

Few Sample Knowledge Distillation for Efficient Network Compression

Tianhong Li^{1*} Jianguo Li² Zhuang Liu³ Changshui Zhang⁴

¹MIT ²Intel Labs ³UC Berkeley ⁴Dept. Automation, Tsinghua University

tianhong@mit.edu, jianguoli@intel.com, zhuangl@berkeley.edu, zsc@tsinghua.edu.cn

Abstract

Deep neural network compression techniques such as pruning and weight tensor decomposition usually require fine-tuning to recover the prediction accuracy when the compression ratio is high. However, conventional fine-tuning suffers from the requirement of a large training set and the time-consuming training procedure. This paper proposes a novel solution for knowledge distillation from **label-free few samples** to realize both data efficiency and training/processing efficiency. We treat the original network as “teacher-net” and the compressed network as “student-net”. A 1×1 convolution layer is added at the end of each layer block of the student-net, and we fit the block-level outputs of the student-net to the teacher-net by estimating the parameters of the added layers. We prove that the added layer can be merged without adding extra parameters and computation cost during inference. Experiments on multiple datasets and network architectures verify the method’s effectiveness on student-nets obtained by various network pruning and weight decomposition methods. Our method can recover student-net’s accuracy to the same level as conventional fine-tuning methods in minutes while using only 1% label-free data of the full training data.

1. Introduction

Deep neural networks have demonstrated extraordinary success in a variety of fields such as computer vision [21, 13], speech recognition [14], and natural language processing [28]. However, their resource-hungry nature greatly hinders their wide deployment in some resource-limited scenarios. To address this limitation, many works have been done to accelerate and/or compress neural networks, among which network pruning [11, 23] and network weight decomposition [7, 18] are particularly popular due to their competitive performance and compatibility.

Network pruning [23, 25] and weight decomposition [37, 19] methods can produce extremely compressed net-

works, but they usually suffer from significant accuracy drops so that fine-tuning is required for possible accuracy recovery. However, current fine-tuning practices are not only time-consuming but also require a fully annotated large-scale training set, which may be infeasible in practice due to privacy or confidential issues. This is **very common** when software or hardware vendors help their customer optimizing deep learning solutions. As a result, when the compression ratio is high, current methods may not recover the dropped accuracy if there are very few training examples (labeled or even unlabeled) for deployment.

To solve this problem, one may ask if it is possible to utilize the knowledge from the original large network to the compressed compact network, since the latter has a similar block-level structure as the former, and already inherits some of the feature representation power from it. A natural solution is to use knowledge distillation (KD) [4, 1, 15], a method for transferring the knowledge from a large “teacher” model to a compact yet efficient “student” model by matching certain statistics between them. Further research introduced various kinds of matching mechanisms [31, 33, 17, 24]. The distillation procedure typically designs a loss function based on the matching mechanisms and optimizes the loss during a full training process. As a result, these methods still require time-consuming training procedure along with fully annotated large-scale training dataset, thus fail to meet our goal of training/processing efficiency and high sample-efficiency.

This paper addresses this issue with a simple and novel method, namely few-sample knowledge distillation (FSKD), for efficient network compression, where “efficient” here means both *training/processing efficiency* and *label-free data-sample efficiency*. As shown in [Figure 1](#), FSKD contains three steps: compressing teacher-net to obtain student-net, aligning student-teacher with the added layers in student-net, and merging the added layers. We first obtain the student-net by pruning or decomposing the teacher-net using existing methods. In this stage, we ensure that both teacher-net and student-net have the same feature map sizes at each corresponding layer block. Second, we add a 1×1 conv-layer at the end of each block of

*This work was done when Tianhong Li was intern at Intel Labs under supervised by Jianguo Li.




Figure 1: Three-step of few-sample knowledge distillation. (1) obtain student-net by compressing teacher-net; (2) add a 1×1 conv-layer at the end of each block of student-net, and align teacher and student by estimating the 1×1 conv-layer’s parameter using least-squared regression; (3) Merge the added 1×1 conv-layer into the previous conv-layer to obtain final student-net.

the student-net. We then forward the few unlabeled data to both the teacher-net and student-net, and align the block-level outputs of the student with the teacher by estimating the parameters of the added layer, using least square regressions. Because there are very few parameters to estimate in the added conv-layers, we could obtain a good estimation with a very small amount of label-free samples. The aligned student-net has the same number of parameters and computation cost as the original one since we prove that the added 1×1 convolution can be merged into the previous convolution layer.

FSKD has many potential usages, especially when full fine-tuning or re-training is infeasible in practice, or the data at hand is only very limited. We name a few concrete cases below. *First*, edge devices have limited computing resources, while FSKD offers the possibility of on-device learning to compress deep models. *Second*, FSKD may help cloud services obtain a compact model when only a few unlabeled data is uploaded by the customer due to privacy or confidential issues. *Third*, FSKD enables fast model convergence if there is a strict time budget for training/fine-tuning. The strong practice requests have been addressed by several recent workshops ¹.

Our major contributions can be summarized as follows:

- To the best of our knowledge, we are the first to show that accuracy recovery from a compressed network can be done with **few unlabeled samples** within minutes using knowledge distillation on desktop PC.
- Extensive experiments show that the proposed FSKD method is widely applicable to deep neural networks compressed by different pruning or decomposition-based methods.
- We demonstrate significant accuracy improvement from FSKD over existing distillation techniques, as

well as compression-ratio and speedup gain over traditional pruning/decomposition-based methods on various datasets and network architectures.

