

ACAM-KD: Adaptive and Cooperative Attention Masking for Knowledge Distillation

Qizhen Lan Qing Tian*

Dept. of Computer Science, University of Alabama at Birmingham, USA

{qlan, qtian}@uab.edu

Abstract

Dense visual prediction tasks, such as detection and segmentation, are crucial for time-critical applications (e.g., autonomous driving and video surveillance). While deep models achieve strong performance, their efficiency remains a challenge. Knowledge distillation (KD) is an effective model compression technique, but existing feature-based KD methods rely on static, teacher-driven feature selection, failing to adapt to the student’s evolving learning state or leverage dynamic student-teacher interactions. To address these limitations, we propose Adaptive student-teacher Cooperative Attention Masking for Knowledge Distillation (ACAM-KD), which introduces two key components: (1) Student-Teacher Cross-Attention Feature Fusion (STCAFF), which adaptively integrates features from both models for a more interactive distillation process, and (2) Adaptive Spatial-Channel Masking (ASCM), which dynamically generates importance masks to enhance both spatial and channel-wise feature selection. Unlike conventional KD methods, ACAM-KD adapts to the student’s evolving needs throughout the entire distillation process. Extensive experiments on multiple benchmarks validate its effectiveness. For instance, on COCO2017, ACAM-KD improves object detection performance by up to 1.4 mAP over the state-of-the-art when distilling a ResNet-50 student from a ResNet-101 teacher. For semantic segmentation on Cityscapes, it boosts mIoU by 3.09 over the baseline with DeepLabV3-MobileNetV2 as the student model.

1. Introduction

Deep learning models have significantly advanced dense visual prediction tasks, such as object detection and segmentation, leading to breakthroughs in applications like autonomous driving, surveillance, and robotics. Despite these successes, deploying high-performance deep models on resource-limited devices remains challenging due to

their high computational complexity and large parameter sizes. Knowledge distillation (KD) has emerged as an effective solution, transferring knowledge from a high-capacity teacher model to a lightweight student model to maintain accuracy while reducing complexity. Among various KD approaches, feature-based knowledge distillation is particularly well-suited for dense prediction tasks like object detection and segmentation. By transferring intermediate feature representations instead of final predictions, it effectively preserves the dense spatial information essential for these tasks.

However, most existing feature-based KD methods rely on fixed or teacher-driven feature selection. They assume that the most critical regions for distillation can be determined solely by the teacher [9, 28, 29] or predefined heuristics, such as bounding boxes, RPN regions, prediction confidence [11, 23, 26]. These methods require the student to passively focus on predetermined regions, disregarding its evolving learning state and unique characteristics. However, the features deemed important by the pre-trained teacher or human heuristics may not align with what the student actually needs at different learning stages. Moreover, in these approaches, the focused areas for distillation remain static for a given image, failing to adapt to the student’s dynamically changing knowledge (i.e., feature representations). Repeatedly emphasizing the same regions in every epoch is inefficient, especially when the student has already acquired sufficient understanding of them. To make things worse, an inaccurate teacher focus likely misleads the student. Fig. 1 provides an example. The student initially exhibits inaccurate attention (Epoch 4) but improves by Epoch 12, focusing better on the foreground plane. However, by Epoch 24, the student’s attention shifts to resemble the teacher’s fixed feature selection instead of further refining its own learning. In addition, while existing methods primarily focus on spatial feature selection—an essential aspect of dense prediction tasks—they often overlook the significance of channel-wise feature selection and fail to consider the varying contributions of different channels.

To address the above challenges in feature-based knowl-

*Corresponding author.

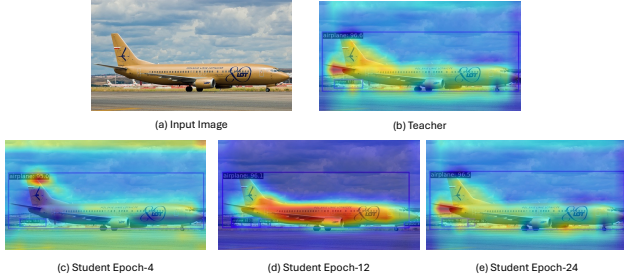


Figure 1. Visualization of attention maps from different models at various training stages. (a) input image, (b) teacher model’s attention, and (c-e) student model’s attention at epochs 4, 12, and 24. The color represents attention intensity, with red indicating the highest focus and blue the lowest. The teacher model’s attention is static and suboptimal. The student’s attention evolves over time, with epoch 12 demonstrating better localization than the teacher.

edge distillation, we propose Adaptive student-teacher Cooperative Attention Masking for enhanced feature-based Knowledge Distillation (ACAM-KD). Specifically, ACAM-KD contains two key modules: (1) Student-Teacher Cross-Attention Feature Fusion (STCA-FF) module, which integrates features from both the pre-trained teacher and the dynamically evolving student models. (2) Building on the cross-attention features from STCA-FF, the second module, Adaptive Spatial-Channel Masking (ASCM), adaptively generates importance masks across both spatial and channel dimensions for more effective distillation. Unlike traditional methods that rely solely on fixed or teacher-defined attention, ACAM-KD leverages cross-attention to facilitate a cooperative knowledge transfer process, allowing the student to dynamically interact with the teacher and refine feature selection based on both teacher-provided knowledge and the student’s own evolving learning state (i.e., feature representations). The attention masks for distillation are dynamically updated throughout the student’s learning process, enabling the distillation to adaptively focus on different tensor regions as the student evolves and its learning needs change. We validate the effectiveness of our approach through extensive experiments on various dense prediction tasks, including object detection and segmentation. Our method consistently outperforms state-of-the-art KD approaches across multiple benchmarks, demonstrating superior performance and adaptability in student models distilled with our framework.

