

# Differentiable Feature Aggregation Search for Knowledge Distillation

Yushuo Guan<sup>1\*</sup>, Pengyu Zhao<sup>1\*</sup>, Bingxuan Wang<sup>1</sup>, Yuanxing Zhang<sup>1</sup>, Cong Yao<sup>2</sup>, Kaigui Bian<sup>1,3</sup>, and Jian Tang<sup>4</sup>  
<sup>1</sup>Peking University, <sup>2</sup>Megvii (Face++) Technology Inc, <sup>3</sup>National Engineering Laboratory for Big Data Analysis and Applications, <sup>4</sup>DiDi AI Labs  
Emails: {david.guan, pengyuzhao, wangbx, longo}@pku.edu.cn, yaocong2010@gmail.com, bkg@pku.edu.cn, tangjian@didiglobal.com

**Abstract.** Knowledge distillation has become increasingly important in model compression. It boosts the performance of a miniaturized student network with the supervision of the output distribution and feature maps from a sophisticated teacher network. Some recent works introduce multi-teacher distillation to provide more supervision to the student network. However, the effectiveness of multi-teacher distillation methods are accompanied by costly computation resources. To tackle with both the efficiency and the effectiveness of knowledge distillation, we introduce the feature aggregation to imitate the multi-teacher distillation in the single-teacher distillation framework by extracting informative supervision from multiple teacher feature maps. Specifically, we introduce DFA, a two-stage Differentiable Feature Aggregation search method that motivated by DARTS in neural architecture search, to efficiently find the aggregations. In the first stage, DFA formulates the searching problem as a bi-level optimization and leverages a novel bridge loss, which consists of a student-to-teacher path and a teacher-to-student path, to find appropriate feature aggregations. The two paths act as two players against each other, trying to optimize the unified architecture parameters to the opposite directions while guaranteeing both expressivity and learnability of the feature aggregation simultaneously. In the second stage, DFA performs knowledge distillation with the derived feature aggregation. Experimental results show that DFA outperforms existing distillation methods on CIFAR-100 and CINIC-10 datasets under various teacher-student settings, verifying the effectiveness and robustness of the design.

**Keywords:** Knowledge Distillation, Feature Aggregation, Differentiable Architecture Search

## 1 Introduction

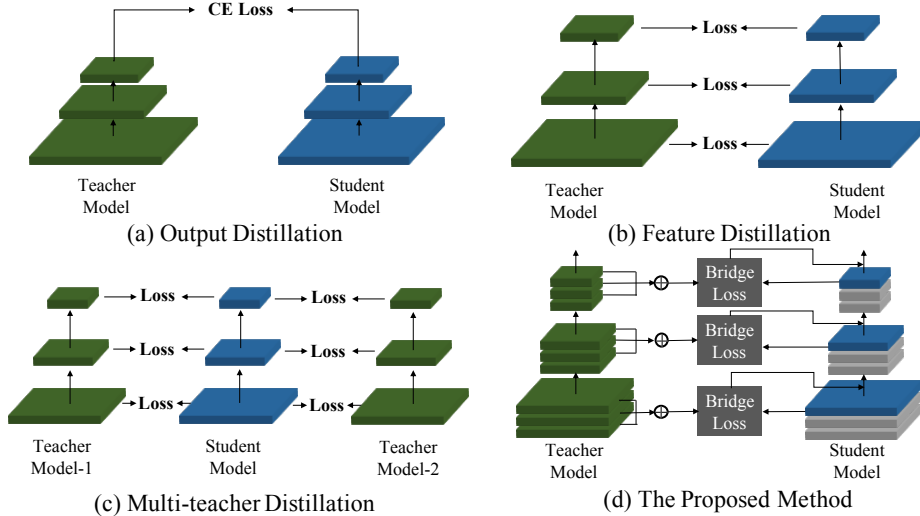
In recent years, visual recognition tasks have been significantly improved by deeper and larger convolutional networks. However, it is difficult to directly deploy such complicated networks on certain computationally limited platforms

---

\* These authors contributed equally to this work.

such as robotics, self-driving vehicles and most of the mobile devices. Therefore, the community has raised increasing attention on model compression approaches such as model pruning [7, 21, 34], model quantization [14, 19, 23] and *knowledge distillation* [13, 30, 33, 37, 38].

Knowledge distillation refers to the methods that supervise the training of a small network (*student*) by using the knowledge extracted from one or more well-trained large networks (*teacher*). The key idea of knowledge distillation is to transfer the knowledge from the teacher networks to the student network. The first attempt of the knowledge distillation for deep neural networks leverages both the correct class labels and the soft targets of the teacher network, i.e., the soft probability distribution over classes, to supervise the training of the student network. The recent advances of knowledge distillation can be mainly divided into two categories: *output distillation* [13] and *feature distillation* [30, 33, 38], as shown in Fig. 1 (a-b). More recent works concentrate on multi-teacher distillation with *feature aggregation* [37], where an ensemble of teacher networks provide richer information from the aggregation of output distributions and feature maps. Although an ensemble of teacher networks could provide richer information from the aggregation of output distributions and feature maps, they require much more computation resources than single-teacher distillation.



**Fig. 1.** Illustrations of different knowledge distillation methods. (a) Output distillation. (b) Feature distillation. (c) Multi-teacher distillation. (d) DFA leverages a novel bridge loss for feature distillation, which takes the advantage of NAS and feature aggregation.

