

DETR with Additional Global Aggregation for Cross-domain Weakly Supervised Object Detection

Zongheng Tang^{1,2,3} Yifan Sun² Si Liu^{1,3*} Yi Yang⁴

¹Institute of Artificial Intelligence, Beihang University ²Baidu Inc

³Hangzhou Innovation Institute, Beihang University ⁴CCAI, Zhejiang University

tzhhhh123@buaa.edu.cn sunyf15@tsinghua.org.cn liusi@buaa.edu.cn yangyics@zju.edu.cn

Abstract

This paper presents a DETR-based method for cross-domain weakly supervised object detection (CDWSOD), aiming at adapting the detector from source to target domain through weak supervision. We think DETR has strong potential for CDWSOD due to an insight: the encoder and the decoder in DETR are both based on the attention mechanism and are thus capable of aggregating semantics across the entire image. The aggregation results, i.e., image-level predictions, can naturally exploit the weak supervision for domain alignment. Such motivated, we propose DETR with additional Global Aggregation (DETR-GA), a CDWSOD detector that simultaneously makes “instance-level + image-level” predictions and utilizes “strong + weak” supervisions. The key point of DETR-GA is very simple: for the encoder / decoder, we respectively add multiple class queries / a foreground query to aggregate the semantics into image-level predictions. Our query-based aggregation has two advantages. First, in the encoder, the weakly-supervised class queries are capable of roughly locating the corresponding positions and excluding the distraction from non-relevant regions. Second, through our design, the object queries and the foreground query in the decoder share consensus on the class semantics, therefore making the strong and weak supervision mutually benefit each other for domain alignment. Extensive experiments on four popular cross-domain benchmarks show that DETR-GA significantly improves cross-domain detection accuracy (e.g., 29.0% \rightarrow 79.4% mAP on PASCAL VOC \rightarrow Clipart_{a11} dataset) and advances the states of the art.

1. Introduction

The cross-domain problem is a critical challenge for object detection in real-world applications. Concretely, there is usually a domain gap between the training and testing

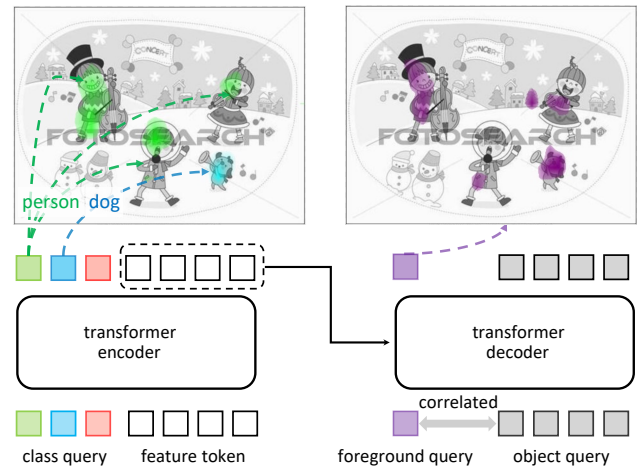


Figure 1. To exploit the weak supervision, DETR-GA aggregates the semantic information across the entire image into image-level predictions. Specifically, DETR-GA adds multiple class queries / a foreground query into the transformer encoder / decoder, respectively. The foreground query is correlated with the object queries but has no position embedding. We visualize the attention score of some queries, e.g., “person” and “dog”. Despite no position supervision, each class query / the foreground query attends to the class-specific / foreground regions for semantic aggregation.

data. This domain gap significantly compromises the detection accuracy when the detector trained on the source domain is directly deployed on a novel target domain. To mitigate the domain gap, existing domain adaptation methods can be categorized into supervised, unsupervised [7, 10, 39], and weakly supervised approaches [17, 21, 33, 58]. Among the three approaches, we are particularly interested in the weakly supervised one because it requires only image-level annotations and achieves a good trade-off between the adaptation effect and the annotation cost. Therefore, this paper challenges the cross-domain weakly supervised object detection (CDWSOD), aiming at adapting the detector from the source to target domain through weak supervision.

We think the DETR-style detector [3, 27, 67] has high potential for solving CDWSOD. In contrast to current CDW-

*Corresponding author: Si Liu

SOD methods dominated by pure convolutional neural network detectors (“CNN detectors”), this paper is the first to explore DETR-style detectors for CDWSOD, to the best of our knowledge. Our optimism for DETR is NOT due to its prevalence or competitive results in generic object detection. In fact, we empirically find the DETR-style detector barely achieves any superiority against CNN detectors for direct cross-domain deployment (Section 4.4). Instead, our motivation is based on the insight, *i.e.*, the DETR-style detector has superiority for combining the strong and weak supervision, which is critical for CDWSOD [17, 21, 58].

Generally, CDWSOD requires using weak (*i.e.*, image-level) supervision on target domain to transfer the knowledge from source domain. Therefore, it is essential to supplement the detector with image-level prediction capability. We argue that this essential can be well accommodated by two basic components in DETR, *i.e.*, the encoder and the decoder. Both the encoder and the decoder are based on the attention mechanism and thus have strong capability to capture long-range dependencies. This long-range modeling capability, by its nature, is favorable for aggregating semantic information to make image-level predictions.

To fully exploit the weak supervision in CDWSOD, this paper proposes DETR with additional Global Aggregation (DETR-GA). DETR-GA adds attention-based global aggregation into DETR so as to make image-level predictions, while simultaneously preserving the original instance-level predictions. Basically, DETR uses multiple object queries in the decoder to probe local regions and gives instance-level predictions. Based on DETR, DETR-GA makes two simple and important changes: for the encoder / decoder, it respectively adds multiple class queries / a foreground query to aggregate semantic information across the entire image. The details are explained below:

1) *The encoder makes image-level prediction through a novel class query mechanism.* Specifically, the encoder adds multiple class queries into its input layer, with each query responsible for an individual class. Each class query probes the entire image to aggregate class-specific information and predicts whether the corresponding class exists in the image. During training, we use the image-level multi-class label to supervise the class query predictions.

Despite NO position supervision, we show these class queries are capable to roughly locate the corresponding position (Fig. 1) and thus exclude the distraction from non-relevant regions. Therefore, our class query mechanism achieves better image-level aggregation effect than the average pooling strategy that is commonly adopted in pure-CNN CDWSOD methods. Empirically, we find this simple component alone brings significant improvement for CDWSOD, *e.g.*, +20.8 mAP on PASCAL VOC \rightarrow Clipart_{test}.

2) *The decoder gives image-level and instance-level predictions simultaneously through correlated object and fore-*

ground queries. To this end, we simply remove the position embedding from an object query and use the remained content embedding as the foreground query. The insight for this design is: in a object query, the position embedding encourages focus on local region [27, 31, 32], while the content embedding is prone to global responses to all potential foreground regions. Therefore, when we remove the position embedding, an object query discards the position bias and becomes a foreground query with global responses (as visualized in Fig. 1). Except this difference (*i.e.*, with or without position embedding), the object queries and the foreground query share all the other elements in the decoder, *e.g.*, the self-attention layer, cross-attention layer. Such correlation encourages them to share consensus on the class semantic and thus benefits the domain alignment along all the classes.

