

Show, Attend and Distill: Knowledge Distillation via Attention-based Feature Matching

Mingi Ji^{1*}, Byeongho Heo², Sungrae Park³

¹ Korea Advanced Institute of Science and Technology (KAIST)

² NAVER AI LAB

³ CLOVA AI Research, NAVER Corp.

qwertgfdcvb@kaist.ac.kr, bh.heo@navercorp.com, sungrae.park@navercorp.com

Abstract

Knowledge distillation extracts general knowledge from a pre-trained teacher network and provides guidance to a target student network. Most studies manually tie intermediate features of the teacher and student, and transfer knowledge through pre-defined links. However, manual selection often constructs ineffective links that limit the improvement from the distillation. There has been an attempt to address the problem, but it is still challenging to identify effective links under practical scenarios. In this paper, we introduce an effective and efficient feature distillation method utilizing all the feature levels of the teacher without manually selecting the links. Specifically, our method utilizes an attention-based meta-network that learns relative similarities between features, and applies identified similarities to control distillation intensities of all possible pairs. As a result, our method determines competent links more efficiently than the previous approach and provides better performance on model compression and transfer learning tasks. Further qualitative analyses and ablative studies describe how our method contributes to better distillation. The implementation code is available at open sourced¹.

Introduction

Knowledge distillation is the technique for transferring knowledge from a source neural network to a target neural network (Hinton, Vinyals, and Dean 2015). The source network, referred to as *a teacher*, indicates a large network that is highly regularized via pre-training, and the target network, referred to as *a student*, is a smaller network for a specific task. The pre-trained teacher directly informs the student of the solution and intermediate process of a problem, and this informative supervision enables fast and effective learning of the student. Based on knowledge distillation, recent studies have shown significant improvements in model compression (Hinton, Vinyals, and Dean 2015; Romero et al. 2014; Yim et al. 2017; Tian, Krishnan, and Isola 2019), cross-domain transfer learning (Orbes-Arteainst et al. 2019; Li et al. 2017; Asami et al. 2017), and continual learning (Li and Hoiem 2017; Hou et al. 2018).

For the success of knowledge distillation, various distillation methods were introduced. Starting from transferring output probability distributions of the teacher (Hinton, Vinyals, and Dean 2015), intermediate features representations (Romero et al. 2014) and their variants (Zagoruyko and Komodakis 2016a; Park et al. 2019; Tian, Krishnan, and Isola 2019) are investigated to identify what knowledge of the teacher helps to build a better student. However, most studies manually links the teacher and student features and perform distillation through the links individually. This manual link selection does not consider the similarity between the teacher and student features, so there is a risk of forcing an incorrect intermediate process to the student. Furthermore, the link selection has a limitation on fully utilizing the whole knowledge of the teacher by choosing a few of all possible links.

To compensate for the limitation, Jang *et al.* (Jang et al. 2019) apply a meta-networks, “learning to transfer (L2T)”, automatically determining the links. In more details, the meta-network consists of individual gates for all possible links, and each gate determines whether distillation through the link contributes to decreasing the classification loss of the student. Their results prove that knowledge distillation with the identified links provides better performance than those with manually selected links. However, the individual gates are not aware of each other although the distillation through the gates simultaneously affect the student. Moreover, their meta-learning scheme requires expensive inner-loop procedures to learn their meta-networks, thus its application can be limited under practical scenarios.

In this paper, we introduce a new feature linking method based on an attention mechanism (Xu et al. 2015; Vaswani et al. 2017), which is called attention-based feature distillation (AFD). Specifically, AFD utilizes an attention-based meta-network that identifies similar features between the teacher and student. The identified similarities are applied to control the distillation intensities for all possible feature pairs. Figure 1 shows an overview of our proposed distillation method to provide graphical descriptions.

When comparing from L2T, our proposed method considers the granularity of the teacher and student features to identify the importance of their links while L2T only uses information for a single pair in a narrow perspective. In addition, AFD learns from feature similarities without any inner-

*This work was done at internship in CLOVA AI Research, NAVER Corp.

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹github.com/clovaai/attention-feature-distillation.