To achieve the same effect as the multi-teacher distillation with less computation overheads, we propose DFA, a two-stage Differentiable Feature Aggregation

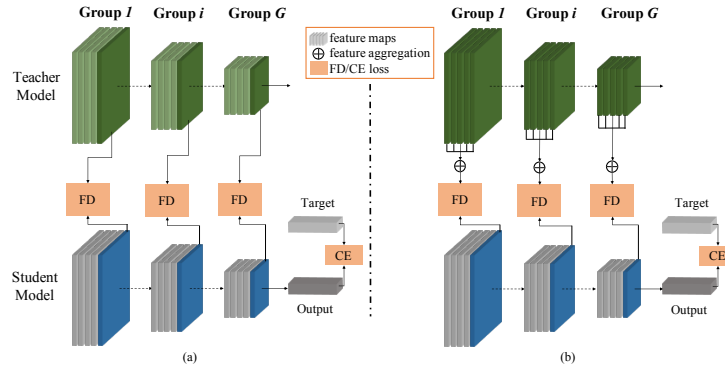
search method in the single-teacher knowledge distillation by coupling features from different layers of a single network as multiple “teachers”, and thus avoids the computation expenses on running several large teacher networks. Specifically, DFA first searches for the appropriate feature aggregation, i.e., the weighted sum of the feature maps, for each layer group in the teacher network by finding the best aggregation weights. Then, it conducts the normal feature distillation with the derived aggregations. Inspired by DARTS [24], DFA leverages the differentiable group-wise search in the first stage, which formulates the searching process as a bi-level optimization problem with feature aggregation weights as the upper-level variable and the model parameters as the lower-level variable. Moreover, as the common distillation loss and cross-entropy loss fail to find the appropriate feature aggregations, a novel *bridge loss* is introduced as the objective function in DFA, where (1) a student-to-teacher path is built for searching the layers that match the learning ability of student network, and (2) a teacher-to-student path is established for finding the feature aggregation with rich features and a wealth of knowledge. Experiments on CIFAR-100 [18] and CINIC [6] datasets show that DFA could outperform the state-of-the-art distillation methods, demonstrating the effectiveness of the feature aggregation search.

The main contributions of this paper are as follows:

- We introduce DFA, a Differentiable Feature Aggregation search method to mimic multi-teacher distillation in the single-teacher distillation framework, which first searches for appropriate feature aggregation weights and then conducts the distillation with the derived feature aggregations.
- We propose a novel bridge loss for DFA. The bridge loss consists of a student-to-teacher path and a teacher-to-student path, which simultaneously considers the expressivity and learnability for the feature aggregation.
- Experimental results show that the performance of DFA surpasses the feature aggregations derived by both hand-crafted settings and random search, verifying the strength of the proposed method.

## 2 Related Work

**Knowledge Distillation:** Knowledge distillation [2] is firstly introduced in model compression. Despite the classification loss, the student network is optimized by an extra cross-entropy loss with the soft target from the teacher network, i.e. the probability distribution softened by temperature scaling. Hinton et al. [13] employ knowledge distillation in the training of deep neural networks. However, with the huge gap of model capacity among the neural networks, it is hard for the student to learn from the output distribution of a cumbersome teacher directly. Thus, several approaches [16, 30, 38] exploit feature distillation in the student training, where the student network mimics the feature maps from the teacher network of different layers. Multi-teacher knowledge distillation [37] takes a further step, which takes full advantage of the feature maps and the class distributions amalgamated from an ensemble of teacher networks. The *feature*



**Fig. 2.** Comparisons between traditional feature distillation and DFA. (a) Traditional methods implement distillation with the last feature map in each layer group of the teacher network. (b) DFA leverages the feature aggregation of the teacher for distillation, which contains rich features and a wealth of knowledge. “FD” and “CE” represent the feature distillation loss  $\mathcal{L}_{fd}$  and cross entropy loss  $\mathcal{L}_{ce}$  respectively.

*aggregation* from multiple teachers helps the student learn from different perspectives. However, compared with single-teacher distillation, more computation resources are demanded for extracting useful information from all the teachers.

**Neural Architecture Search:** With the vigorous development of deep learning, neural architecture search (NAS), an automatic method for designing the structure of neural networks, has been attracting increasing attention recently. The early works mainly sample and evaluate a large number of networks from the search space, and then train the sampled models with reinforcement learning [3, 32, 41, 42] or update the population with the evolutionary algorithm [28, 29]. Though achieving state-of-the-art performance, the above works are all computation expensive. Recent works propose the one-shot approaches [1, 4, 8, 24, 27, 35] in NAS to reduce the computation cost. It models NAS as a single training process for an over-parameterized network covering all candidate sub-networks named supernet, and then selects the network architecture from the trained supernet. Among the one-shot methods, the differentiable architecture search (DARTS) [5, 22, 24, 26, 36, 40] further relaxes the discrete search space to be continuous and couples the architecture parameters with the model parameters in the supernet. Therefore, the architecture parameters can be jointly optimized in the one-shot training along with the model parameters by gradient descent.

There have been several methods designed for combining NAS with knowledge distillation. DNA [20] searches for the light-weight architecture of the student network from a supernet. KDAS [15] builds up the student network progressively based on an ensemble of independently learned student networks. As opposed to these methods, we try to imitate the multi-teacher distillation in the single teacher distillation framework by finding the appropriate feature aggregations in the teacher network with differentiable search strategy.