2. Related Work

Network Pruning methods obtain a small network by pruning weights from a trained larger network, which can keep the accuracy of the larger model if the prune ratio is set properly. [12] proposes to prune the individual weights that are near zero. Recently, filter pruning has become increasingly popular thanks to its better compatibility with off-the-shelf computing libraries, compared with weights pruning. Different criteria have been proposed to select the filters to be pruned, including norm of weights [23], scales of multiplicative coefficients [25], statistics of next layer [27], etc. These pruning based methods usually require iterative loop between pruning and fine-tuning for achieving better pruning ratio and speedup. Meanwhile, **Network Decomposition** methods try to factorize parameter-heavy layers into multiple lightweight ones. For instance, it may adopt low-rank decomposition to fully-connection layers [7], and different kinds of weight decomposition to conv-layers [18, 19, 37]. However, aggressive network pruning or network decomposition usually lead to large accuracy drops, thus fine-tuning is a must to alleviate those drops [23, 25].

Knowledge Distillation (KD) transfers knowledge from a pre-trained large teacher-net (or even an ensemble of networks) to a small student-net, for facilitating the deployment at test time. Originally, this is done by regressing the softmax output of the teacher model [15]. The soft continuous regression loss used here provides richer information than the label based loss, so that the distilled model can be more accurate than training on labeled data with cross-entropy loss. Later, various works have extended this

¹<http://sites.google.com/view/icml2019-on-device-compact-dnn>

approach by matching other statistics, including intermediate feature responses [31, 6], gradient [33], distribution [17], Gram matrix [36], etc. More complicatedly, deep mutual learning [38] trains a cohort of student-nets and teaches each other collaboratively with model distillation throughout the training process. The student-nets in KD are usually designed with random weight initialization, and thus all these methods require a large amount of data (known as the “transfer set”) to transfer the knowledge.

As a result, it is of great interest to start the student-nets with extremely pruned or decomposed networks and explore a KD solution under the few-sample setting. The proposed FSKD has a quite different philosophy on aligning intermediate responses to the closest knowledge distillation method FitNet [31]. FitNet re-trains the whole student-net with intermediate supervision as well as label supervision using a larger amount of fully-annotated data, while FSKD only estimates parameters of the added 1×1 conv-layer with few unlabeled samples. Experiments verify that FSKD is not only more efficient but also more accurate than FitNets.

Learning with few samples has been extensively studied under the concept of one-shot or few-shot learning. One category of methods directly model few-shot samples with generative models [8, 22, 5, 3], while most others study the problem under the notion of transfer learning [2, 30]. In the latter category, meta-learning methods [34, 9] solve the problem in a learning-to-learn fashion, which has been recently gaining momentum due to their application versatility. Most studies are devoted to the image classification task, while it is still less-explored for knowledge distillation from few samples. Recently, some works tried to address this problem. [20] constructs pseudo-examples using the inducing point method and develops a complicated algorithm to optimize the model and pseudo-examples alternatively. [26] records per-layer meta-data for the teacher-net in order to reconstruct a training set, and then adopts a standard training procedure to obtain the student-net. Both are very costly due to the complicated and heavy training procedure. On the contrary, we aim for an efficient solution for knowledge distillation from few unlabeled samples.

3. Method

3.1. Overview

Our FSKD method consists of three steps as shown in [Figure 1](#). *First*, we obtain a student-net either by pruning or by decomposing the teacher-net. *Second*, we add a 1×1 conv-layer at the end of each block of the student-net and align the block-level outputs between teacher and student by estimating the parameters of the added layer from few unlabeled samples. *Third*, we merge the added 1×1 conv-layer into the previous conv-layer so that it will not introduce extra parameters and computation cost into the student-net.

Algorithm 1 Block-coordinate descent algorithm for FSKD

input Student-net s , teacher-net t , input data $\{X_i\}_{i=1}^N$, number of aligned blocks M , number of iterations T

- 1: **for** $k = 1 : T$ **do**
- 2: Random flip input dataset to obtain $\{X'_i\}_{i=1}^N$;
- 3: **for** $j = 1 : M$ **do**
- 4: Feed $\{X'_i\}_{i=1}^N$ to the end of the j -th block of t , obtain response $\{X^t_{ij}\}$;
- 5: Feed $\{X'_i\}_{i=1}^N$ to the end of the j -th block of s , obtain response $\{X^s_{ij}\}$;
- 6: Add 1×1 conv-layer with tensor \mathbf{Q}_j to the end of j -th block of s ;
- 7: Solve \mathbf{Q}_j with least-square regression based on [Equation 1](#);
- 8: Merge \mathbf{Q}_j into previous conv-layer L_j with tensor \mathbf{W}_j to obtain new tensor \mathbf{W}'_j based on [Theorem 1](#) for student-net;
- 9: Update the j -th block of s , to obtain s' ;
- 10: $s = s'$;
- 11: **end for**
- 12: **end for**

output Merged conv-layers $\{\mathbf{W}'_j\}_{j=1}^M$, updated student-net s' ;

Three reasons make the idea works efficiently. *First*, the compressed student-net inherits partial representation power from the teacher network, so adding 1×1 conv-layer is enough to calibrate the student-net and restore the accuracy. *Second*, the 1×1 conv-layers have relatively fewer parameters, which do not require too many data for the estimation. *Third*, the block-level output from teacher-net provides rich information as shown in FitNet [31]. Below, we will first describe our algorithm for block-level output alignment, and then prove why the added 1×1 conv-layer can be merged into the previous conv-layer.

