

# Rethinking Centered Kernel Alignment in Knowledge Distillation

Zikai Zhou<sup>1</sup>, Yunhang Shen<sup>3</sup>, Shitong Shao<sup>1</sup>, Linrui Gong<sup>4</sup>, Shaohui Lin<sup>1,2\*</sup>

<sup>1</sup>School of Computer Science and Technology, East China Normal University, Shanghai, China

<sup>2</sup>Key Laboratory of Advanced Theory and Application in Statistics and Data Science, MOE, China

<sup>3</sup>YouTu Lab, Tencent, Shanghai, China

<sup>4</sup>Hunan University, China

{choukai003,shenyunhang01,1090784053sst,wslgqq277a,shaohuilin007}@gmail.com

## Abstract

Knowledge distillation has emerged as a highly effective method for bridging the representation discrepancy between large-scale models and lightweight models. Prevalent approaches involve leveraging appropriate metrics to minimize the divergence or distance between the knowledge extracted from the teacher model and the knowledge learned by the student model. Centered Kernel Alignment (CKA) is widely used to measure representation similarity and has been applied in several knowledge distillation methods. However, these methods are complex and fail to uncover the essence of CKA, thus not answering the question of how to use CKA to achieve simple and effective distillation properly. This paper first provides a theoretical perspective to illustrate the effectiveness of CKA, which decouples CKA to the upper bound of Maximum Mean Discrepancy (MMD) and a constant term. Drawing from this, we propose a novel Relation-Centered Kernel Alignment (RCKA) framework, which practically establishes a connection between CKA and MMD. Furthermore, we dynamically customize the application of CKA based on the characteristics of each task, with less computational source yet comparable performance than the previous methods. The extensive experiments on the CIFAR-100, ImageNet-1k, and MS-COCO demonstrate that our method achieves state-of-the-art performance on almost all teacher-student pairs for image classification and object detection, validating the effectiveness of our approaches. Our code is available in <https://github.com/Klayand/PCKA>.

## 1 Introduction

Tremendous efforts have been made in compressing large-scale models into lightweight models. Representative methods include network pruning [Frankle and Carbin, 2019], model quantization [Wu *et al.*, 2016], neural architecture search [Wan *et al.*, 2020] and knowledge distilla-

tion (KD) [Hinton *et al.*, 2015]. Among them, KD has recently emerged as one of the most flourishing topics due to its effectiveness [Liu *et al.*, 2021b; Huang *et al.*, 2022; Gong *et al.*, 2023; Shao *et al.*, 2023a] and wide applications [Chong *et al.*, 2022; Chen *et al.*, 2023; Shao *et al.*, 2023b]. Particularly, the core idea of KD is to transfer the acquired representations from a large-scale and high-performing model to a lightweight model by distilling the learning representations in a compact form, achieving precision and reliable knowledge transfer. There are two mainstream approaches for distilling knowledge. The first approach is the logit-based distillation, which aims to minimize the probabilistic prediction (response) scores between the teacher and student by leveraging appropriate metrics [Zhao *et al.*, 2022; Hinton *et al.*, 2015]. The other is feature-based distillation, which investigates the knowledge within intermediate representations to further boost the distillation performance [Yang *et al.*, 2022b; Liu *et al.*, 2021a; Chen *et al.*, 2021; Ahn *et al.*, 2019]. Among them, the design of metrics is essential in knowledge transfer and has been attractive from academic research. Specifically, Kornblith *et al* [2019] proposes the Centered Kernel Alignment (CKA) for the quantitative understanding of representations between neural networks. CKA not only focuses on model predictions but also emphasizes high-order feature representations within the models, providing a comprehensive and enriched knowledge transfer.

Recent studies [Qiu *et al.*, 2022; Saha *et al.*, 2022] introduce CKA to quantitatively narrow the gap of learned representations between the teacher model and the student model, which have achieved significant success. However, their designs are excessively complex and need a large amount of computational resources, making it challenging to achieve fine-grained knowledge transfer and leading to low scalability. Moreover, these methods fail to uncover the essence of CKA, lacking an in-depth analysis of CKA in knowledge distillation. The reason why CKA is effective has not been explored. Therefore, we focus on the theoretical analysis of CKA and rethink a more reasonable architecture design that ensures simplicity and effectiveness while generalizing well across various tasks.

In this paper, we provide a novel perspective to illustrate the effectiveness of CKA, where CKA is regarded as the upper bound of Maximum Mean Discrepancy (MMD) with a constant term, specifically. Drawing from this, we propose

\*Corresponding author

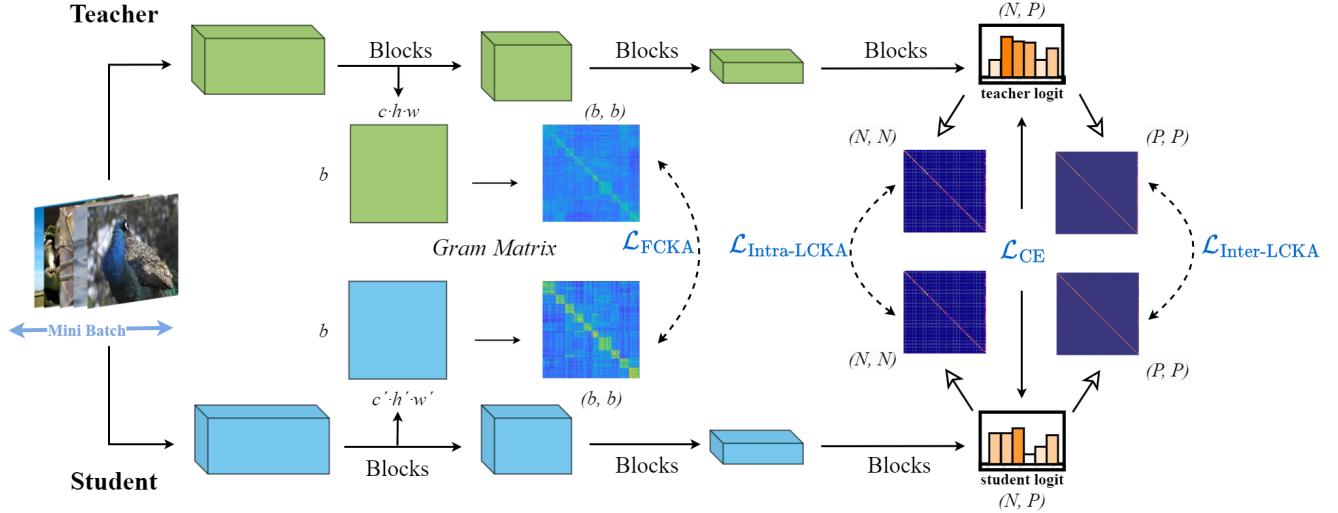


Figure 1: The overall framework of the proposed Relation-based Centered Kernel Alignment (RCKA). We first transform the feature map from the shape of  $(B, C, HW)$  into  $(B, CHW)$  and then compute the CKA similarity of feature maps between the teacher and the student. Besides, we compute the inter-class and intra-class CKA similarity of logits between teacher and student. Here,  $N$  refers to the number of samples, and  $P$  refers to the corresponding probability of class to which this sample belongs.

a Relation-Centered Kernel Alignment (RCKA) framework, which practically establishes a connection between CKA and MMD. Besides, we dynamically customize the application of CKA on instance-level tasks, and introduce Patch-based Centered Kernel Alignment (PCKA), with less computational source yet competitive performance when compared to previous methods. Our method is directly applied not only to logit-based distillation but also to feature-based distillation, which exhibits superior scalability and expansion. We utilize CKA to compute high-order representation information both between and within categories, which better motivates the alleviation of the performance gap between the teacher and student.

