

Knowledge Distillation via the Target-aware Transformer

Anonymous CVPR submission

Paper ID 7432

Abstract

Knowledge distillation becomes a *de facto* standard to improve the performance of small neural networks. Most of the previous works propose to regress the representational features from the teacher to the student in a one-to-one spatial matching fashion. However, people tends to overlook the fact that, due to the architecture differences, the semantic information on the same spatial location usually vary. This greatly undermines the underlying assumption of the one-to-one distillation approach. To this end, we propose a novel one-to-all spatial matching knowledge distillation approach. Specifically, we allow each pixel of the teacher feature to be distilled to all spatial locations of the student features given its similarity, which is generated from a target-aware transformer. Our approach surpasses the state-of-the-art methods by a significant margin on various computer vision benchmarks, such as ImageNet, Pascal VOC and COCOStuff10k.

1. Introduction

Knowledge distillation [18, 28] refers to a simple technique to improve the performance of any machine learning algorithms. One common scenario is to distill the knowledge from a larger teacher neural network to a smaller student one, such that the performance of student model can be significantly boosted comparing to training the student model alone. Concretely, people formulate an external loss function that guides the student feature map to mimic teacher’s. Recently, it has been applied to various downstream applications, such as model compression [39, 43], continual learning [23], and semi-supervised learning [8].

Earlier works only distill the knowledge from the final layer of neural networks, for example, the “logits” in image classification task [1, 18]. Recently, people discover that distilling the intermediate feature maps is a more effective approach to boost the student’s performance. This line of works encourage similar patterns to be elicited in the spatial dimensions [33, 45], and is constituted as state-of-the-art knowledge distillation approach [7, 21].




Figure 1. **Illustration of semantic mismatch.** Suppose that teacher and student are the 3-layers and 2-layers convnets with kernel size 3×3 and stride 1×1 . (a) shows the receptive field of the middle pixel of the final feature map, where the blue box represents the teacher’s receptive field and the orange box is that of the student’s. Since teacher’s model has more convolutional operations, the resulting teacher feature map has a larger receptive field thus contains richer semantic information. (b) Hence, directly regressing the student and teacher’s feature in a one-to-one spatial matching fashion may be suboptimal. (c) We proposed a one-to-all knowledge distillation via a target-aware transformer that can let the teacher’s spatial components be distilled to the entire student feature maps.

To compute the distillation loss of the aforementioned approach, one need to select the source feature map from the teacher and the target feature map from the student, where these two feature maps must have the same spatial dimension. As shown in Figure 1 (b), the loss is computed in a one-to-one spatial matching fashion, that is formulated as a summation of the distance between the source and the target features at each spatial location. One underlying assumption of this approach is the spatial information of each pixel is the same. In practice, this assumption is commonly not valid due to the fact that student model usually has fewer convolutional layers than the teacher. One example

is shown in Figure 1 (a), even at the same spatial location, the receptive field of student feature is often significantly smaller than the teacher’s and thus contains less semantic information. In addition, recent works [5, 10, 38, 44] evidences the importance of receptive field’s influence on the model representation power. Such discrepancy is a potential reason that the current one-to-one matching distillation leads to sub-optimal results.

To this end, we propose a novel one-to-all spatial matching knowledge distillation approach. In Figure 1 (c), our method distills the teacher’s features at each spatial location into all components of the student features through a parametric correlation, *i.e.*, the distillation loss is a weighted summation of all student components. To model such correlation, we formulate a transformer structure that reconstructs the corresponding individual component of the student features and produces an alignment with the target teacher feature. We dubbed this target-aware transformer. As such, we use parametric correlations to measure the semantic distance conditioned on the representational components of student feature and teacher feature to control the intensity of feature aggregation, which address the downside of one-to-one matching knowledge distillation.

As our method computes the correlation between feature spatial locations, it might become intractable when feature maps are large. To this end, we extend our pipeline in a two-step hierarchical fashion: 1) instead of computing correlation of all spatial locations, we split the feature maps into several groups of patches, then performs the one-to-all distillation within each group; 2) we further average the features within a patch into a single vector to distill the knowledge. This reduces the complexity of our approach by order of magnitudes.

We evaluate the effectiveness of our method on two popular computer vision tasks, image classification and semantic segmentation. On the ImageNet classification dataset, the tiny ResNet18 student can be boosted from 70.04% to 72.41% in terms of the top-1 accuracy, and surpasses the state-of-the-art knowledge distillation by 0.8%. As for the segmentation task on COCOstuff10k, comparing to the previous approaches, our approach is able to boost the compact MobilenetV2 architecture by 1.75% in terms of the mean intersection of union (mIoU).

Our contributions can be summarized as follows:

- We propose the knowledge distillation via a target-aware transformer, which enables the whole student to mimic each spatial component of the teacher respectively. In this way, we can increase the matching capability and subsequently improve the knowledge distillation performance.
- We propose the hierarchical distillation to transfer local features along with global dependency instead of the original feature maps. This allows us to apply the proposed method to applications, which are suffered from heavy

computational burden because of the large size of feature maps.

