

CanKD: Cross-Attention-based Non-local operation for Feature-based Knowledge Distillation

Shizhe Sun, Wataru Ohyama
Tokyo Denki University, Tokyo, Japan

yukarari77@outlook.com, w.ohyama@mail.dendai.ac.jp

Abstract

We propose Cross-Attention-based Non-local Knowledge Distillation (CanKD), a novel feature-based knowledge distillation framework that leverages cross-attention mechanisms to enhance the knowledge transfer process. Unlike traditional self-attention-based distillation methods that align teacher and student feature maps independently, CanKD enables each pixel in the student feature map to dynamically consider all pixels in the teacher feature map. This non-local knowledge transfer more thoroughly captures pixel-wise relationships, improving feature representation learning. Our method introduces only an additional loss function to achieve superior performance compared with existing attention-guided distillation methods. Extensive experiments on object detection and image segmentation tasks demonstrate that CanKD outperforms state-of-the-art feature and hybrid distillation methods. These experimental results highlight CanKD’s potential as a new paradigm for attention-guided distillation in computer vision tasks. Code is available at <https://github.com/torihotaru/CanKD>.

1. Introduction

Rapid advances in deep learning have resulted in models of greater depth and width that achieve superior performance in data processing and analysis [16, 22, 34]. Particularly, dense prediction tasks in the field of computer vision, such as semantic segmentation [8, 56] and object detection [32, 43], require models to comprehend image features at the pixel level, and the inherent complexity of these features enables exceptional model performance for complex, heavy-parameter tasks. However, this increased model width and depth pose significant challenges to computational and memory resources.

Knowledge distillation (KD) [3, 17] is a viable method for addressing these challenges, enabling knowledge transfer from a high-performance, heavy-parameter teacher




Figure 1. **Proposed Cross-attention-based Non-local (Can) block.**

model to a moderate-performance, reduced-parameter student model. KD is broadly categorized into two approaches: one based on model output, known as logit distillation [13, 19, 55], and the other based on the intermediate feature layers of the model, known as feature distillation [38]. Instead of directly transferring the final outputs to the student model, the teacher model conveys information in logarithmic form or through intermediate layer features. Compared with distilling the output layer of the teacher model, distilling intermediate layers typically achieves higher accuracy, as the intermediate feature layers of the teacher model are believed to contain more information [46].

Traditional knowledge distillation methods overlook the relationships between different pixels, which are beneficial for further enhancing the model’s distillation performance. To incorporate these relationships and benefit from contextual information, recent studies have employed global attention mechanisms, such as GcBlock [6] and Relation Networks [18]. These mechanisms utilize attention maps to establish connections between individual pixels in feature maps, potentially further enhancing the performance of the student network. Conventional attention-based distillation models independently compute self-attention on the teacher and student feature maps before aligning them. This ap-

proach achieves consistency between pixels at corresponding positions in the student and teacher feature maps.

Inspired by the strong performance of cross-attention in numerous dense prediction models [22, 35], we propose a hypothesis: **Whether, in a dense prediction task, cross-attention capturing relationships between each student pixel at specific locations and all teacher pixels could enhance distillation performance.** Interestingly, in dense prediction distillation tasks, prior attention-based distillation methods have predominantly focused on self-attention, as though an inherent obstacle has impeded the use of cross-attention in this context. We posit that **the limitations of cross-attention performance may stem from the conventional attention computation schemes that have long been employed in Transformers.**

To prove our hypothesis, we propose a novel Cross-Attention-based Non-local distillation method, called **CanKD**. Our method utilizes Cross-Attention Non-local (Can) blocks, as illustrated in Figure 1, to enhance the student feature map and more comprehensively learn from the teacher feature map. Through cross-attention, each pixel in the student feature map can dynamically consider all pixels at various locations within the teacher feature map, thereby capturing and establishing complex pixel-level relationships between the student and teacher models. Meanwhile, we replace the conventional attention computation with a more direct method that expands the attention map values to infinity.

Compared to traditional attention-based distillation modules, such as [50, 53], where both the teacher and student feature maps undergo attention computations, our approach achieves superior performance with fewer parameters.¹ Moreover, compared with other high-performing yet complex knowledge distillation methods [20, 54], our distillation approach introduces only an additional loss function to achieve nearly equivalent or superior performance. Conversely, in FreeKD [54], the student and teacher models include L2 loss in the neck and smooth L1 loss in the head position in addition to the FreeKD loss.

We summarize our main contributions as follows:

1. We present a novel Cross-Attention-based Non-local distillation method (CanKD) to enable the student model to learn from the teacher model more thoroughly, challenging the traditional paradigm in current attention-guided distillation.
2. We evaluate the effectiveness of our method for image detection by testing multiple detector models on the COCO [30] dataset. Our approach outperforms state-of-the-art feature distillation methods. Meanwhile, we evaluate the performance of our method for image segmentation tasks using the Cityscapes dataset [11], further

demonstrating the general applicability of our method to image-related tasks.

3. We evaluate the effectiveness of our method on two recent vision foundation models and one large-scale dataset. Each result exceeds the student baseline model, further confirming the universality of our approach across different model architectures and datasets of varying scales.

2. Related work

2.1. Knowledge distillation on dense prediction tasks

In recent years, several studies applied knowledge distillation to dense prediction tasks. For example, MIMIC [27], as an early study, used L2 loss to force the intermediate layers of the student network to imitate those of the teacher network; CWD [41] processed the activation map of each channel within the teacher and student feature map to derive a soft probability map, rather than conventionally handling activation values at each spatial location; MGD [51] introduced a novel distillation approach that randomly masked the pixels of the student feature map and subsequently used simple convolution blocks to enforce the generation of the teacher’s complete features; DiffKD [20] used a diffusion model to denoise the student feature map, thereby enabling distillation of the refined clean features and the teacher features.

Attention-guided methods are a viable approach to enhance distillation performance in dense prediction tasks with an imbalance between foreground and background pixels. For example, FKD [53] applied a non-local attention block to both the teacher and student feature maps. FGD [50] innovatively partitions the image into foreground and background segments and subsequently applies attention mechanisms to each segment individually. These approaches achieved state-of-the-art (SOTA) performance in dense prediction tasks.

2.2. Cross-attention mechanism

The advent of self-attention mechanisms [45] has significantly advanced the development of artificial intelligence models. Models based on self-attention mechanisms demonstrate exceptional performance in both natural language and image processing tasks [1, 12, 14, 22]. Variants of self-attention are a significant research topic.

Cross-attention was initially introduced as a variant of attention mechanisms in the decoder component of Transformer models [45] and progressively garnered increased popularity [15]. Cross-attention effectively integrates sequences from distinct information sources by combining the queries, keys, and values of two sequences, thereby establishing dynamic associations between them. In computer

¹We compare CanKD with other SOTA methods in parameters and FLOPs, which are shown in Section 10 in supplementary.

vision, cross-attention is extensively employed in text-to-image generation tasks. For instance, in the Stable Diffusion model [37], cross-attention maps various conditional modalities to the intermediate layers of the U-net architecture. Moreover, several models achieved highly favorable results in dense prediction and style transfer by applying cross-attention mechanisms [28, 57].