3.2. Block-level Alignment

In this section, we provide details of our block-level output alignment algorithm. Suppose $X^s, X^t \in \mathbb{R}^{n_o \times d}$ are the block-level output in matrix form for the student-net and teacher-net respectively, where d is the per-channel feature map resolution size. We add a 1×1 conv-layer \mathbf{Q} at the end of each block of student-net before non-linear activation. As \mathbf{Q} is degraded to the matrix form, it can be estimated with least squared regression as

$$\mathbf{Q}^* = \arg \min_{\mathbf{Q}} \sum_{i=1}^N \|\mathbf{Q} * \mathbf{X}_i^s - X_i^t\|, \quad (1)$$

where N is the number of label-free samples used, and “*” means matrix product. The number of parameters of \mathbf{Q} is $n_o \times n_o$, where n_o is the number of output channels in the block, which is usually not too large so that we can estimate \mathbf{Q} with a limited number of samples.

Suppose there are M corresponding blocks in the teacher-net and the student-net required to align, to achieve

our goal, we need minimize the following loss function

$$\mathcal{L}(\mathbf{Q}_j) = \sum_{j=1}^M \sum_{i=1}^N \|\mathbf{Q}_j * X_{ij}^s - X_{ij}^t\|_F, \quad (2)$$

where \mathbf{Q}_j is the tensor for the added 1×1 conv-layer of the j -th block. In practice, we optimize this loss with a block-coordinate descent (BCD) algorithm [35], which greedily handles each of the M blocks in the student-net sequentially as shown in Algorithm-1 (FSKD-BCD). We can also minimize this loss using standard SGD on all added parameters (FSKD-SGD). However, FSKD-BCD has the following advantages over FSKD-SGD: (1) The BCD algorithm processes each block greedily with a sequential update rule, and each \mathbf{Q} can be solved with the same set of few samples by aligning the block-level responses between teacher-net and student-net, while standard SGD considers $\{\mathbf{Q}_j\}$ all together which theoretically requires more data. (2) The BCD algorithm is much more efficient, which can be usually done in less than a minute.

Unless otherwise noted, we use the FSKD-BCD algorithm with one iteration in our experiments. Appendix-A evaluates FSKD with more BCD iterations. Appendix-B makes a comparison between FSKD-BCD and FSKD-SGD.

3.3. Mergeable 1×1 conv-layer

Now we prove that the added 1×1 conv-layer can be merged into the previous conv-layer without introducing additional parameters and computation cost during inference.

Theorem 1. *A pointwise convolution with tensor $\mathbf{Q} \in \mathbb{R}^{n'_o \times n'_i \times 1 \times 1}$ can be merged into the previous convolution layer with tensor $\mathbf{W} \in \mathbb{R}^{n_o \times n_i \times k \times k}$ to obtain the **merged tensor** $\mathbf{W}' = \mathbf{Q} \circ \mathbf{W}$, where \circ is merging operator and $\mathbf{W}' \in \mathbb{R}^{n'_o \times n_i \times k \times k}$ if the following conditions are satisfied.*

- c1. *The output channel number of \mathbf{W} equals to the input channel number of \mathbf{Q} , i.e., $n_o = n'_i$.*
- c2. *No non-linear activation layer like ReLU [29] between \mathbf{W} and \mathbf{Q} .*

The pointwise convolution can be viewed as a linear combination of the kernels in the previous convolution layer. Due to the space limitation, we put the formal proof and the detailed form of the merging operator in Appendix-C. The number of output channels of \mathbf{W}' is n'_o , which is different from that of \mathbf{W} (i.e., n_o). It is easy to have the following corollary.

Corollary 1. *When the following condition is satisfied,*

- c3. *the number of input and output channels of \mathbf{Q} equals to the number of output channel of \mathbf{W} , i.e., $n'_i = n'_o = n_o$, $\mathbf{Q} \in \mathbb{R}^{n_o \times n_o \times 1 \times 1}$,*




Figure 2: Illustration of FSKD on filter pruning and network slimming. At each block, we copy weights of the unpruned part in teacher-net to student-net, and align the feature maps of student-net to those unpruned feature maps of teacher-net by adding a 1×1 conv-layer (red-color) with L2-loss. The added 1×1 layer can be merged into the previous conv-layer in student-net.

the merged convolution tensor \mathbf{W}' has the same parameters and computation cost as \mathbf{W} , i.e. both \mathbf{W}' , $\mathbf{W} \in \mathbb{R}^{n_o \times n_i \times k \times k}$.

This condition is required not only for ensuring the same parameter size and computing cost, but also for ensuring the output-size of current layer matching to the input-size of next layer so that these two layers are connectable.

4. Experiments

We perform extensive experiments on different image classification datasets to verify the effectiveness of FSKD on various student-net construction methods and its advantages over existing distillation methods in terms of both accuracy and speed. Student-nets can be obtained either by pruning based methods such as filter pruning [23] and network slimming [25], or by decomposition-based methods such as network decoupling [10]. We implement the code with PyTorch, and conduct experiments on a desktop PC with Intel i7-7700K CPU and one NVidia 1080TI GPU. The code will be made publicly available. For all experiments, the results are averaged over 5 trials of different randomly selected images.