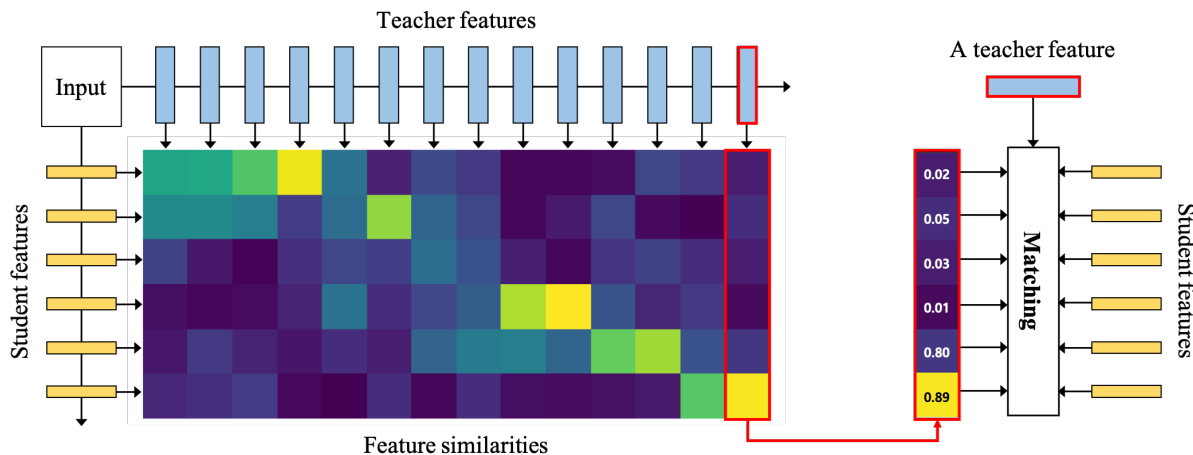


Figure 1: Overview of AFD. An attention-based model determines similarities between the teacher and student features. Knowledge from each teacher feature is transferred to the student with the identified similarities.

loop procedure but L2T learns from the classification loss, which requires expensive Hessian computation. In our experiment, we observe that L2T and ours have distinct linking results from the different objectives, but our method shows better or comparable results on multiple tasks with more efficient computation.

We conduct experiments for model compression on three image classification tasks such as CIFAR-100 (Krizhevsky, Hinton et al. 2009), tinyImageNet, and ImageNet (Deng et al. 2009) and for domain transfer on four specific tasks such as CUB200 (Wah et al. 2011), MIT67 (Quattoni and Torralba 2009), Stanford40 (Yao et al. 2011), and Stanford Dogs (Khosla et al. 2011) with a pre-trained large network on ImageNet. As a result, our method shows performance gains in most of experiments and the analyses on the identified feature links explains how our method works. Further ablation studies and sensitivity analysis provide a guideline to use our method.

Related Work

Hinton *et al.* introduced the concept of knowledge distillation (Hinton, Vinyals, and Dean 2015) by utilizing the output probability distributions of the teacher as a soft label to transfer knowledge. Also, intermediate features from the teacher have been proved to hold additional knowledge that can contribute to improving the student performance. Romero *et al.* (FitNet (Romero et al. 2014)) proposed feature-based distillation that couples the teacher and student features and induce the student to mimic the paired teacher features. However, due to a capacity gap between the teacher and the student, it is challenging for the student to mimic the exact teacher features.

To address the problem, recent works focused on propagating core knowledge from the teacher features. Zagoruyko *et al.* (Zagoruyko and Komodakis 2016a) simplified teacher features by applying channel-wise summations and led the student to learn core knowledge of refined teacher features. Kim *et al.* (Kim, Park, and Kwak 2018) extracted low-dimensional

representations of the features via multiple auto-encoders and transfer them to the student. Relational knowledge distillation aims to transfer relation knowledge between data instances (Park et al. 2019; Tung and Mori 2019; Peng et al. 2019; Liu et al. 2019). In the other direction of feature refinement methods, the distillation regularization terms have been explored to allow the student to accept more knowledge. Ahn *et al.* (Ahn et al. 2019) transferred knowledge by maximizing the mutual information between the feature of the teacher and student. Huang *et al.* (Huang and Wang 2017) utilized the maximum mean discrepancy to propagate knowledge from the teacher features. Tian *et al.* (Tian, Krishnan, and Isola 2019) applied the contrastive learning scheme on relational knowledge distillation.

Although distillation methods on how to refine and propagate knowledge have been continuously advanced as the above, it is still remaining problem how to link intermediate features between the teacher and student. Our method is placed to solve the problem as like L2T (Jang et al. 2019). However, L2T and ours have different properties to identify the links.

Attention-based Feature Distillation

Let $\mathbf{h}^T = \{h_1^T, \dots, h_T^T\}$ be a set of the feature candidates from the teacher and $\mathbf{h}^S = \{h_1^S, \dots, h_S^S\}$ be a set of feature candidates from the student where T and S indicate the numbers of the candidates from the teacher and student, respectively. Each candidate has its own feature map size and channel dimension as $h \in \mathbb{R}^{H \times W \times d}$ where H , W , and d indicate the height, width, and channel dimension, respectively. When two sets of the candidates are given, AFD aims to identify similarities for all possible combinations ($S \times T$ pairs) and transfer knowledge of the teacher candidates to the student with the identified similarities.