2. Related Work

2.1. Knowledge Distillation for Object Detection

A key issue in KD for object detection is to identify which spatial regions and feature representations are most important. FGF [21] introduced pixel-level masks to highlight regions near ground-truth boxes, while GID [6] empha-

sized areas where teacher and student predictions differ significantly. FKD [28] leveraged high-attention features to weigh important regions, and FGD [26] further advanced this by also incorporating ground-truth bounding boxes and capturing global context. Recent works explore learnable feature selection strategies. MasKD [9] proposed learnable receptive tokens, derived from the teacher’s features, to generate pixel-wise masks that capture fine-grained spatial knowledge. FreeKD [29] introduced a frequency-domain semantic prompt to guide the student in focusing on essential feature regions, improving distillation robustness in dense prediction tasks. Despite these advancements, existing KD methods for object detection remain predominantly heuristics-based or teacher-driven, assuming static feature importance throughout training. Although [9] and [29] attempt to learn tokens or prompts for masking, their learning process is conducted offline in one shot, relying solely on a fixed pre-trained teacher. Consequently, the learned tokens or prompts remain static throughout the student’s learning process, failing to account for the student’s evolving state and its dynamic interaction with the teacher model. As a result, for a given input, the student consistently receives the same knowledge across epochs, regardless of its progress. Also, the dynamic interplay between spatial and channel-wise feature selection is often neglected.

2.2. Knowledge Distillation for Semantic Segmentation

Effective knowledge distillation for semantic segmentation is challenging due to high-resolution outputs and complex pixel dependencies. Early KD methods primarily addressed this by enforcing local spatial consistency, aligning teacher and student features at the pixel level. SKD [15] introduced pairwise and holistic distillation strategies, with the former preserving local pixel relationships through pairwise similarity alignment and the latter enforcing high-order structural consistency via adversarial training on score maps. IFVD [22] enhanced knowledge transfer by enforcing intra-class feature consistency, improving the stability of learned representations across spatial regions. CWD [19] introduced a channel-wise feature alignment strategy to help the student model learn the teacher’s feature distribution across channels. CIRKD [24] further advanced relational distillation by modeling both intra-image and cross-image dependencies, promoting structured feature learning beyond direct pixel-wise supervision.

Although some progress has been made, existing KD approaches for segmentation share similar limitations with KD methods for detection, relying on heuristics or teacher-driven, fixed feature importance that fails to adapt to the student’s evolving state. Additionally, they lack a joint optimization of spatial and channel-wise knowledge selection. To address these shortcomings, we propose Adap-

tive student-teacher Cooperative Attention Masking for enhanced feature-based Knowledge Distillation (ACAM-KD). The following section details our methodology.

3. Methodology

Knowledge distillation (KD) transfers knowledge from a high-capacity teacher model to a lightweight student model, aiming to improve efficiency while maintaining performance. Given a teacher model T and a student model S , feature-based KD minimizes the distillation loss $\mathcal{L}_{\text{feat}}$:

$$\frac{1}{M} \sum_{m=1}^M \frac{1}{C \sum_{p=1}^{H \times W} \mathbf{M}_{m,p}} \left\| \mathbf{M}_m \odot (F^T - f_{\text{align}}(F^S)) \right\|_2^2, \quad (1)$$

where F^T and F^S are the teacher’s and student’s feature maps, respectively. The function f_{align} is an adaptation layer applied to F^S to align its feature dimension with F^T . The mask \mathbf{M}_m is applied element-wise to emphasize relevant regions for distillation, where $\mathbf{M}_{m,p}$ represents the selection weight at spatial position p of the m -th mask. C , H , and W denote the number of channels, feature map height, and width, respectively. The loss is averaged over M masks.

Existing feature-based KD methods rely on fixed, teacher-driven mask selection, assuming that key regions can be identified solely from teacher activations or predefined heuristics. This approach disregards the student’s role and evolving knowledge, failing to adapt to its dynamically changing needs at different learning stages. Additionally, maintaining static focus regions throughout distillation is inefficient or even dangerous, and most methods emphasize spatial selection while neglecting channel-wise importance. To address these limitations, we propose Adaptive student-teacher Cooperative Attention Masking for Knowledge Distillation (ACAM-KD), which dynamically refines feature selection through student-teacher interactions. Instead of relying on predefined masks, our approach learns masks adaptively and continuously throughout the KD process. Fig. 2 presents an overview of ACAM-KD, while Sec. 3.1 and 3.2 detail its two main components.