To validate the effectiveness of our approaches, we conduct extensive experiments on image classification (CIFAR-100 [Krizhevsky and Hinton, 2009] and ImageNet-1k [Russakovsky *et al.*, 2015]), and object detection (MS-COCO [Lin *et al.*, 2014]) tasks. As a result, our methods achieve state-of-the-art (SOTA) performance in almost all quantitative comparison experiments with fair comparison. Moreover, following our processing architecture, the performance of the previous distillation methods is further boosted in the object detection task.

Our contribution can be summarized as follows:

- We rethink CKA in knowledge distillation from a novel perspective, providing a theoretical reason for why CKA is effective in knowledge distillation.
- We propose a Relation Centered Kernel Alignment (RCKA) framework to construct the relationship between CKA and MMD, with less computational source yet comparable performance than previous methods, which verifies our theoretical analysis correctly.
- We further dynamically customize the application of CKA for instance-level tasks and propose a Patch-based

Centered Kernel Alignment (PCKA) architecture for knowledge distillation in object detection, which further boosts the performance of previous distillation methods.

- We conduct plenty of ablation studies to verify the effectiveness of our method, which achieves comparable performance on a range of vision tasks. Besides, we visualize the characteristic information of CKA and discover new patterns in it.

## 2 Related Work

Vanilla Knowledge Distillation [Hinton *et al.*, 2015] proposes aligning the output distributions of classifiers between the teacher and student by minimizing the KL-divergence, during training the emphasis on negative logits can be fine-tuned through a temperature coefficient, which serves as a form of normalization during the training process of a smaller student network. Tremendous efforts [Tung and Mori, 2019; Huang *et al.*, 2022; Qiu *et al.*, 2022; Zagoruyko and Komodakis, 2016; Park *et al.*, 2019] have been made on how to design a good metric to align the distribution between the teacher and student.

Designing an suitable alignment method for KD can start from two typical types: Drawing on representations, numerous methods have made significant strides by aligning the intermediate features [Zagoruyko and Komodakis, 2016], the samples' correlation matrices [Tung and Mori, 2019], and the output logits between the teacher and student [Huang *et al.*, 2022]. From a mathematical standpoint, some measure theories are introduced to illustrate the similarity between the teacher and student, such as mutual information [Ahn *et al.*, 2019]. Among these, Centered Kernel Alignment (CKA) is a valuable function for measuring similarity. It simultaneously considers various properties during similarity measures, such as invariance to orthogonal transformations. While the ef-

effectiveness of CKA in KD has been demonstrated in some works [Qiu *et al.*, 2022; Saha *et al.*, 2022], the essence of CKA has not been thoroughly explored, and the unavoidable additional computational costs also limit its application prospect.

In this paper, we will revisit CKA in KD and provide a novel theoretical perspective to prove its effectiveness and analyze how it functions across various distillation settings.

### 3 Methodology

In this section, we first revisit the paradigm of knowledge distillation and then introduce the formula of Centered Kernel Alignment (CKA). Specifically, we derive the formula of the relationship between CKA and Maximum Mean Discrepancy (MMD), where CKA can be decoupled as the upper bound of MMD with a constant term. In light of the above deduction, we outline the methodology of our paper. We apply the proposed methods in image classification and object detection, dynamically customizing CKA for each task.

#### 3.1 The Paradigm of Knowledge Distillation

The existing KD methods can be categorized into two groups. Particularly, the logits-based KD methods narrow the gap between the teacher and student models by aligning the soft targets between them, which is formulated as following loss term:

$$L_{\text{logits}} = D_{\text{logits}}(\mathbf{T}_s(\sigma(z_s; \tau)), \mathbf{T}_t(\sigma(z_t; \tau))), \quad (1)$$

where  $z_s$  and  $z_t$  are the logits from students and teachers, respectively. And  $\sigma(\cdot)$  is the softmax function that produces the category probabilities from the logits, and  $\tau$  is a non-negative temperature hyper-parameter to scale the smoothness of the predictive distribution. Specifically, we have  $\sigma_i(z; \tau) = \text{softmax}(\exp(z_i/\tau))$ .  $D_{\text{logits}}$  is a loss function to capture the discrepancy distributions, *e.g.* Kullback-Leibler divergence. And  $\mathbf{T}_s$  and  $\mathbf{T}_t$  denote the transformation functions in students and teachers, respectively, which usually refer to the identity mapping in Vanilla KD [Hinton *et al.*, 2015].

Similarly, the feature-based KD methods, which aim to mimic the feature representations between teachers and students, are also represented as a loss item:

$$L_{\text{feat}} = D_{\text{feat}}(\mathbf{T}_s(\mathbf{F}_s), \mathbf{T}_t(\mathbf{F}_t)), \quad (2)$$

where  $\mathbf{F}_s$  and  $\mathbf{F}_t$  denote feature maps from students and teachers, respectively. Transformation modules  $\mathbf{T}_s$  and  $\mathbf{T}_t$  align the dimensions of  $\mathbf{F}_s$  and  $\mathbf{F}_t$ .  $D_{\text{feat}}$  computes the distance between two feature maps, such as  $\ell_1$ - or  $\ell_2$  norm. Therefore, the KD methods can be represented by a generic paradigm. The final loss is the weighted sum of the cross-entropy loss  $L_{ce}$ , the logits distillation loss, and the feature distillation loss:

$$L = L_{ce} + \alpha L_{\text{logits}} + \beta L_{\text{feat}}, \quad (3)$$

where  $\alpha$  and  $\beta$  are hyper-parameters controlling the trade-off between these three losses.

### 3.2 Distilling with the Upper Bound

Centered Kernel Alignment (CKA) has been proposed as a robust way to measure representational similarity between neural networks. We first prove that CKA measures the cosine similarity of the Gram matrix between teachers and students.

