yu, Z. Yu, and S. Ramalingam, “Learning strict identity mappings in deep residual networks,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018, pp. 4432–4440.
- [18] J.-H. Luo and J. Wu, “Neural network pruning with residual-connections and limited-data,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020, pp. 1458–1467.
- [19] G.-H. Wang and J. Wu, “Practical network acceleration with tiny sets,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2023, p. in press.
- [20] Y. He, P. Liu, Z. Wang, Z. Hu, and Y. Yang, “Filter pruning via geometric median for deep convolutional neural networks acceleration,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019, pp. 4340–4349.
- [21] P. Molchanov, A. Mallya, S. Tyree, I. Frosio, and J. Kautz, “Importance estimation for neural network pruning,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019, pp. 11 264–11 272.
- [22] Z. Liu, M. Sun, T. Zhou, G. Huang, and T. Darrell, “Rethinking the value of network pruning,” in *Int. Conf. Learn. Represent.*, 2018, pp. 1–21.
- [23] X. Ding, T. Hao, J. Tan, J. Liu, J. Han, Y. Guo, and G. Ding, “ResRep: Lossless cnn pruning via decoupling remembering and forgetting,” in *Int. Conf. Comput. Vis.*, 2021, pp. 4510–4520.
- [24] Y. Li, K. Adamczewski, W. Li, S. Gu, R. Timofte, and L. Van Gool, “Revisiting random channel pruning for neural network compression,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2022, pp. 191–201.
- [25] G.-H. Wang, J. Li, B. Li, and Y. Lu, “Evc: Towards real-time neural image compression with mask decay,” in *Int. Conf. Learn. Represent.*, 2023, pp. 1–23.
- [26] H. Yin, P. Molchanov, J. M. Alvarez, Z. Li, A. Mallya, D. Hoiem, N. K. Jha, and J. Kautz, “Dreaming to distill: Data-free knowledge transfer via DeepInversion,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020, pp. 8715–8724.
- [27] G. Fang, K. Mo, X. Wang, J. Song, S. Bei, H. Zhang, and M. Song, “Up to 100x faster data-free knowledge distillation,” in *AAAI Conf. Arti. Intell.*, vol. 36, no. 6, 2022, pp. 6597–6604.
- [28] Y. Li, F. Zhu, R. Gong, M. Shen, X. Dong, F. Yu, S. Lu, and S. Gu, “MixMix: All you need for data-free compression are feature and data mixing,” in *Int. Conf. Comput. Vis.*, 2021, pp. 4410–4419.
- [29] Y. Liu, W. Zhang, and J. Wang, “Zero-shot adversarial quantization,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021, pp. 1512–1521.
- [30] X. Zhang, H. Qin, Y. Ding, R. Gong, Q. Yan, R. Tao, Y. Li, F. Yu, and X. Liu, “Diversifying sample generation for accurate data-free

- quantization," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021, pp. 15 658–15 667.
- [31] Y. Cai, Z. Yao, Z. Dong, A. Gholami, M. W. Mahoney, and K. Keutzer, "ZeroQ: A novel zero shot quantization framework," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020, pp. 13 169–13 178.
- [32] A. Romero, N. Ballas, S. E. Kahou, A. Chassang, C. Gatta, and Y. Bengio, "FitNets: Hints for thin deep nets," in *Int. Conf. Learn. Represent.*, 2015, pp. 1–13.
- [33] G.-H. Wang, Y. Ge, and J. Wu, "Distilling knowledge by mimicking features," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 11, pp. 8183–8195, 2022.
- [34] C. Shen, X. Wang, Y. Yin, J. Song, S. Luo, and M. Song, "Progressive network grafting for few-shot knowledge distillation," in *AAAI Conf. Arti. Intell.*, vol. 35, no. 3, 2021, pp. 2541–2549.
- [35] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2016, pp. 770–778.
- [36] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "MobileNetV2: Inverted residuals and linear bottlenecks," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018, pp. 4510–4520.
- [37] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet large scale visual recognition challenge," *Int. J. Comput. Vis.*, vol. 115, no. 3, pp. 211–252, 2015.
- [38] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimeshain, L. Antiga *et al.*, "PyTorch: An imperative style, high-performance deep learning library," in *Adv. Neural Inform. Process. Syst.* 32, 2019, pp. 8026–8037.
- [39] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie, "The Caltech-UCSD birds-200-2011 dataset," California Institute of Technology, Tech. Rep. CNS-TR-2011-001, 2011.
- [40] B. Zhou, A. Khosla, A. Lapedriza, A. Torralba, and A. Oliva, "Places: An image database for deep scene understanding," *arXiv preprint arXiv:1610.02055*, 2016.