4.1. Student-net from Pruning teacher-net

Filter Pruning

We first obtain the student-nets using the filter pruning method [23], which prunes out conv-filters according to the L_1 norm of their weights. The L_1 norm of filter weights are sorted and the smallest portion of filters will be pruned to reduce the number of filter-channels in a conv-layer. Figure 2 illustrates how FSKD works for block-level alignment in this case. Note that the number of channels in teacher-net may be different from that in student-net. However, we only match the *un-pruned* part of feature-maps in teacher-net to the feature maps in the student-net so that FSKD is applicable in this case.

	Acc.(%)	FLOPs($\times 10^8$)	Speedup	#Param($\times 10^6$)	Pruned
VGG-16	92.66	3.11	1.00 \times	15	-
Prune-A	85.42	2.06	1.51 \times	5.3	64%
Prune-B	47.90	1.33	2.34 \times	3.4	77%
Prune-C	13.05	1.09	2.85 \times	1.8	88%

Table 1: Prune-A/B/C by **filter pruning** of VGG-16 on CIFAR-10 and their accuracy, FLOPs, #parameters, etc.

	Acc. (%)	#Samples	Time (sec)
VGG-16	92.66	50000	
Prune-A + FSKD	92.37 \pm 0.24	100	4.8
Prune-A + FitNet	91.23 \pm 0.41	100	48.5
Prune-A + FSKD	92.46 \pm 0.15	500	25.5
Prune-A + FitNet	92.13 \pm 0.35	500	139.2
Prune-A + Fine-tuning	90.25 \pm 0.67	500	40.4
Prune-A + Full fine-tuning	92.54 \pm 0.33	50000	1059.6
Prune-B + FSKD	90.17 \pm 0.31	100	3.7
Prune-B + FitNet	88.76 \pm 0.51	100	60.1
Prune-B + FSKD	91.21 \pm 0.23	500	19.3
Prune-B + FitNet	90.68 \pm 0.47	500	157.1
Prune-B + Fine-tuning	83.36 \pm 0.89	500	50.3
Prune-B + Full fine-tuning	91.53 \pm 0.37	50000	1753.4
Prune-C + FSKD	89.55 \pm 0.35	100	7.4
Prune-C + FitNet	85.09 \pm 0.75	100	71.3
Prune-C + FSKD	90.41 \pm 0.31	500	33.5
Prune-C + FitNet	88.31 \pm 0.70	500	180.3
Prune-C + Fine-tuning	78.13 \pm 0.24	500	58.7
Prune-C + Full fine-tuning	90.77 \pm 0.33	50000	2592.3

Table 2: Performance comparison between FitNet, fine-tuning, FSKD by student-nets from **filter pruning** [23] of VGG-16 with pruning scheme A/B/C on CIFAR-10. ‘‘Full fine-tuning’’ uses full training data.

We make a comprehensive study of VGG-16 [32] on CIFAR-10 dataset to evaluate the performance of FSKD along with three different pruning settings. *First*, we follow the original pruning scheme of [23] and obtain **Prune-A**. *Second*, we propose another more aggressive pruning scheme named **Prune-B**, which prunes 10% more filters in the aforementioned layers, and also pruned 20% filters for the remaining layers. *Third*, since previous works show that one time extremely pruning may yield the pruned network unable to recovery from fine-tuning, while the iteratively pruning and fine-tuning procedure is observed effective to obtain extreme model compression [11, 23, 25], we propose **Prune-C** which iteratively runs the pruning and FSKD procedure as described in Appendix-D for 2 iterations to achieve higher compression rate. Table 1 lists the accuracy, FLOPs and #parameter information for three student-nets obtained by these pruning schemes.

For the few-sample setting, we randomly select 100 (10 for each category) and 500 (50 for each category) images from the CIFAR-10 training set, and keep them fixed in all experiments. We evaluate 5 different randomly-selected im-




Figure 3: Accuracy vs #samples on CIFAR-10. Student-net Prune-A (left) Prune-B (right) by **filter pruning** [23].




Figure 4: Left: Layer-level output correlation between teacher-net and student-net before and after FSKD on student-nets (Prune-A) by **filter pruning** [23]. Right: Accuracy change during sequentially block-level alignment.

age set and report the mean and std of accuracy. Table 2 lists the results of different methods of recovering a pruned network, including FitNet [31], fine-tuning with limited data and full training data [23].

As shown in Table 2, our method is much more efficient and provides better accuracy recovery than both FitNet and the fine-tuning procedure adopted in [23], and also more robust with a different set of selected images. For instance, for Prune-B with only 500 samples, our method can recover the accuracy from 47.9% to 91.2% in 19.3s, while FitNet has to take 157.1s to recover the accuracy to 90.7%, and few-sample fine-tuning can only recover the accuracy to 83.4%. When full training set available, it takes about 30 minutes for full fine-tuning to reach similar accuracy as FSKD. This demonstrates the big advantages of FSKD over full fine-tuning based solutions.

Figure 3 further studies the performance versus different amount of training samples. Our method keeps outperforming FitNet under the same training samples. In particular, FitNet experiences a noticeable accuracy drop when the number of samples is less than 100, while FSKD can still recover the accuracy of the pruned network to a high level.