**Theorem 1** (Proof in Appendix<sup>1</sup> C.1). *Let  $X$  and  $Y$  be  $N \times P$  matrices. The CKA similarity  $\|Y^\top X\|_F^2$  is equivalent to the cosine similarity of  $XX^\top$  and  $YY^\top$ , which denote the Gram matrix of  $X$  and  $Y$ , respectively. In other words,*

$$\begin{aligned} S_{\text{CKA}}(X, Y) &= \frac{\|Y^\top X\|_F^2}{\|X^\top X\|_F \|Y^\top Y\|_F} \\ &= \frac{\text{vec}(XX^\top)^\top \text{vec}(YY^\top)}{\|\text{vec}(XX^\top)\|_2 \|\text{vec}(YY^\top)\|_2}, \end{aligned}$$

where *vec* operator represents reshaping the matrix to a vector.

We then derive the formula of the relationship between CKA and MMD, where CKA can be regarded as the upper bound of MMD with a constant term.

**Theorem 2** (Proof in Appendix C.2). *Maximizing CKA similarity is equivalent to minimizing the upper bound of MMD distance:*

$$\begin{aligned} \frac{\|Y^\top X\|_F^2}{\|X^\top X\|_F \|Y^\top Y\|_F} &= -N \mathbb{E}_{i,j} [\langle x_i, x_j \rangle - \langle y_i, y_j \rangle]^2 + 2 \\ &\leq -N (\mathbb{E}_{i,j} [\langle x_i, x_j \rangle] - \mathbb{E}_{i,j} [\langle y_i, y_j \rangle])^2 + 2, \end{aligned}$$

where the inequality is given by Jesen's inequality.

According to Jesen's inequality, CKA can be decoupled as the upper bound of MMD with a constant term. The first term corresponds to minimizing the upper bound of MMD distance with the RKHS kernel. In contrast, the latter constant term acts as a weight regularizer, enhancing the influence of MMD, where it promotes the similarity between features of the same batch, not only instances in the same class but also in different classes. On one hand, optimizing the upper bound of MMD, which has additional stronger constraints, allows it to converge to the optimal solution more quickly and stably. On the other hand, the latter term serves as a weight scaling mechanism, effectively avoiding the challenges of optimization caused by excessively small MMD values, which result in small gradients.

According to the deduction, we successfully transformed our optimization objective from maximizing CKA to minimizing the upper bound of MMD, which makes our method more intuitive and concise. Building upon these findings, we propose our methods, which are more effective than previous methods.

### 3.3 Relation-based Centered Kernel Alignment

As illustrated in Fig. 1, we propose a Relation Centered Kernel Alignment (RCKA) framework in image classification. In this framework, we leverage CKA as a loss function to ensure that the centered similarity matrix is distilled

<sup>1</sup>Appendix can be referred to <https://arxiv.org/abs/2401.11824>

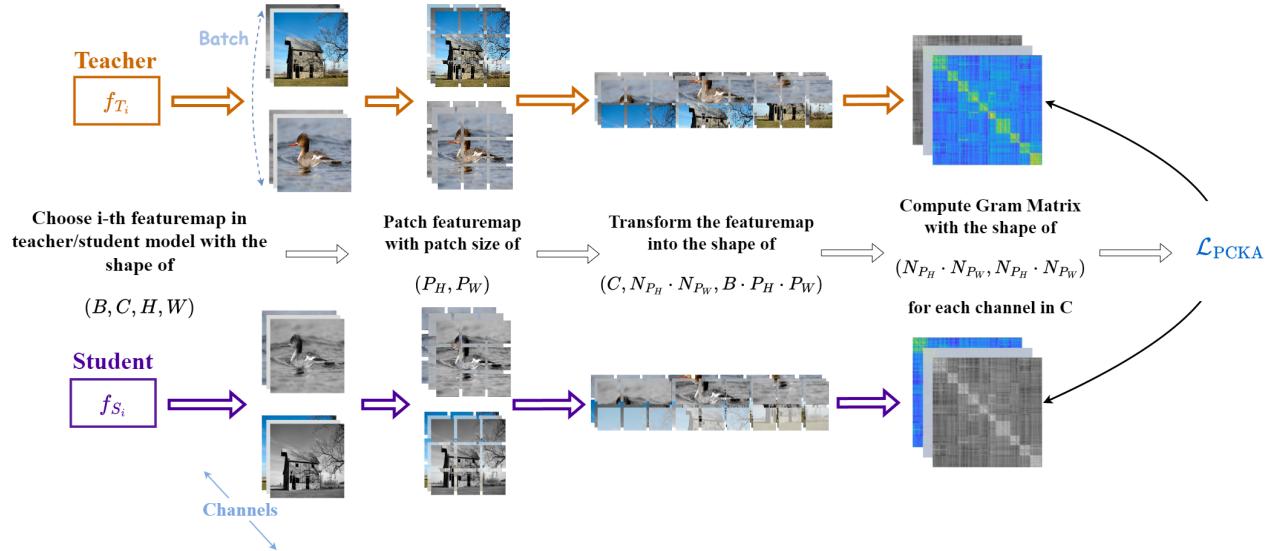


Figure 2: The overall framework of PCKA. We dynamically customize the framework of proposed method based on the characteristics of object detection. In this framework, we first patch the featuremap of the teacher and student with the patchsize  $(P_H, P_W)$ , then transform the featuremap to obtain the Gram matrix between each patch. Finally, we calculate the loss  $L_{PCKA}$ , and get average from dimension  $C$ . Here,  $B, C, H, W$  refer to the batchsize, channels, height and width of the featuremap, respectively.  $N_{P_H}, N_{P_W}$  denote the number of patches cutting along the height and width, respectively.

rather than forcing the student to mimic the teacher’s similarity matrix with a different scale. This is very important because a model’s discriminative capability is dependent on the distribution of its features rather than its scale, which is inconsequential for class separation [Nguyen *et al.*, 2020; Orhan and Pitkow, 2017].

Assume we have a large-scale teacher model  $t$  and a lightweight student model  $s$ . The activation map from layer  $l$  of the teacher is denoted as  $F_t^{(l)} \in \mathbb{R}^{b \times c \times h \times w}$ , whereas the activation map of layer  $l'$  of the student is denoted as  $F_s^{(l')} \in \mathbb{R}^{b \times c' \times h' \times w'}$ .  $c, h$ , and  $w$  denote the channel, height, and width of the teacher, whereas  $c', h'$  and  $w'$  denote that of the student. The mini-batch size is denoted by  $b$ . The logits of the teacher and student are denoted as  $z_t \in \mathbb{R}^{N \times P}$  and  $z_s \in \mathbb{R}^{N \times P'}$ , where  $N$  and  $P$  (or  $P'$ ) refer to the number of samples and the corresponding probability for which class this sample belongs to. Therefore, the formula of our method, similar to Eqn. 3, is represented as:

$$L_{RCKA} = L_{CE} + \alpha L_{FCKA} + \beta (L_{\text{Intra-LCKA}} + L_{\text{Inter-LCKA}}), \quad (4)$$

where  $\alpha$  and  $\beta$  are hyper-parameters controlling the trade-off between the features loss  $L_{FCKA}$  and logits loss  $(L_{\text{Intra-LCKA}} + L_{\text{Inter-LCKA}})$ .