Guo-Hua Wang received his BS degree in the School of Management and Engineering from Nanjing University. He is currently a Ph.D. student in the Department of Computer Science and Technology in Nanjing University, China. His research interests are computer vision and machine learning.



Jianxin Wu (Member, IEEE) received the BS and MS degrees from Nanjing University, and the PhD degree from the Georgia Institute of Technology, all in computer science. He is currently a professor with the Department of Computer Science and Technology and School of Artificial Intelligence, Nanjing University, China, and associated with the State Key Laboratory for Novel Software Technology, China. His research interests include computer vision and machine learning. He was the (senior) area chair of CVPR, ICCV, ECCV,

AAAI, and IJCAI and an associate editor for the IEEE Transactions on Pattern Analysis and Machine Intelligence.

APPENDIX A**PROOF OF CLAIM 1**

We omit the parameters θ in $q_\theta(y|x)$ and $q_\theta(y|f, x)$ for simplicity.

$$KL [p(y, f|x) || q(y, f|x)] \quad (35)$$

$$= \iint p(y, f|x) \ln \frac{p(y, f|x)}{q(y, f|x)} dydf \quad (36)$$

$$= \iint p(f|y, x) p(y|x) \ln \frac{p(f|y, x) p(y|x)}{q(f|y, x) q(y|x)} dydf \quad (37)$$

$$= \int p(y|x) \int p(f|y, x) \left[\ln \frac{p(y|x)}{q(y|x)} + \ln \frac{p(f|y, x)}{q(f|y, x)} \right] dydf \quad (38)$$

$$= \int p(y|x) \int p(f|y, x) \ln \frac{p(y|x)}{q(y|x)} dydf \quad (39)$$

$$+ \int p(y|x) \int p(f|y, x) \ln \frac{p(f|y, x)}{q(f|y, x)} dydf$$

$$= \int p(y|x) \ln \frac{p(y|x)}{q(y|x)} dy \quad (40)$$

$$+ \int p(y|x) KL [p(f|y, x) || q(f|y, x)] dy$$

$$= KL [p(y|x) || q(y|x)] \quad (41)$$

$$+ \int p(y|x) KL [p(f|y, x) || q(f|y, x)] dy$$

$$\geq KL [p(y|x) || q(y|x)]. \quad (42)$$

APPENDIX B**PROOF OF CLAIM 2**

We omit the parameters θ in $q_\theta(y|x)$ and $q_\theta(y|f, x)$ for simplicity.

$$KL [p(y, f|x) || q(y, f|x)] \quad (43)$$

$$= \iint p(y, f|x) \ln \frac{p(y, f|x)}{q(y, f|x)} dydf \quad (44)$$

$$= \iint p(y|f, x) p(f|x) \ln \frac{p(y|f, x) p(f|x)}{q(y, f|x)} dydf \quad (45)$$

$$= \int p(f|x) \int p(y|f, x) \ln \frac{p(y|f, x) p(f|x)}{q(y, f|x)} dydf \quad (46)$$

$$= \mathbb{E}_{f \sim p(f|x)} \int p(y|f, x) \ln \frac{p(y|f, x) p(f|x)}{q(y, f|x)} dy \quad (47)$$

$$= \mathbb{E}_{f \sim p(f|x)} \left[\int p(y|f, x) \ln p(f|x) dy \right. \quad (48)$$

$$\left. + \int p(y|f, x) \ln \frac{p(y|f, x)}{q(y, f|x)} dy \right]$$

$$= \mathbb{E}_{f \sim p(f|x)} \left[\ln p(f|x) + \int p(y|f, x) \ln \frac{p(y|f, x)}{q(y|f, x) q(f|x)} dy \right] \quad (49)$$

$$= \mathbb{E}_{f \sim p(f|x)} \left[\ln p(f|x) - \int p(y|f, x) \ln q(f|x) dy \right. \quad (50)$$

$$\left. + \int p(y|f, x) \ln \frac{p(y|f, x)}{q(y|f, x)} dy \right]$$

$$= \mathbb{E}_{f \sim p(f|x)} \{-\ln q(f|x) + KL [p(y|f, x) || q(y|f, x)]\} + C. \quad (51)$$

For feature mimicking, we let $q_\theta(f|x)$ be a Gaussian whose distribution parameters are computed from x with a deep neural network. So we have

$$q_\theta(f|x) = \frac{1}{\sqrt{2\pi}\sigma_\theta(x)} \exp \left(-\frac{1}{2} \left(\frac{f - \mu_\theta(x)}{\sigma_\theta(x)} \right)^2 \right). \quad (52)$$