**Algorithm 1** Algorithm for the two-stage DFA.

---

1: <b>function</b> AGGREGATION SEARCH 2:   Random initialize $\beta, w$ 3: <b>for</b> group $i = 1, 2, \dots, G$ <b>do</b> 4: <b>for</b> iteration $k = 1, 2, \dots, I_s$ <b>do</b> 5:       Optimize $\beta_i$ by $\mathcal{L}_{\text{val}}(w, \beta)$ 6:       Optimize $w$ by $\mathcal{L}_{\text{train}}(w, \beta)$ 7: <b>end for</b> 8: <b>end for</b> 9:   Reserve the derived $\beta$ . 10: <b>end function</b>	11: <b>function</b> FEATURE DISTILLATION 12:   Get the derived $\beta$ . 13: <b>for</b> iteration $k = 1, 2, \dots, I_v$ <b>do</b> 14: <b>for</b> group $i = 1, 2, \dots, G$ <b>do</b> 15:       Calculate $A_i$ with $\beta_i$ . 16: <b>end for</b> 17:     Update $w$ by minimizing the 18:     loss defined in Eqn. (4). 19: <b>end for</b> 20: <b>end function</b>
---	---

---

### 3 Method

We propose the two-stage Differentiable Feature aggregation (DFA) method for single-teacher knowledge distillation, as outlined in Algorithm 1. In the first stage (i.e., ‘‘AGGREGATION SEARCH’’), DFA searches for appropriate feature aggregation weights. In the second stage (i.e., ‘‘FEATURE DISTILLATION’’), the derived feature aggregations are applied to perform the feature distillation between teacher and student. Details will be described in this section.

#### 3.1 Feature Distillation

DFA is based on feature distillation on  $G$  layer groups, where a layer group denotes the set of layers with the same spatial size in teacher and student networks. The general design schemes for feature distillation are categorized into *teacher transform*, *student transform*, *distillation position* and *distance function* [10]. Teacher transform and student transform extract knowledge from hidden features of the teacher and student networks at the distillation positions respectively. Then, the extracted features are applied to the distance function of distillation. Most approaches [11, 33, 38] adopt  $L_2$  loss as the distance measurement. Let  $N_i^T$  and  $N_i^S$  denote the number of layers in the  $i$ -th layer group of teacher and student network, the distillation loss is defined as:

$$\mathcal{L}_{\text{fd}} = \sum_{i=1}^G L_2(\mathcal{F}_t^i(T_i^{N_i^T}), \mathcal{F}_s^i(S_i^{N_i^S})) \quad (1)$$

where  $T_i^{N_i^T}$  and  $S_i^{N_i^S}$  denote the feature maps of teacher and student networks drawn from the distillation position of the  $i$ -th group. Conforming to the previous work [38], the distillation positions of the teacher and student network lay at the end of each layer group. Besides,  $\mathcal{F}_t^i(\cdot)$  and  $\mathcal{F}_s^i(\cdot)$  in Eqn. (1) represent the teacher transform and student transform respectively, which map the channel numbers of both  $T_i^{N_i^T}$  and  $S_i^{N_i^S}$  to the channel number of teacher feature map. Traditional feature distillation methods are illustrated in Fig. 2 (a).

Different from traditional single-teacher feature distillation methods, DFA utilizes feature aggregation of teacher network as the supervision for student network for each layer group, as shown in Fig. 2 (b). Given the feature aggregation weights  $\alpha_i = \{\alpha_i^1, \dots, \alpha_i^{N_i^T}\}$  of  $i$ -th group in the teacher network, where  $\sum_{j=1}^{N_i^T} \alpha_i^j = 1$ , the feature aggregation of  $i$ -th group  $A_i$  can be computed by:

$$A_i = \sum_{j=1}^{N_i^T} \alpha_i^j T_i^j. \quad (2)$$

The existing feature distillation methods could be seen as a special case of feature aggregation, where the weight of the last layer for each layer group  $i$  of the teacher network, i.e.,  $\alpha_i^{N_i^T}$ , is set to one and the weights of the other layers in the group are set to zero. Given the feature aggregation of different layer groups, the feature distillation loss in Eqn. (1) is changed to:

$$\mathcal{L}_{\text{fd}} = \sum_{i=1}^G L_2(\mathcal{F}_t^i(A_i), \mathcal{F}_s^i(S_i^{N_i^S})). \quad (3)$$

Finally, the student network is optimized by a weighted sum of distillation loss  $\mathcal{L}_{\text{fd}}$  and classification loss  $\mathcal{L}_{\text{ce}}$ :

$$\mathcal{L}_{\text{student}} = \mathcal{L}_{\text{ce}} + \gamma_{\text{fd}} * \mathcal{L}_{\text{fd}} \quad (4)$$

where  $\gamma_{\text{fd}}$  is the balancing hyperparameter.  $\mathcal{L}_{\text{ce}}$  is the standard cross-entropy loss between the ground-truth class label  $gt$  and the output distribution of the student  $p = \{p^1, \dots, p^C\}$ :

$$\mathcal{L}_{\text{ce}}(gt, p) = - \sum_{i=1}^C \mathbb{I}[i = gt] \log(p^i), \quad (5)$$

where  $C$  represents the number of classes and  $\mathbb{I}[\cdot]$  is the indicator function.

### 3.2 Differentiable Group-wise Search

As the feature aggregation weights are continuous and grow exponentially with the number of layer groups, DFA leverages a differentiable architecture search method to efficiently search for the task-dependent feature aggregation weights for better distillation performance. Inspired by previous attempts that divide the NAS search space into blocks [24, 42], DFA implements the feature aggregation search in a group-wise manner, i.e., the weights of other groups keep fixed when searching for the aggregation weights for layer group  $i$ . The group-wise search enables a strong learning capability of the model, leading to only a few epochs to achieve convergence during training. The overall framework of the differentiable group-wise search is shown in Fig. 3.