Overall, DETR-GA utilizes the weak supervision on the encoder and decoder to transfer the detection capability from source to target domain. Experimental results show that DETR-GA improves cross-domain detection accuracy by a large margin. Our main contributions can be summarized as follows:

- As the first work to explore DETR-style detector for CDWSOD, this paper reveals that DETR has strong potential for weakly-supervised domain adaptation because its attention mechanism can fully exploit image-level supervision by aggregating semantics across the entire image.
- We propose DETR-GA, a CDWSOD detector that simultaneously makes “instance-level + image-level” predictions and can utilize both “strong + weak” supervision. The key point of DETR-GA is the newly-added class queries / foreground query in the encoder / decoder, which promotes global aggregation for image-level prediction.
- Extensive experiments on four popular cross-domain benchmarks show that DETR-GA significantly improves CDWSOD accuracy and advances the states of the art. For example, on PASCAL VOC \rightarrow Clipart_{all}, DETR-GA improves the baseline from 29.0% to 79.4% .

2. Related Work

Object detection. Object detection is a fundamental task in computer vision and attract much attention for years. The pure CNN-based detectors can be roughly divided into two-stage methods, *e.g.*, Faster RCNN [36], and single-stage methods, *e.g.*, YOLO [35] and FCOS [48]. Although these CNN-based detectors have made significant advancements, they still require numerous hand-designed components, *e.g.*, non-maximum suppression (NMS), and anchor generation. Recently, DETR [3] propose an end-to-end object detection pipeline without the above hand-designed components. Many DETR-style detectors [27, 31, 32, 67] are then proposed to further improve DETR. For example, DAB-DETR [31] proposes to replace the original position embedding with 4d box anchors. DN-DETR [27] proposes

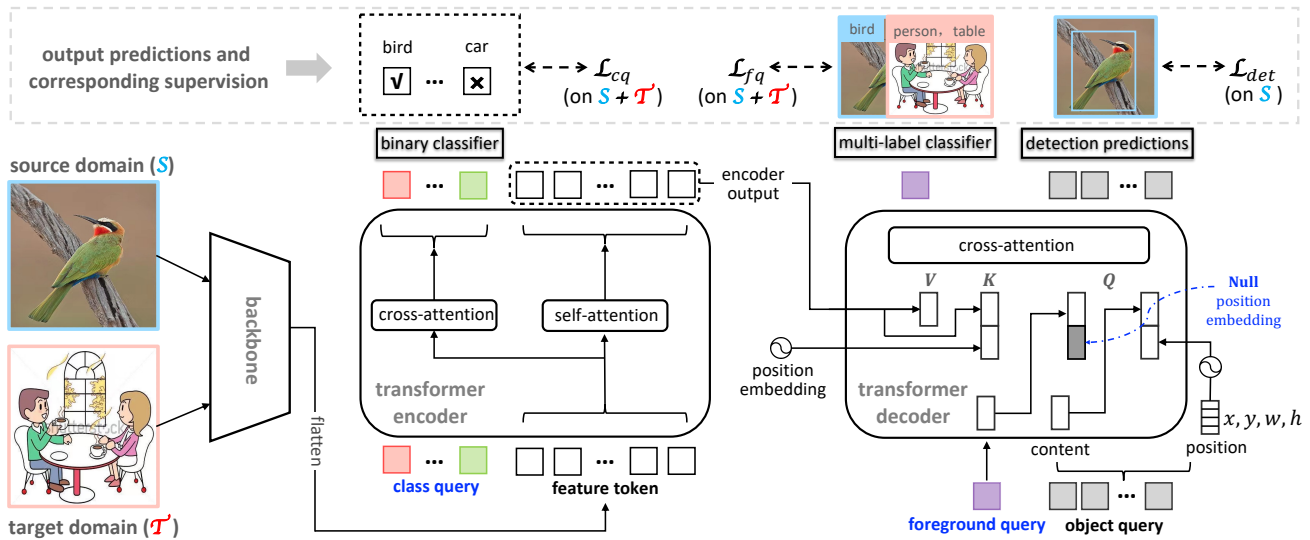


Figure 2. Overview of the proposed DETR-GA. DETR-GA tries to transfer the detection knowledge from the fully-labeled source domain \mathcal{S} to the target domain \mathcal{T} through weak supervision. Based on a DETR-style detector consisting of backbone, encoder and decoder, DETR-GA adds two modules, *i.e.*, **class query** and **foreground query**. In the encoder, multiple class queries aggregate class-specific information from all the feature tokens through cross-attention. In the decoder, the foreground query and the object queries are correlated with each other, because they share most processing except that the foreground query has **null** position embedding. Correspondingly, they make image-level and instance-level predictions, respectively. Overall, we utilize weak supervision / hybrid (weak + strong) supervision for the encoder / decoder, therefore aligning the class semantics between the source and target domain.

a denoising training method to accelerate convergence.

So far as we know, this paper is the first to explore DETR-style detector for CDWSOD. During our preliminary experiments, we find that DETR presents no obvious superiority for direct cross domain deployment, though it usually surpasses the pure-CNN detectors on generic object detection. For example, on PASCAL VOC \rightarrow Clipart_{all}, Faster RCNN and DETR respectively achieve 29.0 and 29.2, which are comparable. However, we still believe that DETR has strong potential for CDWSOD, because its attention mechanism promotes global aggregation and can better exploit weak supervision. Empirically, we show that the proposed DETR-GA achieves significant improvement over direct cross-domain deployment, and surpasses the strongest pure-CNN method by a clear margin(though they achieve similar direct cross-domain deployment accuracy).

Global aggregation in CDWSOD. Aggregating the semantic information into image-level prediction is critical for cross-domain alignment in CDWSOD [17, 58]. Xu [58] uses average pooling for global aggregation on the backbone feature and proposes a softmax-based aggregation to merge all the instance-level predictions from the detection head. ICCM [17] divides each image into semantic clusters and aligns the foreground regions in the target domain with the labeled region in the source domain. In contrast, in DETR-GA, the class and foreground queries utilize attention mechanisms for global aggregation. The query-based global aggregation has an advantage, *i.e.*, the query is

capable of roughly locating the corresponding regions despite no position supervision.

Cross-domain object detection. Cross-domain object detection (CDOD) aims to adapt the detector from the fully-labeled source domain to the label-scarce target domain. Previous works mainly focus on the unsupervised domain adaptation object detection(UDAOD) and related fields [2, 4, 10, 12, 14–16, 20, 30, 41, 42, 45, 46, 50, 52, 60, 62, 66]. The UDAOD methods can be roughly categorized into image translation [19, 25], self-supervision [23, 24, 28, 34, 38], and adversarial training [6–8, 18, 37, 39, 47, 53, 55, 57, 61, 64, 65]. Recently, SFA [53] utilize adversarial training for domain alignment based on the DETR-style detector. Compared with UDAOD, CDWSOD [17, 21, 33] offers more image-level annotations on the target domain. DETR-GA utilizes the additional global aggregation to make image-level domain alignment and achieves better performance.

3. DETR-GA

3.1. Overview

Task definition. Unlike weakly supervised object detection [1, 44, 51, 59] only utilize the weak supervision on the target domain, CDWSOD aims at adapting the detector from the source domain to the target domain through weak supervision. In the source domain, each image has instance-level annotations with object classes and bounding boxes,

while in the target domain, each image has only multi-class labels indicating which classes exist.

