

Knowledge Representing: Efficient, Sparse Representation of Prior Knowledge for Knowledge Distillation

Junjie Liu¹, Dongchao Wen¹, Hongxing Gao¹, Wei Tao¹,

¹Canon Information Technology (Beijing) Co., LTD

{liujunjie, wendongchao, gaohongxing, taowei}@canon-ib.com.cn

Tse-Wei Chen², Kinya Osa², Masami Kato²

²Device Technology Development Headquarters, Canon Inc.

twchen@ieee.org

Abstract

Despite the recent works on knowledge distillation (KD) have achieved a further improvement through elaborately modeling the decision boundary as the posterior knowledge, their performance is still dependent on the hypothesis that the target network has a powerful capacity (representation ability). In this paper, we propose a knowledge representing (KR) framework mainly focusing on modeling the parameters distribution as prior knowledge. Firstly, we suggest a knowledge aggregation scheme in order to answer how to represent the prior knowledge from teacher network. Through aggregating the parameters distribution from teacher network into more abstract level, the scheme is able to alleviate the phenomenon of residual accumulation in the deeper layers. Secondly, as the critical issue of what the most important prior knowledge is for better distilling, we design a sparse recoding penalty for constraining the student network to learn with the penalized gradients. With the proposed penalty, the student network can effectively avoid the over-regularization during knowledge distilling and converge faster. The quantitative experiments exhibit that the proposed framework achieves the state-of-the-arts performance, even though the target network does not have the expected capacity. Moreover, the framework is flexible enough for combining with other KD methods based on the posterior knowledge.

1. Introduction

The deep neural network has achieved the significant improvement in different fields with years, but it also requires higher computational and memory costs. For the purpose to apply these networks to the real-time industrial tasks, the neural network compression [4] is arguably




Figure 1. The pipeline of knowledge representing algorithm: The prior knowledge in teacher network is represented by the knowledge aggregation scheme into higher abstract level. Then the sparse recoding penalty is further used to regularize the gradients in student network for efficient learning these prior knowledge.

the most crucial strategy. As for the network compression problem, the typical solutions are designed to *slim* [27, 37] the network directly, or quantify their parameters distributions [14, 16, 25], and filter the redundant layer dimensions [9, 21].

In contrast to these techniques which aim at directly compressing the network while preserving its performance as much as possible, an alternative solution is to preset a smaller target network as the student, and employ the knowledge from the larger network as teacher to improve student’s performance. Therefore, knowledge distillation [11] (KD) is proposed. The KD mainly assumes the samples distribution is anisotropy [1], but annotations of the samples are not able to represent this intrinsic. Based on the hypothesis, these methods evaluate the samples in the teacher network to produce the decision boundary as a strong posterior distribution, and then use to regularize the gradients optimization of student network. While this helps prevent the student network from being over-fitting, the extra risk of non-convergence is introduced.

A possible solution is to refine the posterior distribution from the teacher network, in order to provide more valuable knowledge for better distilling. The Neuron Selectively Transfer (NST) [12] is proposed to align the distribution selectively with the Maximum Mean Discrepancy (MMD) metric, and the generative adversarial network with KD (KDGAN) [30] is further used to produce a more robust decision boundary for student classifier. However, considering the student network which contains a very limited capacity - the representation ability, this limitation gradually becomes a major bottleneck in network training to further improve the performance of knowledge distillation. In a word, the fine-grained posterior distribution is usually underemployed.

With the constraint from network capacity, an instinctive approach is to introduce the parameters distribution [6] from the teacher network as the prior knowledge [26, 34]. For the typical one, Romero et al. [26] constructs the Hint layer to estimate a parameters distribution with less filter numbers, through using the intermediate features representation of the teacher, and it uses these knowledge to guide the update of student parameters. However, the Hint layer suffers from the over-regularization if the teacher network is too deep.

In this paper, we produce a KD solution mainly focusing on modeling the prior knowledge, while avoiding the negative impacts from over-regularization, and the solution is flexible enough, for combining with other KD methods based on the posterior knowledge. Specially, we propose a knowledge representing (KR) framework, which aims at representing the prior knowledge at more abstract level, and taking full advantage of these knowledge. For answering the question of how to represent the prior knowledge, a knowledge aggregation scheme is firstly suggested. Inspired by the theory of optimal transportation [23, 24], the scheme is designed to alleviate the phenomenon of residual accumulation in the deeper layers. Then, as for the most critical issue of what the dominant prior knowledge is for better distilling, a sparse recoding penalty is proposed. Through employing a learnable threshold in the penalty, it can enhance the gradients of dominant neurons and smooth inactive ones. With these two proposed terms, the proposed framework can prompt the student network to preserve the key features of teacher network, even without a strong representation ability.