- We achieve state-of-the-art performance compared against related alternatives on multiple computer vision tasks by applying our distillation framework.

2. Related Works

The seminal work [18] introduced the idea of knowledge distillation. Specifically, Hinton *et al.* proposed to distill the logits (before softmax layer) from teacher to student by minimizing the KL divergence, where a temperature factor is applied to soften the logits.

Since feature map contains richer representation, Romero *et al.* [33] introduced the intermediate layer transfer between teacher and student. Lately, AT [45] proposed several statistical methods to highlight the dominating area of the feature map and discarded low-response area as noise. Chen *et al.* [3] proposed the semantic calibration which allowed the student to learn from the most semantic-related teacher layer. In [21], the feature similarities between teacher and student were calculated and then were used as weights to balance the feature matching. These early methods intuitively established the links between knowledge source (teacher) and distillation terminal (student) in the one-to-one manner by spatial order. However, they over-estimated the prior of spatial order while neglected the issues of semantic mismatch, *i.e.*, the pixels of teacher feature map often contains richer semantic compared to that of student on the same spatial location. We found that some works [19, 25, 30–32, 40, 43], though unintended, have been proposed to relax the spatial constrain during feature transfer. Typically, they defined the relational graph, and similarity matrix in the feature space of teacher network and transferred it to the student network. For instances, Tung and Mori [40] calculated the similarity matrix where each entry encoded the similarity between two instances. Liu *et al.* [25] measured the correlation between channels by inner-product. They condensed and compressed the entire feature to some properties (often scalar) and thus collapsed the spatial information. On the other, such process damaged the original teacher feature and may lead to sub-optimal solution.

The spread of KD has also driven some methods designed for specific vision tasks including video captioning [29], action recognition [9, 41], object detection [4, 11, 46] and semantic segmentation [16, 26, 42]. Regarding the semantic segmentation, these methods are indeed related to relation knowledge distillation which computes similarity matrix [40]. To investigate the potential of our method, we also adapt the method to semantic segmentation with hierarchical distillation.




Figure 2. Illustration of our framework. (a) **Target-aware Transformer**. Conditioned on the teacher feature and the student feature, the transformation map $Corr.$ is computed and then applied on the student feature to reconfigure itself, which is then asked to minimize the L_2 loss with the corresponding teacher feature. (b) **Patch-group Distillation**. Both teacher and student features are to be sliced and rearranged as groups for distillation. By concatenating the patches within a group, we explicitly introduce the spatial correlation among the patches beyond the patches themselves. (c) **Anchor-point Distillation**. Each color indicates a region. We use average pooling to extract the *anchor* within a local area of the given feature map, forming the new feature map of a smaller size. The generated anchor-point features will participate in the distillation.

3. Method

In this section, we first briefly describe the fundamental elements of feature map knowledge distillation and then introduce the general formulation of our knowledge distillation via a target-aware transformer. As our method computes the point-wise correlation of the given feature maps, the computational complexity becomes intractable on large-scale features, we then introduce the hierarchical distillation approach to address this limitation.

3.1. Formulation

Suppose the teacher and the student are two convolutional neural networks, denoted by T and S . $F^T \in \mathbb{R}^{H \times W \times C}$ and $F^S \in \mathbb{R}^{H \times W \times C'}$ denote the teacher feature and student feature respectively, where H and W are the height and width of the feature map, and C represents the channel numbers. In the pioneer work [18], the distillation loss is formulated by a distance of features that come from the last layer of the network. For example, in the image classification domain, it refers to the “logits” before going in the softmax layer and cross-entropy loss. Concretely, the vanilla distillation loss is defined as:

$$\mathcal{L}_{\text{KL}} = \text{KLD}\left(\sigma\left(\frac{T(x)}{\tau}\right), \sigma\left(\frac{S(x)}{\tau}\right)\right), \quad (1)$$

where $\text{KLD}(\cdot)$ measures the Kullback-Leibler divergence, $\sigma(\cdot)$ is the softmax function, $T(x)$ and $S(x)$ are the output logits given specific input x , and τ is the temperature factor. Without loss of generality, we assume that C' aligns with C

and reshape both F^T and F^S into 2D matrices:

$$\begin{aligned} f^s &= \Gamma(F^S) \in \mathbb{R}^{N \times C}, \\ f^t &= \Gamma(F^T) \in \mathbb{R}^{N \times C}. \end{aligned} \quad (2)$$

Here $\Gamma(\cdot)$ is a function that flattens the 3D feature tensor into the 2D matrix where each row of the matrix is associated with a pixel in the feature tensor by spatial order and $N = H \times W$. We can describe f^s and f^t as two sets of the pixels with cardinality N :

$$\begin{aligned} f^{s\top} &= [f_1^s, f_2^s, f_3^s, \dots, f_N^s], \\ f^{t\top} &= [f_1^t, f_2^t, f_3^t, \dots, f_N^t]. \end{aligned} \quad (3)$$

Previous work [33] simply minimize the discrepancy between two sets f^s and f^t in a one-to-one spatial matching manner, we denote this approach feature matching (FM):

$$\mathcal{L}_{\text{FM}} = \|F^S - F^T\|_2 = \sum_{i=1}^N \|f_i^s - f_i^t\|_2. \quad (4)$$

This formulation assumes that the semantic distributions of the teacher and the student match exactly. However, as mentioned earlier, for the feature maps of the teacher network, which usually encompasses more layers and larger feature channels, the spatial information of the same pixel location contains a richer semantic information compare to the student network. Directly regressing the features in a pixel-wise manner may lead to suboptimal distillation results.