Figure 2 shows the overview of our proposed network. As can be seen, the feature candidates are compared in two directions with the two pooling methods: global average pooling and channel-wise pooling. The similarity identified

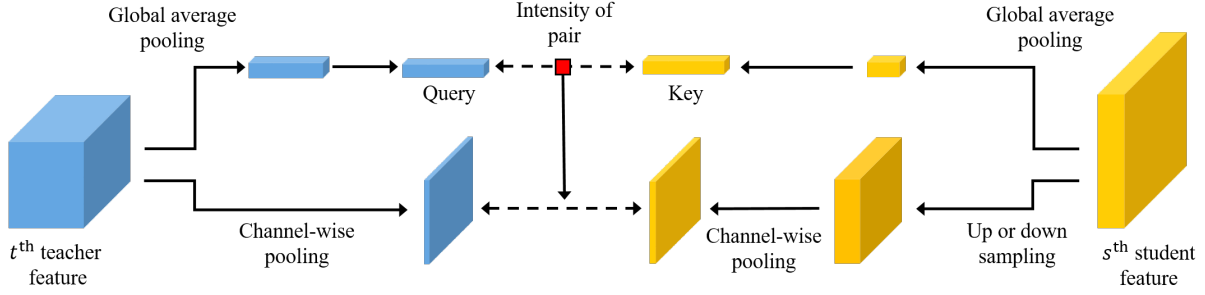


Figure 2: Overview of the proposed meta-network. The globally pooled features are utilized to estimate the similarities and the channel-wisely averaged features are used to calculate the distance between the features.

by two globally pooled features is used as an intensity for transferring knowledge through the distance defined by the channel-wisely averaged features. In order to identify the similarity between h_t^T and h_s^S , AFD adopts a query-key concept of the attention mechanism (Xu et al. 2015; Vaswani et al. 2017). Specifically, each teacher feature generates a query, \mathbf{q}_t , and each student feature identifies a key, \mathbf{k}_s . The followings describe \mathbf{q}_t and \mathbf{k}_s in mathematical expressions;

$$\begin{aligned} \mathbf{q}_t &= f_Q(W_t^Q \cdot \phi^{HW}(h_t^T)), \\ \mathbf{k}_s &= f_K(W_s^K \cdot \phi^{HW}(h_s^S)). \end{aligned} \quad (1)$$

Here, $\phi^{HW}(\cdot)$ indicates a global average pooling. f_Q and f_K are activation function of the query and key. $W_t^Q \in \mathbb{R}^{d \times d_t^T}$ and $W_s^K \in \mathbb{R}^{d \times d_s^S}$ are linear transition parameters for the t -th query and the s -th key. It should be noted that the features have different transition weights since they have different properties through their different levels, *i.e.* a low-level visual feature can represent a line and a high-level visual features can represent an object. Therefore, we apply different transition weights to each features.

By utilizing the queries and keys, attention values that represent relations between teacher and student candidates are calculated with a ‘‘softmax’’ function;

$$\begin{aligned} \alpha_t &= \text{softmax}([(\mathbf{q}_t^T W_1^{Q-K} \mathbf{k}_{t,1} + (\mathbf{p}_t^T)^T \mathbf{p}_1^S) / \sqrt{d}, \\ &\dots, (\mathbf{q}_t^T W_S^{Q-K} \mathbf{k}_{t,S} + (\mathbf{p}_t^T)^T \mathbf{p}_S^S) / \sqrt{d}]). \end{aligned} \quad (2)$$

Here, we introduce additional weight parameters; a bilinear weight, $W_t^{Q-K} \in \mathbb{R}^{d \times d}$, and positional encodings, $\mathbf{p}_t^T \in \mathbb{R}^d$ and $\mathbf{p}_s^S \in \mathbb{R}^d$. The bilinear weight is applied to generalize the attention value from different source ranks since the query and key are identified from different dimensional features (Pirsiavash, Ramanan, and Fowlkes 2009; Kim, Jun, and Zhang 2018). The positional encodings are utilized to share common information over different instances (Vaswani et al. 2017). α_t is the attention vector that capture relation between the t -th teacher feature and whole student features. By utilizing α_t , the teacher feature, h_t^T , enables to transfer its knowledge selectively to student features.

The final distillation term forms as

$$\mathcal{L}_{AFD} = \sum_t \sum_s \alpha_{t,s} \left\| \tilde{\phi}^C(h_t^T) - \tilde{\phi}^C(\hat{h}_s^S) \right\|_2, \quad (3)$$

where $\tilde{\phi}^C$ indicates a combined function of a channel-wise average pooling layer with L2 normalization, $\mathbf{v} / \|\mathbf{v}\|_2$, by

following (Zagoruyko and Komodakis 2016a). In addition, \hat{h}_s^S is up-sampled or down-sampled from h_s^S to match the feature map size to those of the teacher features.

Finally, the regularization term is added to the total loss function as following;

$$\mathcal{L}_{\text{Student}} = \mathcal{L}_{\text{cls}} + \beta \mathcal{L}_{AFD}, \quad (4)$$

where \mathcal{L}_{cls} is the classification loss with ground-truth labels and β is a trade-off parameter controlling the impact of the proposed distillation loss. We use cross entropy for \mathcal{L}_{cls} . Using the loss, the student and the attention-based network are trained simultaneously. It should be noted that the AFD network is trained only with \mathcal{L}_{AFD} that represents the weighted similarities for all possible feature pairs, so AFD does not require expensive Hessian computations to connect its parameters with \mathcal{L}_{cls} .