We further illustrate the per-layer (block) feature responses difference between teacher-net and student-net before and after using FSKD in Figure 4(a). Before applying FSKD, the correlation between teacher-net and student-net is broken due to aggressive compression. However, after FSKD, the per-layer correlations are mostly restored. This verifies the ability of FSKD for recovering lost informa-

Filter-prune-ratio & Method	Acc. before(%)	Acc. after(%)	FLOPs($\times 10^8$)	Speedup	#Param($\times 10^6$)	Pruned
VGG-19	93.38	-	7.97	1.00 \times	20	-
70% + FSKD	15.90	93.41 \pm 0.23	3.91	2.04 \times	2.2	89%
70% + FitNet	15.90	90.47 \pm 0.57	3.91	2.04 \times	2.2	89%
70% + Fine-tuning	15.90	62.86 \pm 2.85	3.91	2.04 \times	2.2	89%
ResNet-164	95.07	-	4.99	1.00 \times	1.7	-
60% + FSKD	54.46	94.19 \pm 0.21	2.75	1.82 \times	1.1	37%
60% + FitNet	54.46	88.94 \pm 0.66	2.75	1.82 \times	1.1	37%
60% + Fine-tuning	54.46	60.94 \pm 3.12	2.75	1.82 \times	1.1	37%
DenseNet-40	94.18	-	5.33	1.00 \times	1.1	-
60% + FSKD	88.24	93.62 \pm 0.09	2.89	1.84 \times	0.5	54%
60% + FitNet	88.24	91.37 \pm 0.17	2.89	1.84 \times	0.5	54%
60% + Fine-tuning	88.24	88.98 \pm 0.79	2.89	1.84 \times	0.5	54%

Table 3: Performance comparison between FSKD, FitNet and fine-tuning on different network structures obtained by **network slimming** [25] with 100 samples randomly selected from CIFAR-10 training set.

Filter-prune-ratio & Method	Acc. before(%)	Acc. after(%)	FLOPs($\times 10^8$)	Speedup	#Param($\times 10^6$)	Pruned
VGG-19	72.08	-	7.97	1.00 \times	20	-
50% + FSKD	9.24	71.98 \pm 0.15	5.01	1.60 \times	5.0	75%
50% + FitNet	9.24	69.52 \pm 0.43	5.01	1.60 \times	5.0	75%
50% + Fine-tuning	9.24	48.75 \pm 2.86	5.01	1.60 \times	5.0	75%
ResNet-164	76.56	-	5.00	1.00 \times	1.7	-
40% + FSKD	46.07	76.11 \pm 0.13	3.33	1.50 \times	1.5	14%
40% + FitNet	46.07	73.87 \pm 0.45	3.33	1.50 \times	1.5	14%
40% + Fine-tuning	46.07	57.45 \pm 1.94	3.33	1.50 \times	1.5	14%
DenseNet-40	73.21	-	5.33	1.00 \times	1.1	-
40% + FSKD	60.62	73.26 \pm 0.07	3.71	1.44 \times	0.71	36%
40% + FitNet	60.62	71.08 \pm 0.33	3.71	1.44 \times	0.71	36%
40% + Fine-tuning	60.62	62.36 \pm 0.97	3.71	1.44 \times	0.71	36%

Table 4: Performance comparison between FSKD, FitNet and fine-tuning on different network structures obtained by **network slimming** [25] with 500 samples randomly selected from CIFAR-100 training set.

tion. We also show the accuracy change during sequentially block-level alignment in [Figure 4\(b\)](#), which clearly demonstrate the effectiveness of our sequentially block-by-block update in the FSKD algorithm.

Network Slimming

We then study the student-net from another filter pruning method named network slimming [25], which removes insignificant filter channels and corresponding feature maps using sparsified channel scaling factors. Network slimming consists of three steps: sparse regularized training, pruning and fine-tuning. Here, we replace the time-consuming fine-tuning step with our FSKD, and follow the original paper [25] to conduct experiments to prune different networks on different datasets. The alignment framework is the same as the filter pruning case as shown in [Figure 2](#).

We apply FSKD on networks pruned from VGG-19, ResNet-164, and DenseNet-40 [16], on both CIFAR-10 and CIFAR-100 datasets. [Table 3](#) lists results on CIFAR-10, while [Table 4](#) lists results on CIFAR-100. Note that the

filter-prune-ratio (like 70% in [Table 3](#)) means the portion of filters that are removed in comparison to the total number of filters in the network. We also apply FSKD on networks pruned from VGG-A (or VGG-11) on ImageNet dataset, as shown in [Table 5](#).

The results show that that FSKD consistently outperforms FitNet and fine-tuning with a notable margin under the few-sample setting on all evaluated networks and datasets. This study demonstrates that FSKD is universally applicable to various network structure and pruning methods, and can recover the accuracy of the pruned network using few unlabeled samples to the same level of fine-tuning using fully annotated training dataset.