3.1. Student-Teacher Cross-Attention Feature Fusion for Distillation Attention

In knowledge distillation, the student should not only passively absorb knowledge from the teacher but also identify the most beneficial learning regions based on its evolving needs. To enable this, we first introduce a student-teacher cross-attention-based feature fusion mechanism (STCA-FF), where the teacher’s feature map defines the query, while the student’s feature map provides both the key and value. Given teacher and student feature maps $F^T, F^S \in \mathbb{R}^{C \times H \times W}$, we apply 1×1 convolutions to project them

into query, key, and value representations:

$$Q = W_q F^T, \quad K = W_k F^S, \quad V = W_v F^S, \quad (2)$$

where $W_q \in \mathbb{R}^{C_q \times C}$ and $W_k \in \mathbb{R}^{C_q \times C}$ reduce the channel dimension to $C_q = C/2$, while $W_v \in \mathbb{R}^{C \times C}$ maintains the original channel dimension. To compute attention over spatial locations, the feature maps are flattened along the spatial dimensions (H, W) , yielding $Q \in \mathbb{R}^{HW \times C_q}$, $K \in \mathbb{R}^{C_q \times HW}$, and $V \in \mathbb{R}^{HW \times C}$. The attention matrix is then computed as:

$$A = \text{softmax} \left(\frac{QK}{\sqrt{C_q}} \right) \in \mathbb{R}^{HW \times HW}. \quad (3)$$

Using the attention weights, we compute the fused features:

$$F_{\text{fused}} = AV \in \mathbb{R}^{C \times H \times W}. \quad (4)$$

The fused features integrate information from both the teacher and student, enabling the distillation process to leverage their joint interactions rather than relying solely on predefined teacher activations. This fused representation is then fed into the next adaptive spatial-channel masking module, where learnable selection units generate spatial and channel-wise masks to prioritize regions for distillation, ensuring knowledge transfer is dynamically adjusted based on both teacher and student features.

3.2. Adaptive Spatial-Channel Masking

In contrast to frozen teacher features, the student-teacher interacted features dynamically evolve throughout the student’s learning process. The distillation mask applied to these fused features should also adapt accordingly. However, the fixed attention masks commonly used in existing methods fail to accommodate this need. To overcome this limitation, we introduce a learnable selection mechanism that dynamically generates feature importance masks during distillation, enabling the student model to selectively focus on the most beneficial spatial and channel-wise regions at a certain learning stage.

Unlike existing spatial-only masking methods, we introduce two complementary sets of learnable selection units for distillation masking: channel-wise selection units $m^c \in \mathbb{R}^{M \times 1}$ and spatial selection units $m^s \in \mathbb{R}^{M \times C}$. Both are learned and updated throughout the distillation. These learnable selection units operate on the dynamically evolving student-teacher fused features, generating adaptive channel-wise and spatial masks to enhance distillation. Given the fused feature representation $F_{\text{fused}} \in \mathbb{R}^{C \times H \times W}$, the selection units generate channel and spatial masks as follows:

$$\mathbf{M}^c = \sigma(m^c \cdot v), \quad \mathbf{M}^s = \sigma(m^s \cdot z), \quad (5)$$

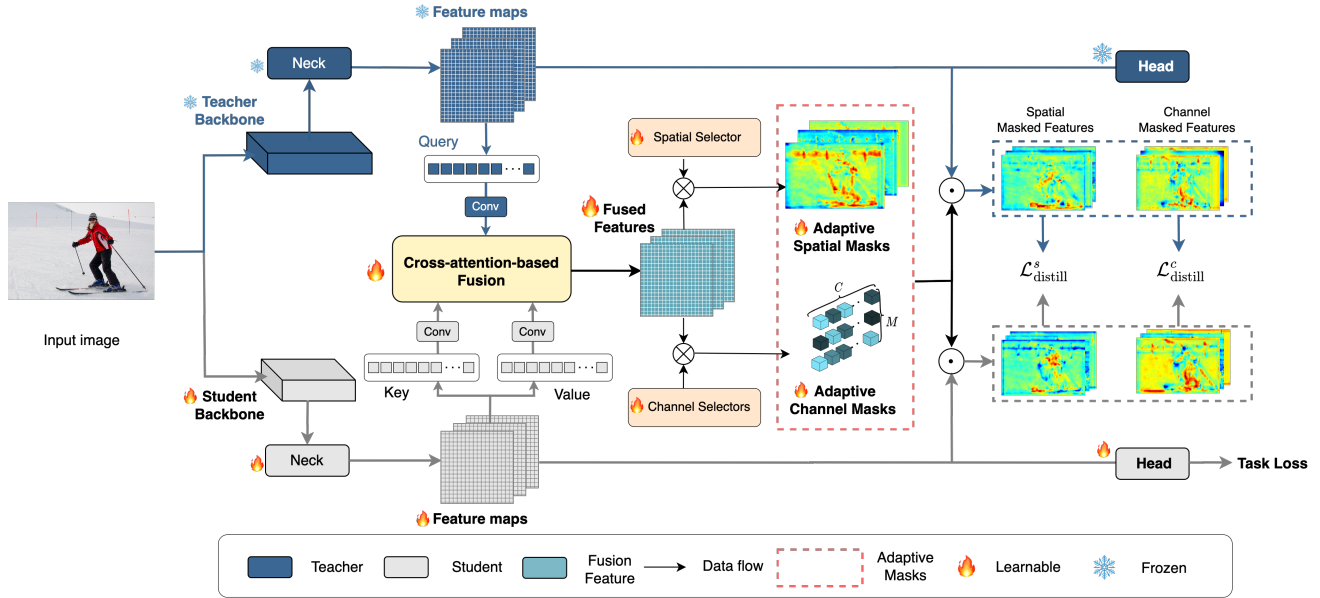


Figure 2. Overview of the proposed Adaptive and Cooperative Attention Masking for Knowledge Distillation (ACAM-KD) framework. ACAM-KD leverages student-teacher cross-attention to dynamically and cooperatively guide the distillation process. Instead of relying on static, teacher-determined attention masking, our distillation masks are updated throughout distillation, adapting to the student’s evolving learning needs and selectively focusing on the most beneficial spatial and channel-wise features at a particular learning stage. Different colors in the adaptive channel masks indicate varying levels of importance.