Experiments

We evaluate our proposed method on model compression tasks that train a smaller, or better model under a limitation of the model capacity and transfer learning tasks that obtain a better model in a specific domain by utilizing a pre-trained model. Following by the quantitative evaluations, we qualitatively analyze how our method works. Finally we provide ablation studies to provide further properties of our methods.

In our experiment, our baseline methods are a traditional knowledge distillation method (KD) firstly introduced by Hinton *et al.* (Hinton, Vinyals, and Dean 2015), three popular feature-level distillation methods (FitNet (Romero et al. 2014), ATT (Zagoruyko and Komodakis 2016a), CRD (Tian, Krishnan, and Isola 2019)) that require a manual feature matching, and one feature-level distillation method (L2T) automatically linking the teacher and student features. For CRD, we set the number of negative samples same as the batch size of each experiment. Note that KD is applied to all baselines to reveal additional gains from the feature distillation methods.

Model Compression

We demonstrate the effectiveness of the proposed distillation method on model compression tasks. The experiments are conducted on three popular benchmark datasets such as *CIFAR-100* (Krizhevsky, Hinton et al. 2009), *tinyImageNet*, and *ImageNet* (Deng et al. 2009). We utilize Residual Network (ResNet) (He et al. 2016) and Wide Residual Network

	Same style					Different style	
Teacher	ResNet56	ResNet110	ResNet110	WRN-40-2	WRN-40-2	WRN-40-2	ResNet34
Student	ResNet20	ResNet20	ResNet56	WRN-16-2	WRN-40-2	ResNet56	WRN-28-2
Teacher	0.7254	0.7409	0.7409	0.7620	0.7620	0.7620	0.7860
Student	0.6940	0.6940	0.7254	0.7289	0.7620	0.7254	0.7532
KD	0.7098	0.7081	0.7483	0.7499	0.7763	0.7497	0.7648
FitNet	0.7005	0.7002	0.7411	0.7522	0.7766	0.7506	0.7644
ATT	0.7054	0.7081	0.7488	0.7520	0.7778	0.7516	0.7720
RKD	0.7043	0.7076	0.7477	0.7459	0.7762	0.7439	0.7632
CRD	0.7095	0.7091	0.7512	0.7515	0.7780	0.7525	0.7697
L2T	0.7037	0.7001	0.7457	0.7486	0.7678	0.7463	0.7640
Ours	0.7153	0.7138	0.7539	0.7547	0.7813	0.7540	0.7747

Table 1: Performance comparison on CIFAR-100. The teachers and the students have same or different architectural style. All experiments are repeated 5 times.

(WRN) (Zagoruyko and Komodakis 2016b) architectural styles. First, we conduct an experiment on CIFAR-100 that consists of 32×32 sized color images for 100 object classes and has 50K training and 10K validation images. For data augmentation, the horizontal flipping and random cropping are applied. We set the batch size as 64 and the maximum iteration as 240 epochs. All models are trained with stochastic gradient descent with 0.9 of momentum, weight decay as 5×10^{-4} , initial learning rate as 0.05, and divide it by 10 at 150, 180, 210 epochs. For the baseline methods, we use their official code and their hyper-parameters. For the distillation loss of our model, we apply $\{30, 50, 100, 200\}$ of beta, β , and choose the best performer. We provide how the performance changes depending on the value of β in *Sensitivity Analysis* section. For the feature candidates of our method, we use all output features of the teacher and student residual blocks as candidate for all experiments, except ResNet110. For ResNet110, we skip one for every two residual blocks and used only half of the entire residual blocks. In detail, the number of the candidates become 9 for ResNet20, 27 for ResNet56, 27 for ResNet110, 8 for ResNet18, 16 for ResNet34, 6 for WRN-16-2, 12 for WRN-28-2, and 18 for WRN-40-2.

Table 1 shows our experiment settings and results on the CIFAR-100 dataset with various network architecture. We pre-train large teacher networks and utilize them to train smaller or same-scaled student networks. The experiments are divided into two groups according to the architectural style of the teacher and student. When considering the student without knowledge distillation, all students shows the worst performance over all experiment settings. With manually linked feature pairs, the baseline models including KD, FitNet, ATT, RKD and CRD show better performances than the vanilla students. When applying L2T that identifies feature links with individual gates, we observed worse performance than other baseline methods even though it identifies beneficial feature pairs to meet their objective. The reason of the degradation is that L2T tends to propagate high-level knowledge of the teacher to the low-level feature of the student (See *Qualitative Studies on Feature Attention* section).

	Tiny ImageNet		ImageNet
Teacher	ResNet34	ResNet34	ResNet34
Student	ResNet18	ResNet34	ResNet18
Teacher	0.6750	0.6750	0.7355
Student	0.6530	0.6750	0.7028
KD	0.6818	0.6971	0.7066
FitNet	0.6779	0.6892	0.7085
ATT	0.6782	0.6961	0.7093
RKD	0.6772	0.6846	0.7137
CRD	0.6819	0.6968	0.7135
Ours	0.6880	0.6981	0.7138

Table 2: Performance comparison on large-scale datasets; Tiny ImageNet and ImageNet.

