

Online Knowledge Distillation with Diverse Peers

Defang Chen,^{1,2} Jian-Ping Mei,^{3*}
Can Wang,^{1,2} Yan Feng,^{1,2} Chun Chen^{1,2}

¹College of Computer Science, Zhejiang University, Hang Zhou, China.

²ZJU-LianlianPay Joint Research Center.

³College of Computer Science, Zhejiang University of Technology, Hang Zhou, China.
{defchern, wcan, fengyan, chenc}@zju.edu.cn, jpmei@zjut.edu.cn

Abstract

Distillation is an effective knowledge-transfer technique that uses predicted distributions of a powerful teacher model as soft targets to train a less-parameterized student model. A pre-trained high capacity teacher, however, is not always available. Recently proposed online variants use the aggregated intermediate predictions of multiple student models as targets to train each student model. Although group-derived targets give a good recipe for teacher-free distillation, group members are homogenized quickly with simple aggregation functions, leading to early saturated solutions. In this work, we propose Online Knowledge Distillation with Diverse peers (OKDDip), which performs two-level distillation during training with multiple auxiliary peers and one group leader. In the first-level distillation, each auxiliary peer holds an individual set of aggregation weights generated with an attention-based mechanism to derive its own targets from predictions of other auxiliary peers. Learning from distinct target distributions helps to boost peer diversity for effectiveness of group-based distillation. The second-level distillation is performed to transfer the knowledge in the ensemble of auxiliary peers further to the group leader, i.e., the model used for inference. Experimental results show that the proposed framework consistently gives better performance than state-of-the-art approaches without sacrificing training or inference complexity, demonstrating the effectiveness of the proposed two-level distillation framework.

Introduction

Modeling with tens of millions of parameters helps deep neural networks achieve great success in various applications. However, over-parameterized models are computationally demanding, making them unsuitable for deployment with limited resources or a stringent requirement on latency (Denil et al. 2013; Han, Mao, and Dally 2016). The distillation technique transfers the knowledge of a teacher model in the form of soft predictions to improve the generalization ability of a less-parameterized student model through regularization (Ba and Caruana 2014; Romero et al. 2015; Yim et al. 2017). Compared to hard ground-truth labels,

the soft predicted distributions carry richer information that helps to optimize small networks more effectively (Hinton, Vinyals, and Dean 2015; Chaudhari et al. 2017).

The vanilla Knowledge Distillation (KD) method is a two-stage process in which a high capacity teacher model is trained and then used for distillation. This will increase both training cost and pipeline complexity. Recent attempts on group-based online knowledge distillation explored less costly and unified models to eliminate the necessity of pre-training a large teacher model. (Zhang et al. 2018; Anil et al. 2018; Lan, Zhu, and Gong 2018; Song and Chai 2018). The main idea is to train a number of student models simultaneously by learning from ground-truth labels as well as distilling from their group-derived soft targets, which is a specific form of aggregation of intermediate peer predictions.

With the absence of a powerful teacher model, the group-derived targets play a key role in transferring group knowledge to each student model. Averaging over predictions of group members is a simple aggregation to derive the targets for representing the group knowledge (Zhang et al. 2018; Anil et al. 2018; Song and Chai 2018). Since the quality of predictions varies among peers, it is important to treat peers unequally (Lan, Zhu, and Gong 2018). Unfortunately, naive aggregation functions tend to cause peers to quickly homogenize, hurting the effectiveness of group distillation (Kuncheva and Whitaker 2003; Zhou 2012). In this paper, we propose a new two-level distillation approach called **Online Knowledge Distillation with Diverse peers** (OKDDip), which involves two types of student models, i.e., multiple auxiliary peers and one group leader. Base distillation is performed among auxiliary peers equipped with a diversity holding mechanism. An ensemble of predictions of these diverse peers is further distilled into the group leader. Unlike naive group-based learning where all peers end up with similar behaviors, trained peer models in our approach could be quite different from each other. It is thus unreasonable to arbitrarily select a single peer for inference. The second-level distillation therefore is necessary to reduce inference cost using only one student model. With this design, OKDDip can be effectively trained without a high capacity teacher network by taking advantages of group distillation and re-

*Corresponding author

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

mains efficient for inference.

A key design of OKDDip is that each auxiliary peer assigns individual weights to all the peers during aggregation to derive their own target distributions. We incorporate an attention-based mechanism (Vaswani et al. 2017) to generate a distinct set of weights for each peer to measure the importance of group members. This allows large variation in derived target distributions and hence boosts peer diversity. Note that asymmetric weights exist in our model, which differs from simple aggregation and allows quality peers to excel.

We performed experimental study to evaluate the classification performance of the proposed approach on CIFAR-10, CIFAR-100 and ImageNet-2012 datasets with a variety of settings based on popular network architectures. Experimental results show that without increasing cost and complexity, our two-level distillation approach consistently generalized better than state-of-the-art online knowledge distillation approaches as well as the classic teacher-guided KD approach. Larger peer diversity and stronger ensemble are observed in our approach when compared to others, demonstrating that the proposed attention-based mechanism works well for diversity holding.