where $\sigma(\cdot)$ is the sigmoid function, $v \in \mathbb{R}^{1 \times C}$ is the spatially average-pooled vector of F_{fused} , and $z \in \mathbb{R}^{C \times HW}$ is the flattened F_{fused} along the spatial dimension. With the masks, we define the channel-wise and spatial feature distillation losses $\mathcal{L}_{\text{distill}}^c$ and $\mathcal{L}_{\text{distill}}^s$ as:

$$\mathcal{L}_{\text{distill}}^c = \frac{1}{M} \sum_{m=1}^M \frac{1}{HW \sum_{k=1}^C \mathbf{M}_{m,k}^c} \left\| \mathbf{M}_m^c \odot (F^T - f_{\text{align}}(F^S)) \right\|_2^2, \quad (6)$$

$$\mathcal{L}_{\text{distill}}^s = \frac{1}{M} \sum_{m=1}^M \frac{1}{C \sum_{p=1}^{H \times W} \mathbf{M}_{m,p}^s} \left\| \mathbf{M}_m^s \odot (F^T - f_{\text{align}}(F^S)) \right\|_2^2, \quad (7)$$

where k indicates the index of channels. By optimizing the two losses, the adaptive masks $\mathbf{M}^c \in \mathbb{R}^{M \times C}$ and $\mathbf{M}^s \in \mathbb{R}^{M \times HW}$ dynamically adjust to emphasize channels and spatial locations that are important from the teacher’s perspective while adapting to the student’s evolving needs at each distillation stage. This dual awareness originates from F_{fused} , which encapsulates the collaborative interaction between teacher and student. The dynamic behavior is driven by both the feature fusion and the subsequent adaptive spatial-channel masking.

Mask Diversity To prevent trivial solutions where all masks converge to similar patterns or degenerate into zero values, we incorporate a Dice coefficient-based diversity

loss that promotes variation by penalizing excessive similarity among masks. Given a set of learnable masks $\{M_i\}_{i=1}^M$, we define the diversity loss as:

$$\mathcal{L}_{\text{div}} = \frac{2 \sum_{i=1}^M \sum_{j=1, j \neq i}^M M_i \cdot M_j}{\sum_{i=1}^M M_i^2 + \sum_{j=1}^M M_j^2}, \quad (8)$$

where the numerator measures the pairwise similarity between different masks, while the denominator normalizes it by the sum of the squared magnitudes of the masks.

3.3. Overall Loss

The total training objective is the sum of the following losses:

$$\mathcal{L} = \mathcal{L}_{\text{task}} + \alpha(\mathcal{L}_{\text{distill}}^c + \mathcal{L}_{\text{distill}}^s) + \lambda \mathcal{L}_{\text{div}}, \quad (9)$$

where $\mathcal{L}_{\text{task}}$ is the task (e.g., detection or segmentation) loss, $\mathcal{L}_{\text{distill}}^s$ and $\mathcal{L}_{\text{distill}}^c$ are the spatial and channel-wise distillation losses, and \mathcal{L}_{div} is the mask diversity loss. α and λ are balancing hyperparameters. In our experiments, we set $\alpha = 1$ and $\lambda = 1$, which already achieves strong performance.

4. Experiments and Results

To demonstrate the effectiveness of our ACAM-KD, we conduct extensive experiments across a variety of model architectures on both object detection and semantic segmentation tasks.

Method	mAP	AP ₅₀	AP ₇₅	AP _s	AP _m	AP _l
<i>Single-Stage Detectors (RetinaNet)</i>						
T: R101	38.9	58.0	41.5	21.0	42.8	52.4
S: R50	37.4	56.7	39.6	20.0	40.7	49.7
FitNet [17]	37.4	57.1	40.0	20.8	40.8	50.9
GID [6]	39.1	59.0	42.3	22.8	43.1	52.3
FRS [7]	39.3	58.8	42.0	21.5	43.3	52.6
FGD [26]	39.6	-	-	22.9	43.7	53.6
MasKD [9]	39.8	59.0	42.5	21.5	43.9	54.0
FreeKD [29]	39.9	-	-	21.2	44.0	53.7
CrossKD [20]	39.7	58.9	42.5	22.4	43.6	52.8
Ours	41.2	60.6	44.1	24.6	45.5	54.1
<i>Two-Stage Detectors (Faster R-CNN)</i>						
T: R101	39.8	60.1	43.3	22.5	43.6	52.8
S: R50	38.4	59.0	42.0	21.5	42.1	50.3
FitNet [17]	38.9	59.5	42.4	21.9	42.2	51.6
GID [6]	40.2	60.7	43.8	22.7	44.0	53.2
FRS [7]	39.5	60.1	43.3	22.3	43.6	51.7
FGD [26]	40.4	-	-	22.8	44.5	53.5
MasKD [9]	40.8	60.7	44.4	23.2	44.6	53.6
FreeKD [29]	40.8	-	-	23.1	44.7	54.0
Ours	41.4	61.7	45.2	24.7	45.0	53.3
<i>Anchor-Free Detectors (RepPoints)</i>						
T: R101	40.5	61.3	43.5	23.4	44.7	53.2
S: R50	38.6	59.6	41.6	22.5	42.2	50.4
FitNet [17]	40.7	61.6	43.9	23.7	44.4	53.3
FGD [26]	41.0	-	-	23.8	45.3	53.6
MasKD [9]	41.1	61.4	44.5	24.1	45.4	53.8
Ours	42.5	63.5	46.0	26.1	46.8	54.2

Table 1. KD performance with ResNet-backed teachers on COCO validation set. T: teacher model, S: student baseline.