3. Feature-based knowledge distillation using Cross-Attention Non-local operation

Feature-based knowledge distillation generally involves optimizing a loss function $\mathcal{L}$ defined by the weighted sum of task loss $\mathcal{L}_{\text{task}}$ and feature distillation loss $\mathcal{L}_{\text{feat}}$. A simple and widely used determination of $\mathcal{L}_{\text{feat}}$ is the L_2 -distance between teacher and student feature maps:

$$\mathcal{L}_{\text{feat}} = \|F_T - F_S\|_2^2, \quad (1)$$

where F_T and F_S denote the teacher and student features maps, respectively. Our proposed CanKD approach transforms the student's feature map F_S in equation (1) before calculating the feature distillation loss. This enhances the guidance from the teacher model to the student model, resulting in improved feature-based knowledge distillation. Figure 2a illustrates the overview of CanKD. The feature map transformation using the Cross-Attention Non-local operation, which is the core of the proposed method, is denoted by the Can block on the student side of the bottleneck layers, between the backbone network and the head layers. The Can block is a specific implementation of the Can operator that introduces the operation anywhere in a neural network model. The following subsections provide a detailed description of the formulation and implementation of the Can operation.

3.1. Cross-Attention-based Non-local operation

The Cross-Attention Non-local operation takes two feature maps as input: an input (student) feature map $\mathbf{X}$ and a reference (teacher) feature map $\mathbf{Y}$, and outputs a single feature map $\mathbf{Z}$. Each feature map is a sequence of spatially or temporally arranged feature vectors:

$$\mathbf{X} = [\mathbf{x}_i], \quad \mathbf{Y} = [\mathbf{y}_i], \quad \mathbf{Z} = [\mathbf{z}_i], \quad (2)$$

where $i = (1, \dots, N)$ is a location index. N denotes the size of feature maps. For image-based tasks, N equals the product of the width W and height H of the feature map. $\mathbf{X}$, $\mathbf{Y}$, and $\mathbf{Z}$ are of equal size. The Can operation captures the relationship between a feature vector at a specific location in the input feature map and the feature vectors corresponding to all locations in the reference feature map. It then transforms the input feature map using the captured relationship information.

The Can operation for a feature vector at the i -th location in the input map is defined by

$$\mathbf{z}_i = \frac{1}{N} \sum_j \xi(\mathbf{x}_i, \mathbf{y}_j) g(\mathbf{y}_j), \quad (3)$$

where $\mathbf{x}_i$ and $\mathbf{y}_j$ are feature vectors at the i -th location in the input (student) and reference (teacher) feature maps, respectively. $\mathbf{z}_i$ is the output feature vector of the same size as $\mathbf{x}$ and $\mathbf{y}$. ξ is a scalar function representing the relationship (similarity) between $\mathbf{x}_i$ and $\mathbf{y}_j$. g is a unary function representing a vector $\mathbf{y}$. This is inspired by the work in [47]. However, while [47] calculates self-attention within the same data source, the Can operation instead assigns feature vectors obtained from different sources to the input and reference, respectively, to compute the relationship.

Some variations exist in the definition of the ξ function. In this study, to facilitate implementation, we use the vector product of the embedded representations of the input and reference vectors:

$$\xi(\mathbf{x}_i, \mathbf{y}_j) = \theta(\mathbf{x}_i)^\top \phi(\mathbf{y}_j). \quad (4)$$

3.2. Implementation of the Can operation

We use the Can operation for feature-based knowledge distillation in dense image-based tasks. We implement this operation as a module, the Can block, to facilitate its incorporation into neural networks. Figure 1 illustrates the structure of the Can block. In dense image-based tasks, feature maps are originally represented as third-order tensors, denoted as $\mathbf{F} \in \mathbb{R}^{H \times W \times C}$. We apply the Can operation by first regarding these third-order tensors as second-order tensors $\mathbf{F} \in \mathbb{R}^{HW \times C}$ and then treating the student and teacher feature maps $\mathbf{F}_S, \mathbf{F}_T$ as the input and reference feature maps $\mathbf{X}, \mathbf{Y}$ in equation (2), respectively. Here, H, W , and C represent the height, width of the feature map, and the number of channels, respectively, which determine the shape of the feature maps.

Following [47], we define the Can block as:

$$\mathbf{F}_S^* = \mathbf{W}_Z \mathbf{Z} + \mathbf{F}_S, \quad (5)$$

thereby combining the Can operation with a residual connection. Each feature vector in the feature map $\mathbf{Z}$ is calculated by equation (3). A residual connection performs an element-wise summation of the original student feature maps and the output feature maps of the Can operation. This enhances the student feature map to effectively incorporate the knowledge from various spatial locations of the teacher feature maps.²

Inspired by the Transformer's decoder in [45], we employ three trainable 1×1 convolutions in space to embed

²The ablation study of residual connection is shown in Section 5 in the supplementary.





Figure 2. Proposed feature-based knowledge distillation using the Cross-Attention-based Non-local operation

functions θ , ϕ , and g in the Can operation. These convolutions perform the cross-attention mechanism to map the student feature map using θ and the teacher feature map using ϕ and g :

$$\begin{aligned}\theta(F_S) &= W_\theta F_S, \\ \phi(F_T) &= W_\phi F_T, \\ g(F_T) &= W_g F_T,\end{aligned}\quad (6)$$

where W_θ , W_ϕ , and W_g are 1×1 convolution weights to be trained. As shown in equation (4), we use dot product similarity as the pairwise function to calculate the affinity. This function can be easily extended to feature maps by utilizing the tensor dot product.

Additionally, we incorporate a down-sampling technique in our block to reduce pairwise computation. We apply a max pooling layer to each of the two embeddings, $\phi(F_T)$ and $g(F_T)$, derived from the teacher feature maps. This technique can reduce pairwise computation.³

3.3. Overall loss

We posit that a discrepancy exists between the distributions of the enhanced student and teacher feature maps. This discrepancy can, to some extent, impact the model's distillation performance. To mitigate this discrepancy, we incorporate instance normalization to standardize both feature maps before computing the L2 loss. Thus, our feature distillation loss function is:

$$\mathcal{L}_{\text{feat}} = \|\Omega(F_T) - \Omega(F_S^*)\|_2^2, \quad (7)$$

where $\Omega()$ is the normalization process.

We introduce a hyperparameter μ to balance the loss of feature distillation and the loss of task. The overall training

loss to optimize is:

$$\mathcal{L}_{\text{student}} = \mathcal{L}_{\text{task}} + \mu \mathcal{L}_{\text{feat}}. \quad (8)$$

Notably, our distillation loss is primarily applied to intermediate feature maps, such as those used in Feature Pyramid Networks (FPN) [31], making it applicable to various architectures. Furthermore, our distillation method merely introduces an additional loss term to the existing loss function, enabling its seamless integration with other distillation approaches.

4. Experiments

We applied the proposed method to two types of dense tasks, namely object detection and semantic segmentation, to evaluate its effectiveness. We followed the standard experimental procedures commonly used in previous studies for both tasks and compared our performance with that of state-of-the-art approaches.⁴

4.1. Object detection

4.1.1. Details

We verify our method on the MSCOCO detection dataset [30], which contains 80 object classes. We train our student models on COCO train2017 and test on COCO val2017. We use average precision (AP) to evaluate performance.

As in [20, 50, 53], we experiment using baseline and enhanced settings. In the baseline setting, we test two-stage detector Faster-RCNN [36], one-stage detector RetinaNet [29], and anchor-free detector FCOS [42]. In the teacher network and student network, the backbone is ResNet-101 (R101) [16] and ResNet-50 (R50), respectively. In the enhanced setting, we replace ResNet-101 with ResNeXt101

³The ablation study of max pooling is available in Section 4 in the supplementary.

⁴We also test classification task with ImageNet-1k, which are shown in Section 9 in supplementary.

Table 1. **Object detection on the COCO validation dataset under baseline setting.** T:Teacher. S:Student. † denotes *inheriting strategy* [21, 26, 54]. F:Feature distillation. L:logit distillation. ‡ denotes introducing the space search algorithm[25, 26]. We train FCOS[42] with techniques similar to those in FGD [50].