Related Work

Knowledge Distillation. KD provides a succinct but effective solution for compressing a pre-trained large teacher model into a smaller student model by steering the student predictions towards teacher predictions (Bucilua, Caruana, and Niculescu-Mizil 2006; Ba and Caruana 2014; Hinton, Vinyals, and Dean 2015; Polino, Pascanu, and Alistarh 2018). Compared to hard ground-truth labels, fine-grained class information in soft predictions helps the small model to reach flatter local minima, which results in more robust performance and improves generalization ability (Pereyra et al. 2017; Keskar et al. 2017). Several recent works attempt to further improve the performance with new formulations of teacher-learned knowledge (Yim et al. 2017; Chen, Zhang, and Dong 2018; Ahn et al. 2019).

Online Knowledge Distillation. Instead of two-stage knowledge transfer, recent work focus on more economic online knowledge distillation without a pre-trained teacher model. Simultaneously training a group of student models by learning from peers’ predictions is an effective substitute for teacher-absent knowledge distillation. Some approaches use individual networks with each one corresponding to a student model (Zhang et al. 2018; Anil et al. 2018), while some others ask all student models to share the same early blocks to further reduce the training cost (Song and Chai 2018; Lan, Zhu, and Gong 2018). The main difference in these approaches is the way that each student model learns from others. In (Zhang et al. 2018), each student model learns from the simple average of predictions of all other group members and requires complex asynchronous updating among different networks. A similar variants called co-distillation investigated the potential benefits for online distillation in the distributed learning (Anil et al. 2018). In (Lan, Zhu, and Gong 2018), all student models share the same target distribution by averaging predictions of all the members

with weights learned by a fully connected layer. Unfortunately, simply treating each peer to be equally important or forcing all the members to learn from the same targets would hurt the diversity among students, which limits the effectiveness of within-group knowledge transfer.

Self-Attention. Attention was introduced in natural language processing for encoding each word with others which are most relevant regarding to the target task (Bahdanau, Cho, and Bengio 2015). It has been also successfully applied to other types of data such as image and graph with state-of-the-art performance. In particular, self-attention or intra-attention refers to the mechanism of capturing global dependencies by calculating the response at a position through attending to all the neighbors (Vaswani et al. 2017). Specifically, the input representation of a position, such as a word in a sequence (Vaswani et al. 2017), a pixel within an image region (Zhang et al. 2019) or a node in a graph (Veličković et al. 2017), is linearly mapped into three vectors called *query*, *key*, and *value*. The output of this position is obtained by averaging the *values* of its neighbors, e.g., words in the same sequence, with different weights, which are calculated by matching the *query* of this position to the *keys* of neighbors.

Online Knowledge Distillation with Diverse Peers

Learning with labels

Suppose we are dealing with classification tasks with a labeled dataset $\mathcal{D} = \{(x_i, y_i)\}_i^n$. The goal of learning is to find a parameterized mapping $f(x, \theta) : \mathcal{X} \mapsto [0, 1]^{|\mathcal{Y}|}$, which can generalize to unseen data. A classifier is typically trained by minimizing the cross entropy between the predicted class probabilities q_i and the one-hot ground-truth label distribution l_i of each training sample

$$\mathcal{L}_{gt} = - \sum_{i,j} l_{ij} \log q_{ij}, \quad (1)$$

where $l_{ij} = 1$ if $y_i = j$, and 0 otherwise. $q_i = \sigma(\mathbf{g}_i, T)$ is calculated with softmax of logits \mathbf{g}_i , i.e., outputs of the last fully connected layer, i.e.,

$$q_{ij} = \frac{\exp(g_{ij}/T)}{\sum_k \exp(g_{ik}/T)}, \quad (2)$$

and parameter T is usually set to 1.

Next, we present the detailed formulation of the proposed approach after a brief review of teacher-required knowledge distillation.

Distillation with a teacher model

In teacher-presented knowledge distillation, a student model is trained by the teacher-predicted soft distribution \mathbf{t} together with hard ground-truth labels. The soft distribution or targets \mathbf{t} are calculated with softmax of logits as Equation (2) with a temperature $T > 1$. A higher T means softer distributions. Following the suggestion in (Lan, Zhu, and Gong 2018), we set T to 3 in this paper for all methods.