4.2. Student-net from Decomposing teacher-net

In this section, we apply FSKD on a decomposition-based method called network decoupling [10], which can decompose a regular convolution layer into the sum of several *depth-wise separable blocks*, where each such block consists of a depth-wise (DW) conv-layer and a point-wise

Filter-prune-ratio & Method	Acc. before(%)	Acc. after(%)	GFLOPs	Speedup	#Param($\times 10^6$)	Pruned
VGG-A	63.3	-	7.74	1.00 \times	132.9	-
50% + FSKD	13.8	62.5 \pm 0.2	5.41	1.43 \times	23.2	83%
50% + FitNet	13.8	58.3 \pm 0.4	5.41	1.43 \times	23.2	83%
50% + Fine-tuning	13.8	19.2 \pm 3.5	5.41	1.43 \times	23.2	83%

Table 5: Performance comparison between FSKD, FitNet and fine-tuning on VGG-A (VGG-11) network structures obtained by **network slimming** [25] with 1000 samples randomly selected from ImageNet training set.

	Acc. before(%)	Acc. after (%)	GFLOPs	Speedup	#Param*($\times 10^6$)	Pruned
VGG-16 (teacher)	68.4	-	15.47	1.00 \times	14.71	-
Decoupled ($T = 2$) + FSKD	0.24	62.7 \pm 0.2	3.76	4.11 \times	3.35	77.2%
Decoupled ($T = 3$) + FSKD	1.57	67.1 \pm 0.1	5.54	2.79 \times	5.02	65.8%
Decoupled ($T = 4$) + FSKD	54.6	67.6 \pm 0.1	7.31	2.12 \times	6.69	54.5%
ResNet-18 (teacher)	67.1	-	1.83	1.00 \times	11.17	-
Decoupled ($T = 2$) + FSKD	0.21	49.5 \pm 0.5	0.55	3.33 \times	2.69	75.9%
Decoupled ($T = 3$) + FSKD	3.99	61.9 \pm 0.3	0.75	2.44 \times	3.95	64.6%
Decoupled ($T = 4$) + FSKD	26.5	65.1 \pm 0.1	0.95	1.92 \times	5.20	53.4%
Decoupled ($T = 5$) + FSKD	53.6	66.3 \pm 0.1	1.15	1.60 \times	6.46	42.2%

Table 6: Performance of FSKD on different student-nets obtained by **network decoupling** [10] VGG-16 and ResNet-18 with different parameters T on **ImageNet** dataset. “*” here means that parameters from FC-layer are not counted, only those from conv-layers are counted, since decoupling only handles the conv-layers.

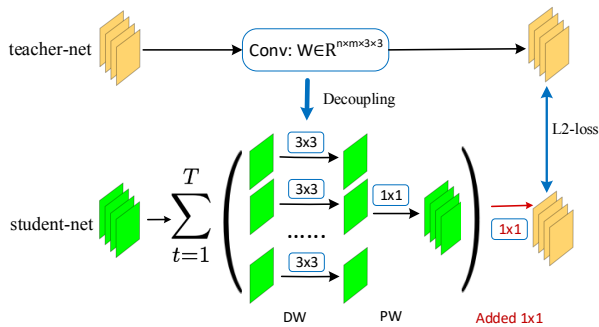


Figure 5: Illustration of FSKD on network decoupling. At each block of teacher-net, we decouple regular-conv into a sum of depthwise + pointwise conv-layers as the block of student-net, and align the feature maps of student-net to that of teacher-net by adding a 1×1 conv-layer (red-color) with L2-loss. The added layer can be merged into previous the pointwise layer in student-net.

(PW, 1×1) conv-layer. The compression ratio increases as the number (T) of such blocks decreases, but the accuracy of the compressed model will also drop. Since each decoupled block ends with a 1×1 convolution, we can apply FSKD at the end of each decoupled block. **Figure 5** illustrates how FSKD works for block-level alignment in this case.

Following [10], we obtain student-nets by decoupling VGG-16 and ResNet-18 pre-trained on ImageNet with different T values. We evaluate the resulted network performance on the validation set of the ImageNet classification task. We randomly select one image from each of the 1000 classes in ImageNet training set to obtain 1000 samples as

our FSKD training set. **Table 6** shows the top-1 accuracy of student-net before and after applying FSKD on VGG-16 and ResNet-18.

It is quite interesting to see that when T is small, we can recover the accuracy of student-net from nearly random guess (0.24%, 0.21%) to a much higher level (62.7% and 49.5%) with only 1000 samples. One possible explanation is that the highly-compressed networks still inherit some representation power from the teacher-net i.e., the depth-wise 3×3 convolution, while lacking the ability to output meaningful predictions due to the degraded and inaccurate 1×1 convolution. The FSKD calibrates the 1×1 convolution by aligning the block-level responses between teacher-net and student-net so that the lost information in 1×1 convolution is compensated, and reasonable recovery is achieved.

In all the other cases, FSKD can recover the accuracy of a highly-compressed network to be comparable with the original network. This shows that FSKD can be applied to student-net compressed by network decomposition, and that FSKD can achieve great performance on large and difficult dataset such as ImageNet.