The  $L_{FCKA}$ ,  $L_{\text{Intra-LCKA}}$  and  $L_{\text{Inter-LCKA}}$  are represented as:

$$L_{FCKA} = S_{CKA}(\mathbf{T}(\mathbf{F}_t), \mathbf{T}(\mathbf{F}_s)), \quad (5)$$

$$L_{\text{Intra-LCKA}} = S_{CKA}(z_t, z_s), \quad (6)$$

$$L_{\text{Inter-LCKA}} = S_{CKA}(z_t^T, z_s^T), \quad (7)$$

where  $\mathbf{T}$  in Eqn. 5 refers to the transformation module  $\mathbb{R}^{b \times c \times h \times w} \rightarrow \mathbb{R}^{b \times c \times h \times w}$ .

Compared with the previous methods, our method has superior scalability and expansion and can be directly applied to both feature and logits distillation. We calculate the Gram matrix to collect high-order inter-class and intra-class representations, encouraging the student to learn more useful knowledge. Also, we provide the relationship between CKA and MMD in Appendix C.2 to better demonstrate the theoretical support of our method.

Because the value of CKA ranges from  $[0, 1]$ , at the beginning of the training process,  $L_{CE}$  plays a more important role than all CKA losses to drive the optimization of the student, which helps the student avoid matching extremely complex representations.

### 3.4 Patch-based Centered Kernel Alignment

In this subsection, we further adapt the proposed RCKA to instance-level tasks such as object detection. However, directly applying RCKA to instance-level tasks may deteriorate performance, as the above tasks are usually trained with a small size of mini-batches (*e.g.* 2 or 4 per GPU), causing the failure of the Gram matrix to collect enough knowledge. Besides, increasing the mini-batch size requires a significant amount of computational resources, making it infeasible in practice. Thus, we dynamically customize our RCKA method for object detection.

Recent works [Shu *et al.*, 2021; Heo *et al.*, 2019] find that distilling the representations of intermediate layers is more effective than distilling the logits in object detection. Therefore, we adjust our method to only target intermediate layers. We follow our core idea in the classification task, which calculates the similarities between different instances by using CKA. Hence, we divide the image feature maps into several patches and compute the similarities between different

Architecture		Same							Different		
Distillation Type	Teacher	RN-110	RN-110	WRN-40-2	WRN-40-2	RN-32×4	VGG-13	WRN-40-2	RN-32×4	VGG-13	
	Student	74.31	74.31	75.61	75.61	79.42	74.64	75.61	79.42	74.64	
		RN-20	RN-32	WRN-40-1	WRN-40-2	RN-8×4	VGG-8	SN-V1	SN-V1	MN-V2	
Feature-based	FitNet [2014]	68.99	71.06	72.24	73.58	73.50	71.02	73.73	73.59	64.14	
	ATKD [2016]	70.22	70.55	72.77	74.08	73.44	71.43	73.32	72.73	59.40	
	SPKD [2019]	70.04	72.69	72.43	73.83	72.94	72.68	74.52	73.48	66.30	
	CCKD [2019]	69.48	71.48	72.21	73.56	72.97	70.71	71.38	71.14	64.86	
	RKD [2019]	69.25	71.82	72.22	73.35	71.90	71.48	72.21	72.28	64.52	
	VID [2019]	70.16	70.38	73.30	74.11	73.09	71.23	73.61	73.38	65.56	
	CRD [2020]	71.46	73.48	74.14	75.48	75.51	73.94	76.05	75.11	69.73	
	OFD [2019]	-	73.23	74.33	75.24	74.95	73.95	75.85	75.98	69.48	
	ReviewKD [2021]	-	71.89	75.09	76.12	75.63	74.84	77.14	76.93	70.37	
	ICKD-C [2021a]	71.91	74.11	74.63	75.57	75.48	73.88	75.19	74.34	67.55	
	DPK [2022]	<b>72.44</b>	<b>74.89</b>	75.27	76.42	-	74.96	74.43	76.00	68.63	
Logit-based	FCKA(ours)	71.49	73.64	74.70	75.53	74.93	74.35	75.98	75.67	68.97	
	KD [2015]	70.67	73.08	73.54	74.92	73.33	72.98	74.83	74.07	67.37	
	DKD [2022]	-	74.11	74.81	76.24	76.32	74.68	76.70	76.45	69.71	
	DIST [2022]	69.94	73.55	74.42	75.29	75.79	73.74	75.23	75.23	68.48	
	IKL-KD [2023]	-	74.26	74.98	76.45	<b>76.59</b>	74.88	77.19	76.64	<b>70.40</b>	
	NKD [2023]	71.26	73.79	75.23	76.37	76.35	74.86	76.59	76.90	70.22	
	LCKA(ours)	70.87	73.64	74.63	75.78	75.12	74.35	76.12	76.43	69.37	
Combined	SRRL [2021]	71.51	73.80	74.75	75.96	75.92	74.40	76.61	75.66	69.14	
	RCKA(ours)	72.26	74.31	<b>75.34</b>	<b>76.51</b>	76.11	<b>74.97</b>	<b>77.21</b>	<b>76.97</b>	70.12	

Table 1: Results on the CIFAR-100 test set. ‘‘Same’’ and ‘‘Different’’ in the first row refer to whether the model architecture is the same for teachers and students. Combined distillation type means that this method transfers the knowledge both on the features and logits. Our methods surpass almost all algorithms with the same distillation type. ‘‘RN’’, ‘‘WRN’’, ‘‘SN’’, and ‘‘MN’’ denote ResNet, Wide ResNet, ShuffleNet, and MobileNet, respectively.

Architecture		Accuracy		Feature-based					Logit-based		Combined		
Teacher	Student	Teacher	Student	OFD	CRD	ReviewKD	ICKD-C	MGD [2022b]	KD	RKD	DKD	DIST	SRRL Ours
ResNet-34	ResNet-18	Top-1	73.31	69.76	71.08	71.17	71.61	72.19	71.80	70.66	70.34	71.70	72.07
		Top-5	91.42	89.08	90.07	90.13	90.51	90.72	90.40	89.88	90.37	90.41	90.42
ResNet-50	MobileNet-V1	Top-1	76.16	70.13	71.25	71.37	72.56	-	72.59	70.68	-	72.05	<b>73.24</b>
		Top-5	92.86	89.49	90.34	90.41	91.00	-	90.74	90.30	-	91.05	<b>91.12</b>

Table 2: Results on the ImageNet validation set. We use ResNet-34 and ResNet-50 released by Torchvision [Marcel and Rodriguez, 2010] as our teacher’s pre-training weight.

patches.