DETR Revisit. Since DETR-GA is based on the DETR pipeline, we first give a brief revisit to DETR. DETR consists of a backbone model, a transformer encoder and a transformer decoder. Given an input image, DETR first uses a backbone model (*e.g.*, ResNet-50 [13]) to extract the feature map. Afterward, DETR flattens the feature map into feature tokens and feeds them into the transformer encoder. In the encoder, all the feature tokens are updated through self-attention block by block. The final output states of these feature tokens flow into the transformer decoder, along with multiple (*i.e.*, 300) learnable object queries. In each decoder block, object queries first perform self-attention between themselves and then cross-attention to the feature tokens. For simplicity, Fig. 2 only presents the cross-attention part. Please refer to the supplementary for a complete view of the DETR decoder. After multiple blocks of self-attention and cross-attention, the output states of the object queries generate instance-level detection predictions (*i.e.*, the object class and bounding box position), which are supervised through a detection loss \mathcal{L}_{det} [3].

Overall structure of DETR-GA. As shown in Figure 2, DETR-GA takes the mixture of source and target images as its input. To fully exploit the weak supervision, DETR-GA supplements the original detector with image-level predictions in both the encoder and decoder. Given the image-level prediction capability of the encoder and decoder, DETR-GA facilitates joint training on the source and target domain. The encoder aligns the class semantics across the source and target domain through weak supervision. Afterward, the decoder further delivers the object detection capability from source to target through hybrid supervision (*i.e.*, strong supervision on the source and weak supervision on the target).

Specifically, for the encoder, we add multiple class queries into its input layer, with each query responsible for an individual class. Each class query searches the entire image and aggregates class-specific information through cross-attention. The aggregated outputs predict whether the corresponding classes exist in the image and are supervised with a binary cross-entropy loss \mathcal{L}_{cq} . For the decoder, we add a foreground query into its input layer. The foreground shares the content embedding with all the object queries and has NO position embedding. Therefore, it can attend to all the potential foreground without positional bias. Its aggregation output predicts which classes exist in the image and uses another multi-class binary cross-entropy loss \mathcal{L}_{fq} for supervision. In the encoder, the aggregated semantic information is naturally used for image-level prediction and can be trained through weak supervision. Importantly, in the decoder, the foreground query is correlated with the object queries, making the image-level and instance-level predic-

tions mutually benefit each other.

Given the additional image-level predictions from the encoder and decoder, DETR-GA is capable to combine weak and strong supervision. Specifically, the encoder receives weak supervision from “source+target” domains (\mathcal{L}_{cq}). The decoder receives weak supervision from “source+target” domains (\mathcal{L}_{fq}), and meanwhile receives strong supervision from the target domain (\mathcal{L}_{det}). Moreover, we further add some common practices [53, 58] on the backbone, *i.e.*, using an adversarial loss \mathcal{L}_{dc} for domain-level alignment and an image-level classification loss \mathcal{L}_{bc} .

During training, we combine all the described loss functions to optimize DETR-GA in an end-to-end manner:

$$\mathcal{L} = \underbrace{\lambda_{bc}\mathcal{L}_{bc} + \lambda_{dc}\mathcal{L}_{dc}}_{\text{backbone}} + \underbrace{\lambda_{cq}\mathcal{L}_{cq}}_{\text{encoder}} + \underbrace{\lambda_{fq}\mathcal{L}_{fq} + \mathcal{L}_{det}}_{\text{decoder}}. \quad (1)$$

During inference, DETR-GA does not need the image-level predictions and thus resumes the standard DETR [3] pipeline without extra computational cost.

3.2. Encoder with Class Queries

We add C (the number of total classes) class queries $\mathbf{Q} = \{\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_C\}$ into the encoder. Each class query \mathbf{q}_i is a learnable vector (*i.e.*, $\mathbf{q}_i \in \mathbb{R}^d$) and corresponds to the i -th class. We will use the output state of \mathbf{q}_i as a binary indicator to identify the presence of the i -th class.

Specifically, let us denote the collection of all the feature tokens as \mathbf{T} , and their position embedding as \mathbf{P}_T . In each transformer block of the encoder, all the class queries \mathbf{Q} are updated through cross-attention to the feature tokens \mathbf{T} , in parallel to the self-attention update of the feature tokens. Formally, in the k -th ($k = 1, 2, \dots, K$) transformer block of the encoder, the self-attention and the cross-attention update are formulated as:

$$\begin{aligned} \mathbf{T}^k &= \mathbf{T}^{k-1} + \text{attn}(\mathbf{T}^{k-1} + \mathbf{P}_T, \mathbf{T}^{k-1} + \mathbf{P}_T, \mathbf{T}^{k-1}) \\ \mathbf{Q}^k &= \mathbf{Q}^{k-1} + \text{attn}(\mathbf{Q}^{k-1}, \mathbf{T}^{k-1} + \mathbf{P}_T, \mathbf{T}^{k-1}), \end{aligned} \quad (2)$$

where $\text{attn}(\text{query}, \text{key}, \text{value})$ denotes the attention layer [49], the superscript k indicates the k -th transformer block. There are a layer norm and a feed-forward network (FFN) [49] following the attention layer and are omitted here for brevity.

Given the output states of the class query, *i.e.*, $\mathbf{Q}^K = \{\mathbf{q}_1^K, \mathbf{q}_2^K, \dots, \mathbf{q}_C^K\}$ (K is the transformer depth), the encoder uses a binary classifier to predict the presence of the i -th class by:

$$p_i = \text{sigmoid}\{\mathbf{q}_i^K \cdot \mathbf{w}_i\}, \quad (3)$$

where \mathbf{w}_i is the weight vector for the i -th class in the classification head, “ \cdot ” is the inner product operation. We note that Eqn. 3 has a minor difference against the canonical

multi-class predictor, which uses a single input to simultaneously predict the presence of multiple classes. In contrast, Eqn. 3 has multiple inputs, *i.e.*, each \mathbf{q}_i^K is responsible for an individual class.

The above classifier is supervised by a standard binary cross-entropy (BCE) loss \mathcal{L}_{cq} , which is formulated as:

$$\mathcal{L}_{cq} = \frac{1}{C} \sum_{i=1}^C [y_i \log p_i + (1 - y_i)(1 - \log p_i)], \quad (4)$$

where y_i is set to 1 if class c_i exists, otherwise 0.

We provide an intuitive review of the domain alignment procedure in the encoder. Through the cross-attention in Eqn. 2, each class query absorbs information from all the feature tokens, therefore facilitating global aggregation. The aggregation results \mathbf{Q}^M are used to predict the presence of each class through Eqn. 3. The corresponding weak supervision (Eqn. 4) is enforced on both the source and target domain and thus aligns each class across two domains.

An important advantage of our query-based aggregation is: though there is no position supervision, each class query is capable of roughly locating its object regions, therefore excluding the distraction from non-relevant regions, as shown in Figure 1 and in the supplementary. In contrast, CNN CDWSOD methods usually use global average pooling [58] for aggregation. It inevitably introduces interference among different classes and thus compromises the domain alignment effect. In an ablation study (Section 4.4), we show that our query-based aggregation is significantly better (*e.g.*, 9.4 mAP higher on PASCAL VOC \rightarrow Clipart_{test}) than the average pooling aggregation.