Method	Distillation	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
<i>One-stage detectors</i>							
T: RetinaNet-R101	N/A	38.9	58.0	41.5	21.0	42.8	52.4
S: RetinaNet-R50	N/A	37.4	56.7	39.6	20.0	40.7	49.7
FGD[50]	F	39.6(2.2†)	-	-	22.9	43.7	53.6
Auto-KD[25]	F+L†	39.5(2.1†)	58.5	42.4	22.8	43.5	53.3
DiffKD[20]	F+L	39.7(2.3†)	58.6	42.1	21.6	43.8	53.3
CanKD	F	39.7(2.3†)	58.6	42.5	21.8	43.8	53.6
CanKD†	F	39.8(2.4†)	58.8	42.6	21.6	44.0	53.7
<i>Two-stage detectors</i>							
T: FasterRCNN-R101	N/A	39.8	60.1	43.3	22.5	43.6	52.8
S: FasterRCNN-R50	N/A	38.4	59.0	42.0	21.5	42.1	50.3
FGD[50]	F	40.4(2.0†)	-	-	22.8	44.5	53.5
PKD[5]	F	40.5(2.1†)	60.9	44.4	22.6	44.8	53.1
Auto-KD[25]	F+L†	40.1(1.7†)	60.6	43.7	22.7	44.0	52.8
DiffKD[20]	F+L	40.6(2.2†)	60.9	43.9	23.0	44.5	54.0
DetKDS†[26]	F†	40.6(2.2†)	60.8	44.1	22.9	44.8	53.5
CanKD	F	40.1(1.7†)	60.4	43.7	22.1	44.3	52.8
CanKD†	F	40.7(2.3†)	60.8	44.4	23.3	44.9	53.6
<i>Anchor-free detectors</i>							
T: FCOS-R101	N/A	40.8	60.0	44.0	24.2	44.3	52.4
S: FCOS-R50	N/A	38.5	57.7	41.0	21.9	42.8	48.6
FGD[50]	F	42.1(3.6†)	-	-	27.0	46.0	54.6
Auto-KD[25]	F+L†	42.0(3.5†)	60.4	45.5	25.9	45.8	54.4
DiffKD[20]	F+L	42.4(3.9†)	61.0	45.8	26.6	44.5	54.0
PCKA+mimic[58]	F+L	40.7(3.3†)	60.5	43.1	23.4	44.8	53.1
DetKDS†[26]	F†	42.3(3.8†)	60.8	45.6	26.1	46.2	54.7
CanKD	F	42.4(3.9†)	60.8	45.9	26.1	46.5	55.0
CanKD†	F	42.4(3.9†)	60.8	45.8	26.8	46.1	55.2

(X101) [48] in the teacher backbone. We test two-stage detector Cascade Mask RCNN [4], one-stage detector RetinaNet [29], and anchor-free detector RepPoints [49]. These methods conduct the distillation in the model’s neck module. As shown in Figure 2b, we distill the output of every layer within the feature pyramid located in the neck (from P3 to P6).

Following the official strategies in MMDetection [7], we use stochastic gradient descent (SGD) with a learning rate of 0.005 and weight decay of 1e-4 with 2× schedule in MMRazor [10].⁵ The loss weights μ in equation (8) are all set to five, as determined by an ablation study. Furthermore, to verify that our method maintains suitable performance with a heterogeneous teacher, we follow the PKD [5] training strategy, using FCOS-X101 as the teacher and RetinaNet-R50 as the student to conduct additional experiments.

4.1.2. Experiment results

We compared existing KD methods in Table 1 using the baseline setting. Our CanKD approach significantly improves the performance of the student model across multiple frameworks, surpassing other state-of-the-art feature distil-

⁵The further details of training strategies are shown in Section 1 in the supplementary.

Table 2. **Object detection on the COCO validation dataset under enhanced setting.** T:Teacher. S:Student. *CM RCNN*: Cascade Mask RCNN. † denotes *inheriting strategy* [21, 26, 54]. F:Feature distillation. L:logit distillation. ‡ denotes introducing the space search algorithm[24, 26].

Method	Distillation	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
<i>One-stage detectors</i>							
T: RetinaNet-X101	N/A	41.0	60.9	44.0	23.9	45.2	54.0
S: RetinaNet-R50	N/A	37.4	56.7	39.6	20.0	40.7	49.7
FKD[53]	F	39.6(2.2†)	58.8	42.1	22.7	43.3	52.5
FGD[50]	F	40.4(3.0†)	-	-	23.4	44.7	54.1
KD-Zero[24]	F+L†	40.9(3.5†)	60.4	43.5	23.2	45.2	54.8
DiffKD[20]	F+L	40.7(3.3†)	60.0	43.2	22.2	45.0	55.2
PCKA+mimic[58]	F+L	40.7(3.3†)	60.4	43.4	23.9	44.7	55.1
DetKDS†[26]	F†	41.0(3.6†)	60.2	43.6	23.0	45.2	55.0
FreeKD†[54]	F+L	41.0(3.6†)	-	-	22.3	45.1	55.7
CanKD	F	41.0(3.6†)	60.6	43.8	24.3	45.9	54.4
CanKD†	F	41.1(3.7†)	60.6	43.8	23.6	45.4	55.3
<i>Two-stage detectors</i>							
T: CM RCNN-X101	N/A	45.6	64.1	49.7	26.2	49.6	60.0
S: FasterRCNN-R50	N/A	38.4	59.0	42.0	21.5	42.1	50.3
FKD[53]	F	41.5(3.1†)	62.2	45.1	23.5	45.0	55.3
FGD[50]	F	42.0(3.6†)	-	-	23.7	46.4	55.5
KD-Zero[24]	F+L†	41.9(3.5†)	62.7	45.5	23.6	45.6	55.6
DiffKD[20]	F+L	42.2(3.8†)	62.8	46.0	24.2	46.6	55.3
PCKA+mimic[58]	F+L	42.4(4.0†)	63.3	46.1	24.3	46.7	56.1
DetKDS†[26]	F†	42.3(3.9†)	62.8	46.2	24.8	46.1	56.0
FreeKD†[54]	F+L	42.4(4.0†)	-	-	24.1	46.7	55.9
CanKD	F	42.0(3.6†)	62.4	45.6	23.8	46.4	55.1
CanKD†	F	42.5(4.1†)	62.9	46.5	23.9	47.0	56.7
<i>Anchor-free detectors</i>							
T: RepPoints-X101	N/A	44.2	65.5	47.8	26.2	48.4	58.5
S: RepPoints-R50	N/A	38.6	59.6	41.6	22.5	42.2	50.4
FKD[53]	F	40.6(2.0†)	61.7	43.8	23.4	44.6	53.0
FGD[50]	F	41.3(2.7†)	-	-	24.5	45.2	54.0
DiffKD[20]	F+L	41.7(3.1†)	62.6	44.9	23.6	45.4	55.9
DetKDS†[26]	F†	42.3(3.7†)	63.1	45.8	24.1	46.4	55.9
FreeKD†[54]	F+L	42.4(3.8†)	-	-	24.3	46.4	56.6
CanKD	F	42.4(3.8†)	62.9	45.6	24.1	46.5	56.4
CanKD†	F	42.6(4.0†)	63.1	45.6	24.1	47.0	56.4

Table 3. **Object detection on the COCO validation dataset with heterogeneous teacher.** T:Teacher. S:Student.

Method	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
T: FCOS-X101	42.7	62.5	45.7	26.0	46.5	54.7
S: RetinaNet-R50	37.4	56.7	39.6	20.0	40.7	49.7
PKD[5]	40.3(2.9†)	59.6	43.0	22.2	44.9	53.7
DetKDS[26]	40.4(3.0†)	59.8	43.0	22.8	44.3	53.8
CanKD	40.5(3.1†)	59.9	43.4	23.2	45.1	54.2

lation methods and even exceeding the performance of specific state-of-the-art hybrid distillation approaches. For instance, CanKD using FasterRCNN [36] with a ResNet-R50 [16] backbone attains a 2.3-point AP improvement over the baseline on COCO.