Hence, the feature mimicking loss becomes

$$\mathbb{E}_{f \sim p(f|x)} [-\ln q_\theta(f|x)] \quad (53)$$

$$= \mathbb{E}_{f \sim p(f|x)} \left[-\ln \frac{1}{\sqrt{2\pi}\sigma_\theta(x)} \exp \left(-\frac{1}{2} \left(\frac{f - \mu_\theta(x)}{\sigma_\theta(x)} \right)^2 \right) \right] \quad (54)$$

$$= \mathbb{E}_{f \sim p(f|x)} \left[\frac{1}{2} \left(\frac{f - \mu_\theta(x)}{\sigma_\theta(x)} \right)^2 - \ln \frac{1}{\sqrt{2\pi}\sigma_\theta(x)} \right]. \quad (55)$$

If we assume the variance is a constant, and let $\sigma(x) = \frac{1}{\sqrt{2\beta}}$, then we have

$$\mathbb{E}_{f \sim p(f|x)} [-\ln q_\theta(f|x)] \quad (56)$$

$$= \mathbb{E}_{f \sim p(f|x)} \left[\beta (f - \mu_\theta(x))^2 - \ln \frac{\sqrt{2\beta}}{\sqrt{2\pi}} \right] \quad (57)$$

$$= \mathbb{E}_{f \sim p(f|x)} \left[\beta (f - \mu_\theta(x))^2 \right] + C. \quad (58)$$

APPENDIX C**PROOF OF CLAIM 3**

Given n training samples x_1, x_2, \dots, x_n and the loss function

$$\mathcal{L} = -\ln q_\theta(f|x), \quad (59)$$

we can obtain local minima solutions by finetuning. By Conjecture 2, the loss landscape is convex w.r.t. parameters, and the local minima will not escape from the current optimization basin. Hence, the solution can be derived by

$$\arg \max_{\theta} \sum_{i=1}^n \ln q_\theta(f_i|x_i). \quad (60)$$

Note that it is equivalent to a Maximum Likelihood Estimate (MLE). The log-likelihood for each sample is

$$l_i(\theta) = \ln q(f_i|x_i), \quad (61)$$

and the total log-likelihood is

$$l(\theta) = \sum_{i=1}^n l_i(\theta). \quad (62)$$

For MLE, the condition for achieving optimal solutions $\hat{\theta}$ is $\frac{\partial l(\theta)}{\partial \theta} \Big|_{\theta=\hat{\theta}} = 0$. With the Taylor expansion on θ and omitting high order terms,

$$0 = \frac{\partial l(\theta)}{\partial \theta} \Big|_{\theta=\hat{\theta}} \approx \frac{\partial l(\tilde{\theta})}{\partial \tilde{\theta}} + \frac{\partial^2 l(\tilde{\theta})}{\partial \tilde{\theta}^2} (\hat{\theta} - \tilde{\theta}). \quad (63)$$

Hence, we have

$$\hat{\theta} \doteq \tilde{\theta} - \frac{\frac{\partial l(\tilde{\theta})}{\partial \tilde{\theta}}}{\frac{\partial^2 l(\tilde{\theta})}{\partial \tilde{\theta}^2}}. \quad (64)$$

Note that this equation holds approximately under our Conjecture 2. For simplicity, we denote $\tilde{\theta}$ as θ in the rest

of this proof. Next, we will analyze the variance of $\hat{\theta}$. For each sample,

$$\frac{\partial l_i(\theta)}{\partial \theta} = \frac{\dot{q}_\theta(f_i|x_i)}{q_\theta(f_i|x_i)}, \quad (65)$$

where $\dot{q}_\theta(f_i|x_i)$ denotes $\frac{\partial q_\theta(f_i|x_i)}{\partial \theta}$.

$$\mathbb{E}_{f \sim q(f|x_i)} \left[\frac{\partial l_i(\theta)}{\partial \theta} \right] = \int \frac{\dot{q}_\theta(f|x_i)}{q_\theta(f|x_i)} q_\theta(f|x_i) df \quad (66)$$

$$= \int \dot{q}_\theta(f|x_i) df \quad (67)$$

$$= \frac{\partial}{\partial \theta} \int q_\theta(f|x_i) df \quad (68)$$

$$= \frac{\partial 1}{\partial \theta} = 0. \quad (69)$$

$$\text{Var}_{f \sim q(f|x_i)} \left[\frac{\partial l_i(\theta)}{\partial \theta} \right] = \int \dot{l}_i(\theta)^2 q_\theta(f|x_i) df. \quad (70)$$