3.3. Decoder with Foreground and Object Queries

In DETR-GA, the transformer decoder has three types of input, *i.e.*, feature tokens (\mathbf{T}^K output from the encoder), a newly-added foreground query $\mathbf{f} \in \mathbb{R}^d$ (which is a vector), and N object queries. The object queries consist of the content embedding $\mathbf{O} = \{\mathbf{o}_1, \mathbf{o}_2, \dots, \mathbf{o}_N\}$ and the position embedding $\mathbf{P}_O \in \mathbb{R}^{N \times d}$.

In the l -th ($l = 1, 2, \dots, L$) transformer block of the decoder, the foreground query and the object queries simultaneously undergo a self-attention update and a following cross-attention update.

The self-attention update is formulated as:

$$[\tilde{\mathbf{O}}^{l-1}, \tilde{\mathbf{f}}^{l-1}] = \text{self}([\mathbf{O}^{l-1} + \mathbf{P}_O^{l-1}, \mathbf{f}^{l-1}]), \quad (5)$$

where “[]” is the concatenation operation, *self* is the self-attention layer. We adopt the same self-attention layer as in DN-DETR [27] and provide a detailed description in the supplementary.

The following cross-attention update is formulated as:

$$\begin{aligned} \mathbf{O}^l &= \mathbf{O}^{l-1} + \text{attn}(H(\tilde{\mathbf{O}}^{l-1}, \mathbf{P}_O^{l-1}), H(\mathbf{T}^M, \mathbf{P}_T), \mathbf{T}^M), \\ \mathbf{f}^l &= \mathbf{f}^{l-1} + \text{attn}(H(\tilde{\mathbf{f}}^{l-1}, \mathbf{0}), H(\mathbf{T}^M, \mathbf{P}_T), \mathbf{T}^M), \end{aligned} \quad (6)$$

where *attn*(query, key, value) denotes the attention layer, *H* is the operation for entangling the content and position embedding (*e.g.*, the concatenation operation as in DAB-DETR [31] and DN-DETR [27]). The foreground query has a null position embedding, as illustrated in Fig. 2 and thus may be viewed as removing the impact of position embedding. The position embedding update is the same as in DN-DETR and is omitted here.

The above self-attention (Eqn. 5) and cross-attention (Eqn. 6) are almost the same as in DAB-DETR and DN-DETR, except for adding a foreground query with no position embedding. Although we adopt the DN-DETR as the direct baseline, we note that Eqn. 5 and Eqn. 6 jointly depict a general DETR decoder and are compatible with other DETR-style detectors. For example, *H* for entangling the content embedding and position embedding can be modified to “+” operation, as in the original DETR.

Albeit simple, there are two good characteristics, *i.e.*, removing the position embedding for global attention and correlating the foreground and object queries:

- The foreground query has no position embedding and thus favors global aggregation. In DETR-style detectors, the position embedding encourages the object query to spatially focus on the local region. In fact, a keynote of recent state-of-the-art DETR progress [27, 31] is to improve the position embedding for localizing the objects. Removing the position embedding removes the potential position bias and encourages \mathbf{f} to consider only the semantic content for information aggregation. Empirically, we find that adding position embedding to the foreground query compromises DETR-GA (Section 4.4).

- The foreground query is strongly correlated with the object queries because they share almost all the update elements (*e.g.*, the self-attention layer, the cross-attention layer, and the same feature tokens). Intuitively, in Eqn. 5 and Eqn. 6, the foreground query may be viewed as a specific object query with ZERO position embedding. This correlation encourages the foreground query and object queries to share consensus on the class semantics and thus benefits class-wise alignment. Empirically, we find removing this correlation compromises DETR-GA (Section 4.4).

Consequently, the encoder simultaneously gives multiple instance-level predictions (from the object query output) and an image-level prediction (from the foreground query output). These predictions are respectively supervised through \mathcal{L}_{det} and \mathcal{L}_{fq} , as shown in Fig. 2. \mathcal{L}_{fq} is a standard BCE loss as in Eqn. 4.

4. Experiments

4.1. Dataset

Following existing works [5, 17, 21, 39, 58], we use four datasets, *i.e.*, PASCAL VOC (VOC) [11], Clipart, Water-

Method	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mean
<i>WS Group</i>																					
PCL [43]	3.4	10.6	2.3	1.7	5.2	3.4	23.3	1.2	5.6	0.4	7.8	3.7	5.6	0.3	24.5	19.7	11.9	3.6	9.2	25.4	8.4
EDRN [40]	2.7	13.5	1.2	4.2	1.8	10.3	25.7	0.4	8.4	0.3	3.2	2.7	1.1	0.7	29.4	17.2	5.2	1.6	2.9	19.1	7.6
<i>UDA Group</i>																					
SWDA [39]	26.2	48.5	32.6	33.7	38.5	54.3	37.1	18.6	34.8	58.3	17.0	12.5	33.8	65.5	61.6	52.0	9.3	24.9	54.1	49.1	38.1
HTD [5]	33.6	58.9	34.0	23.4	45.6	57.0	39.8	12.0	39.7	51.3	21.1	20.1	39.1	72.8	63.0	43.1	19.3	30.1	50.2	51.8	40.3
IIOD [54]	41.5	52.7	34.5	28.1	43.7	58.5	41.8	15.3	40.1	54.4	26.7	28.5	37.7	75.4	63.7	48.7	16.5	30.8	54.5	48.7	42.1
I ³ Net [6]	30.0	67.0	32.5	21.8	29.2	62.5	41.3	11.6	37.1	39.4	27.4	19.3	25.0	67.4	55.2	42.9	19.5	36.2	50.7	39.3	37.8
DBGL [4]	28.5	52.3	34.3	32.8	38.6	66.4	38.2	25.3	39.9	47.4	23.9	17.9	38.9	78.3	61.2	51.7	26.2	28.9	56.8	44.5	41.6
TIA [63]	42.2	66.0	36.9	37.3	43.7	71.8	49.7	18.2	44.9	58.9	18.2	29.1	40.7	87.8	67.4	49.7	27.4	27.8	57.1	50.6	46.3
<i>CDWS Group</i>																					
DT+PL [21]	50.1	75.0	37.0	38.7	58.1	83.4	50.1	38.0	55.2	67.3	51.1	34.8	49.8	89.9	60.2	63.4	28.8	42.4	62.6	70.9	55.3
ICCM [17]	39.8	66.7	37.2	42.5	43.3	48.1	48.1	21.3	46.5	73.0	29.0	29.8	57.3	78.6	67.8	48.7	46.3	19.3	42.8	48.5	46.7
H ² FA [58]	58.1	73.0	56.8	50.4	61.2	98.6	69.5	57.8	66.4	77.1	56.1	84.1	64.3	100.0	78.1	78.2	43.5	65.4	77.3	79.7	69.8
DETR-GA	86.7	85.6	64.2	66.4	67.4	100.0	75.5	70.0	72.2	81.4	83.8	87.0	75.2	100.0	82.8	75.5	67.8	76.1	86.2	85.1	79.4

Table 1. Mean AP performance (%) on Clipart_{all}.