We further investigate the effectiveness of CanKD using an enhanced teacher model. Our results are summarized in Table 2. We observe that CanKD demonstrates even greater superiority when applied with more powerful teacher models. Notably, CanKD achieves a 4.0-point AP improvement over the baseline when using the RepPoints-R50 [49]

student model, surpassing the AP of the following best method, FreeKD, by 0.2 points. With the CanKD method, the RetinaNet-R50 [29] student model even surpasses the RetinaNet-X101 [29] teacher model on the COCO dataset.

Under heterogeneous teacher conditions, our model also demonstrates excellent performance, achieving a 3.1-point AP improvement over the baseline and surpassing the AP of the top current state-of-the-art model DetKDS by 0.1 points, as shown in Table 3.⁶

4.2. Semantic segmentation

4.2.1. Details

We verify our method on the Cityscapes dataset [11], which comprises 5,000 high-quality images (2,975 training images and 500 validation images). In this task, we evaluate performance using the mean Intersection-over-Union (mIoU) metric.

We experiment with two segmentation frameworks, PSPNet [56] and DeeplabV3 [9]. Our teacher backbone is ResNet101. Our student backbone uses two frameworks: ResNet18 and MobileNetV2 [39]. The crop size of all teacher and student models is 512×1024 . The loss weights, μ , of CanKD in equation (6) are set to ten, as determined by an ablation study. In other configurations, the loss weights μ are all set to five following the configuration of the detection task. The training strategies of CWD [41] and MGD [51] follow the original CWD and MGD papers. The distillation position of CanKD and MGD are the model’s backbone output, while the distillation position of CWD is the logit output. To address the channel mismatch between teacher and student feature maps, we follow prior work [50, 51, 54], introducing a 1×1 convolution to align the channel dimensions of the student feature maps with those of the teacher. (We adopt the same strategy for other tasks as well.) Following the settings in MMSegmentation [7], we use SGD with a learning rate of 0.01, weight decay of $5e-4$, and 80K schedule in the MMRazor [10] framework for all methods.

4.2.2. Experiment results

Our semantic segmentation results are recorded in Table 4. Compared with previous state-of-the-art models, CanKD demonstrates outstanding performance for high-resolution images. In homogeneous settings, CanKD achieves a 0.42-point mIoU improvement over the baseline using DeepLabV3-R18, surpassing MGD[51] by 0.17 mIoU points. Furthermore, we observe that CanKD delivers excellent performance when the teacher and student models are heterogeneous. For instance, CanKD improves the mIoU by 0.64 points over the baseline when using the DeepLabV3-R18[9] student model with the PSPNet-R101[56] teacher model.

⁶The model’s confusion matrix and changes in training process are shown in Section 6 and 7 in the supplementary.

Table 4. **Semantic segmentation results on the Cityscapes validation dataset.** FLOPs is measured based on an input image size of 1024×2048 . * indicates results reproduced for this paper.

Method	FLOPs(G)	Params(M)	mIoU(%)	aAcc(%)	mAcc(%)
T: PSPNet-R101	2047.11	65.49	79.74	96.33	86.56
S: PSPNet-R18	433.86	12.63	74.87	95.56	82.07
CWD[41]			75.54	95.75	82.11
MGD*[51]			75.84	95.68	83.07
CanKD			76.24	95.67	85.49
S: DeepLabV3-R18	479.44	13.83	76.70	95.89	84.29
CWD*[41]			77.19	95.93	85.35
MGD*[51]			77.28	95.90	85.11
CanKD			77.34	96.04	85.49
T: DeepLabV3-R101	2777.00	84.63	80.31	96.44	87.16
S: DeepLabV3-R18	479.44	13.83	76.70	95.89	84.29
MGD*[51]			76.95	95.90	84.53
CanKD			77.12	96.01	85.47
S: DeepLabV3-MV2	598.56	18.34	73.84	95.28	83.72
CanKD			75.62	95.94	84.81

Table 5. **Object detection on recent vision foundation models and large-scale dataset.** T:Teacher. S:Student.

Method	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
T: Dino-R101	50.4	68.0	54.8	32.9	53.9	65.2
S: Dino-R50	49.0	66.4	53.3	31.4	52.2	64.0
CWD[41]	49.5(0.5 $\uparrow$)	67.0	54.2	31.5	52.7	64.5
PKD[5]	49.2(0.2 $\uparrow$)	66.7	53.8	32.3	52.4	63.8
QPD[23]	49.5(0.5 $\uparrow$)	67.2	53.9	-	-	-
QRD[23]	49.5(0.5 $\uparrow$)	67.5	54.1	-	-	-
CanKD	49.7(0.7$\uparrow$)	67.1	54.2	32.5	52.8	64.4
T: Dino-Swin-L	58.4	77.1	64.2	41.5	62.2	73.5
S: Dino-Swin-B	56.3	74.7	61.8	40.0	59.8	73.0
CWD[41]	54.7(1.6 $\downarrow$)	72.9	59.9	37.6	57.7	72.4
CanKD	56.5(0.2$\uparrow$)	74.8	61.9	40.4	60.0	72.6
T: GroundingDino-R101	50.0	67.4	54.6	32.3	54.0	64.5
S: GroundingDino-R50	48.9	66.4	53.3	30.7	52.2	64.2
CWD[41]	49.5(0.6 $\uparrow$)	66.7	54.1	32.3	53.0	64.1
PKD[5]	49.4(0.5 $\uparrow$)	66.8	53.9	31.7	52.8	64.0
CanKD	49.8(0.9$\uparrow$)	66.9	54.1	31.5	53.1	65.3
<i>Object365v2 Dataset</i>						
T: RetinaNet-R101	19.5	28.0	20.9	8.7	19.6	26.2
S: RetinaNet-R50	16.7	24.4	17.8	7.2	16.7	22.0
CanKD	18.4(1.7$\uparrow$)	26.6	19.7	7.7	18.3	24.9

4.3. Recent vision foundation models and large-scale dataset

To demonstrate the generality of our approach, we verify CanKD on two popular DETR-based models (Dino [52] and Grounding-Dino [33]) and a large-scale dataset (Object365v2 [40]) with various student backbone (R50, Swin-b[34]) in Table 5. From the results, our CanKD achieves improvements of +0.7, +0.2 and +0.9 over the student baseline on Dino and Grounding-Dino models. Furthermore, on ultra-large Object365v2 dataset, our CanKD outperforms the baseline by +1.7, further demonstrating the generality and advancement of our method.

Table 6. **Ablation study of instance normalization and the Can block.** We use RepPoints-X101 [49] as the teacher and RepPoints-R50 as the student.

Method	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L
w/o KD (baseline)	38.6	59.6	41.6	22.5	42.2	50.4
L2 only	41.4	62.0	44.5	23.3	45.2	54.5
Instance Norm + L2	41.7	62.2	45.0	23.6	46.1	55.6
CAN + Instance Norm + L2(CanKD)	42.4	62.9	45.6	24.1	46.5	56.4

Table 7. **Ablation study of different affinity functions.** We use RepPoints-X101 [49] as the teacher and RepPoints-R50 as the student to test the sensitivity of the affinity function.

