

KD-DETR: Knowledge Distillation for Detection Transformer with Consistent Distillation Points Sampling

Yu Wang, Xin Li, Shengzhao Weng, Gang Zhang, Haixiao Yue,
 Haocheng Feng, Junyu Han, Errui Ding
 Baidu VIS

{wangyu106, lixin68, wengshengzhao, zhanggang03,
 yuehaoxiao, fenghaocheng, hanjunyu, dingerrui}@baidu.com

Abstract

DETR is a novel end-to-end transformer architecture object detector, which significantly outperforms classic detectors when scaling up. In this paper, we focus on the compression of DETR with knowledge distillation. While knowledge distillation has been well-studied in classic detectors, there is a lack of researches on how to make it work effectively on DETR. We first provide experimental and theoretical analysis to point out that the main challenge in DETR distillation is the lack of consistent distillation points. Distillation points refer to the corresponding inputs of the predictions for student to mimic, which have different formulations in CNN detector and DETR, and reliable distillation requires sufficient distillation points which are consistent between teacher and student.

Based on this observation, we propose the first general knowledge distillation paradigm for DETR (KD-DETR) with consistent distillation points sampling, for both homogeneous and heterogeneous distillation. Specifically, we decouple detection and distillation tasks by introducing a set of specialized object queries to construct distillation points for DETR. We further propose a general-to-specific distillation points sampling strategy to explore the extensibility of KD-DETR. Extensive experiments validate the effectiveness and generalization of KD-DETR. For both single-scale DAB-DETR and multi-scale Deformable DETR and DINO, KD-DETR boost the performance of student model with improvements of 2.6% – 5.2%. We further extend KD-DETR to heterogeneous distillation, and achieves 2.1% improvement by distilling the knowledge from DINO to Faster R-CNN with ResNet-50, which is comparable with homogeneous distillation methods.

1. Introduction

In recent years, [2] propose the novel end-to-end detector Detection Transformer (DETR) which eliminates the need

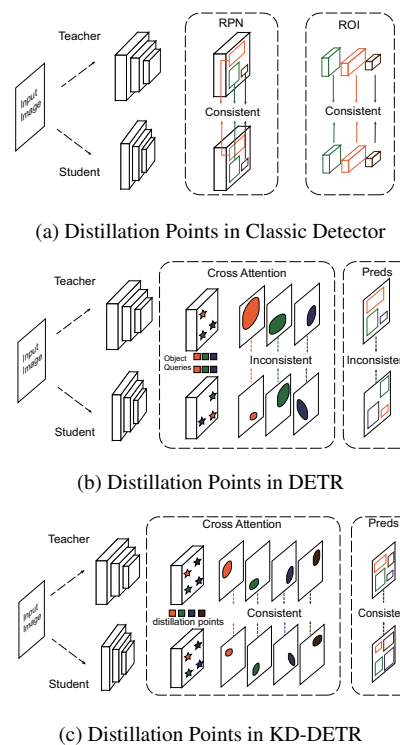


Figure 1. **Schematic Illustrations of Distillation Points in Different Architecture:** (a) Two-Stage Detector: both positive and negative proposals in RPN and ROI are consistent distillation points with strict one-to-one correspondence between teacher and student model; (b) DETR: Object queries lacks spatial or semantic relationship between teacher and student model, resulting in inconsistent distillation points; (c) In KD-DETR, a set of special object queries is introduced to construct consistent distillation points for DETR distillation

for hand-crafted anchors and non-maximum suppression (NMS). [43][14][19][41] further make remarkable stride towards the scalability and potential of DETR, significantly outperforming classical detectors.

Different from classic detectors, DETR interprets object detection as an end-to-end set prediction problem with bipartite matching. A set of learnable object queries are introduced, each responsible for a certain instance. The object queries interact with features extracted from the encoder to make final predictions of box locations and categories. Despite the impressive performance, the growing model scale prevents DETR from being deployed to real-world applications with urgent computation budget requirement.

To address this problem, current works have made efforts in designing efficient DETR architectures, reducing the encoder tokens utilized in cross-attention module to decrease the computation cost[24], and leveraging dense prior from RPN to downsize the decoder layers [36]. In this work, we concentrate on compressing the large-scale DETR model by knowledge distillation[11] approaches, which is a promising technique for model compression and accuracy boosting. Knowledge distillation can transfer the knowledge learned from large and cumbersome DETR models to small and efficient ones by forcing the student to mimic the predictions from teacher, whether logits or internal activations. However, modern knowledge distillation methods are designed under CNN-based detectors, and researches on expanding it to general DETR compression are limited.

We start from experiments of applying classic logit-based distillation[4] to DETR to investigate the key point in DETR distillation. With these experiments, we observe that the critical challenge lies in the different formulations of DETR and classic detectors. Compared with classic detectors, the set-prediction formulation of DETR naturally contains fewer consistent distillation points. We use distillation points to denote the corresponding input of the prediction for mimicking in knowledge distillation, and the sufficiency and consistency of distillation points form the foundation of knowledge distillation. Specifically, abundant distillation points which are kept consistent between teacher and student models are essential for effective distillation. As shown in Figure 1a, classic detectors make predictions for a set of region proposals generated from the sliding window locations on the images with handcraft scales. This pattern naturally ensures a strict spatial correspondence between the large number of proposals made by teacher and student models, even for those negative ones with low confidence, thus providing a sufficient number of consistent distillation points for mimicking. In DETR, as is shown in Figure 1b, the distillation points actually consist of both image and object queries. However, the object queries from teacher and student models are egocentric and even differ in number, thus lacking definite correspondence, especially those redundant negative queries in bipartite matching. As the distillation points in DETR are inconsistent and insufficient, the predictions acquired from teacher are not reliable or informative for student to mimic.