Knowledge is transferred by aligning the student predicted distribution \mathbf{q}' also after softmax with the same temperature (i.e., $T = 3$) to the target distribution. Specifically,

Kullback-Leibler (KL) divergence between \mathbf{t} and \mathbf{q}' may be used to define the distillation loss

$$\mathcal{L}_{dis} = KL(\mathbf{t}, \mathbf{q}') = \sum_{i,j} t_{ij} \log \frac{t_{ij}}{q'_{ij}}. \quad (3)$$

With both hard and soft labels, the total loss for training a teacher-presented distillation model is

$$\mathcal{L}_{KD} = \mathcal{L}_{gt} + T^2 \mathcal{L}_{dis}, \quad (4)$$

where \mathcal{L}_{dis} is multiplied by T^2 before combination to ensure that the contribution of distillation term keeps roughly unchanged if the temperature is changed (Hinton, Vinyals, and Dean 2015). It is worthy to note that the predicted probabilities of the student model \mathbf{q} are computed from the logits with temperature $T = 1$ when aligning with hard ground-truth labels but with a higher temperature when aligning with soft targets in distillation training. To be more clear, we use \mathbf{q} for the $T = 1$ version, and \mathbf{q}' for the high temperature version through out this paper.

Two-level distillation

The proposed two-level distillation framework for group-based knowledge distillation is illustrated in Figure 1. All the m student models including $m - 1$ auxiliary peers and one group leader use the same network architecture, which consists of a feature extractor to produce high-level features followed by a classifier to produce logits. For convenience of denotation, all the student models are indexed from 1 to m , where 1 to $m - 1$ denote the auxiliary peers and m corresponds to the group leader. The predicted distribution (i.e., the high temperature version) of the a th student model is denoted as \mathbf{q}'_a with $a = 1, \dots, m$.

Loss function. For the first-level distillation, each auxiliary peer for $a = 1, 2, \dots, m - 1$ distills from its own group-derived soft targets \mathbf{t}_a , which is computed by aggregating predictions of all peers with different weights

$$\mathbf{t}_a = \sum_{b=1}^{m-1} \alpha_{ab} \cdot \mathbf{q}'_b, \quad (5)$$

where α_{ab} represents the extent to which the b th member is attended in derivation \mathbf{t}_a , and $\sum_b \alpha_{ab} = 1$. We will elaborate on these attention-based weights later. The distillation loss of all auxiliary peers is then given as

$$\mathcal{L}_{dis1} = \sum_{a=1}^{m-1} KL(\mathbf{t}_a, \mathbf{q}'_a), \quad (6)$$

which could be regarded as weighted regularization for the output distribution. As pointed in (Pereyra et al. 2017), penalizing the confident prediction could prevent over-fitting by increasing the probabilities assigned to incorrect classes.

The group knowledge of those auxiliary peers is distilled further to the group leader (i.e., the m th student model) with the second distillation

$$\mathcal{L}_{dis2} = KL(\mathbf{t}_m, \mathbf{q}'_m), \quad (7)$$

which is similar to the classic KD process in the way of transferring the knowledge of an ensemble to a student

model, but in online fashion. With diversity enhancement by the first-level distillation, we simply average the predictions of all auxiliary peers to compute \mathbf{t}_m .

The overall loss of the proposed approach is given as

$$\mathcal{L}_{OKDDip} = \sum_{a=1}^m \mathcal{L}_{gt}(a) + T^2 \mathcal{L}_{dis1} + T^2 \mathcal{L}_{dis2}, \quad (8)$$

where the first term is the total cross entropy loss to the ground-truth labels of all the m student models¹.

Attention-based weights. With different initializations, quality of intermediate predictions varies among peer models, which should contribute different extents in deriving the soft target distributions for each of the auxiliary peers. Simply treating all peers equally would make the distilled model suffer from negative contributions from low quality predictions. We expect the weights to capture the relative importance of peers to a distilled model.

Inspired by the self-attention mechanism (Vaswani et al. 2017), we project the extracted features of each peer model \mathbf{h}_a into two subspaces separately by linear transformation

$$L(\mathbf{h}_a) = \mathbf{W}_L^T \mathbf{h}_a \quad \text{and} \quad E(\mathbf{h}_a) = \mathbf{W}_E^T \mathbf{h}_a, \quad (9)$$

where \mathbf{W}_L and \mathbf{W}_E are the learned projection matrices shared by all auxiliary peers. Similar to self-attention, α_{ab} is calculated as Embedded Gaussian distance with normalization

$$\alpha_{ab} = \frac{e^{L(\mathbf{h}_a)^T E(\mathbf{h}_b)}}{\sum_{f=1}^{m-1} e^{L(\mathbf{h}_a)^T E(\mathbf{h}_f)}}. \quad (10)$$

The two separate transformation matrices may capture different semantic information. Weights generated in this way have the following merits:

- **Asymmetric:** The asymmetric property provides a possible way to suppress negative effect in one direction without stopping positive guidance in the other, which is important for mutual learning between two peers optimized to different levels. It on the one hand reduces the extent to which a well-behaved model is affected by a poorly performed peer by assigning a small weight, and on the other hand allows the less optimized model to learn from the better optimized one with a large weight.
- **Dynamic:** The performance of peer models changes during training, updating weights in iterations allows each model to attend a dynamic set of peers adaptively.

It is shown in our experimental study that the above properties enable OKDDip to outperform state-of-the-art approaches as well as its simplified variants.