To this end, we propose a one-to-all spatial matching knowledge distillation pipeline that allows the each feature

location of the teacher can teach the entire student features in a dynamic manner. To make the whole student mimic a spatial component of the teacher, we propose the **Target-aware Transformer (TaT)** to pixel-wisely reconfigure the semantic of student feature in the certain position. Given a spatial component (alignment target) of the teacher, we use **TaT** to guide the whole student to reconstruct the feature in its corresponding location. Conditioned on the alignment target, **TaT** should reflect the semantic similarity with the components of the student feature. We use a linear operator to avoid changing the distribution of student semantics. The formulation of transformation operator W^i can be defined as:

$$W^i = \sigma(\langle f_1^s, f_i^t \rangle, \langle f_1^s, f_i^t \rangle, \dots, \langle f_N^s, f_i^t \rangle) \quad (5)$$

$$= [w_1^i, w_2^i, \dots, w_N^i],$$

where f_i^t and f_i^s denote the corresponding i -th components of teacher and student, $\langle \cdot, \cdot \rangle$ represents the inner-product and $\|W^i\| = 1$. We use inner-product to measure the semantic distance and softmax function for normalization. Each entry of W^i is like the gate and controls the amount of semantic that will be propagated to the i -th reconfigured point. By aggregating the these related semantic across all the components, we have the result:

$$f_i^{s'} = w_1^i \times f_1^s + w_2^i \times f_2^s + \dots + w_N^i \times f_N^s. \quad (6)$$

The Eq. 5 and Eq. 6 can be combined and rewritten as the form of matrix multiplication: $f_i^{s'} = \sigma(f^s \cdot f_i^t) \cdot f^s$.

Note this is the simple non-parametric method that only depends on the original features. To facilitate the training, we introduce the parametric method with the extra linear transformation applied on the student feature and teacher feature. We observe that parametric version performs better than non-parametric one in ablation study. Guided by the target-aware transformer, the reconfigured student feature can be formulated as:

$$f_i^{s'} = \sigma(\gamma(f^s) \cdot \theta(f_i^t)^\top) \cdot \phi(f^s), \quad (7)$$

where $\theta(\cdot)$, $\gamma(\cdot)$ and $\phi(\cdot)$ are the linear functions consisting of 1×1 conv layer plus the BN layer [20]. We compare the parametric **TaT** to non-parametric one to analyse the effectiveness brought by these linear functions in the Section 4.5. In the case that the channel numbers of F^S do not match with that of F^T , $\gamma(\cdot)$ can help with alignment.

After reconfiguration, each component of $f_i^{s'}$ aggregates the meaningful semantic from the original feature, which enhances the expressivity. We do not require the student to reconstruct the teacher feature in a pixel-to-pixel manner. Indeed, our model allows the student to act as a whole to mimic the teacher. The resulting $f_i^{s'}$ is lately asked to minimize the L_2 loss with the teacher feature. The objective for **TaT** knowledge distillation can be given by:

$$\mathcal{L}_{\text{TaT}} = \|f_i^{s'} - f_i^t\|_2. \quad (8)$$

Finally, the total loss of our proposed method can be defined by:

$$\mathcal{L} = \alpha \mathcal{L}_{\text{Task}} + \beta \mathcal{L}_{\text{KL}} + \epsilon \mathcal{L}_{\text{TaT}}, \quad (9)$$

Here $\mathcal{L}_{\text{Task}}$ can be any loss on the generic machine learning tasks. α , β and ϵ are the weight factors to balance the loss. Empirically, we find that our model benefits from \mathcal{L}_{KL} . However, the model can achieve state-of-the-art without the help of \mathcal{L}_{KL} .

3.2. Hierarchical Distillation

The proposed **TaT** lift the limitation of previous one-to-one spatial matching fashion. However, the computation complexity of **TaT** map will become intractable when it comes to a large feature map. Assuming the spatial dimensions of the feature map are H and W , this means the computation complexity will reach $\mathcal{O}(H^2 \cdot W^2)$. Therefore, we propose a hierarchical distillation approach to address this large feature map limitation. It contains two steps: 1) patch-group distillation that splits the entire feature maps into smaller patches, so to distill local information from the teacher to the student; 2) we further summarize the local patches into one vector and distill this for global information.

3.2.1 Patch-group Distillation

As mentioned above, as the spatial dimension of input feature maps increases, distillation becomes more difficult. A straightforward solution [25] is to divide the feature map into patches and perform distillation within patches individually. However, the correlation between patches is completely ignored, resulting in sub-optimal solutions.