Method	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mean
<i>WS Group</i>																					
WSDDN [1]	1.6	3.6	0.6	2.3	0.1	11.7	4.5	0.0	3.2	0.1	2.8	2.3	0.9	0.1	14.4	16.0	4.5	0.7	1.2	18.3	4.4
CLNet [22]	3.2	22.3	2.2	0.7	4.6	4.8	17.5	0.2	4.8	1.6	6.4	0.6	4.7	0.6	12.5	13.1	14.1	4.1	8.0	29.7	7.8
<i>UDA Group</i>																					
DM [26]	28.5	63.2	24.5	42.4	47.9	43.1	37.5	9.1	47.0	46.7	26.8	24.9	48.1	78.7	63.0	45.0	21.3	36.1	52.3	53.4	41.8
ATF [16]	41.9	67.0	27.4	36.4	41.0	48.5	42.0	13.1	39.2	75.1	33.4	7.9	41.2	56.2	61.4	50.6	42.0	25.0	53.1	39.1	42.1
UMT [10]	39.6	59.1	32.4	35.0	45.1	61.9	48.4	7.5	46.0	67.6	21.4	29.5	48.2	75.9	70.5	56.7	25.9	28.9	39.4	43.6	44.1
AT [29]	33.8	60.9	38.6	49.4	52.4	53.9	56.7	7.5	52.8	63.5	34.0	25.0	62.2	72.1	77.2	57.7	27.2	52.0	55.7	54.1	49.3
<i>CDWS Group</i>																					
DT+PL [21]	51.6	84.0	30.0	41.1	52.3	82.0	50.2	19.0	51.8	58.3	41.3	14.6	47.0	86.2	61.9	58.6	24.9	22.5	47.4	52.8	48.9
PLGE [33]	43.4	52.5	29.4	40.1	30.4	71.9	54.9	3.6	52.4	73.8	53.5	24.0	54.8	89.1	65.1	40.5	32.3	33.8	45.4	61.0	47.6
H ² FA [58]	38.5	70.6	38.9	47.4	59.6	83.5	47.0	29.3	51.5	76.3	44.4	48.1	47.3	79.2	75.7	54.4	53.9	32.0	56.6	51.1	55.3
DETR-GA	64.1	63.0	42.7	52.5	56.6	82.8	57.5	36.3	56.6	67.1	48.2	53.5	41.9	84.0	78.0	62.3	48.9	48.4	75.2	64.4	59.2

Table 2. Mean AP performance (%) on Clipart_{test}.

Training Data	Source Domain	Target Domain		
	BBox	BBox	Tag	None
WSOD			✓	
UDAOD	✓			✓
CDWSOD	✓		✓	

Table 3. The annotations type of different tasks. For example, WSOD methods use the target domain data with image-level tag annotations for training.

color and Comic [21] for method evaluation.

The general object detection dataset VOC is used as the source domain with instance-level annotations. Specifically, we use the `trainval` split of VOC 0712 as the source-domain training data, which provides ~ 16.5 k real-world images of 20 object categories.

The other three artistic painting datasets are used as the target domains with image-level annotations. Clipart has a `train` split and a `test` split, both of which contain 500 images of 20 object categories. There are two popular modes on this dataset: Clipart_{all} merges the two splits for both training and testing [5, 17, 39], while Clipart_{test} uses the former split for training and the latter split for testing [10, 21, 33]. Both Watercolor and Comic contain 2k images of 6 classes. We use 1k images from the `train` split for training and 1k images from the `test` split for testing.

4.2. Implementation details

Baseline. We choose the DETR version of DN-DETR [27], which improves the convergence speed and accuracy

Method	bike	bird	car	cat	dog	person	mean
<i>WS Group</i>							
PCL [43]	1.2	0.4	8.9	2.9	2.3	15.6	5.2
EDRN [40]	1.6	0.5	13.2	7.2	2.5	13.2	6.4
<i>UDA Group</i>							
SWDA [39]	30.3	19.6	28.8	15.2	24.9	46.9	27.6
HTD [5]	35.4	14.8	26.6	13.7	26.9	40.0	26.2
MCAR [64]	47.9	20.5	37.4	20.6	24.5	50.2	33.5
I ³ Net [6]	47.5	19.9	33.2	11.4	19.4	49.1	30.1
DBGL [4]	35.6	20.3	33.9	16.4	26.6	45.3	29.7
<i>CDWS Group</i>							
DT+PL [21]	53.0	23.7	34.4	27.4	27.2	44.0	35.0
ICCM [17]	50.6	23.3	35.4	32.3	33.8	47.1	37.1
PLGE [33]	55.0	21.2	40.0	35.1	37.9	60.9	41.7
H ² FA [58]	55.3	26.6	45.9	38.1	45.6	66.8	46.4
DETR-GA	57.4	36.1	39.9	42.0	49.4	66.0	48.5

Table 4. Mean AP performance (%) on Comic.

of the original DETR [3], as our baseline. We adopt the evaluation metrics in Detectron2 [56] for fair comparisons with previous works [21, 58]. We use ImageNet [9] pre-trained ResNet-101 [13] as the backbone in all experiments unless specified. The loss weights λ_{cq} , λ_{fq} , λ_{bc} and λ_{dc} are empirically set to 10, 10, 1 and 0.5, respectively. Other hyper-parameters are the same as the default setup in DN-DETR.

Training. All models are trained on 4 GPUs with AdamW optimizer. We set the weight decay to 1×10^{-4} and the initial learning rate to 1×10^{-5} . The total batch size is 16, with 8 images per domain. The training has two steps. We first train the source-only model on the source

Method	bike	bird	car	cat	dog	person	mean
<i>WS Group</i>							
PCL [43]	6.7	28.8	20.2	9.5	5.4	27.4	16.3
EDRN [40]	5.2	29.3	15.3	1.4	0.9	34.9	14.5
<i>UDA Group</i>							
SWDA [39]	82.3	55.9	46.5	32.7	35.5	66.7	53.3
HTD [5]	69.2	49.5	49.5	34.9	30.8	61.2	49.2
ATF [16]	78.8	59.9	47.9	41.0	34.8	66.9	54.9
MCAR [64]	87.9	52.1	51.8	41.6	33.8	68.8	56.0
IIOD [54]	95.8	54.3	48.3	42.4	35.1	65.8	56.9
UMT [10]	88.2	55.3	51.7	39.8	43.6	69.9	58.1
I ³ Net [6]	81.1	49.3	46.2	35.0	31.9	65.7	51.5
VDD [55]	90.0	56.6	49.2	39.5	38.8	65.3	56.6
DBGL [4]	83.1	49.3	50.6	39.8	38.7	51.3	53.8
AT [29]	93.6	56.1	58.9	37.3	39.6	73.8	59.9
<i>CDWS Group</i>							
DT+PL [21]	81.0	49.5	39.5	32.3	28.4	62.4	48.8
ICCM [17]	86.6	64.2	52.6	32.4	41.2	67.4	57.4
PLGE [33]	73.7	56.1	50.6	42.5	41.8	74.6	56.5
H ² FA [58]	88.6	52.4	53.6	46.4	44.5	73.8	59.9
DETR-GA	80.5	59.5	55.4	48.9	43.4	72.0	60.0

Table 5. Mean AP performance (%) on Watercolor.

domain dataset. This step takes 50 epochs, and we decrease the learning rate by a factor of 10 after 40 epochs. Then, the source-only model is fine-tuned on both source and target domain data. For the Clipart_{all} datasets, the training takes 30k iterations with a learning rate decay by a factor of 10 at the 20k iterations. For the other datasets, we train 20k iterations with the learning rate divided by 10 at the 10k iterations.