Why learning with distinct target distributions leads to diversity? In the first-level distillation, each auxiliary peer has an individual set of weights to measure the extent to which it attends all group members from its own point of view. Such kind of a personalized aggregation increases

¹Following (Lan, Zhu, and Gong 2018; Laine and Aila 2017), the two distillation terms would be multiplied by an iteration-dependent weighting function during implementation to avoid large contribution of distillation in early stages.

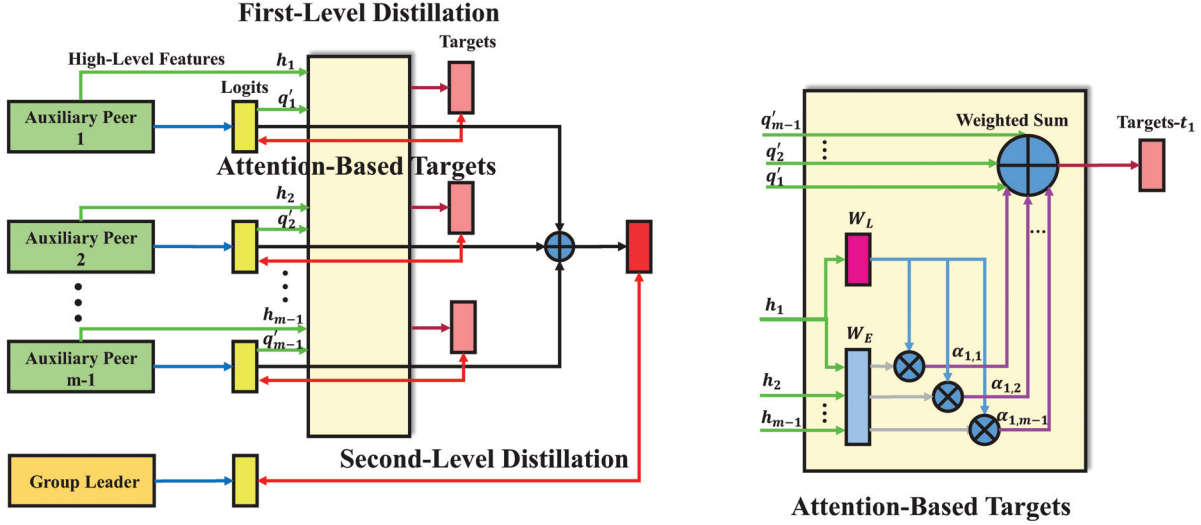


Figure 1: An overview of the proposed Online Knowledge Distillation with Diverse Peers (OKDDip). (Left) Two-level distillation. The first-level distillation performs group-based learning among $m - 1$ auxiliary peers by learning from their own targets. The second-level distillation transfers group knowledge of diverse peers to the group leader, i.e., the final model for deployment. (Right) Attention-based targets derivation. The targets of an auxiliary peer is computed as weighted sum of logits of group members. Each weight α_{ab} for auxiliary peer a to attend member b in deriving its targets is calculated with normalized Embedded Gaussian distance between their mapped features $\mathbf{W}_L^T \mathbf{h}_a$ and $\mathbf{W}_E^T \mathbf{h}_b$, where \mathbf{W}_L and \mathbf{W}_E are two linear projection matrices, \mathbf{h}_a and \mathbf{h}_b are original high-level features.

the independency between soft target distributions of different peers, which is helpful for alleviating diversity degradation during group-based distillation as demonstrated in experiments later. Next, we provide analytical discussions on why \mathcal{L}'_{dis1} leads to peer diversity based on its approximation given below.

If the logits are zero-meaned before computing \mathbf{q}' , the KL-divergence distillation loss in Equation (6) then is approximated by the Mean-Square-Error (Hinton, Vinyals, and Dean 2015):

$$\mathcal{L}'_{dis1} = \frac{1}{2} \sum_{a=1}^{m-1} \|\mathbf{q}'_a - \mathbf{t}_a\|^2 = \frac{1}{2} \sum_{a=1}^{m-1} \|\mathbf{q}'_a - \sum_{b=1}^{m-1} \alpha_{ab} \mathbf{q}'_b\|^2. \quad (11)$$

The above loss function is minimized with $\alpha_{aa} = 1$ and $\alpha_{ab} = 0$ for all $b \neq a$, which means that each peer does not learn from others at this point. This could be regarded as an extreme of our approach where all auxiliary peers are trained independently without any group-based distillation. Incorporated with such kind of loss into the total objective function in Equation (8), the proposed approach is able to keep a proper balance between group sharing and independent learning, which allows it to leverage the information distilled from other peers while preventing quick diversity diminishing.

Training and deployment. The proposed framework may be implemented with branch-based or network-based student models. In a branch-based setting, all student models share the first several layers to use the same low-level features, and separate from each other from a certain layer to have individual branches for high-level feature extraction

and classification. In a network-based setting, student models are individual networks. All auxiliary peers are discarded after training and only the group leader is kept for deployment. There is no additional increase in complexity or cost compared to other group-based approaches given the number and architecture of student models are the same.