Method	$Softmax$	θ, ϕ	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L
Gaussian	✓	×	41.8	62.2	45.0	24.1	46.0	55.9
Embedded Gaussian	✓	✓	41.9	62.3	45.3	24.4	46.4	55.3
Dot product	×	✓	42.4	62.9	45.6	24.1	46.5	56.4

5. Analysis

5.1. Ablation study of Can block and instance normalization

To verify the impact of Can block and instance normalization, we compared CanKD against a variant that removes the Can block and another variant that removes the Can block and instance normalization. Our results are recorded in Table 6. We use RepPoints-X101 as the teacher model and RepPoints-R50 as the student model. To maintain consistency with the original model, the weight for all methods is set to five. Our results successfully demonstrate the effectiveness of both the Can block and instance normalization. Compared with a distillation method that uses only L2 loss, jointly employing instance normalization and the Can block improves the AP by 1.0 points.

5.2. Sensitivity study of the affinity function

We replace the dot product affinity function in equation (4) with two alternative functions: the Gaussian function [2, 44] and the embedded Gaussian function [45].⁷ We use RepPoints-X101 as the teacher model and RepPoints-R50 as the student model. As shown in Table 7, the dot product approach shows a clear advantage over the other two Gaussian methods. We conjecture that this superiority may stem from the dot product’s positive influence on cross-attention.

In Figure 3, we present the heatmaps for each method from intermediate layers. These visualizations indicate that, compared with the dot product approach, the other two Gaussian methods are less effective in capturing the high-activation regions (the warmer-colored areas in the heatmaps) in the teacher’s feature map. We conjecture that introducing the *Softmax* layer in the first two methods normalizes the attention scores obtained from the affinity function into the range [0,1], thereby reducing the difference in

⁷The details of these functions are available in Section 2 in the supplementary.

Table 8. **Ablation study of different hyperparameters for the object detection task.** We use RepPoints-X101 [49] as the teacher and RepPoints-R50 as the student.

Weight	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L
Baseline	38.6	59.6	41.6	22.5	42.2	50.4
$\mu = 2$	42.3	63.0	45.5	24.1	46.5	56.2
$\mu = 5$	42.4	62.9	45.6	24.1	46.5	56.4
$\mu = 8$	41.8	62.0	45.2	24.3	46.4	54.5
$\mu = 10$	41.6	61.7	45.0	23.3	46.1	55.5

Table 9. **Ablation study of different hyperparameters for the semantic segmentation task.** We use PSPNet-R101 [56] as the teacher and DeepLabV3-R18 [9] as the student.

Weight	mIoU(%)	aAcc(%)	mAcc(%)
Baseline	76.70	95.89	84.29
$\mu = 2$	76.96	96.33	85.03
$\mu = 5$	77.28	96.02	85.08
$\mu = 8$	77.00	96.03	84.97
$\mu = 10$	77.34	96.04	85.49

attention values across pixels. Meanwhile, the dot product approach computes attention more directly, facilitating backpropagation. Consequently, we believe these factors account for the superior performance of the dot product strategy within our distillation framework.

5.3. Sensitivity study of weight hyperparameter

In equation (8), we use a hyperparameter μ to balance the magnitudes of the distillation loss and the task loss in the student model. Here, we explore alternative parameter values to investigate the impact of different hyperparameter settings on the model’s performance.

As shown in Table 8, in the object detection task with this teacher-student pairing, setting $\mu = 5$ yields the highest performance, improving the AP by 3.8 points over the baseline. In contrast, $\mu = 10$ gives the lowest performance with an improvement of only 3.0 points. As shown in Table 9, $\mu = 10$ achieves the best performance for the semantic segmentation task, improving mIoU by 0.64 points. Overall, the semantic segmentation task model exhibits less sensitivity to this parameter. Furthermore, according to the log data, when the value of the distillation loss is close to that of the task loss, our model easily achieves peak performance.

5.4. Visualization

In Figures 4 and 5, we show the heatmaps and prediction maps of CanKD for object detection and semantic segmentation. The student model distilled with CanKD effectively inherits the teacher’s knowledge for dense prediction tasks. Compared with the original student feature maps, the distilled model’s activations better align with those of




Figure 3. **Visualization heatmaps for different affinity functions.** We use RepPoints-X101 [49] as the teacher and RepPoints-R50 as the student. These heatmaps are generated from P6 in the FPN layers




Figure 4. **Visualization of student heatmaps on P6 in FPN layers, where the student is distilled with CanKD heatmaps and teacher heatmaps for the COCO validation dataset.** These heatmaps are selected from T:FasterRCNN-R101-S:FasterRCNN-R50 and T:RepPoints-X101-S:RepPoints-R50.




Figure 5. **Visualization of prediction results for semantic segmentation task in the Cityscapes validation dataset.** The teacher is PSPNet-R101 and the student is PSPNet-R50.

the teacher. Moreover, when the student produces correct outputs but the teacher errs, our distilled model preserves

the student’s accurate predictions.⁸ We attribute this to the cross-attention mechanism in the Can module, which builds robust latent relations between student and teacher pixels—unlike conventional attention methods that enforce pixel-level, one-to-one correspondences.

6. Conclusion

In this work, we proposed Cross-Attention-based Non-local Knowledge Distillation (CanKD). By leveraging cross-attention mechanisms, this novel distillation method enhances knowledge transfer in teacher-to-student models. CanKD allows each pixel in the student feature map to dynamically consider all pixels in the teacher feature map. This approach captures complex pixel-level relationships, leading to more effective feature learning and improved model performance.

Through extensive experiments on dense prediction tasks, we demonstrated that CanKD significantly outperforms state-of-the-art feature distillation methods and hybrid distillation approaches. Our method achieves superior performance while introducing only an additional loss function, rendering it computationally efficient compared with existing attention-based distillation techniques. The results highlight CanKD’s potential as a new paradigm for attention-guided knowledge distillation.

Finally, we discuss the limitations of our proposed method. Attention-based distillation methods share a common drawback: they are highly sensitive to pixel-level variations. Consequently, the widespread use of bilinear interpolation for H×W alignment can degrade their performance. Although CanKD performs slightly worse than complex multi-loss distillation methods on large models, its modular design makes it a promising component for integration into such frameworks, potentially enhancing their overall distillation effectiveness.⁹

⁸Additional visualization examples are provided in Section 8 of the supplementary.

⁹A more detailed discussion of the limitations is provided in Section 11 of the supplementary material.

