

# MTTrans: Cross-Domain Object Detection with Mean Teacher Transformer

Jinze Yu<sup>1</sup>, Jiaming Liu<sup>2</sup>, Xiaobao Wei<sup>2</sup>, Haoyi Zhou<sup>1</sup>, Yohei Nakata<sup>3</sup>,  
Denis Gudovskiy<sup>3</sup>, Tomoyuki Okuno<sup>3</sup>, Jianxin Li<sup>1</sup>,  
Kurt Keutzer<sup>4</sup>, and Shanghang Zhang<sup>2\*</sup>

<sup>1</sup>Beihang University    <sup>2</sup>Peking University

<sup>3</sup>Panasonic Holdings Corporation    <sup>4</sup>University of California, Berkeley  
yujinze@buaa.edu.cn, shanghang@pku.edu.cn

**Abstract.** Recently, DETection TRansformer (DETR), an end-to-end object detection pipeline, has achieved promising performance. However, it requires large-scale labeled data and suffers from domain shift, especially when no labeled data is available in the target domain. To solve this problem, we propose an end-to-end cross-domain detection Transformer based on the mean teacher framework, **MTTrans**, which can fully exploit unlabeled target domain data in object detection training and transfer knowledge between domains via pseudo labels. We further propose the comprehensive multi-level feature alignment to improve the pseudo labels generated by the mean teacher framework taking advantage of the cross-scale self-attention mechanism in Deformable DETR. Image and object features are aligned at the local, global, and instance levels with domain query-based feature alignment (DQFA), bi-level graph-based prototype alignment (BGPA), and token-wise image feature alignment (TIFA). On the other hand, the unlabeled target domain data pseudo-labeled and available for the object detection training by the mean teacher framework can lead to better feature extraction and alignment. Thus, the mean teacher framework and the comprehensive multi-level feature alignment can be optimized iteratively and mutually based on the architecture of Transformers. Extensive experiments demonstrate that our proposed method achieves state-of-the-art performance in three domain adaptation scenarios, especially the result of Sim10k to Cityscapes scenario is remarkably improved from 52.6 mAP to 57.9 mAP. Code will be released here.


**Keywords:** Unsupervised Domain Adaptation, Object Detection, Mean Teacher Transformer

## 1 Introduction

Object detection is one of the fundamental computer vision tasks which has been improved dramatically in the last decades. Methods based on Convolutional Neural Networks (CNN) [24,23,21] achieve satisfying results but rely heavily on certain hand-crafted operations and are not fully end-to-end. Recently,

---

\* Corresponding Author



**Fig. 1.** The performance of our proposed MTTrans compared with previous works on different datasets. Our method consistently outperforms the SOTA baseline with a large margin on the scene adaptation and the synthetic to real adaptation, and the performance of each individual category.

Transformer-based approaches [3,47] have been introduced as a promising one-stage detector. Although Transformer-based detectors have shown superior detection and generalization performance compared with CNN-based ones [1,41], they still suffer significant performance degradation caused by domain shift or variation of data distribution when tested in scenarios with domain gaps. These motivate the study of unsupervised domain adaptive (UDA) on transformer object detectors [33,42].

The pivot of UDA is to deal with domain shifts between source and target domains. It empowers the learned model to be transferred from the source to the unlabeled target domain. Most previous works are based on cross-domain feature alignment techniques, which rarely use target data in object detection training due to the lack of labels, thus causing insufficient use of data. In this paper, inspired by the prevalent mean teacher mechanism in semi-supervised learning (SSL) tasks [29,35], we innovatively design MTTrans, a cross-domain detection Transformer based on the mean teacher framework. By extending the mean teacher framework to UDA, we can sufficiently exploit unlabeled target data and transfer the knowledge between domains via pseudo labels.

However, directly incorporating the mean teacher framework to UDA object detection often results in degressive performance: the framework is proposed initially to solve the SSL task, in which the labeled and unlabeled data are in the same data distribution. The goal of UDA is to remove domain-specific components from the extracted features, and the mean teacher does not have such an ability. Such domain-specific features increase the divergence between source and target features and degrade the quality of pseudo labels generated by the mean teacher. Therefore, it is crucial to exploit cross-domain feature alignments to learn domain-invariant features and improve the quality of pseudo labels in the target domain, which is the critical factor for better domain adaptation.