Experiments

We provide experimental results in this section to evaluate the performance of the proposed approach for image classification. In addition to the overall generalization ability, we also study the diversity maintenance ability for group-based distillation and conduct several ablation studies on the attention mechanism as well as the two-level strategy. Finally, we analyze the impact of group size, i.e., the number of student models, and extend our method with an additional pre-trained teacher model. All evaluations are made in comparison with state-of-the-art approaches. More results are provided in supplementary materials.

Datasets and Architectures. Three image classification datasets are used in the following evaluations. CIFAR-10 and CIFAR-100 (Krizhevsky and Hinton 2009) both contain 50,000/10,000 training/testing colored natural images with 32×32 pixels, which are drawn from 10/100 classes. ImageNet-2012 (Russakovsky et al. 2015) is a more challenging dataset consisting of about 1.3 million training images and 50 thousand validation images from 1000 classes. We adopted a standard augmentation procedure as (He et al. 2016; Huang et al. 2017; Zhang et al. 2018). For pre-processing, we normalized all images by channel means

Table 1: Error rates (Top-1, %) on CIFAR-10. OKDDip: network-based (1st column) and branch-based (2nd column).

Network	Baseline	Ind	DML	CL-ILR	ONE	OKDDip
DenseNet-40-12	6.87 ± 0.02	6.97 ± 0.03	6.50 ± 0.02	7.02 ± 0.08	6.85 ± 0.15	5.94 ± 0.05 6.48 \pm 0.12
ResNet-32	6.34 ± 0.03	5.99 ± 0.15	6.18 ± 0.05	6.06 ± 0.07	5.94 ± 0.06	5.62 \pm 0.07 5.58 ± 0.08
VGG-16	6.12 ± 0.15	6.03 ± 0.01	5.94 ± 0.04	6.22 ± 0.10	6.16 ± 0.08	5.88 ± 0.04 5.87 ± 0.03
ResNet-110	5.46 ± 0.02	4.95 ± 0.02	5.68 ± 0.03	4.88 ± 0.12	5.02 ± 0.04	4.54 ± 0.07 4.56 ± 0.11
WRN-20-8	5.27 ± 0.06	5.35 ± 0.02	5.04 ± 0.08	5.12 ± 0.16	5.29 ± 0.02	4.84 ± 0.07 5.06 \pm 0.04

and standard deviations. The size of each training sample is 32×32 for CIFAR-10/100 and 224×224 for ImageNet-2012. Six network architectures are used in our experiments, namely VGG-16 (Simonyan and Zisserman 2015), ResNet-32, ResNet-34, ResNet-110 (He et al. 2016), WRN-20-8 (Zagoruyko and Komodakis 2016), and DenseNet-40-12 (Huang et al. 2017).

Settings. We use stochastic gradient descent with Nesterov momentum for optimization and set the initial learning rate to 0.1, momentum to 0.9. For CIFAR-10/CIFAR-100 dataset, we set the mini-batch size to 128 and weight decay to 5×10^{-4} . The learning rate is divided by 10 at 150 and 225 of the total 300 training epochs for these two datasets. For ImageNet-2012 dataset, we set the mini-batch size to 256, the weight decay to 1×10^{-4} , and the learning rate is divided by 10 at 30 and 60 of the total 90 training epochs. All results are reported in means (standard deviations) over 3 runs. Codes will be released once the paper is accepted.

Approaches compared. We compare the proposed OKDDip to several recently proposed online knowledge distillation approaches, including network-based DML (Zhang et al. 2018), branch-based CL-ILR (Song and Chai 2018), and ONE (Lan, Zhu, and Gong 2018). The “Baseline” approach trains a model by ground-truth labels only and “Ind” refers to the degenerated branch-based approach that trains each student model individually without any group distillation. The classic teacher-guided KD approach (Hinton, Vinyals, and Dean 2015) and a high capacity teacher model are also included for comparison.

For branch-based student models, all student models share the first several blocks of layers and separate from the last block for CIFAR-10/100 and the last two blocks for ImageNet-2012 to form a multi-branch structure as (Lan, Zhu, and Gong 2018). All group-based online knowledge distillation approaches are compared with student models of the same architecture and the same number. Results of group-based approaches are generated with 4 student models except those which are regarding the impact of group size. Since we keep one student model to serve as group leader in OKDDip, the number of peers in group-based aggregation is one less than that of DML, CL-ILR and ONE.

Comparison of classification error rates

Table 1 and Table 2 give the Top-1 classification error rates (%) on CIFAR-10 and CIFAR-100 based on five different network architectures with parameters range from

0.18M (DenseNet-40-12) to 15.30M (VGG-16). Following the original papers, the results of compared methods are the averaged accuracy of different student models, and the results of OKDDip are generated by the group leader. Two groups of results are reported for OKDDip with network-based (1st column) and branch-based (2nd column) implementation.