4.1. Object Detection

4.1.1. Experimental Setup

We conduct experiments on the widely used Microsoft COCO2017 detection dataset [12], which includes 118,287 training images and 5,000 validation images from 80 object classes. We evaluate ACAM-KD on various detectors, including one-stage [13], two-stage [16], and anchor-free detectors [25]. We perform feature distillation at the FPN neck, implementing our approach within the MMDetection framework [2] and following its official training strategies. All models are optimized using SGD with a momentum of 0.9 and a weight decay of 0.0001. In all experiments, we set M to 6 for both spatial and channel cases. Additionally, we adopt the inheritance strategy from previous knowledge distillation works [10, 27] to stabilize early-stage training.

4.1.2. Experimental Results

Tab. 1 and Tab. 2 present object detection performance on COCO under different teacher configurations. With a ResNet-50 (R50) student, our method consistently sur-

Method	mAP	AP ₅₀	AP ₇₅	AP _s	AP _m	AP _l
<i>Single-Stage Detectors (RetinaNet)</i>						
T: X101	41.2	62.1	45.1	24.0	45.5	53.5
S: R50	37.4	56.7	39.6	20.0	40.7	49.7
COFD [8]	37.8	58.3	41.1	21.6	41.2	48.3
FKD [28]	39.6	58.8	42.1	22.7	43.3	52.5
FRS [7]	40.1	59.5	42.5	21.9	43.7	54.3
FGD [26]	40.4	-	-	23.4	44.7	54.1
MasKD [9]	40.9	60.1	43.6	22.8	45.3	55.1
FreeKD [29]	41.0	-	-	22.3	45.1	55.7
Ours	41.5	61.3	44.5	25.3	45.5	54.3
<i>Two-Stage (S: Faster R-CNN, T: Cascade Mask R-CNN)</i>						
T: X101	45.6	64.1	49.7	26.2	49.6	60.0
S: R50	38.4	59.0	42.0	21.5	42.1	50.3
LED [1]	38.7	59.0	42.1	22.0	41.9	51.0
FGFI [21]	39.1	59.8	42.8	22.2	42.9	51.1
COFD [8]	38.9	60.1	42.6	21.8	42.7	50.7
FKD [28]	41.5	62.2	45.1	23.5	45.0	55.3
FGD [26]	42.0	-	-	23.7	46.4	55.5
MasKD [9]	42.4	62.9	46.8	24.2	46.7	55.9
FreeKD [29]	42.4	-	-	24.1	46.7	55.9
Ours	42.6	63.0	46.6	25.2	46.8	55.6
<i>Anchor-Free Detectors (RepPoints)</i>						
T: X101	44.2	65.5	47.8	26.2	48.4	58.5
S: R50	38.6	59.6	41.6	22.5	42.2	50.4
FKD [28]	40.6	61.7	43.8	23.4	44.6	53.0
FGD [26]	41.3	-	-	24.5	45.2	54.0
MasKD [9]	41.8	62.6	45.1	24.2	45.4	55.2
FreeKD [29]	42.4	-	-	24.3	46.4	56.6
Ours	42.8	63.9	46.4	26.7	47.1	54.5

Table 2. KD performance with ResNeXt-backed teachers on COCO validation set. T: teacher model, S: student baseline.

passes previous distillation approaches across single-stage (RetinaNet [13]), two-stage (Faster R-CNN [16]), and anchor-free (RepPoints [25]) detectors. When distilled from a ResNet-101 (R101) teacher, our method improves the student’s mAP by 1.3 for single-stage, 0.6 for two-stage, and 1.4 for anchor-free detectors over the second-best KD method. With a ResNeXt-101 (X101) teacher, we achieve even greater gains, boosting the student’s mAP by 4.1 for single-stage, 4.2 for two-stage, and 4.2 for anchor-free detectors. These results demonstrate the effectiveness of our method in deriving efficient student detectors that outperform the baseline and state-of-the-art competitors.

4.2. Segmentation

4.2.1. Experimental Setup

We validate our method using the Cityscapes dataset [5], which contains a total of 5,000 finely annotated images and 19,998 coarsely annotated images, captured from 50

cities in different weather conditions and seasons. Like most previous works in this area, we only use the finely annotated images. We evaluate all segmentors using mean Intersection-over-Union (mIoU). In all experiments, we adopt the DeepLabV3 framework [3] with a ResNet-101 backbone as the teacher network. For student models, we explore various configurations, combining frameworks such as DeepLabV3 and PSPNet [30] with backbones like ResNet-18 and MobileNetV2 [18].

For the semantic segmentation task, we distill features from the predicted segmentation maps. The models are trained using the MMSegmentation [4] framework with a 40K-iteration schedule and an input resolution of 512×1024. Optimization is performed using SGD with a weight decay of 5e-4, while a polynomial annealing scheduler is applied to adjust the learning rate, starting at 0.02. To ensure that each learnable mask corresponds to a specific semantic category, we set M to 19, matching the number of classes in the Cityscapes dataset.