Our paper makes the following contributions:

- A new penalty is proposed to constrain the optimization of knowledge distillation. It helps the student network to avoid the over-regularization and converge faster. Moreover, the penalty can be further applied on other network optimization problems.
- A new scheme is suggested for aggregating the prior

knowledge. It is able to produce more abstract features and alleviate the phenomenon of residual accumulation.

- According to the proposed framework, the more flexible architecture is allowable for both teacher and student network, without the constraints from model depths or filter scales.

2. Related Work

The latest deep networks are usually accompanied with carefully designed modules [7, 8] and enormous parameters. Though the performance of targeted tasks is obviously being improved, the computation and memory cost gradually become the challenge to employ these networks in real-life applications [16, 28]. Comparing to the traditional neural network compression methods [4] which focus on compressing the original network directly, a solution with the knowledge distillation to compress the deep network attracts more attention from research community in recent years, such as in the tasks of image recognition [34], object detection [2, 22], or recommender systems [39], as the flexibility to obtain an arbitrary architecture of target network. In summary, the KD methods can be categorized into two main groups:

1) *Distilling the posterior distribution from training data*: Considering the possibility to extract the knowledge in an ensemble (teacher) into a single model (student), Hinton et al. [11] introduces the idea of knowledge distillation as a regularizer. Through employing a penalized version [10, 13, 33] of final features of the teacher network, a joint learning is processed with the knowledge from posterior distribution. For refining the posterior distribution to provide more valuable knowledge, the Neuron Selectively Transfer (NST) [12] is proposed to align the distribution selectively with the Maximum Mean Discrepancy (MMD) metric. Furthermore, considering the sample bias is unavoidable, the generative adversarial networks for knowledge distillation (KDGAN) [30] is further used to produce a more robust posterior distribution for student classifier. However, these methods haven't take the capacity of student network into consideration, so the fine-grained posterior distribution is underemployed.

2) *Distilling the prior distribution from model parameters*: An alternative approach is to introduce the parameters distribution from teacher network as the prior knowledge [26, 34, 36]. Romero et al. [26] designs the Hint layer to estimate the parameters distribution by using the intermediate hidden layers from the teacher, and used the Hint layer to guide the distillation. Net2Net [3] suggests a function-preserving transform for extracting the prior knowledge from teacher network to initialize the parameters of the student network. And Yim et al. [34] suggests a representation

operator named FSP matrix. It uses not only the parameters distribution but the intermediate features from the neighbor layers. However, these methods either are constrained by the depth of teacher network, or suffer from the over-regularization.

3. Method

For obtaining a student network that faithfully preserves the key representation ability of the teacher, Sec. 3.1 presents the objective function of the knowledge representing framework. Accordingly, we firstly answer the key problem of what the most important prior knowledge is for distilling in Sec. 3.2, through introducing the mathematical expression of the sparse recoding penalty. Then, we suggest how to represent the prior knowledge from the teacher network, with a knowledge aggregation scheme in Sec. 3.3. Finally, Sec. 3.4 shows the optimization procedure of the objective function.

3.1. Knowledge Representing

As one of the most typical feature representation technique, the deep model produces the decision boundary through modeling the data distribution with the parameters in layers. Given a trained decision boundary $y^t(x, W^t)$, where y^t is generated by teacher network with data distribution x and the parameters W^t , the objective of knowledge distillation is to find the parameters W^s for the student network. Specially, with the W^s and x , the y^s from student network is jointly optimized with the y^t . Through minimizing the dissimilarity of two decision boundaries, the objective function of knowledge distillation is defined as:

$$\arg \min_{W^s} \sum_{i=1}^N \mathcal{L}(y_i^t(x_i, W^t), y_i^s(x_i, W^s)) + \lambda \Phi(W^s) \quad (1)$$

Algorithm 1: Training the Knowledge Representing Algorithm

Input: Weights W^t of teacher network and W^s of student network

Output: Aggregation knowledge \tilde{W}^t and optimized weights of student network W^s