References

- [1] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. In *2020 Advances in Neural Information Processing Systems*, pages 1877–1901, 2020. 2
- [2] Antoni Buades, Bartomeu Coll, and J-M Morel. A non-local algorithm for image denoising. In *2005 IEEE Conference on Computer Vision and Pattern Recognition*, pages 60–65, 2005. 7
- [3] Cristian Buciluă, Rich Caruana, and Alexandru Niculescu-Mizil. Model compression. In *12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 535–541, 2006. 1
- [4] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: Delving into high quality object detection. In *2018 IEEE Conference on Computer Vision and Pattern Recognition*, pages 6154–6162, 2018. 5
- [5] Weihao Cao, Yifan Zhang, Jianfei Gao, Anda Cheng, Ke Cheng, and Jian Cheng. Pkd: General distillation framework for object detectors via pearson correlation coefficient. In *2022 Advances in Neural Information Processing Systems*, pages 15394–15406, 2022. 5, 6
- [6] Yue Cao, Jiarui Xu, Stephen Lin, Fangyun Wei, and Han Hu. Gnet: Non-local networks meet squeeze-excitation networks and beyond. In *2019 IEEE/CVF International Conference on Computer Vision Workshops*, pages 1971–1980, 2019. 1
- [7] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, Zheng Zhang, Dazhi Cheng, Chenchen Zhu, Tianheng Cheng, Qijie Zhao, Buyu Li, Xin Lu, Rui Zhu, Yue Wu, Jifeng Dai, Jingdong Wang, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019. 5, 6
- [8] Liang-Chieh Chen. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017. 1
- [9] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *2018 European Conference on Computer Vision*, pages 801–818, 2018. 6, 7
- [10] MMRazor Contributors. Openmmlab model compression toolbox and benchmark. <https://github.com/open-mmlab/mmrazor>, 2021. 5, 6
- [11] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *2016 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3213–3223, 2016. 2, 6
- [12] Jacob Devlin. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2
- [13] Qianggang Ding, Sifan Wu, Hao Sun, Jiadong Guo, and Shu-Tao Xia. Adaptive regularization of labels. *arXiv preprint arXiv:1908.05474*, 2019. 1
- [14] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2
- [15] Mozhdeh Gheini, Xiang Ren, and Jonathan May. Cross-attention is all you need: Adapting pretrained transformers for machine translation. *arXiv preprint arXiv:2104.08771*, 2021. 2
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. 1, 4, 5
- [17] Geoffrey Hinton. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 1
- [18] Han Hu, Jiayuan Gu, Zheng Zhang, Jifeng Dai, and Yichen Wei. Relation networks for object detection. In *2018 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3588–3597, 2018. 1
- [19] Tao Huang, Shan You, Fei Wang, Chen Qian, and Chang Xu. Knowledge distillation from a stronger teacher. In *2022 Advances in Neural Information Processing Systems*, pages 33716–33727, 2022. 1
- [20] Tao Huang, Yuan Zhang, Mingkai Zheng, Shan You, Fei Wang, Chen Qian, and Chang Xu. Knowledge diffusion for distillation. In *2023 Advances in Neural Information Processing Systems*, pages 65299–65316, 2023. 2, 4, 5
- [21] Zijian Kang, Peizhen Zhang, Xiangyu Zhang, Jian Sun, and Nanning Zheng. Instance-conditional knowledge distillation for object detection. In *2021 Advances in Neural Information Processing Systems*, pages 16468–16480, 2021. 5
- [22] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *2023 IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023. 1, 2
- [23] Gang Li, Wenhao Wang, Xiang Li, Ziheng Li, Jian Yang, Jifeng Dai, Yu Qiao, and Shanshan Zhang. Distilling knowledge from large-scale image models for object detection. In *European Conference on Computer Vision*, pages 142–160. Springer, 2024. 6
- [24] Lujun Li, Peijie Dong, Anggeng Li, Zimian Wei, and Ya Yang. Kd-zero: Evolving knowledge distiller for any teacher-student pairs. *Advances in Neural Information Processing Systems*, 36:69490–69504, 2023. 5
- [25] Lujun Li, Peijie Dong, Zimian Wei, and Ya Yang. Automated knowledge distillation via monte-carlo tree search. In *2023 IEEE/CVF International Conference on Computer Vision*, pages 17413–17424, 2023. 5
- [26] Lujun Li, Yufan Bao, Peijie Dong, Chuanguang Yang, Anggeng Li, Wenhan Luo, Qifeng Liu, Wei Xue, and Yike Guo. Detkds: Knowledge distillation search for object detectors. In *2024 International Conference on Machine Learning*, 2024. 5
- [27] Quanquan Li, Shengying Jin, and Junjie Yan. Mimicking very efficient network for object detection. In *2017 IEEE*

- Conference on Computer Vision and Pattern Recognition*, pages 6356–6364, 2017. 2
- [28] Hezheng Lin, Xing Cheng, Xiangyu Wu, and Dong Shen. Cat: Cross attention in vision transformer. In *2022 IEEE International Conference on Multimedia and Expo*, pages 1–6, 2022. 3
- [29] T Lin. Focal loss for dense object detection. *arXiv preprint arXiv:1708.02002*, 2017. 4, 5, 6
- [30] 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 *2014 European Conference on Computer Vision*, pages 740–755, 2014. 2, 4
- [31] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *2017 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2117–2125, 2017. 4
- [32] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollar. Focal loss for dense object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(02):318–327, 2020. 1
- [33] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *European Conference on Computer Vision*, pages 38–55. Springer, 2024. 6
- [34] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *2021 IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021. 1, 6
- [35] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. 2
- [36] Shaoqing Ren. Faster r-cnn: Towards real-time object detection with region proposal networks. *arXiv preprint arXiv:1506.01497*, 2015. 4, 5
- [37] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. 3
- [38] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*, 2014. 1
- [39] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *2018 IEEE Conference on Computer Vision and Pattern Recognition*, pages 4510–4520, 2018. 6
- [40] Shuai Shao, Zeming Li, Tianyuan Zhang, Chao Peng, Gang Yu, Xiangyu Zhang, Jing Li, and Jian Sun. Objects365: A large-scale, high-quality dataset for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 6
- [41] Changyong Shu, Yifan Liu, Jianfei Gao, Zheng Yan, and Chunhua Shen. Channel-wise knowledge distillation for dense prediction. In *2021 IEEE/CVF International Conference on Computer Vision*, pages 5311–5320, 2021. 2, 6
- [42] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9627–9636, 2019. 4, 5
- [43] Zhi Tian, Xiangxiang Chu, Xiaoming Wang, Xiaolin Wei, and Chunhua Shen. Fully convolutional one-stage 3d object detection on lidar range images. In *2022 Advances in Neural Information Processing Systems*, pages 34899–34911, 2022. 1
- [44] Carlo Tomasi and Roberto Manduchi. Bilateral filtering for gray and color images. In *Sixth International Conference on Computer Vision*, pages 839–846, 1998. 7
- [45] A Vaswani. Attention is all you need. In *2017 Advances in Neural Information Processing Systems*, 2017. 2, 3, 7
- [46] Guo-Hua Wang, Yifan Ge, and Jianxin Wu. Distilling knowledge by mimicking features. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11):8183–8195, 2021. 1
- [47] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *2018 IEEE Conference on Computer Vision and Pattern Recognition*, pages 7794–7803, 2018. 3
- [48] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1492–1500, 2017. 5
- [49] Ze Yang, Shaohui Liu, Han Hu, Liwei Wang, and Stephen Lin. Reppoints: Point set representation for object detection. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9657–9666, 2019. 5, 7, 8
- [50] Zhendong Yang, Zhe Li, Xiaohu Jiang, Yuan Gong, Zehuan Yuan, Danpei Zhao, and Chun Yuan. Focal and global knowledge distillation for detectors. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4643–4652, 2022. 2, 4, 5, 6
- [51] Zhendong Yang, Zhe Li, Mingqi Shao, Dachuan Shi, Zehuan Yuan, and Chun Yuan. Masked generative distillation. In *2022 European Conference on Computer Vision*, pages 53–69, 2022. 2, 6
- [52] Hao Zhang, Feng Li, Shilong Liu, Lei Zhang, Hang Su, Jun Zhu, Lionel M Ni, and Heung-Yeung Shum. Dino: Detr with improved denoising anchor boxes for end-to-end object detection. *arXiv preprint arXiv:2203.03605*, 2022. 6
- [53] Linfeng Zhang and Kaisheng Ma. Improve object detection with feature-based knowledge distillation: Towards accurate and efficient detectors. In *2020 International Conference on Learning Representations*, 2020. 2, 4, 5
- [54] Yuan Zhang, Tao Huang, Jiaming Liu, Tao Jiang, Kuan Cheng, and Shanghang Zhang. Freekd: Knowledge distillation via semantic frequency prompt. In *2024 IEEE/CVF*