The Fisher information is $\mathcal{I}(\theta) = \int \dot{l}_i(\theta)^2 q_\theta(f|x_i) df$, and

$$\mathcal{I}(\theta) = \mathbb{E}_{x \sim p(x)} \int \dot{l}_i(\theta)^2 q_\theta(f|x) df. \quad (71)$$

Note that $\frac{i(\theta)}{n} = \frac{1}{n} \sum_{i=1}^n \dot{l}_i(\theta)$, and with the central limit theorem (CLT), we have

$$\frac{\dot{l}(\theta)}{n} \sim \mathcal{N}\left(0, \frac{\mathcal{I}(\theta)}{n}\right). \quad (72)$$

For $\frac{\partial^2 l_i(\theta)}{\partial \theta^2}$, we have

$$\frac{\partial^2 l_i(\theta)}{\partial \theta^2} = \frac{\ddot{q}_\theta(f_i|x_i)}{q_\theta(f_i|x_i)} - \left(\frac{\dot{q}_\theta(f_i|x_i)}{q_\theta(f_i|x_i)} \right)^2. \quad (73)$$

Then,

$$\mathbb{E}_{f \sim q(f|x_i)} \left[\frac{\partial^2 l_i(\theta)}{\partial \theta^2} \right] \quad (74)$$

$$= \mathbb{E}_{f \sim q(f|x_i)} \left[\frac{\ddot{q}_\theta(f|x_i)}{q_\theta(f|x_i)} \right] - \mathbb{E}_{f \sim q(f|x_i)} \left(\frac{\dot{q}_\theta(f|x_i)}{q_\theta(f|x_i)} \right)^2 \quad (75)$$

$$= -\mathcal{I}(\theta). \quad (76)$$

With the law of large numbers, we have

$$\frac{\ddot{l}(\theta)}{n} = \frac{1}{n} \sum_{i=1}^n \ddot{l}_i(\theta) \xrightarrow{P} -\mathcal{I}(\theta). \quad (77)$$

With Eqs. 72, 77 and 64, we have

$$\hat{\theta} \sim \mathcal{N}\left(\tilde{\theta}, \frac{1}{n\mathcal{I}(\tilde{\theta})}\right). \quad (78)$$

APPENDIX D PROOF OF CLAIM 4

To finetune with the loss in Eq. 18, we have assumed $q(f|x)$ is a Gaussian and $\sigma(x) = \frac{1}{\sqrt{2\beta}}$, hence

$$q(f|x) = \sqrt{\frac{\beta}{\pi}} \exp\left(-\beta(f - \mu_\theta(x))^2\right). \quad (79)$$

We assume $\mu(x) = w \cdot x + b$, where w and b are parameters θ . First, we consider w :

$$\dot{q}_w(f|x) = \sqrt{\frac{\beta}{\pi}} \exp\left(-\beta(f - \mu(x))^2\right) 2\beta(f - \mu(x))x. \quad (80)$$

And,

$$\dot{l}_i(w) = \frac{\dot{q}_w(f_i|x_i)}{q_w(f_i|x_i)} \quad (81)$$

$$= 2\beta(f_i - \mu(x_i))x_i \quad (82)$$

$$= 2\beta x_i(f_i - w \cdot x_i - b), \quad (83)$$

$$\ddot{l}_i(w) = -2\beta x_i^2. \quad (84)$$

$$\mathcal{I}(w) = -\mathbb{E}_x \ddot{l}_i(w) = 2\beta \mathbb{E}[x^2]. \quad (85)$$

By Claim 3, we achieve

$$\hat{w} \sim \mathcal{N}\left(\tilde{w}, \frac{1}{2n\beta \mathbb{E}[x^2]}\right). \quad (86)$$

Then, for b , we also have

$$\dot{q}_b(f|x) = \sqrt{\frac{\beta}{\pi}} \exp\left(-\beta(f - \mu(x))^2\right) 2\beta(f - \mu(x)). \quad (87)$$

$$\dot{l}_i(b) = \frac{\dot{q}_b(f_i|x_i)}{q_b(f_i|x_i)} = 2\beta(f_i - \mu(x_i)) = 2\beta(f_i - w \cdot x_i - b). \quad (88)$$

$$\ddot{l}_i(b) = -2\beta. \quad (89)$$

$$\mathcal{I}(b) = -\mathbb{E}_x \ddot{l}_i(b) = 2\beta. \quad (90)$$

By Claim 3, we achieve

$$\hat{b} \sim \mathcal{N}\left(\tilde{b}, \frac{1}{2n\beta}\right). \quad (91)$$