**Search Space:** Given the teacher and student networks, DFA aims to find the appropriate feature aggregation weights  $\alpha_i$  for each group  $i$ . Different from the DARTS-based methods [24], the search space for the feature aggregation is continuous since the combination of different layers could provide richer information than the individual feature map obtained from the discrete search space. Besides, as only one teacher network is utilized in the aggregation search, the training speed and computation overhead are similar to the standard feature distillation. For a stable training process, we represent the feature aggregation weights  $\alpha$  as a softmax over a set of architecture parameters  $\beta$ :

$$\alpha_i^j = \frac{\exp(\beta_i^j)}{\sum_{j'=1}^{N_i^T} \exp(\beta_i^{j'})}. \quad (6)$$

**Optimization of Differentiable Group-wise Search:** The goal of the differentiable group-wise search is to *jointly optimize* the architecture parameters  $\beta$  and the model parameters  $w$  of the student network. Specifically, the differentiable search tries to find the  $\beta^*$  that minimizes the validation loss  $\mathcal{L}_{\text{val}}(w^*(\beta), \beta)$ , where the weights of the architecture parameters  $w^*$  are obtained by minimizing the training loss  $\mathcal{L}_{\text{train}}(w, \beta)$  for a certain architecture parameter  $\beta$ . Thus, the joint optimization could be viewed as a bi-level optimization problem with  $\beta$  as the upper-level variable and  $w$  as the lower-level variable:

$$\min_{\beta} \mathcal{L}_{\text{val}}(w^*(\beta), \beta) + \lambda \mathcal{R}(\beta) \quad (7)$$


$$\text{s.t. } w^*(\beta) = \operatorname{argmin}_w \mathcal{L}_{\text{train}}(w, \beta), \quad (8)$$

where  $\mathcal{L}_{\text{train}}$  and  $\mathcal{L}_{\text{val}}$  are the training and validation loss respectively.  $\mathcal{R}(\cdot)$  denotes the regularization on the architecture parameters  $\beta$  that could slightly boost the performance of DFA. To solve the bi-level optimization problem,  $\beta$  and  $w$  are alternately trained in a multi-step way by gradient descent to reach a fixed point of architecture parameters and model parameters.

An intuitive option of training and validation loss is to use  $\mathcal{L}_{\text{student}}$  in Eqn. (4). Though it seems that directly learning from  $\mathcal{L}_{\text{student}}$  could result in the appropriate architecture parameters for knowledge distillation, actually, training architecture parameters with  $\mathcal{L}_{\text{student}}$  is equivalent to minimizing the distillation loss between feature aggregation and student feature map:

$$\operatorname{argmin}_{\beta} \mathcal{L}_{\text{student}}(w^*(\beta), \beta) = \operatorname{argmin}_{\beta} \mathcal{L}_{\text{fd}}(w^*(\beta), \beta), \quad (9)$$

as the cross-entropy loss  $\mathcal{L}_{\text{ce}}$  is irrelevant to the architecture parameters. Since the distillation loss only characterizes the distance between the student feature maps and the combinations of the teacher feature maps, the architecture parameters tend to be more inclined to choose teachers that are close to the student. In depth, suppose that after training several epochs, the student feature map has learnt some knowledge from the data distribution through cross-entropy loss and teacher network through distillation loss. As the student network is always



**Fig. 3.** Differentiable group-wise search of DFA. (a) The differentiable search for group  $i$ . (b) The teacher-to-student (TS) path. (c) The student-to-teacher (ST) path. The ST and TS connectors are implemented with  $1 \times 1$  convolutional layers for matching the channel dimensions between teacher and student.

shallower than the teacher network, the knowledge in the deep layers is hard to learn such that the architecture parameters would prefer the teachers in the shallow layers matching the depth and expressivity of the student, other than selecting deep layers with rich semantics and strong expressivity. Therefore, the feature aggregation learnt from the  $\mathcal{L}_{\text{student}}$  deviates from the original target that learning a good teacher for the knowledge distillation. Besides, once the student network finds a matching layer in the teacher group, i.e., the weight of an architecture parameter  $\beta_i^j$  is relative larger than the others, the student transform will learn more about the mapping function from  $S_i^{N_i^S}$  to  $T_i^j$ . Then,  $\beta_i^j$  will grow much faster than other competitors due to the biased training of transform function, and the corresponding feature map will gradually dominant the feature aggregation under the exclusive competition of the architecture parameters. Notice again that the network is more likely to pick shallow layers in the early training stage. Therefore, the student network would unavoidably suffer from the performance collapse using the search results derived from  $\mathcal{L}_{\text{student}}$ .

**Bridge Loss for Feature Aggregation Search:** To search for an appropriate feature aggregation for the knowledge distillation, we introduce the *bridge loss* to connect the teacher and student networks, where the original information flow of the student network is split into two paths.

In the first teacher-to-student (TS) path as illustrated in Fig. 3 (b), DFA takes the feature aggregation of the group  $i$  in the teacher network, i.e.,  $A_i$ , as the input of its  $(i+1)$ -th group of the student network, and then computes the teacher-to-student (TS) loss  $\mathcal{L}_{\text{TS}}$  with standard cross entropy. The TS loss in group  $i$  can be expressed as:

$$\mathcal{L}_{\text{TS}}^i = \mathcal{L}_{\text{ce}}(gt, f_G(\dots(f_{i+1}(\mathcal{F}_t^i(A_i)))))) \quad (10)$$



where  $gt$  denotes the ground-truth class label and  $f_i$  denotes the convolutional layers of group  $i$  in the student network. Different from Eqn. (3), the teacher transform  $\mathcal{F}_t^i$  is now served as a TS connector that converts the channel dimension of the teacher feature map to the student feature map.