The observation above raises the issue: how to obtain sufficient and consistent distillation points for DETR distillation. Previous work[3] explicitly alleviate this issue by utilizing the bipartite matching between the object queries from teacher and student. However, the bipartite matching is not stable[14], and the matched object queries are just similar but not consistent, lacking sufficiency and extensibility. To directly address this issue, we propose a general knowledge distillation paradigm for DETR (KD-DETR) with consistent distillation points sampling. In KD-DETR, as illustrated in Figure 1c, we decouple detection and distillation task by introducing a set of specialized object queries to construct distillation points. The distillation points are unlearnable and shared between teacher and student models, probing the “dark knowledge” in teacher model. In this way, consistent distillation points with customized quantities become available. With the paradigm of KD-DETR, we propose a general-to-specific distillation points sampling strategy to probe comprehensive knowledge in teacher model. We further propose a coordination-based distillation points sampling strategy to extend KD-DETR to heterogeneous distillation between DETR and CNN detector.

To the best of our knowledge, this is the first work to propose a general knowledge distillation paradigm for DETR for both homogeneous and heterogeneous distillation. In this paper, we first provide a thorough analysis of the key points in DETR distillation. Based on the analysis, we design a novel KD-DETR which significantly improves the performance of DETR distillation. KD-DETR has both flexibility to different DETR architectures and potential for scalability, even for heterogeneous distillation between DETR and CNN detectors. We conduct extensive experiments on the MS COCO2017[17] dataset on both homogeneous and heterogeneous distillation, and significantly boosts the performance of student models. DAB-DETR with ResNet-18 and ResNet-50 backbone achieves 41.4%, 45.7% mAP, respectively, which are 5.2%, 3.6% higher than the baseline. Deformable DETR with ResNet-18 and ResNet-50 reaches 43.7% and 48.3% mAP, 3.6% and 3.8% higher than the baseline, and outperform DETRDistill[3] with 1.7% improvement. DINO with ResNet-18 and ResNet-50 also gains 4.4% and 2.6% improvement than baseline. We further extend KD-DETR to heterogeneous distillation, and achieves 2.1% improvement by distilling the knowledge from DINO to Faster R-CNN with ResNet-50, which is comparable with homogeneous distillation methods.

2. Related Work

2.1. Classic Object Detection

Classic detectors with CNN view object detection as a verification task with a sliding window on the image to gen-

erate anchors. The mainstream detectors can be divided into one-stage detectors[28][16] and two-stage detectors[23][1]. One-stage detectors, such as Retinanet[18], YOLO[22], and FCOS[28], directly predict the category and regression of anchors on each pixel of the feature maps. While two-stage detectors such as Faster-RCNN[23] and its variants[21][32][39] introduce a Region Proposal Networks (RPN) to generate proposals, and a ROIpool or ROIAlign[29] to extract features of each region proposal for further classification and regression refinement. Both one-stage and two-stage detectors require post-processing, such as NMS, to remove duplicate predictions.

2.2. Detection Transformer

[2] first propose an end-to-end transformer-based detector without any post-processing. Different from classic object detection, DETR interprets object detection as a set-prediction problem with bipartite matching. Lots of follow-up focus on the slow convergence of DETR[5][26][7][40]. Deformable DETR[43] introduces a deformable attention module by generating reference points for query elements, each of which only concentrates on a small number of locations on the whole feature map. An alternative way is to add more prior information to the object queries in the decoder. Conditional-DETR[20] decouples the context and position features in object queries and generates position features by spatial location. DAB-DETR[19] further introduces the width and height information to the positional features. Anchor DETR[31] also encodes the anchor points as object queries with multiple patterns, and further designs a row-column decouple attention to reduce memory cost. The recent work of DINO[41] draws the existing novel techniques, and further exerts the potential of DETR by enlarging the scale of model and datasets.

Besides, another problem in DETR is the scale and computation cost of the model. Current works solving this problem by designing more efficient DETR architecture. Sparse DETR[24] reduces the computation cost by sparsifying encoder tokens. Efficient DETR[36] otherwise introduce RPN to generate object queries and eliminate the cascading decoder layers in DETR. PnP DETR[30] shorten the length of sampled feature with a poll and pool sampling module.

2.3. Knowledge Distillation

Knowledge distillation is a widely-used method for model compression and accuracy boosting by transferring the knowledge in a large cumbersome teacher model to a small student. [11] first propose the concept of knowledge distillation, where the student mimics the soft predictions from teacher. Knowledge distillation has been utilized in various fields[27][34][33]. According to the objective of mimicking, knowledge distillation can be divided into three categories: response-based[42], feature-based[10][35] and

relation-based[37][38], which distill with logits, intermediate activations and the relation of features in different layers respectively.

Several works focus on applying knowledge distillation to object detection[8][4][12]. [4] successfully distills the features on the neck, the classification head, and the regression head, while [15] chooses to distill the logits and features from the RPN head. To overcome the imbalance of foreground and background, [29] introduces fine-grained mask to focus on the regions close to ground-truth bounding boxes, [6] pays more attention to the regions where teacher and student are divided in predictions.

However, the modern knowledge distillation methods for object detection are built upon the architecture of CNN-based detectors, and are not suitable for DETR due to the completely different transformer architecture. [3] directly introduces response-based and feature-based distillation to DETR with Hungarian-matching. Different from the previous work, we analyze the limitation of the set-prediction formulation in knowledge distillation, and propose a general paradigm for both homogeneous and heterogeneous DETR distillation.