Initialization: $\tilde{W}^t, W^s, \text{MaxIter}$;

while less than the *MaxIter* **do**

- Optimizing** \tilde{W}^t **with** W^t ;
Knowledge aggregation for teacher in Eq. 9 ;
- Optimizing** W^s **with** \tilde{W}^t ;
Sparse recoding for student in Eq. 10 ;

end

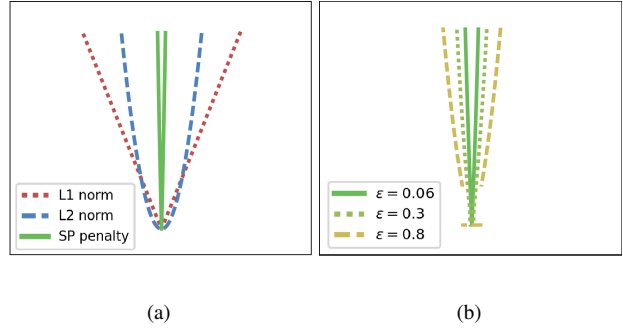


Figure 2. The properties of proposed sparse recoding penalty (best viewed in color): (a) the sparse recoding penalty is able to approximate more strict sparseness; (b) the properties with different setting of ϵ .

where \mathcal{L} represents the metric for evaluating the similarity between the y^t and y^s , and the cross entropy, KDGAN [30], or NST [12] are allowable. Different from the KD methods only evaluating the decision boundary, we further introduce a penalty $\Phi(\cdot)$ in Eq. 1, in order to measure the representation ability of student network. However, if the representation ability of student network is weak, the fine-grained posterior distribution will be underemployed. Then, we extend the objective function Eq. 1 through further introducing the prior knowledge W^t from the teacher network, and the objective function is:

$$\arg \min_{W^s, \tilde{W}^t} \sum_{i=1}^N \mathcal{L}(y_i^t(x_i, W^t), y_i^s(x_i, W^s)) + \mathcal{L}(\tilde{W}^t, W^s) + \gamma \Psi(\tilde{W}^t, W^t) + \lambda \Phi(W^s) \quad (2)$$

Instead of directly employing the parameter distributions W^t from the teacher network as prior knowledge, we firstly represent these distributions as more abstract level, and a knowledge aggregation scheme $\Psi(\cdot)$ is suggested to aggregate W^t into the \tilde{W}^t . With the prior knowledge \tilde{W}^t , the $\mathcal{L}(\tilde{W}^t, W^s)$ is used to guide the update of parameters distributions W^s for the student. Moreover, we propose a sparse recoding penalty to specify the $\Phi(\cdot)$. Through enhancing the magnitude of dominant gradients and filtering the inactive ones, the optimizer no longer requires the parameters distribution W^s of student network to strictly close to the teacher one, and prompts the student network to firstly learn with the most valuable knowledge. In summary, the optimization procedure is represented in Algorithm 1, and we leave over the details in following sections.

3.2. Sparse Recoding Penalty

As demonstrated by previous works [31, 32, 38], prompting the neurons connection being sparse is beneficial for obtaining a well generalization ability. However, such penalties are designed to directly clip the parameters distribution, and the extra risk of over-regularization is introduced. After we analyze the distribution of prorogated gradients in the previous KD methods, we found that major reason for the convergence of oscillatory is that the gradients are not discriminative enough, especially in the student network with a weak representation ability.

Therefore, we propose a sparse recoding penalty $\Phi(\cdot)$, which can penalize the prorogated gradients during the training of deep network. Given an input parameters tensor W , it enhances the high gradients g_j of dominant neurons, and filters the low gradients of inactive neurons. The function is defined as:

$$\Phi(W) = \sum_j \Phi_0(g_j) \quad (3)$$

where

$$\Phi_0(g) = \begin{cases} \frac{1}{\varepsilon}(|g| + g^2), & \text{if } |g| \geq \varepsilon \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

where $\Phi_0(\cdot)$ is a piecewise function that enhances the gradients when $|g| \geq \varepsilon$, and smooths the $|g|$ by zero in others. The ε is a learnable threshold within the update of gradient optimization, and it is initialized with the mean value of parameters distribution. For fairly comparing with other penalties, the Fig.2 shows the curves of $\Phi(\cdot)$ by comparing with the L_1 and L_2 norms. It exhibits that $\Phi(\cdot)$ is a more strict sparse constraint. Moreover, with different parameter setting, properties of the sparse recoding penalty are shown in the figure, and we leave over the further discussion in experiments.