- Conference on Computer Vision and Pattern Recognition*, pages 15931–15940, 2024. [2](#), [5](#), [6](#)
- [55] Borui Zhao, Quan Cui, Renjie Song, Yiyu Qiu, and Jiajun Liang. Decoupled knowledge distillation. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11953–11962, 2022. [1](#)
- [56] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *2017 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2881–2890, 2017. [1](#), [6](#), [7](#)
- [57] Xinyue Zhou, Mingyu Yin, Xinyuan Chen, Li Sun, Changxin Gao, and Qingli Li. Cross attention based style distribution for controllable person image synthesis. In *2022 European Conference on Computer Vision*, pages 161–178, 2022. [3](#)
- [58] Zikai Zhou, Yunhang Shen, Shitong Shao, Linrui Gong, and Shaohui Lin. Rethinking centered kernel alignment in knowledge distillation. In *2024 International Joint Conference on Artificial Intelligence*, pages 5680–5688, 2024. [5](#)

Supplementary for *CanKD: Cross-Attention-based Non-local operation for Feature-based Knowledge Distillation*

Shizhe Sun, Wataru Ohyama
Tokyo Denki University
Tokyo, Japan

yukarari77@outlook.com, w.ohyama@mail.dendai.ac.jp

1. Details about training strategy

In this paper, we verify the superior performance of CanKD by conducting experiments with various model architectures on both object detection and semantic segmentation tasks. We train all models on NVIDIA RTX 6000 Ada 48GB with 2 GPUs. In the object detection task, following the guidelines provided by the mmdetection[3] documentation and adhering to the linear scaling rule[7], we set the learning rate to 0.005. Meanwhile, following the mmdetection 2× training strategy, we employ a multi-step learning rate decay at epochs 16 and 22 with a decay factor of 0.1. In the semantic segmentation task, following the mmsegmentation[4] official training strategy of an 80K schedule, we set the learning rate to 0.01 with a weight decay of 0.0005 and use a polynomial function to decay our learning rate. The number of parameters and flops in the segmentation task results are all calculated by *thop*.

2. Details about affinity function

In this section, we provide a detailed analysis of two alternative affinity functions that deviate from our original method: the Gaussian method[1, 17] and the Embedding Gaussian method[18].

2.1. Gaussian affinity function

From the perspective of early non-local methods, directly processing two feature maps with a Gaussian function represented a straightforward choice. We omit two 1×1 modules from θ and ϕ , and directly compute the affinity between the student’s feature map and the teacher’s feature map. The Gaussian affinity function is determined by:

$$\xi(\mathbf{x}_i, \mathbf{y}_j) = \exp(\mathbf{x}_i^\top \mathbf{y}_j). \quad (1)$$

where, $\mathbf{x}_i$ and $\mathbf{y}_j$ represent the feature vectors at the position i in the student feature map and the position j in the teacher feature map, respectively. The normalization factor

for the Gaussian affinity function is $C = \sum_j \xi(\mathbf{x}_i, \mathbf{y}_j)$. In terms of implementation, this operation can be easily carried out by applying a *Softmax* layer.

2.2. Embedded Gaussian affinity function

Analogous to self-attention in Transformers, the Embedded Gaussian method introduces an embedding space. Specifically, the student and teacher feature maps are projected into the embedding space via two 1×1 convolutional layers, denoted as θ and ϕ . The affinity is then calculated using a Gaussian function, as shown in equation (2).

$$\xi(\mathbf{x}_i, \mathbf{y}_j) = \exp(\theta(\mathbf{x}_i)^\top \phi(\mathbf{y}_j)). \quad (2)$$

The normalization factor for the embedded Gaussian affinity function is also $C = \sum_j \xi(\mathbf{x}_i, \mathbf{y}_j)$.

3. Natural corrupted augmentation analysis

Following the [13], we evaluate the student RetinaNet-R50 detector, which CanKD trained in the COCO-C dataset. The COCO-C dataset is derived from the COCO validation dataset by applying four types of corruption, i.e., transformations—noise, blurring, weather, and digital corruption, to evaluate the model’s robustness. Each corruption category consists of multiple corruption methods, each with six levels of severity. The test results are summarized in Table 1. The results show that CanKD exhibits substantially higher robustness than the other methods. Specifically, it achieves a 2.2 improvement in mPC and a 2.7 improvement in rPC compared to the benchmark.

4. Analysis about maxpooling layer

In this section, we examine whether different scales of the teacher feature maps after max pooling influence the performance of the CanKD. We experimented with max pooling at 4×4 and 8×8 scales. It is important to note that while reducing the pooling scale can facilitate the student

Table 1. **Result of robust object detection via CanKD on COCO-C dataset.**

Method	AP_{clean}	mPC	rPC
RetinaNet-R50	37.4	18.3	48.9
FGD	39.6	20.3	51.3
DiffKD	39.7	20.3	51.1
CanKD	39.8	20.5	51.6

Table 2. **Ablation study on maxpooling scales.** We use RepPoints-X101[19] as teacher and RepPoints-R50 as student.

Scaled	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L
2×2	42.4	62.9	45.6	24.1	46.5	56.4
4×4	42.0	62.3	45.5	24.0	46.2	55.0
8×8	42.1	62.3	45.6	23.9	46.7	55.7

Table 3. **Ablation study on residual connection.** We use RepPoints-X101[19] as teacher and RepPoints-R50 as student.

Scaled	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L
W/O residual connection	41.4	61.8	44.6	24.4	45.6	54.3
W/ residual connection	42.4	62.9	45.6	24.1	46.5	56.4

model’s more comprehensive acquisition of the teacher feature maps, it simultaneously increases the memory footprint on the hardware. The experimental results are presented in Table.2.

From the results, we observe that employing a large-scale max pooling operation eliminates certain crucial information within the teacher feature maps. This reduction hampers the student feature maps from fully assimilating the teacher’s knowledge. Meanwhile, we suggest that employing a max pooling layer to select critical information from the teacher feature maps may not be the optimal choice. A more refined strategy could better preserve essential information while reducing the spatial resolution of the teacher feature maps.

5. Analysis about residual connection

In this section, we examine the role of the residual connection within the Can module, which is in function 5 in the original paper. We believe that directly comparing the feature map produced by the attention component of the Can module with the teacher’s original feature map would cause issues because we do not apply any additional operations to the teacher’s feature map. Therefore, we retain the residual connection so that the student’s original feature map remains involved. The ablation study is in table.3.

According to the experimental results, CanKD with a

residual connection outperforms CanKD without a residual connection by a significant margin. Therefore, we conclude that the residual connection is indispensable in CanKD.

6. Confusion matrix

Here, we present the confusion matrices for the teacher models, the student models, and our proposed CanKD method. Here, Figure 1 is from the FasterRCNN-R50 [14] student model, and Figure 2 is from the FasterRCNN-R50 model distilled with CanKD.

7. Detail for training

By analyzing the output logs from each model, we generated the changes in mAP for each model throughout the training process. The figure is shown in Figure.3.

8. Visualization

In Figure 5, we present a series of visualization figures generated from the student model, CanKD, and the teacher model on the COCO val 2017 dataset[12] to demonstrate the performance and effectiveness of our method in object detection. Meanwhile, in Figure 4, we also present a series of visualization figures generated from the student model, CanKD, and the teacher model on the Cityscapes validation dataset[5]. In Figure 6, we present the heatmaps from different FPN layers for the student model, the distilled student model, and the teacher model. Compared to the original student model, these examples demonstrate that our method has better performance and effectiveness.

9. Experiments on classification task.

To complement the evaluation of CanKD on fundamental tasks, following the official schedule, we conducted experiments on the ImageNet-1K [6] dataset with two teacher–student model pairs, ResNet-34→ResNet-18, ResNet-50→MobileNetv1 in 100e. We choose the last layer output in backbone as our distillation position. The results are shown in table.4. The balanced weight μ in CanKD are all set to 5. However, CanKD is specifically designed for dense prediction tasks, where pixel-level feature alignment is critical. In contrast, classification tasks do not involve pixel-level alignment; therefore, the use of cross-attention may be unnecessary compared to traditional logit distillation methods that already achieve strong results. Meanwhile, CanKD outperforms w/o student, KD [9], and AT [23]. This demonstrates that the performance gains of CanKD in dense prediction do not come at the expense of classification performance. The classification heat map images are shown in Figure 7 which generated by *Grad-CAM* [15].