3. A Closer Look at DETR Distillation

In this section, we first revisit the DETR architecture briefly. Then we conduct a series of classic knowledge distillation experiments on DETR to reveal that the core of DETR distillation is to obtain sufficient and consistent distillation points.

3.1. Revisiting DETR

DETR is built upon the encoder-decoder architecture of transformer. The encoder takes pixels of the feature map from backbone as input for multi-head self-attention to extract context features $X \in \mathbb{R}^{HW \times D}$, where HW denotes the resolution of the feature map, and D denotes the feature dimension. The decoder takes the features from encoder and a set of learnable object queries $\mathbf{Q} = \{q_i \in \mathbb{R}^D | i = 1, \dots, N\}$ as input, where N denotes the number of queries. Each object query is an abstract feature describing a certain instance, and will probe and pool the features from encoder through cross-attention to make predictions of category $\mathbf{C} = \{c_i \in \mathbb{R}^K | i = 1, \dots, N\}$ and location $\mathbf{B} = \{b_i = (bx_i, by_i, bw_i, bh_i) | i = 1, \dots, N\}$, where K denotes the number of categories. Finally, the Hungarian algorithm is used to find a bipartite matching between ground truth and predictions of object queries.

3.2. Consistent Distillation Points

The core idea of knowledge distillation is forcing the student to mimic the prediction of teacher, which can be interpreted as matching the mapping function of student and teacher with a set of distillation points. Distillation points

Strategy	AP	AP_{50}	AP_{75}
Baseline	36.2	56.1	37.9
Inconsistent	35.1	55.2	36.7
Similar Foreground	37.2	57.4	39.9
Similar General	37.4	58.0	40.6

Table 1. Distillation with Different Distillation Points. Inconsistent distillation points are harmful for distillation.

refer to the corresponding input \mathbf{x} of the predictions, as $\mathbf{y} = f(\mathbf{x})$, where f represents the model. In this view, the distillation points should be kept sufficient in quantity and consistent between teacher and student models for effective and reliable matching. However, comparing the formulation of CNN detector and DETR, a critical challenge of DETR distillation lies in the lacking of consistent distillation points.

Classic detector degrades object detection to a verification problem which combines classification and regression, and introduces a set of anchors to specify the region for verification. In this way, the formulation of distillation points consists of the image and the location and scale of anchor $\mathbf{x} = (\mathbf{I}, \mathbf{anchor})$. As the anchors are generated through the sliding window strategy with handcrafted shapes, the locations and scales of anchors are implicitly embedded in the model architecture as prior information. Since student and teacher models share the same or similar architectures, a large number of object proposals generated by teacher and student models naturally have strict spatial correspondence, even for those background regions with low confidence. With the spatial correspondence which can be viewed as an inductive bias of CNN, classic detector naturally guarantees a sufficient number of consistent distillation points.

In contrast, the DETR formulates object detection as a set-prediction problem. The distillation points, therefore, become the combination of the image and the object queries $\mathbf{x} = (\mathbf{I}, \mathbf{q})$. However, the object queries in different models are egocentric, as they are initialized and optimized by themselves independently. Since object queries play the role of probing and pooling the features of certain instances, they have inconsistent concentration preferences in different models. Consequently, the formulation of DETR naturally lacks the ability to provide sufficient distillation points with strict consistency between teacher and student models, and the predictions acquired from teacher are not informative or reliable for student to mimic.

3.3. Distillation with Inconsistent Distillation Points

To validate the analysis above that the sufficiency and consistency of distillation points is the essential challenge in DETR distillation, we start with applying the original logit-based distillation method[4] in classic detector to DETR, which mimics the category and box location logits predictions of teacher model.

We examine three distillation points strategies: *Inconsistent*, *Similar Foreground*, and *Similar General*. In *Inconsistent*, all the object queries are viewed as distillation points with their original permutation; In *Similar Foreground*, only object queries matched to ground truth in bipartite matching will be used as distillation points, and will permute with the same order of ground truth label; *Similar General* further increase the number of distillation points by viewing the average of negative object queries in bipartite matching as a general background distillation point. Experiments are conducted on DAB-DETR and evaluated on MS COCO2017, with ResNet18 as student and ResNet-50 as teacher.

As shown in Table 1, *Inconsistent* distillation points results in great degradation of the student model with unreliable knowledge from teacher model; *Similar Foreground* with semantic-similar foreground distillation points alleviate the problem; and *Similar General* achieves further improvement by increasing the number of distillation points with general background features. The preliminary experiments validate that the sufficiency and consistency of distillation points are of prime importance to improve the performance of student model in DETR distillation.

4. KD-DETR

To address the lack of consistent distillation points in DETR, we propose a general knowledge distillation paradigm for DETR (KD-DETR) with consistent distillation points sampling. As illustrated in Figure 2, KD-DETR introduces a set of specialized object queries $\tilde{\mathbf{q}}$ shared between teacher and student models to construct distillation points. Decoupling the distillation task and detection task, KD-DETR can provide sufficient and consistent distillation points. We denote the original input and sampled distillation points as $\mathbf{x} = \{\mathbf{I}, \mathbf{q}\}$, $\tilde{\mathbf{x}} = \{\mathbf{I}, \tilde{\mathbf{q}}\}$ respectively.

For detection task, the student model is first optimized by its original detection loss: the original input \mathbf{x} is fed into student model to make category and box location predictions, which will be assigned to the ground truth with bipartite matching and calculate detection loss \mathcal{L}_{det} .