To address the issues above, we propose **MTTrans**, an end-to-end cross-domain detection Transformer framework, and the first Transformer-based study that utilizes the mean teacher structure to transfer knowledge between do-

mains. MTTrans transfer domain-invariant task-relevant knowledge between domains and better utilize the unlabeled data via **pseudo label generation by the mean teacher framework** and the **comprehensive multi-level cross-domain feature alignment** based on the cross-scale attention mechanism in Deformable DETR [47]. With pseudo label generation, the unlabeled target domain data are pseudo-labeled and available in the object detection training. In contrast, in previous works, such data are rarely used due to the lack of labels. The multi-level feature alignment aligns both images and object proposals features at the global, instance, and local levels to enhance the reliability and the quality of the generated pseudo labels under domain shift. On the other hand, with target domain images in object detection training, better feature extraction and alignment can be achieved. Therefore, the mean teacher framework and the feature alignment act as a whole and can be optimized iteratively and mutually.

The comprehensive feature alignment strategy includes **domain query-based feature alignment (DQFA)** to align image and object features at the global level, **bi-level graph-based prototype alignment (BGPA)** which builds prototypes with the object features, and a simple **token-wise image feature alignment (TIFA)** technique for the local-level image features. Inspired by SFA [33], DQFA extracts a feature for each of the entire image and all the object proposals via cross-scale attention and shrinks global domain gaps in scene layout. However, it still suffers from the domain shift caused by object visual layout changing. Therefore, we then propose the novel BGPA technique to build and aggregate prototypes for object features based on their visual similarity, which can better learn categorical and spatial correlations (since objects in the same category and spatially connected tend to be visually similar [25]) and achieve more accurate alignment. Finally, TIFA is performed on the local-level image tokens. The quality of pseudo labels is largely improved with the comprehensive multi-level feature alignment, as shown in Fig.5. A more detailed visualization analysis will be performed in Section 4.4.

Finally, the proposed MTTrans framework achieves promising performance on three challenging domain adaptation benchmarks, as shown in Fig. 1. In the weather adaptation scenario, MTTrans outperforms both end-to-end and two-stage detection algorithms by improving the result to 43.4 mAP (Cityscapes to Foggy Cityscapes). In the scene adaptation scenario, the proposed method surpasses the previous state-of-the-art result [33] by +3.7 mAP, improving the result to 32.6 mAP (Cityscapes to BDD100k). In the synthetic to real adaptation scenario, we achieve an improvement of +5.3 mAP (Sim10k to Cityscapes) over the previous state-of-the-art (SFA [33]), achieving 57.9 mAP.

## 2 Related Work

### 2.1 Object Detection

Object detection is one of the fundamental tasks of computer vision [12,19,43]. Recently, CNN-based methods with large-scale labeled training data have become the mainstream object detection approaches, categorized into stronger two-

stage methods [24,20,39] and faster, lighter one-stage methods [23,21,31]. However, these methods extremely depend on handcrafted components, notably the non-maximum suppression (NMS) post-processing, and thus cannot be trained end-to-end. It was recently achieved by DETR [3] and its follow-up work Deformable DETR [47], with the vision Transformers [32]. Deformable DETR proposes a novel deformable multi-head attention mechanism to provide sparsity in attention and multi-scale feature aggregation without feature pyramid structure, allowing for faster training and better performance. In this work, we choose Deformable DETR as the base detector for its simple yet powerful working flow and the great potential of the attention mechanism for cross-domain feature extraction.

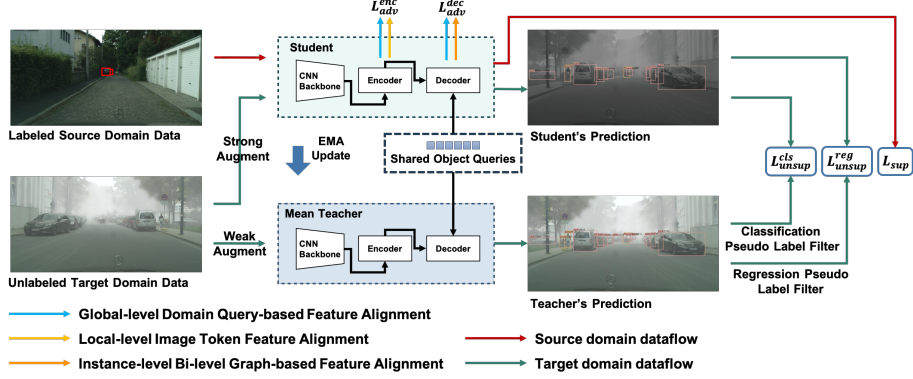
## 2.2 Unsupervised domain adaptive object detection