3.3. Deep Knowledge Aggregation

For representing the prior knowledge as more abstract level, we design a deep knowledge aggregation scheme through stacking the neighbor layers in a very deep network. Specially, with the analysis of prior knowledge distilling in previous methods, we notice that the optimization errors between two networks will be accumulated from layers, since the higher layer in teacher network usually contains a strong representation ability. However, the situation is simply regarded as the phenomenon of gradient vanishing, and cause an over-regularization if the teacher network is too deep. So we name this phenomenon as the residual accumulation, and the proposed scheme will mainly considers this phenomenon. Based on the theory of optimal transportation

[23, 24], the scheme try to reduce the residual accumulation during gradient optimization, through minimizing the inter-domain transportation cost. Given a P_1 and P_2 being two distribution space with probability measures μ and ν respectively, the transportation T preserving $P_1 \rightarrow P_2$ has equal total measure

$$\mu(T(p_1)) = \nu(p_2) \quad (5)$$

where p_1 and p_2 is any measurable subset of P_1 and P_2 . Then the total transportation cost for sending $p_1 \subset P_1$ to $p_2 \subset P_2$ by transportation cost $\tau(p_1, p_2)$ can be defined by

$$\min_{T: P_1 \rightarrow P_2} \int_{P_1} \tau(p_1, T(p_1)) d\mu(p_1) \quad (6)$$

With minimizing the total transportation cost, the distribution P_2 progressively approximates P_1 on measures μ . Assuming a series of neighbor layers ℓ_k, \dots, ℓ_n as set L_k^n , for sending parameter distribution $W_{L_k^n}$ to \tilde{W} with measurable subset $w \subset W_{L_k, \dots, \ell_n}$, the deep knowledge aggregation scheme merges the neighbouring layers to form the higher abstract parameters knowledge. In this case, the function $\Psi(\cdot)$ is formulated as

$$\Psi(\tilde{W}, W_{L_k^n}) = \min_{T: W_{L_k^n} \rightarrow \tilde{W}} \int_{W_{L_k^n}} \tau(w, T(w)) d\mu(w) \quad (7)$$

3.4. Optimization

Instead of directly optimizing the proposed objective function, we design an joint optimization method as the alternative solution. In details, our method uses two stages optimization to alternately solve the Eq. 2.

Optimizing \tilde{W}^t with $W_{L_k^n}^t$ Given an elaborate teacher network with parameter distribution $W_{L_k^n}^t$, we first aggregate the knowledge \tilde{W}^t with $T(W_{L_k^n}^t)$ in here as:

$$\begin{aligned} \arg \min_{\tilde{W}^t} \sum_{i=1}^N \mathcal{L}(y_i^t(x_i, W^t), y_i^s(x_i, W^s)) \\ + \gamma \int_{W_{L_k^n}^t} \tau(w^t, T(w^t)) d\mu(w^t) \end{aligned} \quad (8)$$

As the Eq. 8 involves a transportation cost and the definition of probability measures, it is difficult to directly integrate with gradient descent optimizer. In this case, we use the feature representation F_W as an approximation probability measures, which means the set of features maps F generated by parameters set W . If the transportation cost $\tau(\cdot)$ is defined as the simple L_2 distance, we revise the Eq. 8 as:

$$\arg \min_{W^t} \sum_{i=1}^N \mathcal{L}(y_i^t(x_i, W^t), y_i^s(x_i, W^s)) + \gamma \mu(W_{L_k^n}^t) \|F_{W_{L_k^n}^t} - F_{\tilde{W}^t}\|_2 \quad (9)$$

where γ is a predefined parameter to control the penalty from optimal transportation. The $\mu(W_{L_k^n}^t)$ as a measures function is used to penalize more on the layer with higher accumulation error, and the standard deviation is employed here. Moreover, we remove the part of terms during the derivation for Eq. 8 for fast computation. Then, the solution of \tilde{W}^t can be obtained by gradient descent optimization.

Optimizing W^s with \tilde{W}^t Given an aggregate knowledge \tilde{W}^t , our goal here is further to solve the W^s on student network with sparse recoding penalty, as:

$$\arg \min_{W^s} \mathcal{L}(\tilde{W}^t, W^s) + \lambda \sum_j \Phi_0(g_j^s) \quad (10)$$

where $\Phi_0(g_j^s)$ is designed for prompting the student network to firstly learn with the penalized gradients, and the parameter λ is predefined to control the importance of the sparse recoding penalty.

Instead of directly solving the global optimum for objective function Eq. 2, the two sub-objective functions Eq. 9 and Eq. 10 are designed to overcome the conflict between optimizing the prior knowledge and posterior knowledge simultaneously. Through alternatively minimizing the distribution dissimilarity $\mathcal{L}(\tilde{W}^t, W^s)$ and $\mathcal{L}(y^t, y^s)$, the optimization for Eq. 2 is regarded as an joint optimization procedure. Once the posterior knowledge is dominant during optimization, the optimizer for prior knowledge will penalize the total loss more, and the opposite is also. The gradient is only allowed to descend on the direction that makes both two optimizers are optimal.