For distillation task, the sampled distillation points $\tilde{\mathbf{x}}$ will be fed into both student and teacher to make category and location predictions \mathbf{c}, \mathbf{b} :

$$\mathbf{c}^s, \mathbf{b}^s = f^s(\mathbf{I}, \tilde{\mathbf{q}}), \quad (1)$$

$$\mathbf{c}^t, \mathbf{b}^t = f^t(\mathbf{I}, \tilde{\mathbf{q}}), \quad (2)$$

where f^s, f^t refer to student and teacher model respectively. The distillation loss is calculated with following form:

$$\mathcal{L}_{distill} = \sum_{i=1}^M [\lambda_{cls} \mathcal{L}_{KL}(\hat{\mathbf{c}}_i^t \| \hat{\mathbf{c}}_i^s) + \lambda_{L1} \mathcal{L}_{L1}(\mathbf{b}_i^s, \mathbf{b}_i^t)], \quad (3)$$

$$+ \lambda_{GIoU} \mathcal{L}_{GIoU}(\mathbf{b}_i^s, \mathbf{b}_i^t),$$

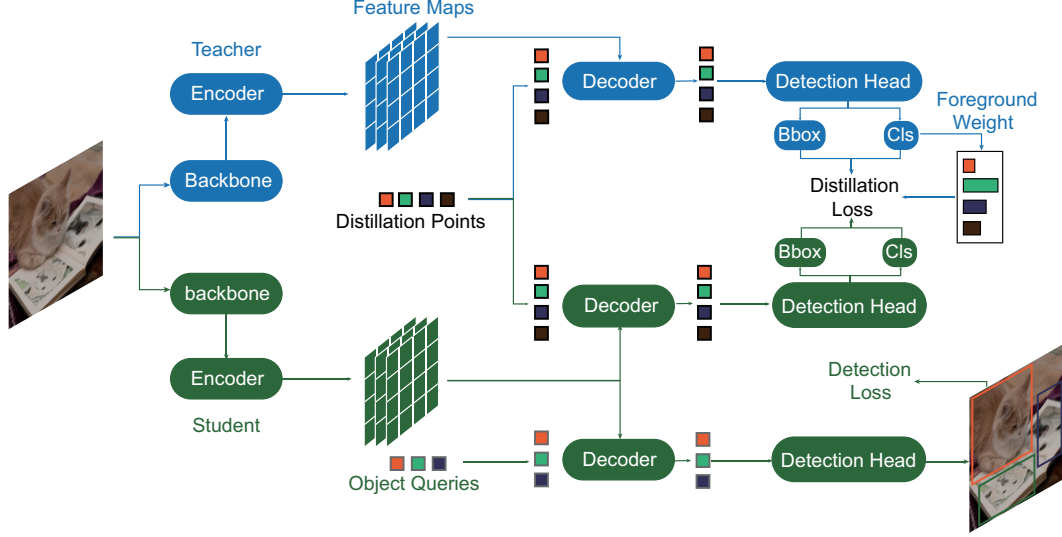


Figure 2. **KD-DETR architecture** KD-DETR decouples the distillation and detection tasks by introducing a set of distillation points shared between teacher and student models. For distillation, the distillation points are served as the query of the transformer decoders for both teacher and student models, and the student will mimic teacher’s classification and box location predictions. The weight of each distillation point is measured by its foreground probability predicted by teacher. For detection task, the original object queries are processed by the student decoder for final prediction.

where M denotes the number of distillation points. For classification, we choose the KL-divergence \mathcal{L}_{KL} as distillation loss with temperature T : $\hat{c} = \text{SoftMax}(\frac{c}{T})$. For box regression, \mathcal{L}_{L1} and \mathcal{L}_{GIoU} represent the L1 and GIoU loss for location distillation, which have the same formulation with detection loss. $\lambda_{cls}, \lambda_{L1}, \lambda_{GIoU}$ represent the coefficient of corresponding loss.

The total Loss is calculated with following form:

$$\mathcal{L} = \mathcal{L}_{distill} + \mathcal{L}_{det} \quad (4)$$

4.1. Distillation Points Sampling

Generally, object queries are a set of abstract features responsible for certain objects by probing and pooling the context features from encoder. Existing works interpret object queries as anchors or reference points, revealing that each object query is sensitive to a particular region on the feature maps. Following this perspective, we provide a comprehensive general-to-specific sampling strategies for distillation points sampling: $\tilde{\mathbf{q}} = \{\mathbf{q}_g, \mathbf{q}_s\}$.

General Sampling with Random Initialized Queries. In general sampling, we hope to probe teacher’s general responses on different locations of the features by sparsely scanning the whole feature maps. Therefore, we randomly initialize a set of object queries, which are uniformly distributed on the features, to construct the general distillation points: $\mathbf{q}_g = \{\mathbf{q}_i \sim \mathcal{U}(0, 1) | i = 1, \dots, M_g\}$, where M_g denotes the number of general distillation points. To learn more general knowledge from teacher, we leave these distillation points unlearnable during training, and re-sample

them every iteration.

Specific Sampling with Teacher Queries. While the general sampling provides a global retrieval of the features, we further propose a specific sampling strategy, focusing on those regions where teacher pays more attention. An intuitive way for specific sampling is to directly reuse the well-optimized object queries in teacher model: $\mathbf{q}_s = \mathbf{q}_{teacher}$. While teacher model is learned to concentrate more on these object queries, the predictions in these areas are more precise and informative.