Domain Adaptive Faster R-CNN [4] is the first work to study domain-adaptive object detection. Most later works follow the cross-domain feature alignment with the adversarial training approach. These works propose to aggregate image or instance features based on their categorical predictions [34,36] or spatial correlations [36,2], and the features are aligned hierarchically at one or some levels in global, local, instance and category levels [34,26,36,2,22]. More recent works include PICA [44], which focuses on Few-shot Domain Adaptation, and Visually Similar Group Alignment (ViSGA) [25], which proposes to aggregate instance features based on their visual similarity with similarity-based hierarchical agglomerative clustering, which achieves state-of-the-art performance on some benchmarks. Other methods focus on applying other domain adaptation techniques or other base detectors, such as Mean Teacher with Object Relations (MTOR) [2] and Unbiased Mean Teacher (UMT) [8]. Regarding the Transformer-based models, Sequence Feature Alignment (SFA) [33] proposes a domain adaptive end-to-end object detector based on Deformable DETR. DA-DETR [42] proposes an alignment technique based on convolutions and spatial and channel attention for DETR. CDTrans [37] and TVT [38] are proposed for cross-domain classification. In this paper, we propose cooperating the mean teacher framework with multi-level cross-domain feature alignment based on Deformable DETR to achieve end-to-end UDA object detection.

## 3 Method

This section introduces our MTTrans framework for domain adaptive detection Transformer. In unsupervised domain adaptation (UDA), the data include the labeled source images and the unlabeled target images. Our goal is to transfer object detection task-specific domain-invariant knowledge from the source to the target domain and make the target domain pseudo-labeled and thus available for object detection training with the mean teacher framework and the multi-level source-target feature alignment.

In Section 3.1, based on the mean teacher, we propose an end-to-end cross-domain detection Transformer framework, **MTTrans**, which learns similarities between the source and target domain with pseudo labels. We then propose multi-level source-target feature alignment in Section 3.2 to address the domain shift problem and further improve the reliability of the pseudo labels generated by the mean teacher for the target domain. Finally, in Section 3.3, we elaborate on the training policy of our framework. The overall pipeline is shown in Fig. 2.




**Fig. 2.** The framework of MTTrans, composed of a *student model* which is the actual task model, and a temporal ensemble version of the student model called the *teacher model*. The teacher model is updated by EMA of the student model to generate pseudo labels for the target domain. To improve the quality of the pseudo labels on the target domain and reduce the source-target discrepancy, we additionally design multi-level feature alignment strategies.

### 3.1 Mean Teacher-based Knowledge Transfer Framework

MTTrans is based on Deformable DETR [47] and is derived from the mean teacher framework, consisting of two models with the same architecture and identical initialization weights. The student model is updated with back-propagation, while the teacher model is updated by student's weights with exponential moving average (EMA). Thus, the teacher model can be considered as multiple temporal ensemble student models: for weights of the teacher model  $\theta'_t$  at time step  $t$ , it is the EMA of successive student's weights  $\theta_t$ :

$$\theta'_t = \alpha \theta'_{t-1} + (1 - \alpha) \theta_t \quad (1)$$

where  $\alpha$  is a smoothing coefficient hyperparameter. All of the teacher's weights are updated according to Eq. 1, except for object query embeddings, which are kept the same between the two models to further enhance consistency between



**Fig. 3.** The detailed structure of the student model, which is composed of a detection Transformer backbone, and the proposed multi-level feature alignment.


them. The object queries are trainable embeddings, initialized with the normal distribution at the start of the training procedure. Then we use the temporally ensembled teacher model to guide the student model’s training in the target domain via pseudo labels.

As a crucial factor for the teacher-student framework, pseudo labels are tactfully generated in our work. Object proposals with any foreground category score higher than a pre-defined threshold are assigned as pseudo labels. In addition, the unlabeled target images are strongly augmented for the student model (denoted as  $I_{tgt}$ ) and weakly augmented for the teacher model (denoted as  $I'_{tgt}$ ) [28]. In such a way, the teacher model’s prediction can be more accurate, enabling the student model to learn from the generated pseudo labels. Although directly applying the mean teacher to the UDA task improves the results, the pseudo labels generated on the target domain are still of low quality because of the data distribution shift. The temporal ensembled model tends to accumulate errors and collapse.

### 3.2 Multi-level Cross-Domain Adversarial Feature Alignment

To address the issue above, we design the comprehensive multi-level cross-domain feature alignment based on the cross-scale self-attention mechanism of Deformable DETR [47] to strengthen the reliability of the pseudo labels generated when faced with domain shift. The alignment is performed on different parts and at three levels of the student model, including **domain query-based feature alignment (DQFA)** for global-level image and instance features on encoder and decoder outputs, **bi-level graph-based prototype alignment (BGPA)** for instance-level object proposal features on decoder outputs, and **token-wise image feature alignment (TIFA)** for local-level image features on encoder outputs, as shown in Fig. 3.