4. Experiments

In this section, we evaluate the proposed knowledge distillation framework with several benchmark datasets. For the base of experiments, we use the deep residual network [8] as the network architecture, and the excerpt of the proposed framework in this architecture is shown in Fig. 3. The c in residual module means the number of aggregated convolution layers. For the problem of optimizing these layers with different spatial scales, the identity mapping (ID) layer [35] is employed also. To ensure a fair comparison, the same data augment strategies are used. Moreover, we employ the similar settings of learning rates, optimization iterations and computation precision (32 float points). The implementation details will be shown in corresponding subsections.

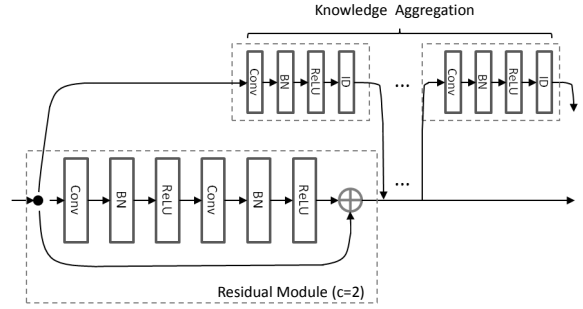


Figure 3. The excerpt of proposed framework on the residual network.

In Sec. 4.1, through comparing with the typical penalties, the property of the sparse recoding penalty is analysed. Then, through comparing with the state-of-the-arts, we evaluate the performance of student networks in general image recognition tasks, and further explore their generalization ability in a revised dataset TCIFAR-100, as described in Sec. 4.2. Finally, the discussions about the optimization procedure of the proposed framework is shown in Sec. 4.3.

4.1. Analysis of Proposed Penalty

As for the sparse recoding penalty, its property through comparing with typical methods is analysed, and we further explore the reason of why the proposed penalty is able to boost the convergence of knowledge distilling. Based on the experiment result, we address that the proposed penalty can be applied on other network optimization problems if the gradients distribution is not discriminative enough.

Penalty Property Given a specific parameters distribution, the traditional penalties [31, 38] form a convex function and obtain the maximal reward in the unique extreme. It penalizes the parameter with higher value to reduce the total loss, for encouraging the value of parameter to close to 0. In contrast to these methods, the sparse recoding penalty is designed to penalize the gradients directly. For the propagated gradients, it filters the gradients with an equal reward within the learnable threshold, in order to slow down the update of inactive neurons. For the gradients out of the threshold, it boosts the update to highlight the dominant neurons. For validating our hypothesis, we visualize the convolutional kernels with the constraint by different penalties in image recognition tasks. The Fig. 4 shows that the sparse recoding penalty can prompt the parameters distribution of the network to be more sparse, through directly regularizing the optimized gradients.

Convergence We have observed fast convergence in our experiment result. In Fig. 5, it illustrates the training loss




Figure 4. With the penalized gradients, the sparse recoding penalty is also to produce more discriminative parameters distribution (best see in color).

on MNIST over the beginning 20,000 iterations. The student network with sparse recoding penalty is better than the traditional penalties. We think one possible reason is that the proposed penalty is designed to penalize the gradients firstly, so it can produce a bigger step for gradients descent in the beginning of network training. Moreover, we evaluate the different types for initializing the parameters distribution in the experiment, and we also found the similar conclusion.

4.2. Performance Analysis

In this section, we firstly conduct the experiments in the image recognition task on CIFAR-10, CIFAR-100 [18] and ILSVRC 2012 [5], in order to evaluate the performance of the proposed knowledge representing framework with the state-of-the-arts. Then, we design a TCIFAR-100 based on CIFAR-100, for further verifying their generalization ability. As the focus of this experiment is analysing the performance of student network with a small capacity, so we reserve the comparison on different tasks as future works.




Figure 5. Convergence speed; the traditional penalties [31, 38] and the sparse recoding penalty (best see in color).