Our redesigned method is illustrated in Fig. 2. In this framework, we first patch the feature maps of the teacher and student with a patch size of  $(P_H, P_W)$ , then transform the feature maps to get the Gram matrix between each patch. Finally, we calculate the loss  $L_{PCKA}$  and get the average from dimension  $C$ . Here,  $N_{P_H}$  and  $N_{P_W}$  denote the number of patches cutting along the height and width, respectively. Therefore, the Patch-based CKA loss is represented as:

$$L_{PCKA} = \gamma S_{CKA}(N_{P_H}^s \cdot N_{P_W}^s, N_{P_H}^t \cdot N_{P_W}^t), \quad (8)$$

where  $N_{P_H}^s \cdot N_{P_W}^s$  and  $N_{P_H}^t \cdot N_{P_W}^t$  are denoted as the number of the student patches and the teacher patches, respectively. Usually,  $N_{P_H}^s \cdot N_{P_W}^s = N_{P_H}^t \cdot N_{P_W}^t$ .  $\gamma$  refers to the loss weight factor.

## 4 Experiments

We conduct extensive experiments on *image classification* and *object detection* benchmarks. The image classification datasets include CIFAR-100 [Krizhevsky and Hinton, 2009] and ImageNet-1k [Russakovsky *et al.*, 2015] and the object detection dataset includes MS-COCO [Lin *et al.*, 2014].

Moreover, we present various ablations and analyses for the proposed methods. More details about these datasets are in Appendix A. We apply a batch size of 128 and an initial learning rate of 0.1 for the SGD optimizer on CIFAR-100. And we follow the settings in [Huang *et al.*, 2022] for the ResNet34-ResNet18 pair and the ResNet50-MobileNet pair on ImageNet-1k. The settings of other classification and detection tasks are in Appendix B.

### 4.1 Image Classification

**Classification on CIFAR-100.** We compare state-of-the-art (SOTA) feature-based and logit-based distillation algorithms on 9 student-teacher pairs. Among them, 6 pairs have the same structure for teachers and students, and the rest of them have different architectures. The results are presented in Tab. 1. Our proposed method outperforms all other algorithms on 4 student-teacher pairs and achieves comparable performance on the rest of them, meanwhile requiring extremely less computational resources and time consumption than the SOTA methods DPK [Qiu *et al.*, 2022] and ReviewKD [Chen *et al.*, 2021]. The comparisons of computational cost are in Appendix F.

T→S Type Method	CM RCNN-X101 [2017]→Faster RCNN-R50 [2015]						RetinaNet-X101→RetinaNet-R50 [2017]						T→S Type Method	FCOS-R101→FCOS-R50 [2019]					
	Two-stage detectors						One-stage detectors							Anchor-free detectors					
	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>	
Teacher	45.6	64.1	49.7	26.2	49.6	60.0	41.0	60.9	44.0	23.9	45.2	54.0	Teacher	40.8	60.0	44.0	24.2	44.3	52.4
Student	38.4	59.0	42.0	21.5	42.1	50.3	37.4	56.7	39.6	20.0	40.7	49.7	Student	38.5	57.7	41.0	21.9	42.8	48.6
KD [2015]	39.7	61.2	43.0	23.2	43.3	51.7	37.2	56.5	39.3	20.4	40.4	49.5	KD [2015]	39.9	58.4	42.8	23.6	44.0	51.1
COFD [2019]	38.9	60.1	42.6	21.8	42.7	50.7	37.8	58.3	41.1	21.6	41.2	48.3	FitNet [2014]	39.9	58.6	43.1	23.1	43.4	52.2
FKD [2021]	41.5	62.2	45.1	23.5	45.0	55.3	39.6	58.8	42.1	22.7	43.3	52.5	GID [2021]	42.0	60.4	45.5	25.6	45.8	54.2
DIST [2022]	40.4	61.7	43.8	23.9	44.6	52.6	39.8	59.5	42.5	22.0	43.7	53.0	FRS [2021]	40.9	60.3	43.6	25.7	45.2	51.2
DIST+mimic [2022]	41.8	62.4	45.6	23.4	46.1	55.0	40.1	59.4	43.0	23.2	44.0	53.6	FGD [2022a]	<b>42.1</b>	-	-	<b>27.0</b>	<b>46.0</b>	<b>54.6</b>
Ours	41.4	62.1	45.2	23.5	45.6	54.9	40.3	59.9	43.0	23.3	44.2	54.9	Ours	39.8	59.0	42.4	22.2	43.6	52.5
Ours + mimic	<b>42.4</b>	<b>63.3</b>	<b>46.1</b>	<b>24.3</b>	<b>46.7</b>	<b>56.1</b>	<b>40.7</b>	<b>60.4</b>	<b>43.4</b>	<b>23.9</b>	<b>44.7</b>	<b>55.1</b>	Ours + mimic	40.7	<b>60.5</b>	43.1	23.4	44.8	53.1

Table 3: Results on the COCO validation set (T→S refers to the distillation from T to S). Here, the content in brackets to the right of “Ours” refers to the methods applied in the distillation process. In addition, CM RCNN-X101 stands for Cascade Mask RCNN-X101.

T→S Method	GFL-R101→GFL-R50 [2020]					
	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>
Teacher	44.9	63.1	49.0	28.0	49.1	57.2
Student	40.2	58.4	43.3	23.3	44.0	52.2
FT [2014]	40.7	58.6	44.0	23.7	44.4	53.2
Inside GT Box	40.7	58.6	44.2	23.1	44.5	53.5
DeFeat	40.8	58.6	44.2	24.3	44.6	53.7
Main Region	41.1	58.7	44.4	24.1	44.6	53.6
FGFI [2019]	41.1	58.8	44.8	23.3	45.4	53.1
FGD [2022a]	41.3	58.8	44.8	24.5	45.6	53.0
GID [2021]	41.5	59.6	45.2	24.3	45.7	53.6
SKD [2022]	42.3	60.2	45.9	24.4	46.7	55.6
Our	<b>42.8</b>	<b>61.2</b>	<b>46.3</b>	<b>24.8</b>	<b>47.1</b>	<b>55.4</b>

Table 4: Results on the COCO validation set (T→S refers to the distillation from T to S). Here, the content in brackets to the right of “Ours” refers to the methods applied in the distillation process. In addition, “Inside GT Box” means we use the GT boxes with the same stride on the FPN layers as the feature imitation regions. “Main Region” means we imitate the features within the main distillation region.