Intuitively, the low-level of a small network cannot mimic the high-level of a large network, which adversely affects performance. Our method that utilizes an attention mechanism to identify similar features between the teacher and student shows the best performance over all experiment settings. In particular, our method shows an improvement over ATT which uses the same feature distance for distillation. In other words, the proposed linking method significantly contributes to distillation performance.

In order to validate our method on more real world environment, we compare our method from other baseline methods on tinyImageNet and ImageNet (Deng et al. 2009). The tinyImageNet dataset consists of 64×64 sized 100K training and 10K validation images for 200 object classes and the ImageNet dataset includes 1.2M training and 50K validation large-scale images for 1K object classes.

For tinyImageNet, we pad the images to 72×72 and then randomly cropped to 64×64 and flipped for data augmentation. We set the batch size as 128 and the maximum iteration as 200 epochs. All models are trained with stochastic gradient descent with momentum 0.9. We set weight decay as 5×10^{-4} , initial learning rate as 0.1, and we divide the learning rate by 5 at 60, 120, 150 and 180 epochs. We adopt ResNet34 for

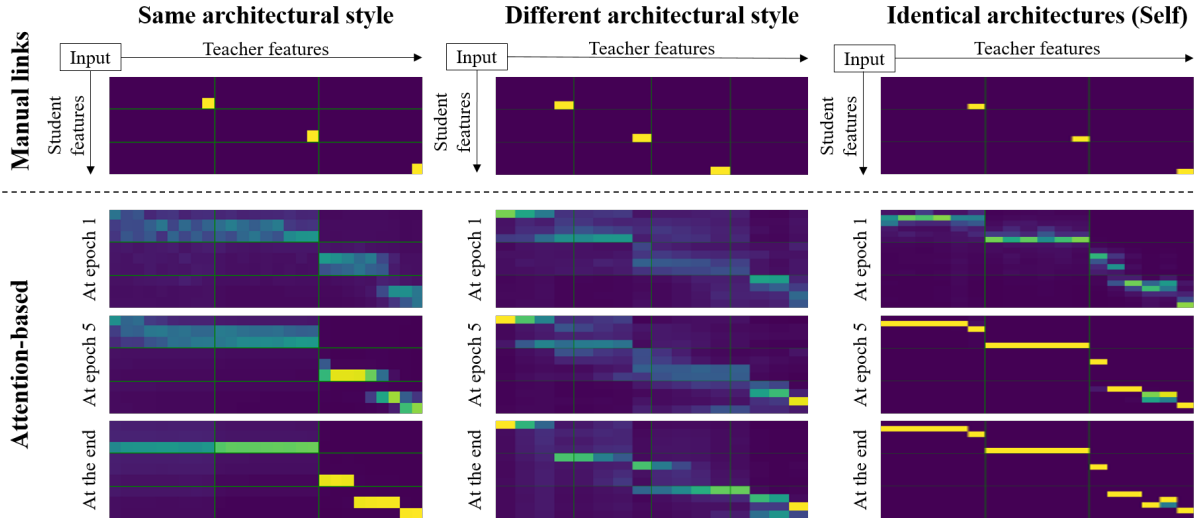


Figure 3: Manual and attention-based feature links for knowledge distillation. Rows and columns of matrices indicate the student and teacher features, respectively. Each matrix is the average overall α at the corresponding training epoch. The pairs are compared on multiple distillation settings; ResNet56 \rightarrow ResNet20 (Same architectural style), ResNet34 \rightarrow WRN-28-2 (Different architectural style), and WRN-40-2 \rightarrow WRN-40-2 (Self). The manual feature links are not changed once they are selected, but AFD are adaptively selected during the distillation.

the teacher and utilize ResNet34 and ResNet18 both for the student. Unlike the ResNet34 and ResNet18 architecture for large image classification, we resize the first convolutional filter size from 7 to 3. For ImageNet, we randomly crop size of 224×224 of each of images and flipped for data augmentation. We optimize the student with initial learning rate as 0.1, divide it by 10 at 30, 60, 90 epochs, and set the maximum iteration as 100 epochs. We set weight decay as 10^{-4} and the batch size as 256. We utilize ResNet34 for the teacher and ResNet18 for the student. For tinyImageNet and ImageNet, we set the hyperparameter, β , as 50.

Table 2 shows the experiment results on tinyImageNet and ImageNet. The proposed method achieve better performance on large-scale datasets than other baseline knowledge distillation methods. It should be noted that we do not conduct L2T on these large-scaled datasets due to the heavy time complexity to update their meta-network. In addition, our method shows constantly better performances than ATT that holds the same distillation loss for the manually selected feature pairs.