In contrast to Liu *et al.* [25], we propose the patch-group distillation (See Figure 2 (b)) that allows the student to learn the local feature from patches and retain the correlation among them to some extent. Given the original student feature F^S and teacher feature F^T , they are partitioned into $n \times m$ patches of size $h \times w$, where $h = H/n$, $w = W/m$. They are further arranged as g groups sequentially where each group contains $p = n \cdot m/g$ patches. Specifically, the patches in a group will be concatenated channel-wisely, forming a new tensor of size $h \times w \times c \cdot g$ that would be used for distillation lately. In this way, each pixel of the new tensor contains the features from p positions of the original feature, which explicitly includes the spatial pattern. Therefore, during the distillation, the student can learn not only the single pixel but the correlation among them. Intuitively, a larger group will introduce richer correlation but complex correlation will turn to difficult to learn. We study the effectiveness of different group sizes in the experiments.

Similar to the formulation presented in Section 3.1, the patch-group distillation can be given by simply replacing

the original input with the reorganized one, and the variant is denoted as \mathcal{L}_{TaT}^P . To relax the strict constraints of the spatial pattern in the patch-group, we set the $\theta(\cdot)$ as linear transformation in our experiments.

3.2.2 Anchor-point Distillation

The patch-group distillation can learn the fined-grained feature on the patch level and retain the spatial correlation among the patches to some extent. However, it is not capable of perceiving the long-range dependency. As will see in the ablation study, the attempt to preserve the global correlation through concatenating all the patches would fail. Furthermore, a large feature map will hinder the inductive bias since it may encourage the student to integrate the less relevant semantic from remote positions by mistake, which may deteriorate the subsequent distillation performance. For complex scenes, long-range dependency is important to capture the relation (*e.g.* layout) of different objects.

We address the conundrum by the proposed anchor-point distillation. As shown in Figure 2 (c), we summarize the local area to compact representation, referred to *anchor*, within a local area that is representative to describe the semantic of the given area, forming the new feature map of smaller size. Since the new feature map consists of the summary of the original feature, it can approximately substitute the original one to obtain the global dependency. We simply use average pooling to extract the anchor points. Then all the anchors are scattered back to the associated position to form a new feature map. The anchor-point feature is used for distillation as described in Section 3.1 and the objective is denoted as \mathcal{L}_{TaT}^A . The patch-group distillation enables the student to mimic the local feature while the anchor-point distillation allows the student to learn the global representation over the coarse anchor-point feature, which are complementary to each other. Therefore, the combination of these two objectives can bring the best of two worlds. Our objective designed for semantic segmentation can be written by:

$$\mathcal{L}_{Seg} = \alpha\mathcal{L}_{CE} + \delta\mathcal{L}_{TaT}^P + \zeta\mathcal{L}_{TaT}^A \quad (10)$$

4. Experiment

In this section, we empirically evaluate the effectiveness of the proposed method through extensive experiments. On image classification, we leverage the commonly used benchmark in knowledge distillation such as Cifar-100 [22] and ImageNet [12], and show our model can improve the student performance by a significant margin compared to many state-of-the-art baselines. In addition, we extend our method to another popular computer vision task, semantic segmentation to further demonstrate the generalization abil-

ity of our method. We nonetheless provide a detailed ablation study in the end of this section.

4.1. Datasets

Cifar-100 [22]. This benchmark contains 100 categories including 600 samples each. For each category, there are 500 images for training while 100 images for testing. We report top-1 accuracy as evaluation metric.

ImageNet [12]. This is a challenging benchmark for image classification including more than one million training samples with 1,000 categories. Similarly, we report the top-1 accuracy to measure the model performances.

Pascal VOC [13]. This benchmark contains 20 foreground classes with a background class. It provides 1,464 training, 1,499 validation, and 1,456 testing samples. Apart from the fine annotated samples, we also use additional coarse annotated images from [14] for training, resulting in 10,582 training samples. We report the mean Intersection over Union (mIoU) on the validation set to measure the proposed method.

COCOStuff10k [2]. The challenging dataset is developed on MSCOCO [24] by adding dense pixel-wise stuff label, resulting in 172 classes: 80 for thing, 91 for stuff, and 1 for unlabeled. It contains 9k training samples and 1k validation samples. We report the mIoU to evaluate our method.

4.2. Implementation Details

Image classification. For the experiments on Cifar-100, we use SGD optimizer [37] and the total running epoch is set to 240. The initial learning rate is 0.05 with a decay rate 0.1 at epoch 150, 180, and 210. In terms of data augmentation, the input images will be randomly cropped and flipped horizontally. We use Bayesian optimization [36] for hyper-parameters (*i.e.* α and ϵ in Eq. 9) searching. We report the exact values in the supplementary materials. For the ImageNet experiments, we use the AdamW optimizer [27] and train all of the models for 100 epochs with a batch size of 2048. The initial learning rate is set to 1.6e-4 and decay by 0.1 at epoch 30, 60, and 90. We apply standard data augmentations including random crop and horizontal flip. We use a simple grid search on the hyper-parameters, and set $\alpha=0.5$, $\beta=0.5$ and $\epsilon=0.1$ in Eq. 9.