4.2.2. Experimental Results

Tab. 3, Tab. 4, and Tab. 5 report the semantic segmentation results across different student architectures evaluated on the Cityscapes dataset. FLOPs are calculated based on an input size of 512×1024 . As shown, our method consistently outperforms previous knowledge distillation approaches when distilling from a DeepLabV3-R101 teacher to various lightweight student models, including DeepLabV3 with ResNet-18 and MobileNetV2 [18] backbones. Furthermore, we investigate heterogeneous distillation by transferring knowledge from DeepLabV3-R101 to PSPNet [30] with a ResNet-18 backbone.

For DeepLabV3-MBV2 as the student, our method improves mIoU by 3.09 over the baseline and 0.79 over the best competing KD method. With DeepLabV3-R18, our approach achieves 77.53 mIoU, surpassing the strongest competitor by 0.53 mIoU. Similarly, for PSPNet-R18, our method boosts mIoU by 3.44 over the baseline and 0.65 over the best prior method. These results underscore the effectiveness and generalizability of our approach across different segmentation models, consistently improving student performance in low-computation settings.

4.3. Runtime Efficiency Analysis

In addition to the number of parameters (Params) and FLOPs, Tab. 6 compares the runtime efficiency of different object detectors in terms of memory footprint and inference speed (frames per second, FPS). The large two-stage teacher model with X101 as its backbone incurs the highest computational cost, requiring more than 2000G FLOPs and consuming significant CUDA memory, making it less suitable for resource-constrained applications. In contrast, the student models with R50 backbones have much

Method	Params FLOPs mIoU		
	(M)	(G)	(%)
T: DeepLabV3-R101	84.74	695	78.07
S: DeepLabV3-R18			72.96
SKD [15] CVPR'19			75.42
IFVD [22] ECCV'20			75.59
CWD [19] ICCV'21			75.55
CIRKD [24] CVPR'22	13.84	120	76.38
LAD [14] WACV'24			76.78
MasKD [9] ICLR'23			77.00
FreeKD [29] CVPR'24			76.45
MGD [27] ECCV'22			76.02
Ours			77.53

Table 3. Comparison of knowledge distillation methods for semantic segmentation on Cityscapes with the DeepLabV3-R18 student model. T: teacher model, S: student baseline model.

Method	Params FLOPs mIoU		
	(M)	(G)	(%)
T: DeepLabV3-R101	84.74	695	78.07
S: DeepLabV3-MBV2			73.12
SKD [15] CVPR'19			73.82
IFVD [22] ECCV'20	3.2	31.14	73.50
CWD [19] ICCV'21			74.66
CIRKD [24] CVPR'22			75.42
MasKD [9] ICLR'23			75.26
Ours			76.21

Table 4. Comparison of knowledge distillation methods for semantic segmentation on Cityscapes with the DeepLabV3-MBV2 student model. T: teacher model, S: student baseline model.

Method	Params FLOPs mIoU		
	(M)	(G)	(%)
T: DeepLabV3-R101	84.74	695	78.07
S: PSPNet-R18			72.55
SKD [15] CVPR'19			73.29
IFVD [22] ECCV'20			73.71
CWD [19] ICCV'21	12.61	109	74.36
CIRKD [24] CVPR'22			74.73
MasKD [9] ICLR'23			75.34
Ours			75.99

Table 5. Comparison of knowledge distillation methods for semantic segmentation on Cityscapes with the PSPNet-R18 student model. T: teacher model, S: student baseline model.

lower complexity while maintaining competitive performance. Inference speed is generally correlated with model size, with lighter models (e.g., RetinaNet-R50) achieving over 40 FPS, while heavier models with X101 backbones

Model	Params (M)	FLOPs (G)	Mem (MB)	FPS
<i>Single-Stage Detectors (RetinaNet)</i>				
T: X101	95.86	424	367	29.4
T: R101	56.96	283	220	30.7
S: R50	37.97	215	148	41.9
<i>Two-Stage Detectors (Faster R-CNN)</i>				
T: X101	135.0	2014	528	20.6
T: R101	60.75	255	244	31.1
S: R50	41.75	187	171	42.1
<i>Anchor-Free Detectors (RepPoints)</i>				
T: X101	94.74	380	230	16.6
T: R101	55.84	239	224	24.5
S: R50	36.85	171	151	31.4

Table 6. Efficiency comparison of different object detectors. In addition to the number of parameters (Params) and FLOPs, we report CUDA memory usage (Mem) and inference speed (FPS), measured on an NVIDIA A100 GPU (80GB). For the two-stage teacher detector with an X101 backbone, Cascade R-CNN is employed. All models are evaluated with an input resolution of 1088×800. T: teacher model, S: student model.

(e.g., RepPoints-X101) operate at only around 20 FPS.

Additionally, Tab. 7 presents hardware-aware efficiency metrics for segmentation models, including CUDA memory usage and FPS, complementing the earlier tables that include comparisons of parameter counts and FLOPs. The teacher model DeepLabV3-R101 has the highest memory consumption (1828 MB) while achieving only 11.9 FPS. In contrast, the DeepLabV3-R18 and PSPNet-R18 student models achieve over 59 FPS (nearly 6 times faster than the teacher) while consuming significantly less memory (568MB and 621MB, respectively), making them better suited for real-time applications. Notably, DeepLabV3-MBV2, despite having fewer FLOPs and parameters, runs at a lower FPS (52.9) than DeepLabV3-R18 (59.2). This is likely due to the inefficiency of depthwise separable convolutions on the specific CUDA-based GPU used for evaluation. While they reduce FLOPs, their lower computational density results in suboptimal parallelization, making them less efficient than standard dense convolutions, as used in ResNet-18.