Transfer Learning

Transfer learning with knowledge distillation utilizes the teacher pre-trained on a source domain task to train the student for a target domain task. We investigate the effectiveness of our method on transfer learning tasks. We adopt a ResNet34 pre-trained on ImageNet as the teacher network and transfer its knowledge into the students for four tasks, such as Caltech-UCSD Bird (CUB 200) (Wah et al. 2011), MIT Indoor Scene Recognition (MIT67) (Quattoni and Torralba 2009), Stanford 40 Actions (Stanford40) (Yao et al. 2011) and Stanford Dogs (Khosla et al. 2011). CUB 200 consists of 5k training and 6k validation images with 200 bird species. MIT67 contains 5k training and 1k validation images

	CUB200	MIT67	Stanford40	Stanford Dogs
Scratch	0.4215	0.4891	0.3693	0.5808
FitNet	0.4893	0.5488	0.4450	0.6725
ATT	0.5774	0.5918	0.5929	0.6970
L2T	0.6505	0.6485	0.6308	0.7808
Ours	0.6829	0.6647	0.6792	0.7606

Table 3: Performance comparison on multi-domain transfer learning tasks; from a ResNet34 model pre-trained with ImageNet (source domain) to a ResNet18 models for CUB200, MIT67, Stanford40, and Stanford Dogs (target domains). Scratch, ATT and L2T are referred from (Jang et al. 2019). All experiments are repeated 3 times.

with 67 type of indoor scenes. Stanford40 has 4k training and 5k validation images with 40 human actions. Stanford Dogs consists of 12k training and 8k validation images with 120 dog species. The images of transfer learning datasets consists of large scale images. The student architecture for the specific target tasks are set as ResNet18. All experiment settings are followed by those of L2T (Jang et al. 2019) and our hyper-parameter, β , is set as 1,000.

Table 3 summarizes the experiment results for the multi-domain transfer learning tasks. For most datasets, our method shows better performance than L2T and ATT. In transfer learning, dataset from target tasks give limited information. Therefore, training without any knowledge transfer (scratch) shows the worst performance with large margin. Comparing ATT with L2T and our method, it can be seen that identifying the feature links is effective in transfer learning tasks.




Figure 4: Activation map corresponding to each distillation methods and the teacher. The teacher is ResNet56 and the students are ResNet20. The colored boxes indicate the links from the teacher to the student.

Qualitative Studies on Feature Attention

Feature Links. Here, we analyze the attention values, α , learned from our proposed method in order to provide its further properties. First, we observe the attention maps of various architecture pairs, such as the same architectural style, different architectural style and identical architecture (self).

Figure 3 shows the feature links α during the training phases. As can be seen, the traditional knowledge distillation methods manually link the teacher and student features and transfer teacher’s knowledge only through the pre-defined links. In contrast to the manual feature links, our method identifies feature links in a data-driven way and thus the feature links α are changed over training steps and converged at the end. In the case of the same architectural style, low and mid-levels of the teacher features are linked with a low-level student feature. This indicates our method affects the student to use more layers to learn high-level of the teacher features. In the case of the different architectural style, the lowest-level and the highest-level features are linked among themselves and mid-level features are smoothly connected. These results show that it is difficult to manually create feature links in different architecture styles. In the case of the identical architectures, the features are linked in the order of the levels but the student tends to use more layers to extract high-level features. Based on observations, we can infer that AFD tends to link features in the same level. However, when the teacher and student have different architectures, the connection is extended to other levels. It is an advantage of our method that the feature links can be changed and extended according to the difference between the teacher and the student’s architecture, and improves the distillation performance in various architecture settings.

Activation Map Figure 4 shows activation maps of a teacher and a students trained with knowledge distillation methods such as ATT with manual feature pairs, L2T, and ours. When comparing links between the teacher and student features, ATT has manually set ordered links through the levels of the features. In the L2T, the identified links switch the order of the levels; the last feature of the teacher is connected with the mid-level feature of the student (green box), and the mid-level feature of the teacher is linked to the last feature of the student (blue box). The identified links of the




Figure 5: Sensitivity analysis for β . Red line indicates accuracy of model compression task (WRN-40-2 \rightarrow WRN-16-2). Blue line indicates accuracy of transfer learning task (ResNet34 \rightarrow ResNet18)

L2T may hinder the student training by transferring different order of features learned by the teacher, see green and blue boxes. In contrast to both, our proposed method spreads the high-level features of the teacher to various feature levels of the student (blue, red, and green boxes) and identify links for the student to train while keeping the order of the features from the teacher. More interestingly, compared with ATT, the low-level and the mid-level features of the AFD student tends to mimic activated regions of the high-level features of the teacher.

Sensitivity Analysis

To investigate the impact of the hyperparameter of AFD, β , we evaluate our model by varying the value of β . β is used to train the attention map, α , which determines the links of the AFD network, and to decide the degree of how much the student mimic the teacher features. We perform the sensitivity analysis for β with the model compression and the transfer learning tasks. Figure 5 shows the accuracy of each task.