Semantic segmentation. We choose the DeepLabV3+ [6] as the base architecture, where it contains a backbone to extract feature and a head to generate the segmentation results. For the teacher, we follow [6] to use the ResNet101 as the backbone model. For the student, we select two networks, ResNet18 which shares a similar architecture design as the backbone, and MobileNetV2 [34] which is drastically different.

We use random flip and Gaussian blur for data augmentation. The samples are randomly cropped and rescaled to

Table 1. **Top-1 accuracy(%) on Cifar-100**. The loss term \mathcal{L}_{KL} in Eq. 9 is removed in this experiment.

Method	Network Architecture						
	WRN-40-2 WRN-16-2	WRN-40-2 WRN-40-1	ResNet56 ResNet20	ResNet110 ResNet20	ResNet110 ResNet32	ResNet32×4 ResNet8×4	VGG13 VGG8
Teacher	75.61	75.61	72.34	74.31	74.31	79.42	74.64
Vanilla	73.26	71.98	69.06	69.06	71.14	72.50	70.36
KD [18]	74.92	73.54	70.66	70.67	73.08	73.33	72.98
FitNet [33]	73.58	72.24	69.21	68.99	71.06	73.50	71.02
AT [45]	74.08	72.77	70.55	70.22	72.31	73.44	71.43
SP [40]	73.83	72.43	69.67	70.04	72.69	72.94	72.68
CC [32]	73.56	72.21	69.63	69.48	71.48	72.97	70.71
RKD [30]	73.35	72.22	69.61	69.25	71.82	71.90	71.48
PKT [31]	74.54	73.45	70.34	70.25	72.61	73.64	72.88
FSP [43]	72.91	NA	69.95	70.11	71.89	72.62	70.20
NST [19]	73.68	72.24	69.60	69.53	71.96	73.30	71.53
CRD [39]	75.48	74.14	71.16	71.46	73.48	75.51	73.94
ICKD [25]	75.64	74.33	71.76	71.68	73.89	75.25	73.42
Ours w/o \mathcal{L}_{KL}	76.06	74.97	71.59	71.70	74.05	75.89	74.39

513×513 during the training and are resized to the same resolution during the testing. The student backbone ResNet18 is trained for 100 epochs with an initial learning rate $7e-3$ on the Pascal VOC and $1e-2$ on COCOStuff10k respectively. For MobileNetV2, the learning rate is set to $7e-3$ for all datasets. We incorporate the cosine learning rate scheduler for all experiments.

4.3. Image Classification

Results on Cifar-100. To show the generalization ability of our method, we applied our distillation approach to various network architectures, including ResNet [15], VGG [35] and WideResNet [44]. And in these experiments, we set $\theta(\cdot)$ as an identical function instead of linear transformation, e.g., Conv+BN. As shown in Table 1, our method surpasses all baselines on six out of seven teacher-student settings, often by a significant margin. This evidences the effectiveness and generalization ability of our approach. Compared to the closest baseline, FitNet [33], which directly computes the distillation loss in one-to-one fashion, our approach improves on average 2.72%. The results when distilling to ResNet20 is interesting. In this case, using a less powerful teacher, ResNet56, results a better student performance on average comparing to using ResNet110. In particular, directly distilling the feature in one-to-one fashion deteriorates the student’s performance compared to vanilla training. Our distillation approach addressed such mismatch and achieves 71.70% which is 2.64% better than the baseline.

Results on ImageNet. Since Cifar-100 only contains 50,000 training images, we further evaluate our approach

on a more challenging dataset. Here, we choose ResNet34 and ResNet18 as teacher and student model respectively. We show the Top-1 accuracy of the student and teacher model in Table 2. Our method outperforms the state-of-the-art methods by a significant margin. Notice that, even without the help of \mathcal{L}_{KL} , our model can reach 72.07% on a tiny ResNet18, comparing to some methods which rely on the \mathcal{L}_{KL} by more than 1%. When enabling \mathcal{L}_{KL} , the proposed method can further improve the Top-1 accuracy of the student to 72.41%. Compared to SCKD [3] which uses an attention mechanism to re-allocate the most semantic-related teacher layers to the student, our method has a significant improvement. That means even matching two layers of teacher and student with similar semantics, the student may not be able to catch up with the teacher in the pixel-to-pixel manner due to semantic mismatch. In contrast, our method leverages a target-aware transformer, to address the semantic mismatch in a more efficient manner.