The second student-to-teacher (ST) path has the similar effect as the original  $\mathcal{L}_{\text{student}}$  in Eqn. (4), i.e., exploring the feature aggregation weights that match the learning ability of the student network. As shown in Fig. 3 (c), the information flow starts from the student input and ends at the  $i$ -th group in the student network. Then, the student network produces  $S_i^{N_i^S}$ , the last feature map of group  $i$ , and compares  $S_i^{N_i^S}$  with the feature aggregation  $A_i$  by a student-to-teacher (ST) loss  $\mathcal{L}_{\text{ST}}$ :

$$\mathcal{L}_{\text{ST}}^i = L_2\left(\frac{\text{vec}(\mathcal{F}_s^i(S_i^{N_i^S}))}{\|\text{vec}(\mathcal{F}_s^i(S_i^{N_i^S}))\|_2}, \frac{\text{vec}(A_i)}{\|\text{vec}(A_i)\|_2}\right) \quad (11)$$

where  $\text{vec}(\cdot)$  represents the vectorization of the tensor that converts the tensor into a column vector. Same as the distillation loss in Eqn. (3), the student transform  $\mathcal{F}_s^i$  is served as a ST connector to map the channel numbers from student feature map to teacher feature map, as opposed to  $\mathcal{F}_t^i$ .

For each group  $i$ , the bridge loss  $\mathcal{L}_{\text{Bridge}}^i$  integrates ST loss  $\mathcal{L}_{\text{ST}}$  and TS loss  $\mathcal{L}_{\text{TS}}$  in a single training process for both training loss  $\mathcal{L}_{\text{train}}$  and validation loss  $\mathcal{L}_{\text{val}}$ :

$$\mathcal{L}_{\text{Bridge}}^i = \gamma_{\text{ST}}\mathcal{L}_{\text{ST}}^i + \gamma_{\text{TS}}\mathcal{L}_{\text{TS}}^i. \quad (12)$$

where  $\gamma_{\text{ST}}$  and  $\gamma_{\text{TS}}$  are balancing hyperparameters.

Different from  $\mathcal{L}_{\text{student}}$ , both the model parameters and architecture parameters can be trained towards the ground truth through bridge loss. For the model parameters, the student network before  $(i + 1)$ -th group tries to imitate the teacher feature aggregation  $A_i$ , while the student after  $(i + 1)$ -th group learns to optimize the cross entropy given  $A_i$ . Hence, the joint optimization of ST loss and TS loss for the model weights can be regarded as an approximation of the cross entropy loss of the student network. On the other hand, the architecture parameters are trained by directly optimizing the cross-entropy loss  $\mathcal{L}_{\text{TS}}$  given the input of the aggregation of teacher feature maps, analogous to the common differentiable architecture search in [24, 35]. The deep layers with rich features will be assigned by higher weights in the architecture search, as they contribute to the reduction of the validation loss. In this case, DFA achieves the feature aggregation which helps the student learn a wealth of knowledge from the teacher. Besides, we still keep ST loss in the architecture training as a regularization. Specifically, though the rich features in the deep layers contribute better performance in the teacher network, they are not always suitable for the student to learn due to the mismatch of expressive power, e.g., it is impractical for a student of three layers to learn from the teacher of ten layers. Introducing ST loss in the validation loss can help the network select the shallow teachers that match the student expressivity, such that the derived feature could provide more knowledge to the student in the early training stage and accelerate the model

convergence. In this way, the ST loss and TS loss act as two players against each other, trying to optimize the unified architecture parameters to the opposite directions while guaranteeing both expressivity and learnability of the feature aggregation simultaneously. Hence, the derived feature is more likely to achieve better performance in the final knowledge distillation.

After searching for the architecture parameters with the differentiable group-wise search, DFA trains the student network thoroughly with the derived feature aggregation weights by Eqn. (4).

### 3.3 Time Complexity Analysis

The two-stage design of DFA would not increase the time complexity compared with other feature distillation methods. Let  $t_T^i, t_S^i$  denote the computing time of layer group  $i$  in the teacher and student network, the differentiable search in the first stage could be calculated by:

$$\mathcal{T}_1 = \left( \sum_{j=1}^i (t_T^j + t_S^j) \right) + \left( \sum_{j=1}^i t_T^j + \sum_{j=i+1}^G t_S^j \right) = 2 \sum_{j=1}^i t_T^j + \sum_{j=1}^G t_S^j. \quad (13)$$

As the second stage is a usual feature distillation, the overall time complexity of DFA could be derived as:

$$\begin{aligned} \mathcal{T} &= \mathcal{T}_1 + \left( \sum_{j=1}^G t_T^j + \sum_{j=1}^G t_S^j \right) = 2 \sum_{j=1}^i t_T^j + \sum_{j=1}^G t_T^j + 2 \sum_{j=1}^G t_S^j \\ &= O\left(\sum_{j=1}^G t_T^j\right) + O\left(\sum_{j=1}^G t_S^j\right) \end{aligned} \quad (14)$$

which is competitive with other feature distillation methods.

### 3.4 Implementation Details

The ST connector (student transform) and TS connector (teacher transform) are both implemented with the one-layer convolution networks in order to reconcile the channel dimensions between student and teacher feature maps. The weight of ST loss and TS loss are set to  $\gamma_{ST} = 1e - 3$  and  $\gamma_{TS} = 1$  in Eqn. (12). In both stage of DFA, the pre-ReLU features are extracted in the student and teacher networks for the knowledge distillation, where values no smaller than -1 are preserved in the feature maps to retain knowledge from both positive values and negative values while avoiding the exploding gradient problem. The model parameters are initialized by He initialization [9]. The feature aggregation weights in each layer group are initialized in the way that only the last feature map has weight 1 and all other feature maps are allocated zero weight.