Figure 1. Confusion matrix from FasterRCNN-R50.

Table 4. **Classification task on ImageNet-1K.** We use ResNet-34, ResNet-50 [8] as teachers and Resnet-18, MobileNet-v1 [10] students. All baselines results are from [21].

Method	<i>Top1</i>	<i>Top5</i>	Method	<i>Top1</i>	<i>Top5</i>
T:ResNet-34	73.62	91.59	T:ResNet-50	76.55	93.06
T:ResNet-18	69.90	89.43	S:MobileNetv1	69.21	89.02
KD [9]	70.68	90.16	KD [9]	70.68	90.30
AT [23]	70.59	89.73	AT [23]	70.72	90.03
CanKD	70.73	89.81	CanKD	70.89	89.90

10. Analysis about parameter and computational complexity

In this section, we analyze the parameter count and computational complexity of CanKD. We compare its param-

ters and FLOPs with several strong feature distillation methods in table.5, and our results show that CanKD has fewer parameters and lower complexity than existing attention-based and mask-based distillation methods. Moreover, since CanKD is applied only during training, it does not affect the parameter count or complexity of the student model at inference time, further demonstrating that CanKD is a lightweight yet effective attention-based distillation approach.

11. Limitation of CanKD

Attention-based distillation methods share a common drawback: when the teacher and student feature maps differ in spatial resolution ($H \times W$), performance degradation occurs. We demonstrate this phenomenon using a simple teacher–student pair in Table.6. Furthermore, from the dis-




Figure 2. Confusion matrix from FasterRCNN-R50 distilled with CanKD.

Table 5. **Analysis on parameter and FLOPs.** We use *thop* library to calculate params and FLOPs in same student and teacher’s input with $2 \times 256 \times 64 \times 64$. Because FGD [20] including bounding box distillation, it’s hard to clarify the complexity.

Method	KD method	FLOPs(G)	Params(M)
FKD [24]	Attention	3.23	0.46
MGD [21]	Mask	9.66	1.18
FGD [20]	Attention	-	0.14
CanKD	Attention	1.09	0.13

tillation results of Dino-Swin-L and Dino-Swin-B, CanKD shows only marginal improvement for this pair, which previous research attribute to the substantial architectural gap between Transformer-based backbones, where a sim-

Table 6. The experiment result with mismatch $H \times W$

Method	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L
T:MaskRCNN-SwinS	48.2	69.8	52.8	32.1	51.8	62.7
S:Retina-R50	37.4	56.7	39.6	20.0	40.7	49.7
PKD [2]	41.5	60.6	44.1	22.9	45.2	56.4
DetKDS [11]	41.4	60.8	44.4	23.4	45.1	55.5
CanKD	40.5	59.0	43.3	22.5	44.7	54.2

ple CNN-based approach cannot be effectively applied [22]. The observation that CWD [16] performs even worse than the baseline further corroborates this finding.

References

- [1] Antoni Buades, Bartomeu Coll, and J-M Morel. A non-local algorithm for image denoising. In *2005 IEEE computer society conference on computer vision and pattern recognition*




Figure 3. Line chart about all student models, distilled student models, and teacher models.




Figure 4. Additional sampling from PSPNet-R18[25], distilled PSPNet-R18 with CanKD and PSPNet-R101. All of these figures are selected from Cityscapes val dataset.

(CVPR'05), pages 60–65. Ieee, 2005. 1

- [2] Weihang Cao, Yifan Zhang, Jianfei Gao, Anda Cheng, Ke Cheng, and Jian Cheng. Pkd: General distillation framework for object detectors via pearson correlation coefficient. In *2022 Advances in Neural Information Processing Systems*, pages 15394–15406, 2022. 4
- [3] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, Zheng Zhang, Dazhi Cheng, Chenchen Zhu, Tianheng Cheng, Qijie Zhao, Buyu Li, Xin Lu, Rui Zhu, Yue Wu, Jifeng Dai, Jingdong Wang, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. MMDetection: Open

mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019. 1

- [4] MMSegmentation Contributors. MMSegmentation: Openmmlab semantic segmentation toolbox and benchmark. <https://github.com/open-mmlab/mms Segmentation>, 2020. 1
- [5] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. 2




Figure 5. **Additional sampling from RepPoints-R50, distilled RepPoints-R50 with CanKD and RepPoints-X101.** All of these figures are selected from COCO val 2017 dataset.

- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 2
- [7] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large minibatch sgd: Training imagenet in 1 hour, 2018. 1
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 3
- [9] Geoffrey Hinton. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 2, 3
- [10] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. 3
- [11] Lujun Li, Yufan Bao, Peijie Dong, Chuanguang Yang, Anggeng Li, Wenhan Luo, Qifeng Liu, Wei Xue, and Yike Guo. Detkds: Knowledge distillation search for object detectors. In *2024 International Conference on Machine Learning*, 2024. 4
- [12] 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 *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. 2




Figure 6. **Heatmaps from FasterRCNN-R50, distilled FasterRCNN-R50 with CanKD and FasterRCNN-R101.** These figures are generated from P3 to P5 in FPN layers




Figure 7. **Heatmaps from ResNet-18, distilled ResNet-18 with CanKD and ResNet-32.** These figures are generated from Layer.1 to Layer.4 in backbone.

[13] Claudio Michaelis, Benjamin Mitzkus, Robert Geirhos, Evgenia Rusak, Oliver Bringmann, Alexander S. Ecker, Matthias Bethge, and Wieland Brendel. Benchmarking robustness in object detection: Autonomous driving when win-

ter is coming. *arXiv:1907.07484*, 2019. 1

[14] Shaoqing Ren. Faster r-cnn: Towards real-time object detection with region proposal networks. *arXiv preprint arXiv:1506.01497*, 2015. 2

[15] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017. 2

[16] Changyong Shu, Yifan Liu, Jianfei Gao, Zheng Yan, and Chunhua Shen. Channel-wise knowledge distillation for dense prediction. In *2021 IEEE/CVF International Conference on Computer Vision*, pages 5311–5320, 2021. 4

[17] Carlo Tomasi and Roberto Manduchi. Bilateral filtering for gray and color images. In *Sixth international conference on computer vision (IEEE Cat. No. 98CH36271)*, pages 839–846. IEEE, 1998. 1

[18] A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017. 1

[19] Ze Yang, Shaohui Liu, Han Hu, Liwei Wang, and Stephen Lin. Reppoints: Point set representation for object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9657–9666, 2019. 2

[20] Zhendong Yang, Zhe Li, Xiaohu Jiang, Yuan Gong, Zehuan Yuan, Danpei Zhao, and Chun Yuan. Focal and global knowledge distillation for detectors. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4643–4652, 2022. 4

[21] Zhendong Yang, Zhe Li, Mingqi Shao, Dachuan Shi, Zehuan

- Yuan, and Chun Yuan. Masked generative distillation. In *2022 European Conference on Computer Vision*, pages 53–69, 2022. [3](#), [4](#)
- [22] Zhendong Yang, Zhe Li, Ailing Zeng, Zexian Li, Chun Yuan, and Yu Li. Vitkd: Feature-based knowledge distillation for vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1379–1388, 2024. [4](#)
- [23] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. *arXiv preprint arXiv:1612.03928*, 2016. [2](#), [3](#)
- [24] Linfeng Zhang and Kaisheng Ma. Improve object detection with feature-based knowledge distillation: Towards accurate and efficient detectors. In *2020 International Conference on Learning Representations*, 2020. [4](#)
- [25] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017. [5](#)