5. Analysis and Discussion

FSKD with Arbitrary Data

In this section, we try to answer the following question: is FSKD totally label-free? For example, is FSKD still valid if the available few samples are arbitrary images and the teacher network never sees these images before? To answer

	Acc. (%)	#Samples
VGG-16	92.66	50000 (CIFAR-10)
Prune-B + FSKD	90.17±0.31	100 (CIFAR-10)
Prune-B + FSKD	90.15±0.31	100 (CIFAR-100)
Prune-B + FSKD	91.21±0.23	500 (CIFAR-10)
Prune-B + FSKD	91.20±0.25	500 (CIFAR-100)

Table 7: Performance comparison between FSKD using samples from CIFAR-10 and CIFAR-100. Student-nets from **filter pruning** [23] of VGG-16 with pruning scheme B on CIFAR-10.

this question, we evaluate FSKD’s performance on VGG-16 model trained on CIFAR-10 and compressed using filter pruning (prune-B), with the few samples for FSKD are randomly selected from CIFAR-100 instead of CIFAR-10.

As shown in Table 7, there is no statistical difference in accuracy between FSKD using data from CIFAR-10 or CIFAR-100. This shows that FSKD aligns the student-net with the teacher-net without any information about the labels of the data. Even if the input images are of classes it has never seen before (CIFAR-100 does not include classes in CIFAR-10), FSKD can still recover the student network to the same accuracy level. This further demonstrates FSKD’s potential in situations where only a few samples of unlabeled data are available.

What if student-nets are hand-designed?

Our previous experiments construct student-nets by pruning or decomposing teacher-net, and then apply FSKD to boost their performance. People may be interested in the problem “what if student-nets are hand-designed with random initialization”. In fact, there are two existing works [20, 26] making some pioneer trials on this topic with specific methods for “pseudo” examples generation. Here we conduct experiments to compare our method to these two methods under the same few-sample setting on the same dataset MNIST, for a fair comparison.

Due to different network structures used in these two methods, we make a separate comparison. For [20], the teacher-net has 3 conv-layers followed by 2 fully-connected layers. For [26], the teacher-net is a standard LeNet-5. For both cases, the student-net is the “half-sized” to that of the corresponding teacher-net in terms of the number of feature map channels per conv-layers. As the channel number between student-net and teacher-net is different, we adopt the same strategy as in Figure 2 for filter pruning. That means, the student-net only corresponds to the un-pruned part of the teacher-net, which is obtained the same as [23]. One difference is that we did not copy the weight from un-pruned part of teacher-net to the student-net, while keeping the weight of student-net randomly initialized. For both cases, we compared our FSKD with (1) standard SGD trained on few samples with labeled loss; (2) method from [20] or [26] under the same setting; (3) the FitNet method trained on few

#labeled data	10	20	50	100	200	all-meta-data
SGD	37.91	46.0	66.0	78.3	86.7	-
[20]	44.1	53.9	70.4	80.0	86.6	-
FitNet	86.1	92.3	94.5	96.0	96.5	-
FSKD	94.4	96.5	97.0	97.5	97.8	-
SGD	57.1	68.3	81.3	85.8	89.7	-
[26]	-	-	-	-	-	92.5
FitNet	90.3	94.2	96.1	96.7	97.3	-
FSKD	95.5	97.2	97.6	98.0	98.1	-

Table 8: Performance of FSKD on hand designed student-nets with random initialization, compared with previous works [20, 26].

samples. In order to better simulate the few-sample setting, we do not apply data augmentation to the training set. We randomly pick 10, 20, 50, 100 and 200 samples from the MNIST training set and keep these few-sample sets fixed across this study. Table 8 lists the comparison results. It shows that SGD with few samples performs the worst, while [20] performs better on the same settings than SGD (still worse in the case of 200 samples). The data-free method [26] performs better than SGD. On both cases, FitNet shows much better performance than SGD and the two compared methods, while our FSKD further outperform FitNet with a noticeable gap, where the gap becomes smaller and smaller when the number of samples increases. This may be due to the following reason. FSKD can be viewed as a special case of FitNet. FitNet optimizes all the weights between teacher-net and student-net using standard SGD algorithm, while FSKD optimizes only the weights from added 1×1 conv-layers in student-net with the BCD algorithm. The BCD algorithm is more sample-efficient than the SGD based algorithm so that FSKD performs both much more efficient and accurate than FitNet on few-sample settings. When training samples used are increased, FSKD will converge to FitNet in the end.

6. Conclusion

We proposed a novel yet simple method, namely few-sample knowledge distillation (FSKD) for efficient network compression, while “efficient” lies in both *training/processing efficiency* and *label-free sample efficiency*. FSKD works for student-nets constructed by either pruning or decomposing teacher-nets with different methods. It demonstrates great efficiency over fine-tuning based solution and advantages over traditional knowledge distillation methods like FitNet by a large margin in the few-sample setting, with extra merits that FSKD is totally label-free in the optimization.

Acknowledgements: Changshui Zhang is funded by NSFC under Grant No. 61876095 & 61751308, and Beijing Academy of Artificial Intelligence (BAAI).