5. Ablation Studies

5.1. Spatial and Channel Masking

Tab. 8 illustrates the impact of spatial and channel masking on knowledge distillation performance. With our student-teacher feature fusion module in place, both spatial-only

Model	Mem (MB)	FPS
T:DeepLabV3-R101	1828	11.9
S:DeepLabV3-R18	568	59.2
S:DeepLabV3-MBV2	470	52.9
S:PSPNet-R18	621	59.4

Table 7. Efficiency comparison of different segmentation models. CUDA memory usage and FPS are measured on an NVIDIA A100 GPU (80GB). Parameter counts and FLOPs are provided in earlier tables. All models are evaluated with an input resolution of 1088×800. T: teacher model, S: student model.

and channel-only masking improve mAP over the baseline student (R50). Spatial masking yields a slightly higher mAP gain (40.9 vs. 40.4), likely due to its ability to preserve spatial dimensions, which is essential for dense prediction tasks. Notably, combining both spatial and channel masking further enhances performance, achieving the highest mAP of 41.2. AP_s benefits the most from spatial masking, while AP_m and AP_l see the greatest gains when both masking strategies are used together. These results highlight the complementary nature of spatial and channel masking, validating the design of our approach.

Method	mAP	AP_{50}	AP_{75}	AP_s	AP_m	AP_l
T: R101	38.9	58.0	41.5	21.0	42.8	52.4
S: R50	37.4	56.7	39.6	20.0	40.7	49.7
Spatial (Sp)	40.9	60.4	43.4	25.4	44.3	52.3
Channel (Ch)	40.4	60.1	43.2	24.5	44.1	52.3
Ours (Sp+Ch)	41.2	60.6	44.1	24.6	45.5	54.1

Table 8. Ablation study on spatial and channel masking strategies. Model: RetinaNet, Dataset: COCO. T: teacher model, S: student model.

5.2. Query Selection in Cross-Attention

Method	mAP	AP_s	AP_m	AP_l
Query from S	41.0	24.4	45.1	54.0
Query from T (Ours)	41.2	24.6	45.5	54.1

Table 9. Comparison of query selection strategies in our student-teacher cross-attention feature fusion (STCA-FF) module. Model: RetinaNet, Dataset: COCO. T: teacher model, S: student model.

Tab. 9 compares different query selection strategies in our student-teacher cross-attention feature fusion (STCA-FF) module. Our default approach, where the teacher generates the query while the student provides the key and value, achieves 41.2 mAP, outperforming the alternative configuration in which the student defines the query (41.0

Method	mAP	AP ₅₀	AP ₇₅	AP _s	AP _m	AP _l
T: R101	38.9	58.0	41.5	21.0	42.8	52.4
S: R50	37.4	56.7	39.6	20.0	40.7	49.7
No Masking	37.4	57.1	40.0	20.8	40.8	50.9
Fixed Masking from Teacher	39.8	59.0	42.5	21.5	43.9	54.0
Adaptive Masking from Teacher	39.9	59.1	42.9	21.7	43.7	53.9
Our ACAM-KD	41.2	60.6	44.1	24.6	45.5	54.1

Table 10. Ablation study on fixed and adaptive masking strategies for knowledge distillation. Model: RetinaNet, Dataset: COCO. T: teacher model, S: student model.

mAP). This is because queries from the more knowledgeable teacher can effectively guide the student on where to attend within its evolving key-value space. In contrast, student-generated queries sometimes yield unstable masks in early stages due to unclear attention targets.

5.3. Fixed and Adaptive Masking

Tab. 10 compares fixed and adaptive masking strategies in knowledge distillation. The no-masking baseline achieves 37.4 mAP, indicating that the student model benefits only marginally from standard distillation without strategic feature selection. Fixed distillation masking learned by the teacher offline (as in [9]) improves mAP to 39.8. Allowing the teacher-driven masking to adapt during distillation (i.e., adaptive masking from the teacher) refines the selection process but results in only a slight improvement to 39.9 mAP. This suggests that making the mask learnable alone has limited impact when the knowledge source remains strictly from the teacher. Our ACAM-KD, which combines student-teacher cross-attention feature fusion with adaptive masking, achieves the highest performance of 41.2 mAP. Notably, our method is particularly effective in handling small objects, as AP_s shows the largest improvement, reaching 24.6. The superiority of ACAM-KD stems from the cooperative interaction between the teacher and student in identifying the most suitable knowledge for the student’s evolving needs at different distillation stages.

5.4. Effect of Diversity Loss

Fig. 3 shows that different spatial masks focus on distinct regions, emphasizing their complementary roles in knowledge transfer. This variation is encouraged by the diversity loss, which prevents multiple masks from collapsing into similar patterns. By fostering diversity, the loss leads to masks that capture a wider range of informative features, enabling the student model to acquire richer and more comprehensive knowledge.

6. Conclusion

Existing feature-based knowledge distillation methods rely on static, teacher-driven supervision, lacking adaptability




Figure 3. Visualization of spatial attention masks associated with different learnable selection units. Variations in highlighted regions, encouraged by our diversity loss, ensure complementary feature learning for effective knowledge distillation. Warmer colors (red/yellow) indicate higher attention, while cooler colors (blue) denote lower attention.

to the student’s evolving learning state. These approaches often overlook the value of dynamic student-teacher interactions and joint spatial-channel feature selection. In this paper, we propose Adaptive and Cooperative Attention Masking for Knowledge Distillation (ACAM-KD), which integrates student-teacher cross-attention feature fusion and adaptive spatial-channel masking. Our method grants the student greater autonomy and fosters a more interactive and adaptive distillation process, resulting in consistent performance gains in dense prediction tasks (e.g., object detection and semantic segmentation). Experimental results highlight the effectiveness of dynamically tailoring knowledge transfer to the student’s needs, offering a promising direction for efficient model compression.