Foreground Rebalance Weight. The imbalance between foreground and background regions is one critical problem in object detection distillation, not special in DETR. An intuitive way is to utilize the classification scores of distillation points which are predicted by teacher model to rebalance the distillation loss. Concretely, those distillation points with higher classification scores are regarded as foreground distillation points, containing more useful information for detection, and should be given more attention.

$$w_i = \max_{c \in [0, K]} p^t(y_c | \mathbf{q}_i), \quad (5)$$

where $p^t(y_c | \mathbf{q}_i)$ denotes the probability of \mathbf{q}_i assigned to category c predicted by teacher model, and w_i denotes the foreground rebalance weight of \mathbf{q}_i . In this way, the distillation loss in Eq. (3) will be written as follow:

$$\begin{aligned} \mathcal{L}_{distill} = & \sum_{i=1}^M w_i [\lambda_{cls} \mathcal{L}_{KL}(\hat{\mathbf{c}}_i^t || \hat{\mathbf{c}}_i^s) \\ & + \lambda_{L1} \mathcal{L}_{L1}(\mathbf{b}_i^s, \mathbf{b}_i^t) + \lambda_{GIoU} \mathcal{L}_{GIoU}(\mathbf{b}_i^s, \mathbf{b}_i^t)] \end{aligned} \quad (6)$$

Models		Epochs	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L	GFlops	Params
DAB-DETR	ResNet-50(T)	50	42.1	63.1	44.7	21.5	45.7	69.3	90	44M
	ResNet-18(S)	50	36.2	56.1	37.9	16.9	39.0	53.5	49	23M
	Ours	50	41.4	61.4	44.2	21.2	44.7	58.7	49	23M
	Gains		+5.2	+5.3	+6.3	+4.3	+5.7	+5.2		
DAB-DETR	ResNet-101(T)	50	43.5	63.9	44.6	23.6	47.3	61.5	157	63M
	ResNet-50(S)	50	42.1	63.1	44.7	21.5	45.7	60.3	90	44M
	Ours	50	45.7	66.3	49.4	26.4	49.8	62.7	90	44M
	Gains		+3.6	+3.2	+4.7	+4.9	+4.1	+2.4		
Deformable-DETR	ResNet-50(T)	50	44.5	63.6	48.7	27.1	47.6	59.6	171	40M
	ResNet-18(S)	50	40.1	58.1	43.7	22.4	42.8	54.2	127	24M
	Ours	50	43.7	62.1	47.7	25.9	46.8	57.6	127	24M
	Gains		+3.6	+4.0	+4.0	+3.5	+4.0	+3.4		
Deformable-DETR	ResNet-101(T)	50	48.0	66.7	52.6	30.5	52.3	62.3	238	59M
	ResNet-50(S)	50	44.5	63.6	48.7	27.1	47.6	59.6	171	40M
	Ours	50	48.3	66.7	52.9	30.8	52.1	62.5	171	40M
	Gains		+3.8	+3.1	+4.2	+3.7	+4.5	+2.9		
DINO	ResNet-50(T)	36	50.9	69.0	55.3	34.6	54.1	64.6	245	47M
	ResNet-18(S)	12	44.0	61.2	48.1	27.4	46.9	56.9	200	30M
	Ours	12	48.4	65.5	53.0	31.6	51.7	62.3	200	30M
	Gains		+4.4	+4.3	+4.9	+4.2	+4.8	+5.4		
DINO	ResNet-101(T)	36	51.3	69.5	55.8	34.8	54.8	65.8	311	66M
	ResNet-50(S)	12	49.0	66.6	53.5	32.0	52.3	63.0	245	47M
	Ours	12	51.6	69.6	56.6	34.2	54.8	66.9	245	47M
	Gains		+2.6	+3.0	+3.1	+2.2	+2.5	+3.9		

Table 2. Results of the proposed KD-DETR with different DETR detectors and backbones with various scale.

4.2. Generalization to Heterogeneous Distillation

To further extend the generalization of the KD-DETR paradigm, we apply the idea to heterogeneous distillation between DETR and CNN detector. Intuitively, both the anchor in the CNN detector and the object query in DETR represent certain locations on the image and share spatial consistency. We first construct a set of distillation points \tilde{q} with the coordinate of anchors in CNN detector, and then convert them to the formulation of object queries in DETR. The distillation loss is in the formulation of Eq. (3).

We further propose a simple but effective strategy for distillation points sampling in heterogeneous distillation with Intersection over Union (IoU). Specifically, we calculate the IoU between anchors and ground truth grounding boxes, and select the top k anchors as distillation points. Details about heterogeneous distillation are in Appendix 7.

5. Experiment

To validate the effectiveness and generalization of KD-DETR, we evaluate it on different DETR architectures including DAB-DETR, Deformable DETR and DINO, with two scales of backbones: ResNet-50 and ResNet18. We also extend KD-DETR to heterogeneous distillation, and

evaluate on distillation between DINO-Res50 and Faster-RCNN Res50. To support our analysis of consistent distillation points, we further conduct an extensive ablation study.

5.1. Experimental Settings

Datasets: All the proposed experiments are evaluated on MS COCO2017[17] spanning 80 categories, with the default split of 117k training images for training and 5k validate images for testing. Standard COCO evaluation metrics are adopted for validation.

Implementation Details: As KD-DETR is a plug-and-play distillation module, we follow the original settings of hyper-parameters and optimizer of all the student model for the training of detection part. We choose ResNet[9] as backbone, which are pre-trained on ImageNet1K[13]. We propose the inheriting strategy[12] to initialize students' level embeddings on multi-scale DETR (details in Appendix 8.2). For distillation task, we set hyper-parameters of the coefficient of the distillation loss as $\lambda_{kl} = 1$, $\lambda_{L1} = 5$, $\lambda_{GIoU} = 2$. The number of General distillation points is 300, 300, 900 for DAB-DETR, Deformable DETR, and DINO. For heterogeneous distillation, the hyper-parameter k for IoU sampling is set to 10. We train our models on Nvidia A100 GPUs with batch size set to 16 in total.