From these two tables, it is shown that OKDDip achieved lower error rates than all other approaches in both types of student models. Specifically, for the network-based setting, OKDDip outperformed the “Baseline” and DML by 17% and 20% in the best case (ResNet-110), showing that the two-level distillation strategy with attention-based mechanism is more effective than existing ones for group learning. It is also seen that OKDDip achieved lower error rates than “Ind”, CL-ILR and ONE by 16%, 8% and 9% in the best case (ResNet-32/110), respectively, showing that our framework remains to be effective with the multi-branch setting. Generally, OKDDip gives slightly better results with network-based student models, which have more independent parameters to give a larger room to maintain the diversity of peers. We also found that the compared group-based online approaches consistently achieved better performance than the “Baseline” and “Ind” on CIFAR-100 but sometimes fall behind on CIFAR-10, especially for VGG-16 network architecture, which indicates that the homogenization problem tends to become even more severe for dealing with easier dataset.

Table 3 gives the classification error rates for the ImageNet-2012. We obtain similar observations as above. For the following experiments, results of OKDDip are generated based on the branch-based setting.

Diversity analysis

Next, we evaluate whether OKDDip produces more diverse student models compared with other group-based approaches. For each method, the diversity is measured by the average Euclidean distance between the predictions of each pair of peers.

Figure 2 plots the results of peer diversity comparison with four approaches for CIFAR-100. The diversity of “Ind” can be regarded as an upper bound since the student models are trained independently.


As shown in Figure 2, peer diversity for all the approaches are very small when starting from random initialization and climb rapidly after several optimization steps. However, it

Table 2: Error rates (Top-1, %) on CIFAR-100. OKDDip: network-based (1st column) and branch-based (2nd column).

Network	Baseline	Ind	DML	CL-ILR	ONE	OKDDip
DenseNet-40-12	28.97 ± 0.23	29.20 ± 0.09	26.64 ± 0.17	28.61 ± 0.12	28.76 ± 0.18	26.10 ± 0.03 28.34 ± 0.02
ResNet-32	28.76 ± 0.08	27.84 ± 0.05	26.47 ± 0.26	27.44 ± 0.05	26.50 ± 0.13	25.40 ± 0.08 25.63 ± 0.14
VGG-16	26.19 ± 0.12	25.81 ± 0.18	25.33 ± 0.03	25.62 ± 0.11	25.63 ± 0.03	24.88 ± 0.06 25.15 ± 0.19
ResNet-110	24.12 ± 0.20	23.54 ± 0.15	22.50 ± 0.11	21.56 ± 0.09	21.67 ± 0.12	21.09 ± 0.17 21.14 ± 0.14
WRN-20-8	22.50 ± 0.44	21.85 ± 0.12	20.21 ± 0.11	20.44 ± 0.13	21.19 ± 0.12	19.63 ± 0.07 20.06 ± 0.05

Table 3: Error rates (Top-1, %) for ResNet-34 on ImageNet-2012. OKDDip: network-based (1st column) and branch-based (2nd column).

Baseline	DML	CL-ILR	ONE	OKDDip
26.76	26.03	26.06	25.92	25.42 25.60



(a) ResNet-32

(b) ResNet-110

Figure 2: Peer diversity comparison with branch-based models during training on CIFAR-100.

drops quickly for CL-ILR and ONE until the learning rate is decreased (at the 150 epoch), from which it climb again but with a slower speed. Through the whole epochs, the diversity of peers trained by OKDDip is significantly larger than CL-ILR and ONE for both the two ResNet architectures, approaching to those of “Ind”, especially for ResNet-110. This demonstrates that the self-attention mechanism in proposed two-level distillation framework works successfully for alleviating homogenization during group-based distillation.

To show whether diverse peers lead to a stronger ensemble, we also evaluate the accuracy of averaged predictions of trained peers, which can represent how powerful the generated group knowledge to some extent. Therefore, only auxiliary peers are counted in for OKDDip.

From Table 4, thanks to large peer diversity, more effective ensemble is observed with OKDDip, which even achieve lower error rates than the ensemble of individually trained student models. Together with Table 2 and Figure 2, we can find that the compared method CL-ILR and ONE improve the accuracy of individual student models ranging from 0.7% to 8% while decreasing the diversity among different peers, which sacrifices the effectiveness of ensemble models. The error rates even increase by 14% and 5% for

Table 4: Error rates (Top-1, %) of ensemble predictions with branch-based student models on CIFAR-100.

Network	CL-ILR	ONE	OKDDip	Ind
VGG-16	25.56	25.54	24.95	25.62
ResNet-32	27.01	24.90	23.45	23.74
ResNet-110	20.19	20.14	19.54	20.18

ResNet-32. But OKDDip can learn more distinct targets for each auxiliary peer and thus maintain the diversity, which further benefit the learning for group leader, leading to the success for both individual and ensemble models. It is also shown that diversity only is not enough to ensure a good ensemble when comparing with the results of “Ind”.