**Classification on ImageNet-1k.** We also conduct experiments on the large-scale ImageNet to evaluate our methods. Our RCKA achieves comparable results with other algorithms, even outperforms them, as shown in Tab. 2. We find that with the increasing of categories and instances, it is more challenging for the student to mimic the high-order distribution of the teacher. Moreover, in Appendix Tab. 8, we explore the feature distillation for ViT-based models on ImageNet-1k. It is noted that our method outperforms other methods, which indicates that our method achieves better scalability and performance.

## 4.2 Object Detection

**Detection on MS-COCO.** Comparison experiments are run on three kinds of different detectors, *i.e.*, tow-stage detectors, one-stage detectors, and anchor-free detectors. As shown in Tab. 3, PCKA outperforms the precious methods almost on all three kinds of metrics, by aligning the high-order patch-wise presentations. We believe that aligning feature maps of the student and teacher in low-order could also improve the performance of PCKA, driven by mimicking low-order representations in the early stage and then learning high-order and complex representations gradually. Thus, we follow [Huang *et al.*, 2022] by adding auxiliary mimic loss, *i.e.*, translating the student feature maps from the teacher feature map by a convolution layer and supervising them

utilizing  $\mathcal{L}_{MSE}$ , to the detection distillation task. We conclude from Tab. 3 that PCKA-based mimic loss achieves the best performance on Cascade RCNN-X101-Cascade RCNN-R50 and RetinaNet-X101-RetinaNet-R50 pairs. We also conduct experiments on the other four architectures, as shown in Tab. 4 and Appendix Tab. 9. These results further validate the effectiveness of our proposed method.

## 4.3 Ablations and Visualizations

We conduct ablation studies in three aspects: **(a)** the effect of hyper-parameters. **(b)** effectiveness of the proposed modules. **(c)** unexplored phenomenon during training.

**Ablation studies on hyperparameters.** As shown in Tab. 16, Tab. 14 and Tab. 17 in Appendix, we conduct the ablation studies on the size of mini-batch, loss scaling factor  $\gamma$  on  $L_{PCKA}$  and the number of intermediate layers for distilling. We find the local optima values are the mini-batch size 12, loss scaling factor 10, and 3 layers of distillation.

**The upper bound of MMD.** In Theorem 2, we derive the relationship between CKA and MMD, where CKA is the upper bound of MMD with a constant term. To validate this, We conduct the experiment, which is shown in Tab. 6, we notice that CKA, which is the upper bound of MMD, has additional stronger constraints. Because of this, CKA converges to the optimal solution more quickly and stably, compared with MMD.

**The dimension to average.** In PCKA framework, we cut the activations of the teacher and student in the shape of  $(C, N_{P_H} \cdot N_{P_W}, B \cdot P_H \cdot P_W)$ . We also carry out the experiments of averaging on different dimensions, shown in Tab. 5. We find averaging on channel dimension is the optima.

**Patch distillation.** We explore the effectiveness of cutting activations into patches. As shown in Tab. 7, several standard distillation methods [Hinton *et al.*, 2015; Zagoruyko and Komodakis, 2016; Huang *et al.*, 2022] all perform well with patch cutting, validating the effectiveness of cutting patches. With the smaller representation distribution in patches, it is easier to align the teacher and student. Thus, the proposed PCKA architecture amazingly boosts the previous methods.

**Visualize the CKA value.** We present some visualizations to show that our method does bridge the teacher-student gap in logit-level. In particular, we visualize the logit similarity for 6 teacher-student pairs in Appendix G. We find that our method significantly improves the logit-similarity.

T→S		RetinaNet-X101→RetinaNet-R50					
Method		AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>
Teacher		41.0	60.9	44.0	23.9	45.2	54.0
Student		37.4	56.7	39.6	20.0	40.7	49.7
Batch avg.		38.5	57.9	40.8	20.7	41.5	52.5
Spatial avg.		39.3	58.7	41.9	21.4	41.3	50.9
Mix-up avg.		38.2	58.1	40.4	21.3	41.9	50.9
Channel avg.		<b>40.3</b>	<b>59.9</b>	<b>43.0</b>	<b>23.3</b>	<b>44.2</b>	<b>54.9</b>

Table 5: The ablation study on the COCO validation set (T→S refers to the distillation from T to S). Here, we explore which dimension we should choose to get better results. "Mix-up" means the 1st distilling layer uses Batch avg. method, the 2nd distilling layer uses Spatial avg. method and the final distilling layer uses Channel avg. method

**Visualize the training process.** We further visualize the training process of different detectors and the patch effect on the RetinaNet-X101-RetinaNet-R50 pair. The results are in Tab. 4 and Tab. 3 in the Appendix.

T→S		RetinaNet-X101→RetinaNet-R50					
Method		AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>
Teacher		41.0	60.9	44.0	23.9	45.2	54.0
Student		37.4	56.7	39.6	20.0	40.7	49.7
MMD w/ patch		38.5	57.7	40.9	22.2	42.8	51.3
PCKA		<b>40.3</b>	<b>59.9</b>	<b>43.0</b>	<b>23.3</b>	<b>44.2</b>	<b>54.9</b>

Table 6: Experiments on the upper bound of MMD. We derive the formula that CKA is the upper bound of MMD with a constant term. From these experiments, we can prove that optimizing the upper bound of MMD can better improve the performance, compared with MMD.

**Visualize the inference outputs.** We first visualize the confusion matrix of the proposed method in Fig. 6, and then visualize the annotated images of training with/without patches averaging different dimensions in Fig. 7, respectively. These figures reveal that our method can collect the similarities between different classes, and also show the effectiveness of our method on the object detection task.

## 5 Discussion

**PCKA in image classification.** We apply PCKA to the classification task, and it also outperforms well on the methods with the same distillation type, as shown in Tab. 15 in Appendix. However, PCKA performs worse on the teacher-student pairs with different architectures. As cutting activations of different architectures contain more dissimilar and harmful representations, bringing difficulty in transferring knowledge to the student.