For the model compression task, we use WRN-40-2 as the teacher network and WRN-16-2 as the student network with CIFAR-100 dataset. As can be seen from the red line in Figure 5, the accuracy decreases when β is more than 1,000 compared to the interval between 30 to 200. For the transfer learning task, we use ResNet34 as the teacher network and ResNet18 as the student network with MIT67 dataset. In contrast with model compression task, the transfer learning task shows better performance when the β value is relatively large. However, we observe that the different result from the model compression task using the same network architecture of transfer learning setting with ImageNet dataset (we use β as 50 for ImageNet). The rationale behind the gap of the hyperparameter lies in the degree of reliance on teachers’ knowledge caused by the size of the dataset and the relevance between source and target tasks (Tung and Mori 2019; Ahn et al. 2019; Park et al. 2019)

Ablation Studies

Linking methods. In order to reveal the benefits of the similarity-based links, we compare ours from other linking

Teacher candi.	Linking method	ResNet56 \rightarrow ResNet20
Random	Random link	0.6945 ± 0.002
	Ordered	0.6999 ± 0.003
	AFD (co-train)	0.7105 ± 0.002
Equal interval	Random link	0.6941 ± 0.005
	Ordered	0.7025 ± 0.002
	AFD (pre-trained)	0.7121 ± 0.001
	AFD (co-train)	0.7135 ± 0.002

Table 4: Ablation studies on selecting the teacher candidates and linking them to the student features. All experiments are repeated 5 times.

# of candi. (teacher)	# of candi.(student)		
	3	6	9
3	0.7140	0.7097	0.7115
9	0.7109	0.7149	0.7135
27	0.7121	0.7152	0.7153

(a) ResNet56 \rightarrow ResNet20

# of candi. (teacher)	# of candi.(student)		
	3	6	12
4	0.7698	0.7704	0.7730
8	0.7714	0.7717	0.7724
16	0.7706	0.7733	0.7747

(b) ResNet34 \rightarrow WRN-28-2

Table 5: Ablation studies on the numbers of the candidates for both the teacher and the student. The candidates are set as the output features of the residual blocks that are sequentially equidistant between themselves.

methods in Table 4. In this experiment, we set the numbers of the teacher and student candidates same as 9 for the same architectural style (ResNet56 \rightarrow ReNet20) to use all possible student features. We choose the teacher candidates in two ways; “Random” that randomly selects the candidates from all features upon all residual blocks and “Equal interval” that selects the features that are sequentially equidistant between themselves. For the linking method, we evaluate three linking methods; “Random link”, “Ordered”, “AFD”.

As shown in Table 4, “Ordered” shows better performance than “Random link”. It should be noted that the combination of “Equal interval” and “Ordered” is usually used when manually selecting links between the teacher and student features. Interestingly, when applying the links identified from the pre-trained AFD, we observe the performance improvement although we only change links from the usually link setting. The result proves that there is more effective way to set the links than manually decided links. In addition, when training AFD together, AFD shows the best performances. This experiments prove the superiority of AFD on identifying links of the feature pairs and transferring teacher’s knowledge to the student.

Number of candidates. Here, we analyze the impact of the numbers of the teacher and student candidates. For this

L1	L2	KL	Cosine
0.7513	0.7547	0.7523	0.7541

Table 6: Accuracy according to distance metrics.

A1	A2	Max
0.7520	0.7547	0.7528

Table 7: Accuracy according to pooling methods.

experiment, the candidates are set as the output features of the residual blocks that are sequentially equidistant between themselves. Table 5 shows distillation performances over varying numbers of the candidates. When the student candidates are more than half of the total features, and using all teacher features (27 for ResNet56 and 16 for ResNet34) provides better student performances. In the other hand, with the small number of the student candidates, using all teacher features causes a information bottleneck and degrades the performances. This experiments provide a guidance to choose the number of the candidates.

Distance Metrics and Pooling Methods We use L2 distance for distilling the teacher’s feature to the student according to the trained link α , see equation 3. However, other distance metric can be used for distillation (Huang and Wang 2017; Ahn et al. 2019; Zagoruyko and Komodakis 2016a) and it may affect the behavior of the student. Therefore, we explore four distance metrics, L1, L2, KL divergence, and cosine similarity, on model compression task with WRN-40-2 as the teacher network and WRN-16-2 as the student network. Table 6 shows that L2 distance is the optimal metric for our experiments, so we use L2 distance for the whole experiments.

Also, channel-wise pooling method, $\tilde{\phi}^C$, applied to feature in equation 3 may affect the performance of distillation. Therefore, we compare three channel-wise pooling methods including max-pooling ($\max_i |h_i|$) and average-pooling ($\frac{1}{q} \sum_i |h_i|^p$). We denote A1 and A2 as average pooling methods with $p = 1$ and $p = 2$, respectively. As can be seen in the table 7, A2 shows the best performance. Therefore, we use A2 for all experiment in this paper.

Conclusion

In this paper, we have proposed an attention-based distillation method adaptively transferring knowledge of teacher features to multiple levels of the student layers. With the proposed method, the teacher features are linked with the student features with the attention map and the student learns from the teacher through the identified links. The proposed method is efficiently learned simultaneously during the student’s training phase while the previous feature linking method requires an additional inner-loop procedure. Our experiment proves the benefits of the proposed method on two knowledge distillation applications such as model compression and transfer learning. Our further analysis shows that our method adjusts the feature levels of the student regardless the architectural styles of the teacher and student and provides better performance than the baseline methods.