Ablation Study

To further show the benefit of each individual OKDDip components, especially for the self-attention (SA) mechanism, we perform various ablation studies on CIFAR-100 based on ResNet-32. Specifically, we compare the performance of OKDDip with the following five ways of ablations.

(1) w/o SA (random). A random attention matrix with normalization is used to put randomly generated belief among peers. This leads to higher error rates by 2.61% (28.24%-25.63%).

(2) w/o SA (entropy). In order to validate the effectiveness of dynamic weight modeling of SA, we completely remove \mathcal{L}_{dis1} from objective function and let the peers only learn from \mathcal{L}_{gt} with an entropy term, i.e., each peer attends only to itself, which reduces the performance by 1.08%.

(3) w/o SA (mean). Simple average is applied to aggregate the predictions of peers in the first-level distillation. This causes 0.72% performance drop due to quick homogenization.

(4) w/o SA (asymmetry). Another special case that ablates asymmetry of SA by forcing \mathbf{W}_E and \mathbf{W}_L as identity matrices. The weaker performance (0.42% error rate increase) indicates that the asymmetry merit of OKDDip indeed helps to alleviate more the negative effect from the poor-optimized models to the well-behaved models during training.

(5) w/o two-level. The second-level distillation is ablated by removing the group leader from training and inference with a randomly chosen student model. This increases the error rate by 2.16%, which confirms the usefulness of the second-level distillation.

Table 5: Ablation study: Error rates (Top-1, %) for ResNet-32 on CIFAR-100

w/o SA (random)	w/o SA (entropy)	w/o SA (mean)	w/o SA (asymmetry)	w/o two-level	OKDDip
28.24 ± 0.16	26.71 ± 0.19	26.35 ± 0.14	26.05 ± 0.17	27.79 ± 0.14	25.63 ± 0.14

Table 6: Error rates (Top-1, %) for ResNet-32 with an additional teacher

DataSet	Baseline	KD	OKDDip	OKDDip+KD	Teacher
CIFAR-10	6.34 ± 0.03	6.08 ± 0.11	5.58 ± 0.08	5.36 ± 0.06	5.27 ± 0.23
CIFAR-100	28.76 ± 0.08	26.51 ± 0.14	25.63 ± 0.14	24.92 ± 0.08	24.12 ± 0.20




Figure 3: Impact of group size with branch-based ResNet-32 on CIFAR-100.

Impact of the group size

In this section, we evaluate the impact of the group size on the performance of group-based distillation approaches. We compare OKDDip with ONE and CL-ILR using ResNet-32 with the branch-based setting. Figure 3 plots the error rates of the three approaches on CIFAR-100 with respect to the total number of branches changing from 3 to 8.

It is seen that OKDDip performs the best among the three in all the cases. OKDDip has a curve that sloped down more sharply, which means that there is still relative large room for further improvement if a larger group size is allowed in deployment. Although ONE gave better results than CL-ILR, further improvement by adding more branches indicates that increasing the group size does not help much due to the homogenization problem.

When a teacher is available

Although our approach is mainly designed for teacher-free deployment, it is still interesting to know that whether it can be further improved if a pre-trained teacher is indeed available. Here, we use ResNet-110 as the teacher model, the multi-branch ResNet-32 as student models of OKDDip. To make a teacher-guided OKDDip, denoted as OKDDip+KD, the KL-divergence losses between the predictions of each student and the teacher are added into the original loss function of OKDDip for optimization.

Table 6 gives the results of OKDDip+KD compared with the “Baseline”, the classic KD approach, and the teacher-free OKDDip. We also included the results of the teacher

model. It is observed that among all the four approaches with a small student model for inference, OKDDip+KD becomes the most competitive one, which approaches the level of the teacher model and outperforms “Baseline” by 15% and 13% on CIFAR-10 and CIFAR-100 respectively. This shows that a powerful teacher is still helpful to further improve the generalization ability of OKDDip. The effectiveness of using a teacher is also demonstrated by comparing KD and the “Baseline”. More importantly, it is observed that with group-based two-level online distillation, the teacher-absent OKDDip already outperformed the teacher-assisted KD with significant improvement.

Conclusion

Group-based knowledge distillation is a good substitute for knowledge transfer when a pre-trained high capacity model is not easily accessible. It is critical but challenging for group learning without too much diversity diminishing among peers. We proposed a novel two-level framework for effective online distillation. The base distillation works as diversity maintained group distillation with several auxiliary peers, which are discarded after training. The second-level distillation transfers the diversity enhanced group knowledge to the ultimate student model called group leader. Experimental results show that by distilling from distinct target distributions derived with weights from an attention-based component, peer diversity is maintained to a relative large extent throughout group learning, leading to effective online knowledge transfer. This finally makes the proposed approach outperforms the state-of-the-art online knowledge distillation approaches without additional training and inference cost. Our results also show that although a teacher model is still helpful, our teacher-free OKDDip already achieves higher accuracy by a large margin than the teacher-presented KD model, making it a promising competitive choice for deployment.