**Average on the channel, boosting the performance.** The results in Tab. 7 reveal an interesting phenomenon, where the performance of the previous distillation methods is boosted by averaging the loss on channel dimension after the activations are cut into patches. Instead of directly matching the whole representation distribution in the activations, cutting patches makes the alignment between the teacher and student easier with a smaller representation distribution in

T→S		RetinaNet-X101→RetinaNet-R50					
Method		AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>
Teacher		41.0	60.9	44.0	23.9	45.2	54.0
Student		37.4	56.7	39.6	20.0	40.7	49.7
KD		37.2	<b>56.5</b>	39.3	20.4	40.4	49.5
KD w/ patch		39.3	58.7	41.9	21.4	41.3	50.9
AT		34.4	52.3	36.4	17.7	37.2	47.8
AT w/ patch		37.4	56.6	39.9	20.8	40.6	49.8
DIST		39.8	59.5	42.5	22.0	43.7	53.0
DIST w/ patch		40.2	59.6	<b>43.2</b>	22.7	<b>44.8</b>	53.9
PCKA w/o patch		36.4	55.8	38.7	20.6	39.8	48.7
PCKA(ours)		<b>40.3</b>	<b>59.9</b>	43.0	<b>23.3</b>	44.2	<b>54.9</b>

Table 7: Ablation study of distillation methods with(w/) or without(w/o) patch on the COCO validation set (T→S refers to the distillation from T to S). Here, we surprisingly notice that previous distillation method performance can be improved by image patching.

patches. Besides, cutting into patches follows the idea in the classification framework, thus PCKA calculates the inter-class similarities and intra-class similarities in patches. Moreover, due to the superiority of cosine similarity over distance-based losses [Boudiaf *et al.*, 2020] and high-order distribution representations collected by the Gram matrix, PCKA outperforms DIST and AT.

**Positional information loss.** In PCKA, we cut the activation of the teacher and student into patches, and then flatten them into a vector. Although this operation damages the original positional information, performance does not deteriorate. We suppose that CKA ensures the focus of the optimization is the shape of the distribution, rather than the raw values in the Gram matrix, which is vital because a model's discriminative capability is dependent on the distribution of its features rather than its scale. Besides, at the beginning, the effect brought by PCKA is smaller, compared with CE loss. Therefore, CE loss motivates the optimization of the student model steadily, and starting from a certain moment, PCKA drives the student model to align complex and high-order representations, improving the generalization ability.

## 6 Conclusion

In this paper, we provide a novel theoretical perspective of CKA in knowledge distillation, which can be simplified as the upper bound of MMD with a constant term. Besides, we dynamically customize the application of CKA based on the characteristics of each task, with less computational source yet comparable performance than previous methods. Furthermore, we propose a novel processing architecture for knowledge distillation in object detection task, which can further boost the performance of previous distillation methods. Our experimental results, including both qualitative and quantitative ones, demonstrate the effectiveness of our methods. In future research, we will further explore the relationship between all similarity metric-based distillation methods, and explore the theoretical reason why averaging on the channel dimension with patches can boost the performance of previous methods.

## Acknowledgements

This work is supported by the National Natural Science Foundation of China (NO. 62102151), Shanghai Sailing Program (21YF1411200), CCF-Tencent Rhino-Bird Open Research Fund, the Open Research Fund of Key Laboratory of Advanced Theory and Application in Statistics and Data Science, Ministry of Education (KLATASDS2305), the Fundamental Research Funds for the Central Universities.

## References

[Ahn *et al.*, 2019] Sungsoo Ahn, Shell Xu Hu, et al. Variational information distillation for knowledge transfer. In *CVPR*, pages 9163–9171, Long Beach, CA, USA, Jun. 2019. IEEE.

[Boudiaf *et al.*, 2020] Malik Boudiaf, Jérôme Rony, Imtiaz Masud Ziko, Eric Granger, Marco Pedersoli, Pablo Piantanida, and Ismail Ben Ayed. A unifying mutual information view of metric learning: Cross-entropy vs. pairwise losses. In *ECCV*, 2020.

[Cai and Vasconcelos, 2017] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: Delving into high quality object detection. *CVPR*, pages 6154–6162, 2017.

[Chen *et al.*, 2021] Pengguang Chen, Shu Liu, et al. Distilling knowledge via knowledge review. In *CVPR*, pages 5008–5017, Virtual Event, Jun. 2021. IEEE.

[Chen *et al.*, 2023] Huanran Chen, Yinpeng Dong, Zhengyi Wang, Xiao Yang, Chengqi Duan, Hang Su, and Jun Zhu. Robust classification via a single diffusion model. *arXiv preprint arXiv:2305.15241*, 2023.

[Chong *et al.*, 2022] Zhiyu Chong, Xinzhu Ma, Hong Zhang, Yuxin Yue, Haojie Li, Zhihui Wang, and Wanli Ouyang. Monodistill: Learning spatial features for monocular 3d object detection. *ArXiv*, abs/2201.10830, 2022.

[Cui *et al.*, 2023] Jiequan Cui, Zhuotao Tian, Zhisheng Zhong, Xiaojuan Qi, Bei Yu, and Hanwang Zhang. Decoupled kullback-leibler divergence loss. *ArXiv*, abs/2305.13948, 2023.

[Dai *et al.*, 2021] Xing Dai, Zeren Jiang, et al. General instance distillation for object detection. In *CVPR*, pages 7842–7851, Virtual Event, Jun. 2021.

[de Rijk *et al.*, 2022] Philip de Rijk, Lukas Schneider, Marius Cordts, and Dariu M. Gavrila. Structural knowledge distillation for object detection. *ArXiv*, abs/2211.13133, 2022.

[Frankle and Carbin, 2019] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *ICLR*, New Orleans, LA, USA, May 2019.

[Gong *et al.*, 2023] Linrui Gong, Shaohui Lin, Baochang Zhang, Yunhang Shen, Ke Li, Ruizhi Qiao, Bohan Ren, Muqing Li, Zhou Yu, and Lizhuang Ma. Adaptive hierarchy-branch fusion for online knowledge distillation. In *AAAI*, 2023.

[Heo *et al.*, 2019] Byeongho Heo, Jeesoo Kim, and Sangdoo et al. A comprehensive overhaul of feature distillation. In *ICCV*, pages 1921–1930, Seoul, Korea (South), Oct.-Nov. 2019. IEEE.

[Hinton *et al.*, 2015] Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2(7), 2015.

[Huang *et al.*, 2022] Tao Huang, Shan You, et al. Knowledge distillation from a stronger teacher. In *NeurIPS*, New Orleans, LA, USA, Nov.-Dec. 2022. NIPS.

[Kornblith *et al.*, 2019] Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey E. Hinton. Similarity of neural network representations revisited. *ArXiv*, abs/1905.00414, 2019.

[Krizhevsky and Hinton, 2009] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. *Handbook of Systemic Autoimmune Diseases*, 2009.

[Li *et al.*, 2020] Xiang Li, Wenhui Wang, Lijun Wu, Shuo Chen, Xiaolin Hu, Jun Li, Jinhui Tang, and Jian Yang. Generalized focal loss: Learning qualified and distributed bounding boxes for dense object detection. *ArXiv*, abs/2006.04388, 2020.