4.4. Semantic Segmentation

As the feature map size is fairly small when performing distillation on image classification, we plan to further investigate the generalization ability of our method on semantic segmentation, where the feature size is drastically larger. As in Section 3.2, we adapt our TaT method with the patch-group and anchor-point scheme. We select two popular benchmarks, Pascal VOC and COCOStuff10k, and present the results in Table 3 and Table 5 respectively. Our method clearly surpasses all baselines by a clear margin. For instance, on Pascal VOC, the proposed model can improve the MobileNetV2 by more than 5%, which shows great po-

Table 2. Top-1 Accuracy(%) on ImageNet validation set. The ResNet34 is employed as the teacher backbone and the ResNet18 is selected as the student backbone. Our method can boost the performance of the tiny ResNet18 beyond 72% and outperforms other methods without \mathcal{L}_{KL} .

Method	Vanilla	AT [45]	CRD [39]	SAD [21]	ICKD [25]	KR [7]	Ours	KD [18]	SCKD [3]	CC [32]	RKD [30]	Ours	Teacher
w/ \mathcal{L}_{KL}								✓	✓	✓	✓	✓	-
Top-1	70.04	70.59	71.17	71.38	71.59	71.61	72.07	70.68	70.87	70.74	71.34	72.41	73.31

Table 3. Comparing the semantic segmentation results (in mIoU%) of different methods on Pascal VOC. We can observe that our method surpasses all previous baselines by a significant margin, specifically, on the popular compact architecture MobilenetV2, our method improves the student by 5.39% comparing to the stand-alone training, and by 1.06% comparing to the state-of-the-art method ICKD. † indicates reproducing by training 100 epochs, using the official released code.

	ResNet18	MobilenetV2
Student	72.07	68.46
KD [18]	73.74	71.73
AT [45]	73.01	71.39
FitNet [33]	73.31	69.23
Overhaul† [17]	73.98	72.30
ICKD [25]	75.01	72.79
Ours	75.76	73.85

Table 4. Non-parametric vs. parametric implementation of target-aware transformer on ImageNet, where check mark indicates applying linear function.

$\theta(\cdot)$	$\gamma(\cdot)$	Top-1 Acc.
		72.22
	✓	72.41
✓	✓	72.35

Table 5. Comparing the semantic segmentation results (in mIoU%) of different methods on COCOStuff10k. As most baselines do not provide the code on the COCO dataset except KR, we only compare our method to KR in this case. We reproduce the baseline using the official code with the same training procedure. Our method surpasses the baseline by nearly 2%, and further demonstrate the effectiveness of our approach.

	Student	KR [7]	Our	Teacher
ResNet18	26.33	26.73	28.75	33.10
MobilenetV2	26.29	26.63	28.05	33.10

tential to unlock the hardware limitation. On the challenging benchmark COCOStuff10k, the model can improve the ResNet18 and MobileNetV2 by 2.42% and 1.76%.

Table 6. Impact of function $\theta(\cdot)$ on a variety of network architectures. We report the top-1 accuracy on Cifar-100. *id* indicates for identity mapping.

Teacher	Student	Conv+BN	<i>id</i>
ResNet56	ResNet20	71.45	71.59
ResNet110	ResNet20	71.68	71.70
ResNet110	ResNet32	73.75	74.05
ResNet32×4	ResNet8×4	75.30	75.89
VGG13	VGG8	73.48	74.39




Figure 3. The performance of our model under different ϵ on ImageNet. Here the loss \mathcal{L}_{KL} is removed and α is set to 0.1.

4.5. Ablation Study

Here, we provide detailed ablation study to validate each component of our approach.

Linear transformation functions. We first study the impact of function $\theta(\cdot)$. We are interested in that if the learning target (*i.e.* teacher feature) is fixed, can the student adapt itself through target-aware transformer better? We compare different settings of $\theta(\cdot)$ including identical mapping against Conv+BN. The result on Cifar100 is presented on Table 6. Surprisingly, the identical mapping for $\theta(\cdot)$ always performs better.

We further investigate the non-parametric implementation by setting both $\theta(\cdot)$ and $\gamma(\cdot)$ as identical mapping on ImageNet (Table 4). The result shows that the semi-parametric version performs best, where the fixed teacher and the linear transformation applying to student feature can facilitate the student to reconfigure itself.

Validating ϵ . To investigate the efficacy brought by the proposed Eq. 8, we then further explore the different settings of the coefficient ϵ used in Eq. 9 (See Figure 3). When increased from 0.05 to 0.25, the objective \mathcal{L}_{TaT} can bring positive and stable effect.

Table 7. Contribution of patch-group and anchor-point distillation. We observe that patch-group distillation presents more efficacy.

Patch-group	Anchor-point	mIoU
✓	✓	75.76
✓		75.63
	✓	75.37

Table 8. Performance (%) of anchor-point distillation on Pascal VOC under different pooling kernel sizes.

Pooling kernel	2×2	4×4	8×8	16×16
Speed up (\times)	16	256	4096	65536
mIoU	75.37	75.27	74.79	74.56

We also conduct the thorough experiments to understand the contribution brought by the proposed patch-group distillation \mathcal{L}_{TaT}^P and anchor-point distillation \mathcal{L}_{TaT}^A . As discussed previously, \mathcal{L}_{TaT}^A is proposed to learn the global representation to capture long-range dependency while \mathcal{L}_{TaT}^P is designed to concentrate on local feature. By covering each one of them, the individual effectiveness of the two components can be examined. As shown in Table. 7, both objectives can improve the vanilla student significantly while \mathcal{L}_{TaT}^P presents more efficacy. The combination of both components achieves the best performance, demonstrating that the two proposed objectives are complementary.