Method	AP	AP_S	AP_M	AP_L
Deformable DETR Res50	44.5	27.1	47.6	59.6
FGD[34]	44.1	25.9	47.7	58.8
FitNet[25]	44.9	27.2	48.4	59.6
DETRDistill[3]	46.6	28.5	50.0	60.4
Ours	48.3	30.8	52.1	62.5

Table 3. Comparison with state-of-the-art on Deformable DETR.

5.2. Distillation on Different DETR benchmarks

We evaluate our method on three typical DETR architectures: DAB-DETR[19], a single-scale DETR with special object query settings; Deformable-DETR[43], a multi-scale DETR with deformable attention module; and DINO[41], which combines a series of novel techniques including deformable attention, two-stage object queries settings, and a de-noising module, to evaluate our effectiveness on well-optimized model with high accuracy. Distillation results of DINO with Swin Transformer backbone and compressing the layers of DETR are in Appendix 8.2, 8.3

The results are illustrated in Table 2. For DAB-DETR, KD-DETR significantly boosts the performance of ResNet-18, and ResNet-50 with 5.2%, and 3.6% mAP improvement respectively. Note that student with ResNet-50 even surpass the teacher by 2.2% mAP. For Deformable DETR with ResNet-18 and ResNet-50, KD-DETR achieve 3.6% and 3.8% mAP gains, and student with ResNet-50 also exceed the teacher with 0.3% mAP. For DINO with ResNet-18 and ResNet-50, KD-DETR also improves 4.4% and 2.6% mAP of the student models on 12-epoch training scheduler.

Table 3 shows the comparison with state-of-the-art distillation method. With simple logit-based distillation alone, KD-DETR outperforms former methods, with an improvement of 1.7% on Deformable DETR with ResNet-50 compared with DETRDistill[3]

5.3. Generalization to Heterogeneous Distillation

Table 4 presents the results with Faster RCNN ResNet-50 as student and DINO ResNet-50 as teacher. KD-DETR works well on heterogeneous distillation task, improving the student models with 2.1% mAP, achieving the performance of state-of-the-art homogeneous methods. This validate our analysis on the effect of distillation points on knowledge distillation, and bridge the gap of transferring knowledge between detectors with different architectures.

5.4. Ablation Study

The ablation study is conducted on MS COCO2017 with DAB-DETR with backbone of ResNet50 as teacher and ResNet18 as student.

Analysis on the Benefit of Knowledge Distillation

Method	AP	AP_S	AP_M	AP_L
RCNN-Res50	38.4	21.5	42.1	50.3
FGFI[29]	39.3	22.5	42.3	52.2
GID[6]	40.2	22.7	44.0	53.2
FGD[34]	40.4	22.8	44.5	53.5
Ours	40.5	22.7	44.6	53.8

Table 4. **Heterogeneous Distillation:** Distillation between DINO and Faster RCNN

Object Queries	Distillation Points	AP	AP_{50}	AP_{75}
300	-	36.2	56.1	37.9
300+300	-	37.7	58.0	41.1
300	300	40.2	60.7	42.8

Table 5. **Benefit of Knowledge Distillation:** Simply increasing the set of object queries leads to trivial improvement

In KD-DETR, we introduce a set of specialized distillation points in the training process to the original DETR architecture. [41][14] have proven that the number of object queries will affect the performance of model. To validate that the benefit is from knowledge distillation rather than the increase of object queries, we conduct an ablation experiment by applying an additional set of object queries with the same number of distillation points.

As shown in Table 5, simply increasing the number of object queries only achieves trivial gains, and the main contribution of boosting the model’s performance is from the knowledge transferred from teacher with distillation points.

Analysis on the Distillation Points Sampling Strategy

A comprehensive general-to-specific distillation points sampling scheme is introduced in this paper, including three strategies: General Sampling, Specific Sampling, and Foreground Rebalance Weight. The ablation results are illustrated in Table 6. The general sampling strategy with randomly initialized queries can boost the performance of student model for 2.5% mAP, while the specific sampling strategy with teacher queries achieves an improvement of 3.8%. When refining the general distillation points with foreground rebalance weight, the general sampling strategy and specific sampling strategy yield larger gains for 3.7% and 4.0%, respectively. The combination of the three strategies further promotes the performance for 5.2%.

It is also important to note that the performance of general sampling with foreground rebalance weight is fairly close to the specific sampling. Since the teacher queries in specific sampling are well-optimized and concentrate more on the foreground regions, such phenomena validate that the general sampling with randomly initialized queries can

General	Specific	FRW	AP	AP_{50}	AP_{75}
			36.2	56.1	37.9
✓			38.7	59.1	40.7
✓		✓	39.9	60.2	42.5
	✓		40.0	60.5	42.4
	✓	✓	40.2	60.7	42.8
✓	✓	✓	41.4	61.4	44.2

Table 6. Distillation with Different Distillation Points Sampling Strategies. FRW refers to Foreground Rebalance Weight.

Number	10	50	100	300	900
AP	38.6	39.3	39.5	39.9	39.5
AP_{50}	58.3	59.2	59.3	60.2	60.4
AP_{75}	41.0	41.3	41.8	42.5	41.7

Table 7. Distillation with Different Number of General Distillation Points

probe the whole feature maps evenly, and foreground rebalance weight can help the student focus more on the foreground regions. In addition, the further improvement brought by the combination of three strategies indicates that the specific distillation points bring more information from the concentration of teacher model.