APPENDIX E PROOF OF CLAIM 5

The classification loss in Eq. 13 can be written into

$$\mathbb{E}_{x \sim p(x)} KL[p(y|x) || q(y|x)]. \quad (92)$$

With the one-hot coding \mathbf{y} , this loss is equal to

$$KL[p(y|x) || q(y|x)] = \sum_{j=1}^C p(y_j|x) \ln \frac{p(y_j|x)}{q(y_j|x)} \quad (93)$$

$$= -\ln q(y_c|x), \quad (94)$$

where c denotes the ground-truth class. This loss follows the canonical form in Claim 3. In a deep neural network, x is first transformed into the feature f , then $q(y_c|f)$ is calculated as

$$q(y_c|f) = \sigma(\mathbf{w} \cdot f + \mathbf{b})_c = \frac{\exp(w_c \cdot f + b_c)}{\sum_{j=1}^C \exp(w_j \cdot f + b_j)}. \quad (95)$$

If the i -th sample belongs to class c , the log-likelihood is

$$l_i(\mathbf{w}, \mathbf{b}) = \ln \sigma(\mathbf{w} \cdot f_i + \mathbf{b})_c \quad (96)$$

$$= w_c \cdot f_i + b_c - \ln \sum_{j=1}^C \exp(w_j \cdot f_i + b_j). \quad (97)$$

First, we consider \mathbf{w} . We have

$$\dot{l}_i(w_c) = f_i - \frac{\exp(w_c \cdot f_i + b_c) f_i}{\sum_{j=1}^C \exp(w_j \cdot f_i + b_j)} = f_i - \sigma(\mathbf{w} \cdot f_i + \mathbf{b})_c f_i, \quad (98)$$

$$\dot{l}_i(w_{c'}) = -\frac{\exp(w_{c'} \cdot f_i + b_{c'}) f_i}{\sum_{j=1}^C \exp(w_j \cdot f_i + b_j)} = -\sigma(\mathbf{w} \cdot f_i + \mathbf{b})_{c'} f_i, \quad (99)$$

where $c' \neq c$. Then,

$$\ddot{l}_i(w_j) = \frac{\partial \sigma(\mathbf{w} f_i + \mathbf{b})_j f_i}{\partial w_j} \quad (100)$$

$$= -\sigma(\mathbf{w} f_i + \mathbf{b})_j [1 - \sigma(\mathbf{w} f_i + \mathbf{b})_j] f_i^2. \quad (101)$$

If we denote $\epsilon_j = \sigma(\mathbf{w} \cdot f_i + \mathbf{b})_j [1 - \sigma(\mathbf{w} \cdot f_i + \mathbf{b})_j]$, then ϵ_j always approach zero in practice. We have $\dot{l}_i(w_j) = -\epsilon_j f_i^2$, and

$$\mathcal{I}(w_j) = -\mathbb{E}_f \ddot{l}_i(w_j) = \epsilon_j \mathbb{E}[f^2]. \quad (102)$$

Finally, we achieve

$$\hat{w}_j \sim \mathcal{N}(\tilde{w}_j, \frac{1}{n \epsilon_j \mathbb{E}[f^2]}). \quad (103)$$

Now, we consider \mathbf{b} .

$$\dot{l}_i(b_c) = 1 - \frac{\exp(w_c \cdot f_i + b_c)}{\sum_{j=1}^C \exp(w_j \cdot f_i + b_j)} = 1 - \sigma(\mathbf{w} \cdot f_i + \mathbf{b})_c, \quad (104)$$

$$\dot{l}_i(b_{c'}) = -\frac{\exp(w_{c'} \cdot f_i + b_{c'})}{\sum_{j=1}^C \exp(w_j \cdot f_i + b_j)} = -\sigma(\mathbf{w} \cdot f_i + \mathbf{b})_{c'}, \quad (105)$$

where $c' \neq c$. And,

$$\ddot{l}_i(b_j) = \frac{\partial \sigma(\mathbf{w} f_i + \mathbf{b})_j}{\partial b_j} = -\sigma(\mathbf{w} f_i + \mathbf{b})_j [1 - \sigma(\mathbf{w} f_i + \mathbf{b})_j]. \quad (106)$$

$$\mathcal{I}(b_j) = -\mathbb{E}_f \ddot{l}_i(b_j) = \epsilon_j. \quad (107)$$

Finally, we achieve

$$\hat{b}_j \sim \mathcal{N}(\tilde{b}_j, \frac{1}{n \epsilon_j}). \quad (108)$$