Validating the anchor-point distillation. Then, we give more insight concerning the proposed objectives’ functionality through sensitivity analysis. Specifically, we investigate the hyper-parameters that would influence the behavior of the training process. In terms of the anchor-point distillation, this work utilizes average pooling to extract the anchor in a local area from the original feature, forming the associated anchor-point feature. It is a trade-off between reducing computation overhead and summarizing fine-grained spatial information since a bigger kernel would reduce feature size along with more informative representation, *e.g.*, when feature map size is reduced to 1×1 , it degrades to ignoring the spatial information and posing one-to-one fashion distillation. Thus we study the pooling kernel size that directly yields different feature resolutions. The result exhibited in Table 8 shows that the amount of distillation calculation is greatly reduced with the increasing pooling size. On the other hand, excessive pooling range would omit useful and informative representation and damage the performance.

Validating the patch-group distillation. Next we analyze the two key factors of patch-group distillation, *i.e.* patch size $h \times w$ and groups g . In Table 9, we found that generally, smaller patch size is advantageous to patch-group distillation and overlage patch size, however, may be un-

Table 9. Performance (%) of patch-group distillation on Pascal VOC under different settings of patch size ($h \times w$). Groups is equal to patches $g = n \times m$.

Patch size	32×32	16×16	8×8	4×4
mIoU	75.33	75.45	75.50	75.47

Table 10. Performance (%) of patch-group distillation on Pascal VOC under different settings of groups where patch size is 8×8 and patch numbers is 256.

Groups	1	32	64	128	256
mIoU	75.26	75.57	75.63	75.62	75.50

favourable since it approaches the original feature. Regarding the groups, it merges the patches as a group for joint distillation. In the experiment shown in Table 10, the patch size is set to 8×8 , which divides the original feature map into $128/8 * 128/8 = 256$ patches. There are two extreme situations. When only one group is used, it indicates that all of the patches will be distilled jointly. On the contrary, using 256 groups means each patch is distilled individually. In this example, we found that 4 patches as a group can reach the best performance.

5. Conclusion

This work develops a framework for knowledge distillation through a target-aware transformation that enables the student to aggregate the useful semantic over itself to enhance the expressivity of each pixel, which allows the student to act as a whole to mimic the teacher rather than minimize each partial divergence in parallel. Our method is successfully extended to semantic segmentation by the proposed hierarchical distillation consisting of patch-group and anchor-point distillation, designed to focus on local feature and long-range dependency. We conduct thorough experiments to validate the effectiveness of the method and advance the state-of-the-art.

6. Discussion

Potential negative societal impact. Our method has no ethical risk on dataset usage and privacy violation as all the benchmarks are public and transparent.

Limitations. There are some issues of interest that we would like to explore in the future: (1) Currently, we only select the last layer of the backbone network for distillation. It would be interesting to see the efficacy when multiple layers are get involved with distillation which has been explored by some works [7, 45]. (2) Also, we didn’t investigate the effectiveness on other applications like object detection, which may need to design the new objective to fit the nature of specific application.

References

- 864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
- inverse detectors. *2011 International Conference on Computer Vision*, pages 991–998, 2011. 5
- [15] Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 6
- [16] Tong He, Chunhua Shen, Zhi Tian, Dong Gong, Changming Sun, and Youliang Yan. Knowledge adaptation for efficient semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 578–587, 2019. 2
- [17] Byeongho Heo, Jeesoo Kim, Sangdoon Yun, H. Park, N. Kwak, and J. Choi. A comprehensive overhaul of feature distillation. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1921–1930, 2019. 7
- [18] Geoffrey E. Hinton, Oriol Vinyals, and J. Dean. Distilling the knowledge in a neural network. *ArXiv*, abs/1503.02531, 2015. 1, 2, 3, 6, 7
- [19] Zehao Huang and Naiyan Wang. Like what you like: Knowledge distill via neuron selectivity transfer. *arXiv preprint arXiv:1707.01219*, 2017. 2, 6
- [20] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. PMLR, 2015. 4
- [21] Mingi Ji, Byeongho Heo, and S. Park. Show, attend and distill: Knowledge distillation via attention-based feature matching. *ArXiv*, abs/2102.02973, 2021. 1, 2, 7
- [22] Alex Krizhevsky. Learning multiple layers of features from tiny images. *tech report*, 2009. 5
- [23] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(12):2935–2947, 2018. 1
- [24] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 5
- [25] Li Liu, Qingle Huang, Sihao Lin, Hongwei Xie, Bing Wang, Xiaojun Chang, and Xiaodan Liang. Exploring inter-channel correlation for diversity-preserved knowledge distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8271–8280, October 2021. 2, 4, 6, 7
- [26] Yifan Liu, Ke Chen, Chris Liu, Zengchang Qin, Zhenbo Luo, and Jingdong Wang. Structured knowledge distillation for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2604–2613, 2019. 2
- [27] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 5
- [28] Asit K. Mishra and Debbie Marr. Apprentice: Using knowledge distillation techniques to improve low-precision network accuracy. In *ICLR*, 2018. 1
- [29] Boxiao Pan, Haoye Cai, De-An Huang, Kuan-Hui Lee, Adrien Gaidon, Ehsan Adeli, and Juan Carlos Niebles. Spatio-temporal graph for video captioning with knowledge
- 918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