We follow the same training scheme as DARTS: only the training set are used to update model parameters, and the validation set are leveraged for better feature aggregation parameters. For the update of the architecture parameters,

DFA adopts Adam [17] as the optimizer with the momentum of (0.5, 0.999), where the learning rate and weight decay rate are both set to 1e-3.

## 4 Experiments

### 4.1 CIFAR-100

CIFAR-100 is a commonly used visual recognition dataset for comparing distillation methods. There are 100 classes in CIFAR-100, and each class contains 500 training images and 100 testing images. To carry out architecture search, the original training images are divided into the training set and validation set with the 7:3 ratio.

We compare our method with eight single-teacher distillation methods: *KD* [13], *FitNets* [30], *AT* [38], *Jacobian* [31], *FT* [16], *AB* [12], *SP* [33] and *Margin* [11]. All the experiments are performed on Wide Residual Network [39]. The feature aggregation is searched for 40 epochs at each layer group. We adopt the same training schemes as [39] to train the model parameters in all methods. Specifically, the model is trained by SGD optimizer of 5e-4 weight decay and 0.9 momentum for 200 epochs on both training and validation set. The learning rate is set to 0.1 initially and decayed by 0.2 at 60, 120 and 160 epochs. The batch size is 128. We utilize random crop and random horizontal flip as the data augmentation.

We explore the performance of our method on several teacher-student pairs, which vary in depth (number of layers), width (number of channels), or both, as shown in Table 1. We conduct the experiments on CIFAR-100 over the above teacher-student pairs, and depict the results in Table 2. For the teacher-student pair (1) of different widths, DFA has a 1.34% improvement

	Teacher	Size	Student	Size
(1)	WRN_28_4	5.87M	WRN_16_4	2.77M
(2)	WRN_28_4	5.87M	WRN_28_2	1.47M
(3)	WRN_28_4	5.87M	WRN_16_2	0.7M

**Table 1.** The configuration of teacher and student networks in experiments on CIFAR-100.

over the output distillation method KD, and also outperforms the other state-of-the-art feature distillation methods. DFA even exhibits a better performance than the teacher network. For the teacher-student pair (2) of different depths, DFA surpasses other feature distillation methods by 0.34%-2.64%. DFA also achieves state-of-the-art results on (3), where the student network compresses both width and depth of the teacher network. The above experiments verify the effectiveness and robustness of DFA in various scenarios.

### 4.2 CINIC-10

CINIC-10 is a large classification dataset containing 270000 images, which are equally split into training, validation and test set with the presence of 10 object categories. The images are collected from CIFAR-10 and ImageNet. Comparing

	Teacher	Student	KD	FitNets	AT	Jacobian	FT	AB	SP	Margin	DFA
(1)	79.17	77.24	78.4	78.17	77.49	77.84	78.26	78.65	78.7	79.11	<b>79.74</b>
(2)	79.17	75.78	76.61	76.14	75.54	76.24	76.51	76.81	77.41	77.84	<b>78.18</b>
(3)	79.17	73.42	73.35	73.65	73.3	73.28	74.17	73.9	74.09	75.51	<b>75.85</b>

**Table 2.** The experiment results on CIFAR-100. We compare the proposed DFA with eight distillation methods. The best results are illustrated in bold. DFA outperforms other state-of-the-art methods.

with CIFAR datasets, CINIC-10 could present a more principled perspective of generalisation performance. We explore the performance of DFA and other three state-of-the-art feature distillation methods on CINIC-10. All the experiments are performed on ShuffleNetV2 [25], an efficient network architecture at mobile platforms. We use several variants of ShuffleNetV2 in the experiments, and the basic configuration is shown in Table 3 conforming to [33]. In the model training process, the SGD optimizer is leveraged, with weight decay of  $5e-4$  and momentum of 0.9. All models are trained with 140 epochs. The learning rate is set to 0.01 initially and decayed by 0.1 at 100, 120 epochs. The batch size is 96. Same as CIFAR dataset, we utilize random crop and random horizontal flip as the data augmentation. We search 5 epochs for each feature group in DFA. The experiment results are shown in Table 4. DFA outperforms the vanilla cross-entropy training as well as the state-of-the-art feature distillation methods.

Group	Block	$k$	$c$	$n$	Output Size
1	Conv-BN-ReLU	3	24	1	$32 \times 32$
2	ShuffleNetV2 block	3	$116x$	4	$16 \times 16$
3	ShuffleNetV2 block	3	$232x$	8	$8 \times 8$
4	ShuffleNetV2 block	3	$464x$	4	$4 \times 4$
5	ShuffleNetV2 block	3	$1024 * \max(1, x)$	1	$4 \times 4$
6	AvgPool-FC	1	10	1	$1 \times 1$

**Table 3.** The configuration of ShuffleNetV2 on CINIC-10 experiments. We leverage standard ShuffleNetV2 blocks in each layer group, where  $k$  denotes the kernel size,  $c$  and  $n$  specify the number of channels and blocks in each layer group. In the end, we add an average pooling layer and a fully connected layer to make final predictions.