References

- [1] Jimmy Ba and Rich Caruana. Do deep nets really need to be deep? In *NIPS*, 2014. 1
- [2] Evgeniy Bart and Shimon Ullman. Cross-generalization: Learning novel classes from a single example by feature replacement. In *CVPR*. IEEE, 2005. 3
- [3] Kartikeya Bhardwaj, Naveen Suda, and Radu Marculescu. Dream distillation: A data-independent model compression framework. *arXiv preprint arXiv:1905.07072*, 2019. 3
- [4] Cristian Bucila, Rich Caruana, Alexandru Niculescu-Mizil, et al. Model compression. In *SIGKDD*. ACM, 2006. 1
- [5] Hanting Chen, Yunhe Wang, Chang Xu, Zhaohui Yang, Chuanjian Liu, Boxin Shi, Chunjing Xu, Chao Xu, and Qi Tian. Data-free learning of student networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3514–3522, 2019. 3
- [6] Tianqi Chen, Ian Goodfellow, Jonathon Shlens, et al. Net2net: Accelerating learning via knowledge transfer. In *ICLR*, 2016. 3
- [7] Emily Denton, Zaremba, Yann Lecun, et al. Exploiting linear structure within convolutional networks for efficient evaluation. In *NIPS*, 2014. 1, 2
- [8] Li Fei-Fei, Rob Fergus, Pietro Perona, et al. One-shot learning of object categories. *IEEE Trans PAMI*, 2006. 3
- [9] Chelsea Finn, Pieter Abbeel, Sergey Levine, et al. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017. 3
- [10] Jianbo Guo, Yuxi Li, Weiyao Lin, Yurong Chen, and Jianguo Li. Network decoupling: From regular to depthwise separable convolutions. In *BMVC*, 2018. 4, 6, 7
- [11] Song Han, Huizi Mao, Bill Dally, et al. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. In *NIPS*, 2016. 1, 5
- [12] Song Han, Jeff Pool, John Tran, William Dally, et al. Learning both weights and connections for efficient neural network. In *NIPS*, 2015. 2
- [13] K. He, X. Zhang, J. Sun, et al. Deep residual learning for image recognition. In *CVPR*, 2016. 1
- [14] Geoffrey Hinton, Li Deng, Dong Yu, et al. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal Processing Magazine*, 29(6), 2012. 1
- [15] G. Hinton, O. Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 1, 2
- [16] Gao Huang, Zhuang Liu, Kilian Q Weinberger, and Laurens van der Maaten. Densely connected convolutional networks. In *CVPR*, 2017. 6
- [17] Zehao Huang and Naiyan Wang. Like what you like: Knowledge distill via neuron selectivity transfer. *arXiv preprint arXiv:1707.01219*, 2017. 1, 3
- [18] M. Jaderberg, A. Vedaldi, A. Zisserman, et al. Speeding up convolutional neural networks with low rank expansions. In *BMVC*, 2014. 1, 2
- [19] Y. Kim, E. Park, S. Yoo, et al. Compression of deep convolutional neural networks for fast and low power mobile applications. In *ICLR*, 2016. 1, 2
- [20] Akisato Kimura, Zoubin Ghahramani, Koh Takeuchi, et al. Few-shot learning of neural networks from scratch by pseudo example optimization. In *BMVC*, 2018. 3, 8
- [21] A. Krizhevsky and G. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012. 1
- [22] Brenden Lake, Ruslan Salakhutdinov, Jason Gross, et al. One shot learning of simple visual concepts. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 33, 2011. 3
- [23] Hao Li, Asim Kadav, Durdanovic I, et al. Pruning filters for efficient convnets. *ICLR*, 2017. 1, 2, 4, 5, 8
- [24] Muyang Li, Ji Lin, Yaoyao Ding, Zhijian Liu, Jun-Yan Zhu, and Song Han. Gan compression: Efficient architectures for interactive conditional gans. *arXiv preprint arXiv:2003.08936*, 2020. 1
- [25] Zhuang Liu, Jianguo Li, Zhiqiang Shen, et al. Learning efficient convolutional networks through network slimming. In *ICCV*, 2017. 1, 2, 4, 5, 6, 7
- [26] Raphael Gontijo Lopes, Stefano Fenu, Thad Starner, et al. Data-free knowledge distillation for deep neural networks. *arXiv preprint arXiv:1710.07535*, 2017. 3, 8
- [27] J. Luo, J. Wu, W. Lin, et al. Thinet: A filter level pruning method for deep neural network compression. In *ICCV*, 2017. 2
- [28] Tomáš Mikolov, Martin Karafiát, Lukáš Burget, et al. Recurrent neural network based language model. In *INTER-SPEECH*, 2010. 1
- [29] Vinod Nair and Geoffrey Hinton. Rectified linear units improve restricted boltzmann machines. In *ICML*, 2010. 4
- [30] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *ICLR*, 2017. 3
- [31] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, et al. Fitnets: Hints for thin deep nets. In *ICLR*, 2015. 1, 3, 5
- [32] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 5
- [33] Suraj Srinivas and Francois Fleuret. Knowledge transfer with jacobian matching. *arXiv preprint arXiv:1803.00443*, 2018. 1, 3
- [34] Oriol Vinyals, Charles Blundell, Tim Lillicrap, et al. Matching networks for one shot learning. In *NIPS*, 2016. 3
- [35] Yangyang Xu and Wotao Yin. A globally convergent algorithm for nonconvex optimization based on block coordinate update. *Journal of Scientific Computing*, 72(2):700–734, 2017. 4
- [36] Junho Yim, Donggyu Joo, Jihoon Bae, et al. A gift from knowledge distillation: Fast optimization, network minimization and transfer learning. In *CVPR*, 2017. 3
- [37] X. Zhang, J. Zou, J. Sun, et al. Accelerating very deep convolutional networks for classification and detection. *IEEE TPAMI*, 38(10), 2016. 1, 2
- [38] Ying Zhang, Tao Xiang, Timothy Hospedales, et al. Deep mutual learning. In *CVPR*, 2018. 3