[Lin *et al.*, 2014] Tsung-Yi Lin, Michael Maire, et al. Microsoft coco: Common objects in context. In *ECCV*, pages 740–755, Zurich, Switzerland, Sept. 2014. Springer.

[Lin *et al.*, 2017] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. *ICCV*, pages 2999–3007, 2017.

[Liu *et al.*, 2021a] Li Liu, Qingle Huang, Sihao Lin, Hongwei Xie, Bing Wang, Xiaojun Chang, and Xiao-Xue Liang. Exploring inter-channel correlation for diversity-preserved knowledge distillation. *ICCV*, pages 8251–8260, 2021.

[Liu *et al.*, 2021b] Songhua Liu, Tianwei Lin, Dongliang He, Fu Li, Ruifeng Deng, Xin Li, Errui Ding, and Hao Wang. Paint transformer: Feed forward neural painting with stroke prediction. In *ICCV*, pages 6598–6607, October 2021.

[Marcel and Rodriguez, 2010] Sébastien Marcel and Yann Rodriguez. Torchvision the machine-vision package of torch. In *ACM MM*, pages 1485–1488, Firenze, Italy, Oct. 2010. ACM.

[Nguyen *et al.*, 2020] Thao Nguyen, Maithra Raghu, and Simon Kornblith. Do wide and deep networks learn the same things? uncovering how neural network representations vary with width and depth. *ArXiv*, abs/2010.15327, 2020.

[Orhan and Pitkow, 2017] Emin Orhan and Xaq Pitkow. Skip connections eliminate singularities. In *ICLR*, 2017.

[Park *et al.*, 2019] Wonpyo Park, Dongju Kim, et al. Relational knowledge distillation. In *CVPR*, pages 3967–3976, Long Beach, CA, USA, Jun. 2019. IEEE.

[Peng *et al.*, 2019] Baoyun Peng, Xiao Jin, et al. Correlation congruence for knowledge distillation. In *ICCV*, pages 5007–5016, 2019.

[Qiu *et al.*, 2022] Zengyu Qiu, Xinzhu Ma, Kunlin Yang, Chunya Liu, Jun Hou, Shuai Yi, and Wanli Ouyang. Better teacher better student: Dynamic prior knowledge for knowledge distillation. *ArXiv*, abs/2206.06067, 2022.

[Ren *et al.*, 2015] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE PAMI*, 39:1137–1149, 2015.

[Romero *et al.*, 2014] Adriana Romero, Nicolas Ballas, et al. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*, 2014.

[Russakovsky *et al.*, 2015] Olga Russakovsky, Jia Deng, et al. Imagenet large scale visual recognition challenge. *IJCV*, 115(3):211–252, 2015.

[Saha *et al.*, 2022] Aninda Saha, Alina Bialkowski, and Sara Khalifa. Distilling representational similarity using centered kernel alignment (cka). In *BMVC*, 2022.

[Shao *et al.*, 2023a] Shitong Shao, Huanran Chen, Zhen Huang, Linrui Gong, Shuai Wang, and Xinxiao Wu. Teaching what you should teach: a data-based distillation method. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pages 1351–1359, 2023.

[Shao *et al.*, 2023b] Shitong Shao, Xu Dai, Shouyi Yin, Lujun Li, Huanran Chen, and Yang Hu. Catch-up distillation: You only need to train once for accelerating sampling. *arXiv preprint arXiv:2305.10769*, 2023.

[Shu *et al.*, 2021] Changyong Shu, Yifan Liu, et al. Channel-wise knowledge distillation for dense prediction. In *ICCV*, pages 5311–5320, Montreal, Canada, Oct. 2021. IEEE.

[Tian *et al.*, 2019] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. *ICCV*, pages 9626–9635, 2019.

[Tian *et al.*, 2020] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive representation distillation. In *ICLR*, Addis Ababa, Ethiopia, Apr. 2020. OpenReview.net.

[Tung and Mori, 2019] Frederick Tung and Greg Mori. Similarity-preserving knowledge distillation. In *ICCV*, pages 1365–1374, Seoul, Korea (South), Oct-Nov. 2019. IEEE.

[Wan *et al.*, 2020] Alvin Wan, Xiaoliang Dai, Peizhao Zhang, Zijian He, Yuandong Tian, Saining Xie, Bichen Wu, Matthew Yu, Tao Xu, Kan Chen, Péter Vajda, and Joseph Gonzalez. Fbnetv2: Differentiable neural architecture search for spatial and channel dimensions. *CVPR*, pages 12962–12971, 2020.

[Wang *et al.*, 2019] Tao Wang, Li Yuan, Xiaopeng Zhang, and Jiashi Feng. Distilling object detectors with fine-grained feature imitation. *CVPR*, pages 4928–4937, 2019.

[Wu *et al.*, 2016] Jiaxiang Wu, Cong Leng, Yuhang Wang, Qinghao Hu, and Jian Cheng. Quantized convolutional neural networks for mobile devices. In *CVPR*, Las Vegas, NV, USA, Jun. 2016. IEEE.

[Yang *et al.*, 2021] Jing Yang, Brais Martinez, et al. Knowledge distillation via softmax regression representation learning. In *ICLR*. OpenReview.net, 2021.

[Yang *et al.*, 2022a] Zhendong Yang, Zhe Li, et al. Focal and global knowledge distillation for detectors. In *CVPR*, pages 4643–4652, New Orleans, LA, USA, Jun. 2022. IEEE.

[Yang *et al.*, 2022b] Zhendong Yang, Zhe Li, Mingqi Shao, Dachuan Shi, Zehuan Yuan, and Chun Yuan. Masked generative distillation. In *ECCV*, 2022.

[Yang *et al.*, 2023] Zhendong Yang, Ailing Zeng, Zhe Li, Tianke Zhang, Chun Yuan, and Yu Li. From knowledge distillation to self-knowledge distillation: A unified approach with normalized loss and customized soft labels. *ArXiv*, abs/2303.13005, 2023.

[Zagoruyko and Komodakis, 2016] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. In *ICLR*, 2016.

[Zhang and Ma, 2021] Linfeng Zhang and Kaisheng Ma. Improve object detection with feature-based knowledge distillation: Towards accurate and efficient detectors. In *ICLR*, pages 1–14. OpenReview.net, 2021.

[Zhao *et al.*, 2022] Borui Zhao, Quan Cui, Renjie Song, Yiyu Qiu, and Jiajun Liang. Decoupled knowledge distillation. In *CVPR*, pages 11953–11962, New Orleans, LA, USA, Jun. 2022. IEEE.

[Zhixing *et al.*, 2021] Du Zhixing, Rui Zhang, Ming Chang, Shaoli Liu, Tianshi Chen, Yunji Chen, et al. Distilling object detectors with feature richness. *NeurIPS*, 34:5213–5224, Dec. 2021.