References

- Ahn, S.; Hu, S. X.; Damianou, A.; Lawrence, N. D.; and Dai, Z. 2019. Variational Information Distillation for Knowledge Transfer. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Asami, T.; Masumura, R.; Yamaguchi, Y.; Masataki, H.; and Aono, Y. 2017. Domain adaptation of DNN acoustic models using knowledge distillation. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 5185–5189. IEEE.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- 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.; Vinyals, O.; and Dean, J. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Hou, S.; Pan, X.; Change Loy, C.; Wang, Z.; and Lin, D. 2018. Lifelong learning via progressive distillation and retrospection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 437–452.
- Huang, Z.; and Wang, N. 2017. Like what you like: Knowledge distill via neuron selectivity transfer. *arXiv preprint arXiv:1707.01219*.
- Jang, Y.; Lee, H.; Hwang, S. J.; and Shin, J. 2019. Learning What and Where to Transfer. In *International Conference on Machine Learning*, 3030–3039.
- Khosla, A.; Jayadevaprakash, N.; Yao, B.; and Li, F.-F. 2011. Novel dataset for fine-grained image categorization: Stanford dogs. In *Proc. CVPR Workshop on Fine-Grained Visual Categorization (FGVC)*, volume 2.
- Kim, J.; Park, S.; and Kwak, N. 2018. Paraphrasing complex network: Network compression via factor transfer. In *Advances in Neural Information Processing Systems*, 2760–2769.
- Kim, J.-H.; Jun, J.; and Zhang, B.-T. 2018. Bilinear attention networks. In *Advances in Neural Information Processing Systems*, 1564–1574.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images .
- Li, J.; Seltzer, M. L.; Wang, X.; Zhao, R.; and Gong, Y. 2017. Large-Scale Domain Adaptation via Teacher-Student Learning. *Proc. Interspeech 2017* 2386–2390.
- Li, Z.; and Hoiem, D. 2017. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence* 40(12): 2935–2947.
- Liu, Y.; Cao, J.; Li, B.; Yuan, C.; Hu, W.; Li, Y.; and Duan, Y. 2019. Knowledge Distillation via Instance Relationship Graph. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7096–7104.
- Orbes-Arteainst, M.; Cardoso, J.; Sørensen, L.; Igel, C.; Ourselin, S.; Modat, M.; Nielsen, M.; and Pai, A. 2019. Knowledge distillation for semi-supervised domain adaptation. In *OR 2.0 Context-Aware Operating Theaters and Machine Learning in Clinical Neuroimaging*, 68–76. Springer.
- Park, W.; Kim, D.; Lu, Y.; and Cho, M. 2019. Relational knowledge distillation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3967–3976.
- Peng, B.; Jin, X.; Liu, J.; Li, D.; Wu, Y.; Liu, Y.; Zhou, S.; and Zhang, Z. 2019. Correlation congruence for knowledge distillation. In *Proceedings of the IEEE International Conference on Computer Vision*, 5007–5016.
- Pirsiavash, H.; Ramanan, D.; and Fowlkes, C. C. 2009. Bilinear classifiers for visual recognition. In *Advances in neural information processing systems*, 1482–1490.
- Quattoni, A.; and Torralba, A. 2009. Recognizing indoor scenes. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 413–420. IEEE.
- Romero, A.; Ballas, N.; Kahou, S. E.; Chassang, A.; Gatta, C.; and Bengio, Y. 2014. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*.
- Tian, Y.; Krishnan, D.; and Isola, P. 2019. Contrastive representation distillation. *arXiv preprint arXiv:1910.10699*.
- Tung, F.; and Mori, G. 2019. Similarity-preserving knowledge distillation. In *Proceedings of the IEEE International Conference on Computer Vision*, 1365–1374.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems*, 5998–6008.
- Wah, C.; Branson, S.; Welinder, P.; Perona, P.; and Belongie, S. 2011. The caltech-ucsd birds-200-2011 dataset .
- Xu, K.; Ba, J.; Kiros, R.; Cho, K.; Courville, A.; Salakhudinov, R.; Zemel, R.; and Bengio, Y. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, 2048–2057.
- Yao, B.; Jiang, X.; Khosla, A.; Lin, A. L.; Guibas, L.; and Fei-Fei, L. 2011. Human action recognition by learning bases of action attributes and parts. In *2011 International Conference on Computer Vision*, 1331–1338. IEEE.
- Yim, J.; Joo, D.; Bae, J.; and Kim, J. 2017. A Gift From Knowledge Distillation: Fast Optimization, Network Minimization and Transfer Learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Zagoruyko, S.; and Komodakis, N. 2016a. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. *arXiv preprint arXiv:1612.03928*.
- Zagoruyko, S.; and Komodakis, N. 2016b. Wide residual networks. *arXiv preprint arXiv:1605.07146*.