Acknowledgments

This work is supported by National Natural Science Foundation of China (Grant No: U1866602), National Key Research and Development Project (Grant No: 2018AAA0101503), and the Zhejiang Provincial Natural Science Foundation (Grant No: LY20F020023). The authors would also like to thank Mr. Gaoqi Chen for his generous support for computing resources.

References

- Ahn, S.; Hu, S. X.; Damianou, A. C.; Lawrence, N. D.; and Dai, Z. 2019. Variational information distillation for knowledge transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 9163–9171.
- Anil, R.; Pereyra, G.; Passos, A.; Ormándi, R.; Dahl, G. E.; and Hinton, G. E. 2018. Large scale distributed neural network training through online distillation. In *International Conference on Learning Representations*.
- Ba, J., and Caruana, R. 2014. Do deep nets really need to be deep? In *Advances in Neural Information Processing Systems*, 2654–2662.
- Bahdanau, D.; Cho, K.; and Bengio, Y. 2015. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations*.
- Bucilua, C.; Caruana, R.; and Niculescu-Mizil, A. 2006. Model compression. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 535–541.
- Chaudhari, P.; Choromanska, A.; Soatto, S.; LeCun, Y.; Baldassi, C.; Borgs, C.; Chayes, J. T.; Sagun, L.; and Zecchina, R. 2017. Entropy-sgd: Biasing gradient descent into wide valleys. In *International Conference on Learning Representations*.
- Chen, S.; Zhang, C.; and Dong, M. 2018. Coupled end-to-end transfer learning with generalized fisher information. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4329–4338.
- Denil, M.; Shakibi, B.; Dinh, L.; De Freitas, N.; et al. 2013. Predicting parameters in deep learning. In *Advances in Neural Information Processing Systems*, 2148–2156.
- Han, S.; Mao, H.; and Dally, W. J. 2016. Deep compression: Compressing deep neural network with pruning, trained quantization and Huffman coding. In *International Conference on Learning Representations*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- Hinton, G. E.; Vinyals, O.; and Dean, J. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Huang, G.; Liu, Z.; Van Der Maaten, L.; and Weinberger, K. Q. 2017. Densely connected convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4700–4708.
- Keskar, N. S.; Mudigere, D.; Nocedal, J.; Smelyanskiy, M.; and Tang, P. T. P. 2017. On large-batch training for deep learning: Generalization gap and sharp minima. In *International Conference on Learning Representations*.
- Krizhevsky, A., and Hinton, G. 2009. Learning multiple layers of features from tiny images. *Technical Report*.
- Kuncheva, L. I., and Whitaker, C. J. 2003. Measures of diversity in classifier ensembles and their relationship with the ensemble accuracy. *Machine Learning* 51(2):181–207.
- Laine, S., and Aila, T. 2017. Temporal ensembling for semi-supervised learning. In *International Conference on Learning Representations*.
- Lan, X.; Zhu, X.; and Gong, S. 2018. Knowledge distillation by on-the-fly native ensemble. In *Advances in Neural Information Processing Systems*, 7528–7538.
- Pereyra, G.; Tucker, G.; Chorowski, J.; Kaiser, Ł.; and Hinton, G. 2017. Regularizing neural networks by penalizing confident output distributions. *arXiv preprint arXiv:1701.06548*.
- Polino, A.; Pascanu, R.; and Alistarh, D. 2018. Model compression via distillation and quantization. In *International Conference on Learning Representations*.
- Romero, A.; Ballas, N.; Kahou, S. E.; Chassang, A.; Gatta, C.; and Bengio, Y. 2015. Fitnets: Hints for thin deep nets. In *International Conference on Learning Representations*.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M. S.; Berg, A. C.; and Li, F. 2015. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision* 115(3):211–252.
- Simonyan, K., and Zisserman, A. 2015. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*.
- Song, G., and Chai, W. 2018. Collaborative learning for deep neural networks. In *Advances in Neural Information Processing Systems*, 1832–1841.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, 5998–6008.
- Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; and Bengio, Y. 2017. Graph attention networks. In *International Conference on Learning Representations*.
- Yim, J.; Joo, D.; Bae, J.; and Kim, J. 2017. A gift from knowledge distillation: Fast optimization, network minimization and transfer learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7130–7138.
- Zagoruyko, S., and Komodakis, N. 2016. Wide residual networks. In *Proceedings of the British Machine Vision Conference*.
- Zhang, Y.; Xiang, T.; Hospedales, T. M.; and Lu, H. 2018. Deep mutual learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4320–4328.
- Zhang, H.; Goodfellow, I.; Metaxas, D.; and Odena, A. 2019. Self-attention generative adversarial networks. In *Proceedings of the International Conference on Machine Learning*, 7354–7363.
- Zhou, Z.-H. 2012. *Ensemble methods: foundations and algorithms*. Chapman and Hall/CRC.