4.3. The effectiveness of DETR-GA

Comparison with the state-of-the-art methods. We compare the proposed DETR-GA with previous state-of-the-art methods. According to the training data and annotations, we classify all the methods into weakly supervised (“WSOD”), unsupervised domain adaptive (“UDAOD”), and CDWSOD methods, as shown in Table 3.

The results on Clipart_{all}, Clipart_{test}, Watercolor and Comic are summarized in Table 1, Table 2, Table 4 and Table 5, respectively. We observe that on all these four datasets, the proposed DETR-GA achieves superior CDWSOD accuracy. For example, on Clipart_{all}, DETR-GA achieves 79.4% mAP, surpassing the strongest competitor (H²FA) by +9.6% mAP. In this paper, we report 79.4%, 59.2%, 60.0% and 48.5% mAP for CDWSOD (with VOC as the source domain) on Clipart_{all}, Clipart_{test}, Watercolor and Comic, respectively. These results outperform all the competing methods (sometimes by a large margin, *e.g.*, on Clipart_{all}), setting new states of the art.

DETR-GA makes good utilization of the weak supervision. Since the DETR-style detectors usually perform favorably against the pure-CNN detectors on the generic object detection tasks, a natural question is: whether the superiority of DETR-GA is because DETR itself has higher base performance than the pure-CNN detectors.

To figure out this question, we investigate the improve-

Based	Method	Clipart _{all}	Clipart _{test}	Watercolor	Comic
CNN	source-only	29.2	29.5	41.4	19.9
	DT+PL [21]	55.3 (+26.1)	48.9 (+19.4)	48.8 (+7.4)	35.0 (+15.1)
	H ² FA [58]	69.8 (+40.6)	55.3 (+25.8)	59.9 (+18.5)	46.4 (+26.5)
DETR	source-only	29.0	30.6	42.2	20.4
	DETR-GA	79.4 (+50.4)	59.2 (+28.6)	60.0 (+17.8)	48.5 (+28.1)
	oracle	—	64.7	60.0	54.7

Table 6. While CNN and DETR detectors perform comparably for direct cross-domain deployment, DETR-GA brings larger improvement and achieves higher CDWSOD accuracy.

ment over direct cross-domain deployment. In Table 6, the “source-only” row indicates the model is trained on the source domain and directly deployed on the target domain. For the pure-CNN detectors, we choose two methods, “DT+PL” [21] and H²FA [58]. “DT+PL is arguably the earliest and the most popular CDWSOD method, while H²FA is the most recent state-of-the-art method. For fair comparison, we report their results based on the ResNet-101 backbone. From Table 6, we draw two observations as below:

First, under the direct cross-domain deployment, the achieved results of DETR-style detector and pure-CN detector are very close. On Clipart_{test}, Watercolor and Comic, DETR (DN-DETR, in particular) is slightly higher than Faster RCNN by 1.1%, 0.8% and 0.5% mAP, respectively, while on Clipart_{all}, DETR is -0.2% mAP lower. We thus infer that DETR does not always promise superiority on the cross-domain object detection task.

Second, DETR-GA achieves larger improvement for CDWSOD. On these datasets, weakly-supervised domain adaptation brings 50.4%, 28.6%, 17.8% and 28.1 mAP improvement for DETR, which are larger than (or comparable to) the improvement on the pure-CNN detector. Therefore, though the direct cross-domain deployment results are close, the proposed DETR-GA achieves higher CDWSOD accuracy than the pure-CNN detectors. It validates our key argument, *i.e.*, DETR can better exploit image-level supervision for CDWSOD.

4.4. Ablation study

Table 7 investigates the major components of DETR-GA through ablation on four benchmarks. We use CQ and FQ to indicate the class query and foreground query for brevity. The baseline adopts only common practices [53, 58] for aligning the backbone features. We draw some important observations below.

DETR-GA gains individual benefit from each component, particularly from the proposed CQ and FQ. Comparing Lines (c)-(d) against the source-only model in Line (a), we observe that each component alone brings noticeable improvement. Moreover, the CQ and FQ make the most prominent improvement. For example, on Clipart_{test}, using the CQ and FQ improve the accuracy by +20.8% and

	baseline	CQ	FQ	Comic	Clipart _{test}
(a)				20.4	30.6
(b)	✓			43.3 (+22.9)	49.3 (+18.7)
(c)		✓		42.5 (+22.1)	51.4 (+20.8)
(d)			✓	44.3 (+23.9)	53.6 (+23.0)
(e)	✓	✓		47.7 (+27.3)	56.6 (+26.0)
(f)		✓	✓	46.1 (+25.7)	54.4 (+23.8)
(g)	✓		✓	47.8 (+27.4)	58.1 (+27.5)
(h)	✓	✓	✓	48.5 (+28.1)	59.2 (+28.6)

Table 7. Effectiveness of different modules in DETR-GA, where mean AP performance (%) over all classes is reported. We use CQ and FQ to indicate the class query and foreground query.

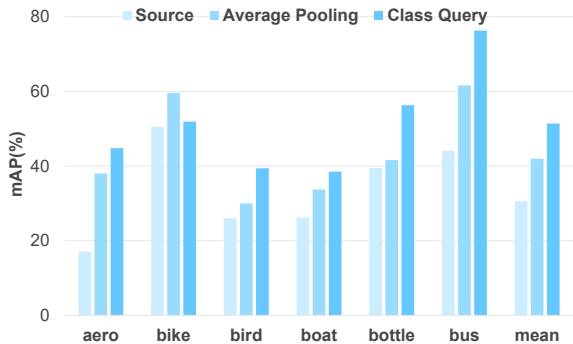


Figure 3. Mean AP (%) performance on Clipart_{test}. We compare the class query with average pooling on the encoder. We only show the first 6 classes and average accuracy for brevity.

Foreground Query	Position Embedding	Weight Sharing	Map
✓			51.5
✓	✓		51.6
✓		✓	53.6

Table 8. Effectiveness of different settings for foreground query. The mean AP performance (%) on Clipart_{test} is reported.

+23.0% mAP, respectively.

Combining these components achieves complementary benefits. For example, combining CQ and FQ (Line (f)) is higher than using only CQ (FQ) by +5.2 (0.8) mAP. Based on “CQ + FQ”, adding the common practices on the backbone further brings 0.9 and 4.8 mAP on Comic and Clipart_{test}, respectively.

Investigation on the class query. We compare two global aggregation methods for the encoder, *i.e.*, the proposed class query and the popular average pooling. The results are shown in Figure 3. It is observed that using average pooling achieves some improvement but is significantly inferior to the proposed class query. The reason is that the class query has some capability to exclude the distraction from non-relevant regions (as shown in Figure 1) and thus facilitates better global aggregation.

Investigation on the foreground query. The foreground query is correlated with the object queries but has no position embedding. We validate these two designs in Table 8. We observe that: 1) canceling the correlation between foreground query and object queries (the 1st row) decreases the mAP by 2.1, and 2) using position embedding




Figure 4. The attention map of (a) the foreground query without position embedding, and (b) the foreground query with position embedding. Adding position embedding compromises the global aggregation capability of the foreground query.

decreases the mAP by 2.0.