	Teacher	Size	Acc	Student	Size	Acc	AT	SP	Margin	DFA	DFA-T
(1)	$x = 2.0$	5.37M	86.14	$x = 1.0$	1.27M	83.28	84.71	85.32	85.29	85.38	<b>85.41</b>
(2)	$x = 2.0$	5.37M	86.14	$x = 0.5$	0.36M	77.34	79.06	79.15	78.79	<b>79.51</b>	79.45
(3)	$x = 1.0$	1.27M	83.28	$x = 0.5$	0.36M	77.34	78.42	79.02	79.69	<b>79.97</b>	79.38

**Table 4.** Experimental results on CINIC-10. The best results are illustrated in bold. “DFA-T” represents the version that searches the feature aggregation weights on CIFAR-100 and implement the feature distillation on CINIC-10.

**Ablation Study on Differentiable Search:** We study the robustness of DFA by searching the feature aggregation weights from a small dataset and then distilling on a larger dataset. Specifically, we build up a variant of DFA, named as DFA-T, by searching the feature aggregation weights on CIFAR-100 and distilling on CINIC-10. It can be observed that DFA-T is only little inferior to DFA, and still achieves state-of-the-art performance.

### 4.3 The Effectiveness of Differentiable Search

**Comparisons with Other Search Methods:** We perform additional experiments on CIFAR-100 to verify the effectiveness of the differentiable feature aggregation search. We compare DFA with the following methods: "Random" represents the method that all feature aggregation weights are randomly selected; "Average" indicates that all feature maps in a layer group share the same feature aggregation weight; "Last" denotes the method that only the last feature map in each layer group is leveraged for the feature distillation, which is widely used in the knowledge distillation. The results are shown in Table 5.


Obviously, "Random" weights or "Average" weights would degrade the performance, indicating the necessity of a decently designed feature selection strategy. Different from "Last", we observe that DFA would allocate positive weights to the shallow layers in the shallow groups to retrieve knowledge rapidly from the teacher network. The weight assignment in DFA reveals that feature aggregation contributes to transferring knowledge from the teacher network to the student network, while the differentiable group-wise search helps achieve the optimal feature aggregation weights. Hence, DFA brings about a remarkable improvement on the knowledge distillation task.

Method	WRN_16.2	WRN_16.4	WRN_28.2
Student	73.42	77.24	75.78
Random	74.25	78.92	77.07
Last	75.51	79.11	77.86
Average	74.11	78.81	76.99
DFA	<b>75.85</b>	<b>79.74</b>	<b>78.18</b>

**Table 5.** Experimental results of DFA and other search methods on CIFAR-100.

**Result Analysis:** In Fig. 4 *Left*, we display the student network’s feature aggregation weights in the configuration (1) of the CIFAR-100 experiment. The feature aggregation weights are initialized with the "Last" scheme and then searched for 40 epochs. It is obvious that the domination of the last feature map in each layer group are weakened as the training continues.

**Sensitivity Analysis:** We study the impact of regularization on the feature aggregation search mentioned in Sec. 3.2. The right figure in Fig. 4 displays the accuracy of the student models with  $\lambda$  ranging from 0 (no regularization) to  $1e - 3$ . Experimental results show that DFA is robust to the regularization in range  $(0, 1e - 3]$ , except a slight decrease without regularization.



**Fig. 4.** *Left.* Heatmap of feature aggregation weights in different layers of WRN\_16-4. *Right.* The accuracy of the student network on CINIC-10,  $\lambda$  ranges from  $[0, 1e-3]$ . We select the three teacher and student pairs in the CINIC-10 experiments.

## 5 Conclusion

In this paper, we propose DFA, a two-stage feature distillation method via differentiable aggregation search. In the first stage, DFA leverages the differentiable architecture search to find appropriate feature aggregation weights. It introduces a bridge loss to connect the teacher and student, where a teacher-to-student loss is built for searching the teacher with rich features and a wealth of knowledge, while a student-to-teacher loss is used to find the aggregation weights that match the learning ability of the student network. In the second stage, DFA performs a standard feature distillation with the derived feature aggregation weights. Experiments show that DFA outperforms several state-of-the-art methods on CIFAR-100 and large-scale CINIC-10 datasets, verifying both the effectiveness and robustness of the design. In-depth analysis also reveals that DFA decently allocates feature aggregation weights on the knowledge distillation task.

## Acknowledgment

This work is partially supported by National Key Research and Development Program No. 2017YFB0803302, Beijing Academy of Artificial Intelligence (BAAI), and NSFC 61632017.