7. Acknowledgment

This work was supported by the National Science Foundation (NSF) under Award No. 2153404 and No. 2412285.

References

- [1] Guobin Chen, Wongun Choi, Xiang Yu, Tony Han, and Manmohan Chandraker. Learning efficient object detection mod-

- els with knowledge distillation. *Advances in Neural Information Processing Systems*, 30, 2017. 5
- [2] 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
- [3] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018. 6
- [4] MMSegmentation Contributors. MMSegmentation: Openmmlab semantic segmentation toolbox and benchmark. <https://github.com/open-mmlab/mms Segmentation>, 2020. 6
- [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. 5
- [6] Xing Dai, Zeren Jiang, Zhao Wu, Yiping Bao, Zhicheng Wang, Si Liu, and Erjin Zhou. General instance distillation for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7842–7851, 2021. 2, 5
- [7] Zhixing Du, Rui Zhang, Ming Chang, Shaoli Liu, Tianshi Chen, Yunji Chen, et al. Distilling object detectors with feature richness. *Advances in Neural Information Processing Systems*, 34:5213–5224, 2021. 5
- [8] Byeongho Heo, Jeessoo Kim, Sangdoon Yun, Hyojin Park, Nojun Kwak, and Jin Young Choi. A comprehensive overhaul of feature distillation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1921–1930, 2019. 5
- [9] Tao Huang, Yuan Zhang, Shan You, Fei Wang, Chen Qian, Jian Cao, and Chang Xu. Masked distillation with receptive tokens. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023. 1, 2, 5, 6, 8
- [10] Tao Huang, Yuan Zhang, Mingkai Zheng, Shan You, Fei Wang, Chen Qian, and Chang Xu. Knowledge diffusion for distillation. *Advances in Neural Information Processing Systems*, 36:65299–65316, 2023. 5
- [11] Quanquan Li, Shengying Jin, and Junjie Yan. Mimicking very efficient network for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6356–6364, 2017. 1
- [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 *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 740–755. Springer, 2014. 5
- [13] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2980–2988, 2017. 5
- [14] Tao Liu, Chenshu Chen, Xi Yang, and Wenming Tan. Re-thinking knowledge distillation with raw features for semantic segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1155–1164, 2024. 6
- [15] Yifan Liu, Ke Chen, Chris Liu, Zengchang Qin, Zhenbo Luo, and Jingdong Wang. Structured knowledge distillation for semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2604–2613, 2019. 2, 6
- [16] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017. 5
- [17] 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. 5
- [18] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018. 6
- [19] Changyong Shu, Yifan Liu, Jianfei Gao, Zheng Yan, and Chunhua Shen. Channel-wise knowledge distillation for dense prediction. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5311–5320, 2021. 2, 6
- [20] Jiabao Wang, Yuming Chen, Zhaohui Zheng, Xiang Li, Ming-Ming Cheng, and Qibin Hou. Crosskd: Cross-head knowledge distillation for object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16520–16530, 2024. 5
- [21] Tao Wang, Li Yuan, Xiaopeng Zhang, and Jiashi Feng. Distilling object detectors with fine-grained feature imitation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4933–4942, 2019. 2, 5
- [22] Yukang Wang, Wei Zhou, Tao Jiang, Xiang Bai, and Yongchao Xu. Intra-class feature variation distillation for semantic segmentation. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VII 16*, pages 346–362. Springer, 2020. 2, 6
- [23] Chenhongyi Yang, Mateusz Ochal, Amos Storkey, and Elliot J Crowley. Prediction-guided distillation for dense object detection. In *European Conference on Computer Vision*, pages 123–138. Springer, 2022. 1
- [24] Chuanguang Yang, Helong Zhou, Zhulin An, Xue Jiang, Yongjun Xu, and Qian Zhang. Cross-image relational knowledge distillation for semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12319–12328, 2022. 2, 6
- [25] Ze Yang, Shaohui Liu, Han Hu, Liwei Wang, and Stephen Lin. Reppoints: Point set representation for object detection.

In *The IEEE International Conference on Computer Vision (ICCV)*, 2019. [5](#)

- [26] Zhendong Yang, Zhe Li, Xiaohu Jiang, Yuan Gong, Zehuan Yuan, Danpei Zhao, and Chun Yuan. Focal and global knowledge distillation for detectors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4643–4652, 2022. [1](#), [2](#), [5](#)
- [27] Zhendong Yang, Zhe Li, Mingqi Shao, Dachuan Shi, Zehuan Yuan, and Chun Yuan. Masked generative distillation. In *European conference on computer vision*, pages 53–69. Springer, 2022. [5](#), [6](#)
- [28] Linfeng Zhang and Kaisheng Ma. Improve object detection with feature-based knowledge distillation: Towards accurate and efficient detectors. In *International conference on learning representations*, 2020. [1](#), [2](#), [5](#)
- [29] Yuan Zhang, Tao Huang, Jiaming Liu, Tao Jiang, Kuan Cheng, and Shanghang Zhang. Freekd: Knowledge distillation via semantic frequency prompt. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15931–15940, 2024. [1](#), [2](#), [5](#), [6](#)
- [30] 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. [6](#)