Number of General Distillation Points

We investigate the influence of the number of general distillation points by varying the number from 10 to 900. As is shown in Table 7, the improvement is significant when increasing the distillation points from 10 to 50, while gradually saturating when continuing to increase to 300. There is even a slight degradation when further increasing the number. These phenomena validate that the general distillation points can effectively probe the knowledge from teacher’s attention in different regions when sparsely distributed on the feature maps. However, dense distillation points will introduce more noise of background information, and will be harmful for distillation.

Visualization of the Distillation Points

In order to understand what knowledge has been transferred from teacher to student better, we visualize the attention map of different kinds of distillation points. Figure 3 illustrates the attention of the same specific distillation points from teacher model, student model and KD-DETR. It can be seen that the specific distillation points focus more on the features of foreground regions. While the attention of the original student model without distillation is different from the teacher model, KD-DETR can align the concentration of teacher and student models. On the contrary, the general distillation points, as illustrated in Figure 4, concentrate more on the background regions, and objects that are not included in the ground truth annotations, which can pro-

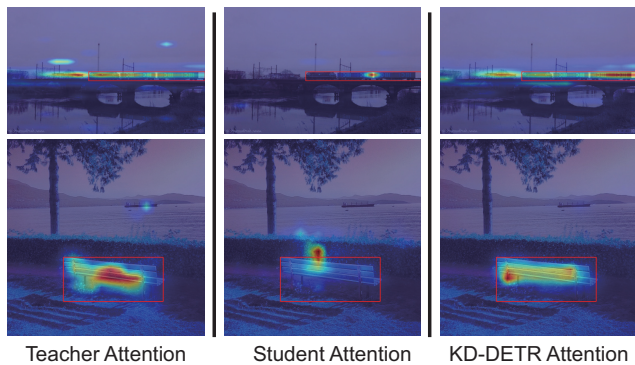


Figure 3. Attention Map Visualization of Specific Distillation Points: Images from left to right are from teacher, original student and student with KD-DETR. The corresponding predicted bounding box of the distillation points are marked with red rectangles.




Figure 4. Attention Map Visualization of General Distillation Points: Teacher model’s attention of general distillation points focuses more on the background regions, providing more information for student to mimic.

vide more additional semantic information to the student. In this way, the combination of general and specific distillation points can provide more comprehensive knowledge.

6. Conclusion

In this paper, we study the compression of DETR with knowledge distillation. We provide thorough experimental and theoretical analysis on the key point in DETR distillation. Based on the analysis, we propose the first general knowledge distillation paradigm for DETR (KD-DETR), together with a comprehensive general-to-specific consistent distillation sampling scheme. We conduct extensive experiments to demonstrate the flexibility, generalization, and extensibility of KD-DETR on various DETR architectures, and for both homogeneous and heterogeneous distillation. For homogeneous distillation, KD-DETR compresses both the scale of backbone and transformer layers, significantly boosting the performance of student model. For heterogeneous distillation, KD-DETR effectively transfers the knowledge from DETR to CNN detector.