## References

1. Bender, G., Kindermans, P.J., Zoph, B., Vasudevan, V., Le, Q.: Understanding and simplifying one-shot architecture search. In: ICML. pp. 550–559 (2018)
2. Bucilu, C., Caruana, R., Niculescu-Mizil, A.: Model compression. In: KDD. pp. 535–541. ACM (2006)
3. Cai, H., Yang, J., Zhang, W., Han, S., Yu, Y.: Path-level network transformation for efficient architecture search. In: International Conference on Machine Learning. pp. 678–687 (2018)
4. Cai, H., Zhu, L., Han, S.: Proxylessnas: Direct neural architecture search on target task and hardware. In: ICLR (2019)
5. Chen, X., Xie, L., Wu, J., Tian, Q.: Progressive differentiable architecture search: Bridging the depth gap between search and evaluation. In: ICCV (2019)
6. Darlow, L.N., Crowley, E.J., Antoniou, A., Storkey, A.J.: Cinic-10 is not imagenet or cifar-10. arXiv preprint arXiv:1810.03505 (2018)
7. Dong, X., Yang, Y.: Network pruning via transformable architecture search. In: Advances in Neural Information Processing Systems. pp. 759–770 (2019)
8. Dong, X., Yang, Y.: One-shot neural architecture search via self-evaluated template network. In: ICCV. pp. 3681–3690 (2019)
9. He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In: Proceedings of the IEEE international conference on computer vision. pp. 1026–1034 (2015)
10. Heo, B., Kim, J., Yun, S., Park, H., Kwak, N., Choi, J.Y.: A comprehensive overhaul of feature distillation. In: ICCV (October 2019)
11. Heo, B., Kim, J., Yun, S., Park, H., Kwak, N., Choi, J.Y.: A comprehensive overhaul of feature distillation. arXiv preprint arXiv:1904.01866 (2019)
12. Heo, B., Lee, M., Yun, S., Choi, J.Y.: Knowledge transfer via distillation of activation boundaries formed by hidden neurons. In: AAAI. vol. 33, pp. 3779–3787 (2019)
13. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531 (2015)
14. Hubara, I., Courbariaux, M., Soudry, D., El-Yaniv, R., Bengio, Y.: Binarized neural networks. In: Advances in neural information processing systems. pp. 4107–4115 (2016)
15. Kang, M., Mun, J., Han, B.: Towards oracle knowledge distillation with neural architecture search. arXiv preprint arXiv:1911.13019 (2019)
16. Kim, J., Park, S., Kwak, N.: Paraphrasing complex network: Network compression via factor transfer. In: Advances in Neural Information Processing Systems. pp. 2760–2769 (2018)
17. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
18. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009)
19. Leng, C., Dou, Z., Li, H., Zhu, S., Jin, R.: Extremely low bit neural network: Squeeze the last bit out with admm. In: Thirty-Second AAAI Conference on Artificial Intelligence (2018)
20. Li, C., Peng, J., Yuan, L., Wang, G., Liang, X., Lin, L., Chang, X.: Blockwisely supervised neural architecture search with knowledge distillation. arXiv preprint arXiv:1911.13053 (2019)

21. Li, H., Kadav, A., Durdanovic, I., Samet, H., Graf, H.P.: Pruning filters for efficient convnets. arXiv preprint arXiv:1608.08710 (2016)
22. Li, W., Gong, S., Zhu, X.: Neural graph embedding for neural architecture search. In: AAAI (2020)
23. Lin, X., Zhao, C., Pan, W.: Towards accurate binary convolutional neural network. In: Advances in Neural Information Processing Systems. pp. 345–353 (2017)
24. Liu, H., Simonyan, K., Yang, Y.: Darts: Differentiable architecture search. In: ICLR (2019)
25. Ma, N., Zhang, X., Zheng, H.T., Sun, J.: Shufflenet v2: Practical guidelines for efficient cnn architecture design. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 116–131 (2018)
26. Nayman, N., Noy, A., Ridnik, T., Friedman, I., Jin, R., Zelnik, L.: Xnas: Neural architecture search with expert advice. In: Advances in Neural Information Processing Systems. pp. 1975–1985 (2019)
27. Pham, H., Guan, M., Zoph, B., Le, Q., Dean, J.: Efficient neural architecture search via parameter sharing. In: ICML. pp. 4092–4101 (2018)
28. Real, E., Aggarwal, A., Huang, Y., Le, Q.V.: Regularized evolution for image classifier architecture search. In: AAAI. vol. 33, pp. 4780–4789 (2019)
29. Real, E., Moore, S., Selle, A., Saxena, S., Suematsu, Y.L., Tan, J., Le, Q.V., Kurakin, A.: Large-scale evolution of image classifiers. In: ICML. pp. 2902–2911 (2017)
30. Romero, A., Ballas, N., Kahou, S.E., Chassang, A., Gatta, C., Bengio, Y.: Fitnets: Hints for thin deep nets. arXiv preprint arXiv:1412.6550 (2014)
31. Srinivas, S., Fleuret, F.: Knowledge transfer with jacobian matching. arXiv preprint arXiv:1803.00443 (2018)
32. Tan, M., Chen, B., Pang, R., Vasudevan, V., Sandler, M., Howard, A., Le, Q.V.: Mnasnet: Platform-aware neural architecture search for mobile. In: CVPR. pp. 2820–2828 (2019)
33. Tung, F., Mori, G.: Similarity-preserving knowledge distillation. In: ICCV. pp. 1365–1374 (2019)
34. Wen, W., Wu, C., Wang, Y., Chen, Y., Li, H.: Learning structured sparsity in deep neural networks. In: Advances in neural information processing systems. pp. 2074–2082 (2016)
35. Xie, S., Zheng, H., Liu, C., Lin, L.: Snas: stochastic neural architecture search. In: ICLR (2019)
36. Xu, Y., Xie, L., Zhang, X., Chen, X., Qi, G.J., Tian, Q., Xiong, H.: Pc-darts: Partial channel connections for memory-efficient differentiable architecture search. In: ICLR (2020)
37. You, S., Xu, C., Xu, C., Tao, D.: Learning from multiple teacher networks. In: KDD. pp. 1285–1294. ACM (2017)
38. Zagoruyko, S., Komodakis, N.: Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. arXiv preprint arXiv:1612.03928 (2016)
39. Zagoruyko, S., Komodakis, N.: Wide residual networks. arXiv preprint arXiv:1605.07146 (2016)
40. Zela, A., Elsken, T., Saikia, T., Marrakchi, Y., Brox, T., Hutter, F.: Understanding and robustifying differentiable architecture search. In: ICLR (2020)
41. Zoph, B., Le, Q.V.: Neural architecture search with reinforcement learning. In: ICLR (2017)
42. Zoph, B., Vasudevan, V., Shlens, J., Le, Q.V.: Learning transferable architectures for scalable image recognition. In: CVPR. pp. 8697–8710 (2018)