4.2.1 CIFAR-10

The CIFAR-10 is an image recognition dataset [18] which includes 50,000 training images and 10,000 test images, and per training class has 5,000 images while test class has 1000 images. For all images, they store in RGB format with size of 32×32 . We use a trained teacher network with 26 layers, which is structured as 5 residual modules. For student network, it contains 8 layers with 2 residual modules, which has roughly 1/3 parameters of the teacher. In details, with the same parameters settings and training strategies, we reduce about 1/3 number of the filters on each layer for the student network, in order to evaluate the case if the target network contains a weak representation ability. And we set the c of knowledge aggregation as 3, which aggregates each three layer of teacher network into higher abstract level for one layer in student network.

	Accuracy	Params
Teacher ResNet-26	91.91	$\sim 0.36M$
Student ResNet-8 (Original)	87.91	$\sim 0.12M$
FitNet [26]	88.57	$\sim 0.12M$
FSP [34]	88.70	$\sim 0.12M$
Proposed-Dense	89.11	$\sim 0.09M$
Proposed	90.65	$\sim 0.09M$
NTS [12]	88.98	$\sim 0.12M$
KDGAN [30]	88.62	$\sim 0.12M$
Proposed + KDGAN [30]	91.35	$\sim 0.09M$

Table 1. ResNet-8 in CIFAR-10 Classification rates(%). Proposed: the KR framework. Proposed-Dense: the KR framework but removing the sparse recoding penalty.

In Tab. 1, it summarizes the obtained results. Based on the proposed framework, the student network which contains less parameters wins the methods [26, 34] focusing on prior knowledge with a significant improvement. For the state-of-the-arts [12, 30] by modeling the posterior knowledge, the proposed framework also achieves the comparable performance. For the self-comparison, we remove the sparse recoding penalty in KR framework and name it as the KR-Dense. And the experiment proves the importance to sparsely penalize the gradients during the distilling optimization, if the student network only has a small capacity. Besides, through combining with the KDGAN [30], a further improvement confirms that our method is flexible for the extension.

4.2.2 CIFAR-100

The CIFAR-100 is an augmented version of CIFAR-10. It contains the same amount of images and size of CIFAR-10, which includes 50,000 training images and 10,000 test images, so only has 100 samples per class. Similar the setting to CIFAR-10, we use a trained teacher network with

	Accuracy	Params
Teacher ResNet-32	64.06	$\sim 0.46\text{M}$
Student ResNet-14 (Original)	58.65	$\sim 0.19\text{M}$
FitNet [26]	61.28	$\sim 0.19\text{M}$
FSP [34]	63.33	$\sim 0.19\text{M}$
Proposed Method	63.95	$\sim 0.17\text{M}$
NTS [12]	63.78	$\sim 0.19\text{M}$
KDGAN [30]	63.96	$\sim 0.19\text{M}$
Proposed Method + KDGAN [30]	63.98	$\sim 0.17\text{M}$

Table 2. ResNet-14 in CIFAR-100 Classification rates(%). With the similar network architecture, we further reduce the output channels in each layer for saving the total parameters.

32 layers as 6 residual modules, and student is composed of 14 layers as 3 residual modules. Besides, the reduction of about 1/3 filter number is still used, and c is set as 3.

Tab. 2 shows results of student network with evaluated methods. Though the proposed method achieves the comparable performance than the state-of-the-arts [12, 30] with less parameters, the improvement for our method is not obvious. We think one possible reason is that the ResNet-14 has a stronger representation ability than the ResNet-8.

4.2.3 ILSVRC 2012

The ILSVRC 2012 classification challenge involves the recognition task to classify one image into 1,000 leaf-node categories in the ImageNet hierarchy [19]. It has about 1.2 million images for training, 50,000 for validation and 100,000 testing images. Although training the very deep network on such enormous datasets to achieve satisfied performance has been a solvable issue, how to obtain the comparable performance with a tiny network by the knowledge distillation still confuses the research community, especially for the methods [22, 26, 34] with prior knowledge. We think the major reason is that the depth of teacher network in ILSVRC 2012 is very deep, so the student network in these methods seriously suffers from the over-regularization.

Tab. 3 shows the errors of Top-1 and Top-5. With the c which is set as 4 in knowledge aggregation scheme, we found the situation of over-regularization is alleviated, and it prompts the KR framework to achieve the better performance.