References

- [1] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: high quality object detection and instance segmentation. *PAMI*, 43(5):1483–1498, 2019. 3
- [2] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, pages 213–229. Springer, 2020. 1, 3
- [3] Jiahao Chang, Shuo Wang, Hai-Ming Xu, Zehui Chen, Chenhongyi Yang, and Feng Zhao. Detrdistill: A universal knowledge distillation framework for detr-families. In *ICCV*, pages 6898–6908, 2023. 2, 3, 7
- [4] Guobin Chen, Wongun Choi, Xiang Yu, Tony Han, and Manmohan Chandraker. Learning efficient object detection models with knowledge distillation. *NIPS*, 30, 2017. 2, 3, 4
- [5] Xiyang Dai, Yinpeng Chen, Jianwei Yang, Pengchuan Zhang, Lu Yuan, and Lei Zhang. Dynamic detr: End-to-end object detection with dynamic attention. In *ICCV*, pages 2988–2997, 2021. 3
- [6] Xing Dai, Zeren Jiang, Zhao Wu, Yiping Bao, Zhicheng Wang, Si Liu, and Erjin Zhou. General instance distillation for object detection. In *CVPR*, pages 7842–7851, 2021. 3, 7
- [7] Peng Gao, Minghang Zheng, Xiaogang Wang, Jifeng Dai, and Hongsheng Li. Fast convergence of detr with spatially modulated co-attention. In *ICCV*, pages 3621–3630, 2021. 3
- [8] Jianyuan Guo, Kai Han, Yunhe Wang, Han Wu, Xinghao Chen, Chunjing Xu, and Chang Xu. Distilling object detectors via decoupled features. In *CVPR*, pages 2154–2164, 2021. 3
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 6
- [10] Byeongho Heo, Jeesoo Kim, Sangdoon Yun, Hyojin Park, Nojun Kwak, and Jin Young Choi. A comprehensive overhaul of feature distillation. In *ICCV*, pages 1921–1930, 2019. 3
- [11] Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2(7), 2015. 2, 3
- [12] Zijian Kang, Peizhen Zhang, Xiangyu Zhang, Jian Sun, and Nanning Zheng. Instance-conditional knowledge distillation for object detection. *NIPS*, 34:16468–16480, 2021. 3, 6, 1
- [13] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *ACM*, 60(6):84–90, 2017. 6
- [14] Feng Li, Hao Zhang, Shilong Liu, Jian Guo, Lionel M Ni, and Lei Zhang. Dn-detr: Accelerate detr training by introducing query denoising. In *CVPR*, pages 13619–13627, 2022. 1, 2, 7
- [15] Quanquan Li, Shengying Jin, and Junjie Yan. Mimicking very efficient network for object detection. In *CVPR*, pages 6356–6364, 2017. 3
- [16] Xiang Li, Wenhai Wang, Lijun Wu, Shuo Chen, Xiaolin Hu, Jun Li, Jinhui Tang, and Jian Yang. Generalized focal loss: Learning qualified and distributed bounding boxes for dense object detection. *NIPS*, 33:21002–21012, 2020. 3
- [17] 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 *ECCV*, pages 740–755. Springer, 2014. 2, 6
- [18] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, pages 2980–2988, 2017. 3
- [19] Shilong Liu, Feng Li, Hao Zhang, Xiao Yang, Xianbiao Qi, Hang Su, Jun Zhu, and Lei Zhang. Dab-detr: Dynamic anchor boxes are better queries for detr. In *ICLR*, 2021. 1, 3, 7
- [20] Depu Meng, Xiaokang Chen, Zejia Fan, Gang Zeng, Houqiang Li, Yuhui Yuan, Lei Sun, and Jingdong Wang. Conditional detr for fast training convergence. In *ICCV*, pages 3651–3660, 2021. 3
- [21] Jiangmiao Pang, Kai Chen, Jianping Shi, Huajun Feng, Wanli Ouyang, and Dahua Lin. Libra r-cnn: Towards balanced learning for object detection. In *CVPR*, pages 821–830, 2019. 3
- [22] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, pages 779–788, 2016. 3
- [23] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *NIPS*, 28, 2015. 3
- [24] Byungseok Roh, JaeWoong Shin, Wuhyun Shin, and Saehoon Kim. Sparse detr: Efficient end-to-end object detection with learnable sparsity. In *ICLR*, 2021. 2, 3
- [25] 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. 7
- [26] Zhiqing Sun, Shengcao Cao, Yiming Yang, and Kris M Kitani. Rethinking transformer-based set prediction for object detection. In *ICCV*, pages 3611–3620, 2021. 3
- [27] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive representation distillation. In *ICLR*, 2019. 3
- [28] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *ICCV*, pages 9627–9636, 2019. 3
- [29] Tao Wang, Li Yuan, Xiaopeng Zhang, and Jiashi Feng. Distilling object detectors with fine-grained feature imitation. In *CVPR*, pages 4933–4942, 2019. 3, 7
- [30] Tao Wang, Li Yuan, Yunpeng Chen, Jiashi Feng, and Shuicheng Yan. Pnp-detr: Towards efficient visual analysis with transformers. In *ICCV*, pages 4661–4670, 2021. 3
- [31] Yingming Wang, Xiangyu Zhang, Tong Yang, and Jian Sun. Anchor detr: Query design for transformer-based detector. In *AAAI*, pages 2567–2575, 2022. 3
- [32] Xiaopeng Yan, Ziliang Chen, Anni Xu, Xiaoxi Wang, Xiaodan Liang, and Liang Lin. Meta r-cnn: Towards general solver for instance-level low-shot learning. In *CVPR*, pages 9577–9586, 2019. 3
- [33] Chuanguang Yang, Helong Zhou, Zhulin An, Xue Jiang, Yongjun Xu, and Qian Zhang. Cross-image relational knowledge distillation for semantic segmentation. In *CVPR*, pages 12319–12328, 2022. 3
- [34] Zhendong Yang, Zhe Li, Xiaohu Jiang, Yuan Gong, Zehuan Yuan, Danpei Zhao, and Chun Yuan. Focal and global

- knowledge distillation for detectors. In *CVPR*, pages 4643–4652, 2022. 3, 7
- [35] Zhendong Yang, Zhe Li, Mingqi Shao, Dachuan Shi, Zehuan Yuan, and Chun Yuan. Masked generative distillation. In *ECCV*, pages 53–69. Springer, 2022. 3
- [36] Zhuyu Yao, Jiangbo Ai, Boxun Li, and Chi Zhang. Efficient detr: improving end-to-end object detector with dense prior. *arXiv preprint arXiv:2104.01318*, 2021. 2, 3
- [37] Junho Yim, Donggyu Joo, Jihoon Bae, and Junmo Kim. A gift from knowledge distillation: Fast optimization, network minimization and transfer learning. In *CVPR*, pages 4133–4141, 2017. 3
- [38] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. In *ICLR*, 2016. 3
- [39] Gongjie Zhang, Shijian Lu, and Wei Zhang. Cad-net: A context-aware detection network for objects in remote sensing imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 57(12):10015–10024, 2019. 3
- [40] Gongjie Zhang, Zhipeng Luo, Yingchen Yu, Kaiwen Cui, and Shijian Lu. Accelerating detr convergence via semantic-aligned matching. In *CVPR*, pages 949–958, 2022. 3
- [41] Hao Zhang, Feng Li, Shilong Liu, Lei Zhang, Hang Su, Jun Zhu, Lionel Ni, and Heung-Yeung Shum. Dino: Detr with improved denoising anchor boxes for end-to-end object detection. In *ICLR*, 2022. 1, 3, 7
- [42] Borui Zhao, Quan Cui, Renjie Song, Yiyu Qiu, and Jiajun Liang. Decoupled knowledge distillation. In *CVPR*, pages 11953–11962, 2022. 3
- [43] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. In *ICLR*, 2020. 1, 3, 7