- 972 distillation. In *Proceedings of the IEEE/CVF Conference*
973 *on Computer Vision and Pattern Recognition*, pages 10870–
974 10879, 2020. 2
- 975 [30] Wonpyo Park, Dongju Kim, Yan Lu, and Minsu Cho. Re-
976 lational knowledge distillation. *2019 IEEE/CVF Conference*
977 *on Computer Vision and Pattern Recognition (CVPR)*, pages
978 3962–3971, 2019. 2, 6, 7
- 979 [31] N. Passalis and A. Tefas. Learning deep representations with
980 probabilistic knowledge transfer. In *ECCV*, 2018. 2, 6
- 981 [32] Baoyun Peng, Xiao Jin, Jiaheng Liu, Shunfeng Zhou, Y. Wu,
982 Y. Liu, Dong sheng Li, and Z. Zhang. Correlation congru-
983 ence for knowledge distillation. *2019 IEEE/CVF Interna-*
984 *tional Conference on Computer Vision (ICCV)*, pages 5006–
985 5015, 2019. 2, 6, 7
- 986 [33] A. Romero, Nicolas Ballas, S. Kahou, Antoine Chassang, C.
987 Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets.
988 *CoRR*, abs/1412.6550, 2015. 1, 2, 3, 6, 7
- 989 [34] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zh-
990 moginov, and Liang-Chieh Chen. Mobilenetv2: Inverted
991 residuals and linear bottlenecks. In *Proceedings of the*
992 *IEEE conference on computer vision and pattern recogni-*
993 *tion*, pages 4510–4520, 2018. 5
- 994 [35] Karen Simonyan and Andrew Zisserman. Very deep convo-
995 lutional networks for large-scale image recognition. *arXiv*
996 *preprint arXiv:1409.1556*, 2014. 6
- 997 [36] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Prac-
998 tical bayesian optimization of machine learning algorithms.
999 *Advances in neural information processing systems*, 25,
1000 2012. 5
- 1001 [37] Ilya Sutskever, James Martens, George Dahl, and Geoffrey
1002 Hinton. On the importance of initialization and momentum
1003 in deep learning. In *International conference on machine*
1004 *learning*, pages 1139–1147. PMLR, 2013. 5
- 1005 [38] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet,
1006 Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent
1007 Vanhoucke, and Andrew Rabinovich. Going deeper with
1008 convolutions. In *Proceedings of the IEEE conference on*
1009 *computer vision and pattern recognition*, pages 1–9, 2015.
1010 2
- 1011 [39] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Con-
1012 trastive representation distillation. *ICLR*, 2020. 1, 6, 7
- 1013 [40] F. Tung and G. Mori. Similarity-preserving knowledge dis-
1014 tillation. *2019 IEEE/CVF International Conference on Com-*
1015 *puter Vision (ICCV)*, pages 1365–1374, 2019. 2, 6
- 1016 [41] Xionghui Wang, Jian-Fang Hu, Jian-Huang Lai, Jianguo
1017 Zhang, and Wei-Shi Zheng. Progressive teacher-student
1018 learning for early action prediction. In *Proceedings of*
1019 *the IEEE/CVF Conference on Computer Vision and Pattern*
1020 *Recognition*, pages 3556–3565, 2019. 2
- 1021 [42] Yukang Wang, W. Zhou, T. Jiang, X. Bai, and Yongchao Xu.
1022 Intra-class feature variation distillation for semantic segmen-
1023 tation. In *ECCV*, 2020. 2
- 1024 [43] Junho Yim, Donggyu Joo, Jihoon Bae, and Junmo Kim. A
1025 gift from knowledge distillation: Fast optimization, network
1026 minimization and transfer learning. *2017 IEEE Conference*
1027 *on Computer Vision and Pattern Recognition (CVPR)*, pages
1028 7130–7138, 2017. 1, 2, 6
- [44] Sergey Zagoruyko and Nikos Komodakis. Wide residual net-
works. *ArXiv*, abs/1605.07146, 2016. 2, 6
- [45] Sergey Zagoruyko and Nikos Komodakis. Paying more
attention to attention: Improving the performance of con-
volutional neural networks via attention transfer. *ArXiv*,
abs/1612.03928, 2017. 1, 2, 6, 7, 8
- [46] Linfeng Zhang and Kaisheng Ma. Improve object detec-
tion with feature-based knowledge distillation: Towards ac-
curate and efficient detectors. In *International Conference*
on Learning Representations, 2020. 2