4.2.4 Generalization Ability

We further explore the generalization ability of previous methods and the proposed framework. Based on the data resource from CIFAR-100, we reproduce the CIFAR-100 as the TCIFAR-100 with the data distortion strategies. In details, each image in training and test set is distorted by

	Top-1	Top-5
Teacher ResNet-101	22.68	6.58
Student Inception-BN [15]	25.74	8.07
FitNet [26]	25.30	7.93
NTS [12]	24.34	7.11
KDGAN [30]	24.11	6.98
Proposed Method	23.47	6.85
Proposed Method + KDGAN [30]	23.18	6.79

Table 3. ImageNet Classification errors (Top-1 and Top-5%).

the artifacts, from a gaussian distribution ($\sigma = 1$) with the random sample. The Fig. 6 shows the examples. In Tab. 4, it shows the proposed framework achieves a significant improvement than state-of-the-arts. We believe the KR framework is able to produce a student network with stronger generalization ability, since the joint optimization prevents the optimizer from being trapped in local extremum.



Figure 6. *left* CIFAR-100; *right* TCIFAR-100

	Accuracy	Params
Teacher ResNet-32	61.25	$\sim 0.46\text{M}$
Student ResNet-14 (Original)	54.37	$\sim 0.19\text{M}$
FitNet [26]	56.77	$\sim 0.19\text{M}$
FSP [34]	57.31	$\sim 0.19\text{M}$
Proposed Method	60.03	$\sim 0.17\text{M}$
NTS [12]	57.88	$\sim 0.19\text{M}$
KDGAN [30]	58.15	$\sim 0.19\text{M}$
Proposed Method + KDGAN [30]	60.33	$\sim 0.17\text{M}$

Table 4. ResNet-14 in TCIFAR-100 Classification rates(%). The transformed CIFAR-100 dataset is reproduced by the CIFAR-100.

4.3. Optimization Discussion

In this section, we further discuss the implementation details of optimizing the proposed framework, and analysis the optimization procedure with different settings.

Implementation Details As for the training on CIFAR-10 and CIFAR-100, the learning rate for Eq. 9 is set as 0.1, and was changed to 0.01, and 0.001 at two steps (30k and 48k) respectively. The optimizer for Eq. 10 started at a smaller learning rate 0.01, but also is reduced according to similar strategies. For the ILSVRC 2012, the learning rate for Eq. 9 is set as 0.1 with a ploy decreasing in each 6 epoch, and the optimizer for Eq. 10 started at learning rate 0.005. The weight decay of 0.00001 and momentum of 0.9

are all used. For the works related to quantization strategies [16, 25], we try to evaluate the performance if combining these works with our framework. Since the quantization techniques transfer the parameters distribution into a discrete space, we found the optimization will be seriously impacted and convergence performance also be influenced. However, this analysis is out of the scope of this paper, so it is left as future work.

Joint Optimization For optimizing the \tilde{W}^t with W^t by Eq. 9 and the W^s with \tilde{W}^t by Eq. 10, we use two different optimizers to separately training these two sub-objective functions. Moreover, we tried different initialization techniques for parameters, and we found the objective function is harder to converge, if the initialization on \tilde{W}^t is very different from W^s . We also consider the types for different optimizers [17, 20, 29]. Through changing the two optimizers as Adam [17] or RMS [29], we found it caused a performance oscillation but less than 1%.

5. Conclusion

In this paper, we propose a knowledge representing (KR) framework mainly focusing on modeling the parameters distribution as prior knowledge. We suggest a knowledge aggregation scheme to represent the parameters knowledge from teacher network into more abstract level, for alleviating the phenomenon of residual accumulation in the deeper layers. We also design a sparse recoding penalty for constraining the student network to learn with the penalized gradients. It helps the student network to avoid the over-regularization during knowledge distilling and converge faster. In conclusion, the proposed framework can prompt the student network to preserve the key features of teacher network, even though the student network does not have a strong representation ability.

Acknowledgements. We thanks all reviewers for providing the constructive suggestions.