Moreover, we visualize the attention map of the foreground query in Figure 4. From Figure 4 (a), it is observed that the foreground query roughly attends to all the potential foreground. However, if we add position embedding to the foreground query (and still maintain the image-level supervision), the foreground query is prone to local focus and misses some important regions, as shown in Figure 4 (b).

5. Conclusion

In this work, we propose DETR with additional Global Aggregation (DETR-GA) for cross-domain weakly supervised object detection (CDWSOD). The main idea of DETR-GA is very straightforward: we add multiple class queries for the encoder and a foreground query for the decoder, respectively, to aggregate the semantics into image-level predictions. The class queries in the encoder are able to aggregate the global semantics corresponding to the classes. And the foreground query in the decoder is correlated with the object queries, thus combining the strong and weak supervision spontaneously benefits domain alignment. Extensive experiments show the superiority of DETR-GA over the competing state-of-the-art.

Limitations. Deformable DETR is another popular DETR-style detector, which replaces the origin attention in DETR [3] with deformable attention. Although the current class query is compatible to the Deformable DETR encoder (in the supplementary material), we think exploring deformable-attention-based class query might fit it better.

Acknowledgements. This work was supported by the Fundamental Research Funds for the Central Universities (No. 226-2022-00051), National Key R&D Program of China (2022ZD0115502), National Natural Science Foundation of China (NO. 62122010), CAAI-Huawei MindSpore Open Fund, Zhejiang Provincial Natural Science Foundation of China under Grant No.LDT23F02022F02, Key Research and Development Program of Zhejiang Province under Grant 2022C01082.

References

- [1] Hakan Bilen and Andrea Vedaldi. Weakly supervised deep detection networks. In *CVPR*, 2016. 3, 6
- [2] Qi Cai, Yingwei Pan, Chong-Wah Ngo, Xinmei Tian, Lingyu Duan, and Ting Yao. Exploring object relation in mean teacher for cross-domain detection. In *CVPR*, 2019. 3
- [3] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. 1, 2, 4, 6, 8
- [4] Chaoqi Chen, Jiongcheng Li, Zebiao Zheng, Yue Huang, Xinghao Ding, and Yizhou Yu. Dual bipartite graph learning: A general approach for domain adaptive object detection. In *ICCV*, 2021. 3, 6, 7
- [5] Chaoqi Chen, Zebiao Zheng, Xinghao Ding, Yue Huang, and Qi Dou. Harmonizing transferability and discriminability for adapting object detectors. In *CVPR*, 2020. 5, 6, 7
- [6] Chaoqi Chen, Zebiao Zheng, Yue Huang, Xinghao Ding, and Yizhou Yu. I³Net: Implicit instance-invariant network for adapting one-stage object detectors. In *CVPR*, 2021. 3, 6, 7
- [7] Yuhua Chen, Wen Li, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Domain adaptive Faster R-CNN for object detection in the wild. In *CVPR*, 2018. 1, 3
- [8] Yuhua Chen, Haoran Wang, Wen Li, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Scale-aware domain adaptive Faster R-CNN. *IJCV*, 2021. 3
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, 2009. 6
- [10] Jinhong Deng, Wen Li, Yuhua Chen, and Lixin Duan. Unbiased mean teacher for cross-domain object detection. In *CVPR*, 2021. 1, 3, 6, 7
- [11] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (VOC) challenge. *IJCV*, 2010. 5
- [12] Qiqi Gu, Qianyu Zhou, Minghao Xu, Zhengyang Feng, Guangliang Cheng, Xuequan Lu, Jianping Shi, and Lizhuang Ma. PIT: Position-invariant transform for cross-FoV domain adaptation. In *ICCV*, 2021. 3
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 4, 6
- [14] Mengzhe He, Yali Wang, Jiayi Wu, Yiru Wang, Hanqing Li, Bo Li, Weihao Gan, Wei Wu, and Yu Qiao. Cross domain object detection by target-perceived dual branch distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9570–9580, 2022. 3
- [15] Zhenwei He and Lei Zhang. Multi-adversarial Faster-RCNN for unrestricted object detection. In *ICCV*, 2019. 3
- [16] Zhenwei He and Lei Zhang. Domain adaptive object detection via asymmetric tri-way Faster-RCNN. In *ECCV*, 2020. 3, 6, 7
- [17] Luwei Hou, Yu Zhang, Kui Fu, and Jia Li. Informative and consistent correspondence mining for cross-domain weakly supervised object detection. In *CVPR*, 2021. 1, 2, 3, 5, 6, 7
- [18] Cheng-Chun Hsu, Yi-Hsuan Tsai, Yen-Yu Lin, and Ming-Hsuan Yang. Every pixel matters: Center-aware feature alignment for domain adaptive object detector. In *ECCV*, 2020. 3
- [19] Han-Kai Hsu, Wei-Chih Hung, Hung-Yu Tseng, Chun-Han Yao, Yi-Hsuan Tsai, Maneesh Singh, and Ming-Hsuan Yang. Progressive domain adaptation for object detection. *computer vision and pattern recognition*, 2019. 3
- [20] Zhengdong Hu, Yifan Sun, and Yi Yang. Suppressing the heterogeneity: A strong feature extractor for few-shot segmentation. In *The Eleventh International Conference on Learning Representations*, 2023. 3
- [21] Naoto Inoue, Ryosuke Furuta, Toshihiko Yamasaki, and Kiyoharu Aizawa. Cross-domain weakly-supervised object detection through progressive domain adaptation. In *CVPR*, 2018. 1, 2, 3, 5, 6, 7
- [22] Vadim Kantorov, Maxime Oquab, Minsu Cho, and Ivan Laptev. Contextlocnet: Context-aware deep network models for weakly supervised localization. In *ECCV*, 2016. 6
- [23] Mehran Khodabandeh, Arash Vahdat, Mani Ranjbar, and William G Macready. A robust learning approach to domain adaptive object detection. In *ICCV*, 2019. 3
- [24] Seunghyeon Kim, Jaehoon Choi, Taekyung Kim, and Changick Kim. Self-training and adversarial background regularization for unsupervised domain adaptive one-stage object detection. In *ICCV*, 2019. 3
- [25] Taekyung Kim, Minki Jeong, Seunghyeon Kim, Seokeon Choi, and Changick Kim. Diversify and match: A domain adaptive representation learning paradigm for object detection. *computer vision and pattern recognition*, 2019. 3
- [26] Taekyung Kim, Minki Jeong, Seunghyeon Kim, Seokeon Choi, and Changick Kim. Diversify and match: A domain adaptive representation learning paradigm for object detection. In *CVPR*, 2019. 6
- [27] Feng Li, Hao Zhang, Shilong Liu, Jian Guo, Lionel M Ni, and Lei Zhang. Dn-detr: Accelerate detr training by introducing query denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13619–13627, 2022. 1, 2, 5, 6
- [28] Shuai Li, Jianqiang Huang, Xian-Sheng Hua, and Lei Zhang. Category dictionary guided unsupervised domain adaptation for object detection. In *AAAI*, 2021. 3
- [29] Yu-Jhe Li, Xiaoliang Dai, Chih-Yao Ma, Yen-Cheng Liu, Kan Chen, Bichen Wu, Zijian He, Kris Kitani, and Peter Vajda. Cross-domain adaptive teacher for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7581–7590, 2022. 6, 7
- [30] Feng Liu, Xiaosong Zhang, Fang Wan, Xiangyang Ji, and Qixiang Ye. Domain contrast for domain adaptive object detection. *IEEE TCSVT*, 2021. 3
- [31] 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. *arXiv preprint arXiv:2201.12329*, 2022. 2, 5
- [32] 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*, 2021. 2
- [33] Shengxiong Ouyang, Xinglu Wang, Kejie Lyu, and Yingming Li. Pseudo-label generation-evaluation framework for cross domain weakly supervised object detection. In *ICIP*, 2021. 1, 3, 6, 7
- [34] Rindra Ramamonjison, Amin Banitalebi-Dehkordi, Xinyu Kang, Xiaolong Bai, and Yong Zhang. SimROD: A simple adaptation method for robust object detection. In *ICCV*, 2021. 3
- [35] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016. 2
- [36] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. *NeurIPS*, 2015. 2
- [37] Farzaneh Rezaeianaran, Rakshith Shetty, Rahaf Aljundi, Daniel Olmeda Reino, Shanshan Zhang, and Bernt Schiele. Seeking similarities over differences: Similarity-based domain alignment for adaptive object detection. In *ICCV*, 2021. 3
- [38] Aruni RoyChowdhury, Prithvijit Chakrabarty, Ashish Singh, SouYoung Jin, Huaizu Jiang, Liangliang Cao, and Erik Learned-Miller. Automatic adaptation of object detectors to new domains using self-training. In *CVPR*, 2019. 3
- [39] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. Strong-weak distribution alignment for adaptive object detection. In *CVPR*, 2019. 1, 3, 5, 6, 7
- [40] Yunhang Shen, Rongrong Ji, Yan Wang, Zhiwei Chen, Feng Zheng, Feiyue Huang, and Yunsheng Wu. Enabling deep residual networks for weakly supervised object detection. In *ECCV*, 2020. 6, 7
- [41] Zhiqiang Shen, Mingyang Huang, Jianping Shi, Zechun Liu, Harsh Maheshwari, Yutong Zheng, Xiangyang Xue, Marios Savvides, and Thomas S Huang. CDTD: A large-scale cross-domain benchmark for instance-level image-to-image translation and domain adaptive object detection. *IJCV*, 2021. 3
- [42] Peng Su, Kun Wang, Xingyu Zeng, Shixiang Tang, Dapeng Chen, Di Qiu, and Xiaogang Wang. Adapting object detectors with conditional domain normalization. In *ECCV*, 2020. 3
- [43] Peng Tang, Xinggang Wang, Song Bai, Wei Shen, Xiang Bai, Wenyu Liu, and Alan Yuille. PCL: Proposal cluster learning for weakly supervised object detection. *IEEE TPAMI*, 2018. 6, 7
- [44] Peng Tang, Xinggang Wang, Xiang Bai, and Wenyu Liu. Multiple instance detection network with online instance classifier refinement. In *CVPR*, 2017. 3
- [45] Zongheng Tang, Yue Liao, Si Liu, Guanbin Li, Xiaojie Jin, Hongxu Jiang, Qian Yu, and Dong Xu. Human-centric spatio-temporal video grounding with visual transformers. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(12):8238–8249, 2021. 3
- [46] Renshuai Tao, Hainan Li, Tianbo Wang, Yanlu Wei, Yifu Ding, Bawei Jin, Hongping Zhi, Xianglong Liu, and Aishan Liu. Exploring endogenous shift for cross-domain detection: A large-scale benchmark and perturbation suppression network. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21157–21167. IEEE, 2022. 3
- [47] Kun Tian, Chenghao Zhang, Ying Wang, Shiming Xiang, and Chunhong Pan. Knowledge mining and transferring for domain adaptive object detection. In *ICCV*, 2021. 3
- [48] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. FCOS: Fully convolutional one-stage object detection. In *ICCV*, 2019. 2
- [49] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 4
- [50] Vibashan VS, Vikram Gupta, Poojan Oza, Vishwanath A Sindagi, and Vishal M Patel. MeGA-CDA: Memory guided attention for category-aware unsupervised domain adaptive object detection. In *CVPR*, 2021. 3
- [51] Fang Wan, Pengxu Wei, Zhenjun Han, Jianbin Jiao, and Qixiang Ye. Min-entropy latent model for weakly supervised object detection. *IEEE TPAMI*, 2019. 3
- [52] Yu Wang, Rui Zhang, Shuo Zhang, Miao Li, YangYang Xia, XiShan Zhang, and ShaoLi Liu. Domain-specific suppression for adaptive object detection. In *CVPR*, 2021. 3
- [53] Wang Wen, Cao Yang, Zhang Jing, He Fengxiang, Zha Zheng-Jun, Wen Yonggang, and Tao Dacheng. Exploring sequence feature alignment for domain adaptive detection transformers. In *ACM MM*, 2021. 3, 4, 7
- [54] Aming Wu, Yahong Han, Linchao Zhu, and Yi Yang. Instance-invariant domain adaptive object detection via progressive disentanglement. *IEEE TPAMI*, 2021. 6, 7
- [55] Aming Wu, Rui Liu, Yahong Han, Linchao Zhu, and Yi Yang. Vector-decomposed disentanglement for domain-invariant object detection. In *ICCV*, 2021. 3, 7
- [56] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019. 6
- [57] Minghao Xu, Hang Wang, Bingbing Ni, Qi Tian, and Wenjun Zhang. Cross-domain detection via graph-induced prototype alignment. In *CVPR*, 2020. 3
- [58] Yunqiu Xu, Yifan Sun, Zongxin Yang, Jiaxu Miao, and Yi Yang. H2fa r-cnn: Holistic and hierarchical feature alignment for cross-domain weakly supervised object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14329–14339, 2022. 1, 2, 3, 4, 5, 6, 7
- [59] Ke Yang, Dongsheng Li, and Yong Dou. Towards precise end-to-end weakly supervised object detection network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8372–8381, 2019. 3
- [60] Bo Zhang, Tao Chen, Bin Wang, Xiaofeng Wu, Liming Zhang, and Jiayuan Fan. Densely semantic enhancement for domain adaptive region-free detectors. *IEEE TCSVT*, 2021. 3
- [61] Yixin Zhang, Zilei Wang, and Yushi Mao. RPN prototype alignment for domain adaptive object detector. In *CVPR*, 2021. 3

- [62] Ganlong Zhao, Guanbin Li, Ruijia Xu, and Liang Lin. Collaborative training between region proposal localization and classification for domain adaptive object detection. In *ECCV*, 2020. 3
- [63] Liang Zhao and Limin Wang. Task-specific inconsistency alignment for domain adaptive object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14217–14226, 2022. 6
- [64] Zhen Zhao, Yuhong Guo, Haifeng Shen, and Jieping Ye. Adaptive object detection with dual multi-label prediction. In *ECCV*, 2020. 3, 6, 7
- [65] Yangtao Zheng, Di Huang, Songtao Liu, and Yunhong Wang. Cross-domain object detection through coarse-to-fine feature adaptation. In *CVPR*, 2020. 3
- [66] Xinge Zhu, Jiangmiao Pang, Ceyuan Yang, Jianping Shi, and Dahua Lin. Adapting object detectors via selective cross-domain alignment. In *CVPR*, 2019. 3
- [67] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020. 1, 2