References

- [1] Anoop Korattikara Balan, Vivek Rathod, Kevin P Murphy, and Max Welling. Bayesian dark knowledge. In *Advances in Neural Information Processing Systems*, pages 3438–3446, 2015.
- [2] Guobin Chen, Wongun Choi, Xiang Yu, Tony Han, and Manmohan Chandraker. Learning efficient object detection models with knowledge distillation. In *Advances in Neural Information Processing Systems*, 2017.
- [3] Tianqi Chen, Ian Goodfellow, and Jonathon Shlens. Net2net: Accelerating learning via knowledge transfer. *Computer Science*, 2015.
- [4] Yu Cheng, Duo Wang, Pan Zhou, and Tao Zhang. A survey of model compression and acceleration for deep neural networks. *arXiv preprint arXiv:1710.09282*, 2017.
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Fei-Fei Li. Imagenet: A large scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2009.
- [6] Misha Denil, Babak Shakibi, Laurent Dinh, Nando De Freitas, et al. Predicting parameters in deep learning. In *Advances in Neural Information Processing Systems*, pages 2148–2156, 2013.
- [7] Huang Gao, Liu Zhuang, and van der Maaten Laurens. Densely connected convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [9] Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1389–1397, 2017.
- [10] Byeongho Heo, Minsik Lee, and Sangdoon Yun. Knowledge distillation with adversarial samples supporting decision boundary. In *Proceedings of Association for the Advancement of Artificial Intelligence*, 2019.
- [11] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *Computer Science*, 14(7):38–39, 2015.
- [12] Zehao Huang and Naiyan Wang. Like what you like: Knowledge distill via neuron selectivity transfer. *arXiv preprint arXiv:1707.01219*, 2017.
- [13] Zehao Huang and Naiyan Wang. Data-driven sparse structure selection for deep neural networks. In *Proceedings of the European Conference on Computer Vision*, pages 304–320, 2018.
- [14] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Quantized neural networks: Training neural networks with low precision weights and activations. *The Journal of Machine Learning Research*, 18(1):6869–6898, 2017.
- [15] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the International Conference on Machine Learning*, 2015.
- [16] Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2704–2713, 2018.
- [17] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations*, 2015.
- [18] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. *Master’s thesis*, 2009.

- [19] Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, 2012.
- [20] Bottou Leon. Stochastic gradient tricks. *Springer*, 7700:430445, 2012.
- [21] Jian-Hao Luo, Jianxin Wu, and Weiyao Lin. Thinet: A filter level pruning method for deep neural network compression. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5058–5066, 2017.
- [22] Ping Luo, Zhenyao Zhu, Ziwei Liu, Xiaogang Wang, and Xiaoou Tang. Face model compression by distilling knowledge from neurons. In *Proceedings of Association for the Advance of Artificial Intelligence*, 2016.
- [23] Arjovsky Martin, Chintala Soumith, and Bottou Leon. Wasserstein generative adversarial networks. In *Proceedings of the International Conference on Machine Learning*, pages 214–223, 2017.
- [24] Lei Na, Kehua Su, Cui Li, Shing Tung Yau, and David Xianfeng Gu. A geometric view of optimal transportation and generative model. 2017.
- [25] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In *Proceedings of the European Conference on Computer Vision*, pages 525–542. Springer, 2016.
- [26] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. *Computer Science*, 2014.
- [27] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4510–4520, 2018.
- [28] Hao Shen. Towards a mathematical understanding of the difficulty in learning with feedforward neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [29] Tijmen Tieleman and Geoffrey Hinton. Rmsprop: Divide the gradient by a running average of its recent magnitude. 2012.
- [30] Xiaojie Wang, Rui Zhang, Yu Sun, and Jianzhong Qi. Kdgan: Knowledge distillation with generative adversarial networks. In *Advances in Neural Information Processing Systems*, pages 783–794, 2018.
- [31] Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Learning structured sparsity in deep neural networks. In *Advances in Neural Information Processing Systems*, pages 2074–2082, 2016.
- [32] Mike Wu, Michael C Hughes, Sonali Parbhoo, Maurizio Zazzi, Volker Roth, and Finale Doshi-Velez. Beyond sparsity: Tree regularization of deep models for interpretability. *Proceedings of Association for the Advance of Artificial Intelligence*, 2018.
- [33] Zheng Xu, Yen-Chang Hsu, and Jiawei Huang. Learning loss for knowledge distillation with conditional adversarial networks. In *Proceedings of the International Conference on Learning Representations*, 2017.
- [34] Junho Yim, Donggyu Joo, Jihoon Bae, and Junmo Kim. A gift from knowledge distillation: Fast optimization, network minimization and transfer learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, 2017.
- [35] Xin Yu, Zhiding Yu, and Srikumar Ramalingam. Learning strict identity mappings in deep residual networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [36] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. In *Proceedings of the International Conference on Learning Representations*, 2017.
- [37] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6848–6856, 2018.
- [38] Yuchen Zhang, Jason D Lee, and Michael I Jordan. 11-regularized neural networks are improperly learnable in polynomial time. In *Proceedings of the International Conference on Machine Learning*, pages 993–1001, 2016.
- [39] Guorui Zhou, Ying Fan, Runpeng Cui, Weijie Bian, Xiaoqiang Zhu, and Gai Kun. Rocket launching: a universal and efficient framework for training well-performing light net. In *Proceedings of Association for the Advance of Artificial Intelligence*, 2018.