**Global-level Domain Query-based Feature Alignment (DQFA).** Inspired by the global memory [11,9] and the sparse attention mechanism [45,5], we adopt DQFA in SFA [33] to extract and align global-level context features



**Fig. 4.** The proposed bi-level graph-based prototype alignment (BGPA) for instance-level object proposal features.  $M$  prototypes (set as 9 in our work) for instance features are first generated with decoder output by an MLP. The prototype features are then aggregated before alignment by constructing a bi-level graph and a GCN.

for images or object proposals. Besides, the domain query provides a global link between any other local feature tokens, through which any two tokens can attend to each other. The two global-level feature tokens generated by the two domain queries are further classified by two domain discriminators, respectively. The domain discriminators' optimization goal is to make the extracted features more domain-specific and distinguishable, which is opposite to the purpose of the cross-domain model, which tries to learn more domain-invariant features. Therefore, the domain adversarial training method [10] is adopted to insert gradient reversal layers and reverse back-propagated gradients from domain discriminators to optimize the detection model to extract domain-invariant features.

#### Instance-level Bi-level Graph-based Prototype Alignment (BGPA).

The above DQFA can effectively shrink global domain gaps in scene layout, but it still suffers from domain shift caused by object visual layout changing. Previous work proposes aggregating and further aligning instance features based on their visual similarities [25]. In MTTrans, we propose BGPA to build and aggregate prototypes for object features based on their feature similarities with prototyping and graph construction, which can learn from categorical and spatial correlations (since objects in the same category and spatially connected tend to be visually similar [25]) and achieve a more accurate alignment. As illustrated in Fig. 4 (a),  $M$  prototypes (set as 9 in this work) for object proposals are generated by an MLP with the output feature tokens of the decoder. In Fig. 4 (b), we construct an undirected bi-level graph with the prototypes and the decoder outputs. The prototypes are connected to each other and the decoder outputs, while edges' weights are calculated by the nodes' cosine similarity. As shown in Fig. 4 (c), the created graph is then processed with a graph convolutional network (GCN) [18,30]. Finally, the graph-aggregated prototypes are aligned with a domain discriminator to alleviate domain shift at the instance level.

**Local-level Token-wise Image Feature Alignment (TIFA).** Though adopting the above two alignments, the mean teacher framework still suffers from the domain shift caused by detailed changes at the local level. Thus, we additionally add a simple yet effective TIFA that aligns the multi-scale image local feature tokens with domain adversarial training. As shown in Fig. 3, image feature maps of multiple spatial scales are retrieved from different CNN backbone layers and then flattened and concatenated, forming a sequence of local-level image feature tokens. The feature tokens are aggregated with each other by the cross-scale self-attention mechanism in the encoder, obtaining context-enriched local-level features. TIFA aligns these tokens one by one with the domain discriminator. There exist some more complex feature alignment methods at the local level, and BGPA proposed above can also be applied to these feature tokens. However, these methods may result in performance degradation compared with the simple TIFA, as shown in Section 4.3.

### 3.3 Progressive Cross-Domain Knowledge Transfer with Mean Teacher and Adversarial Feature Alignment

Compared with previous mean teacher Transformer works that focus on semi-supervised learning, we propose not only the comprehensive multi-level feature alignment described above to strengthen the reliability of the cross-domain pseudo labels but also a two-step progressive transfer training policy.

First of all, the model needs to learn from the labeled source domain data first and then be able to generate informative pseudo labels. Moreover, it is difficult to simultaneously train a randomly initialized model to perform well on both two distinctly different datasets. Therefore, we divide the mean teacher training procedure into a burn-in and a transfer training step. In the first burn-in step, we train the student model in the labeled source domain with object detection and feature alignment tasks. The loss function is a combination of both tasks:

$$\mathcal{L}_{sup} = \mathcal{L}_{det}(I_{src}, y_{src}) - \mathcal{L}_{adv} \quad (2)$$

Given a source domain image  $I_{src} \in \mathcal{D}_{src}$ , the backbone  $G$  and the encoder  $Enc$  of the Transformer produce the image features  $f_{src}$ , then the decoder  $Dec$  produces the features for the object proposals  $g$ , and finally, the detection head  $H$  obtains the bounding boxes and the object category predictions. The supervised detection loss is defined as follows:

$$\mathcal{L}_{det}(I_{src}, y_{src}) = l(H(g), y_{src}) \quad (3)$$

where  $l(\cdot)$  denotes the supervised object detection loss which keeps the same with Deformable DETR [47].

The cross-domain adversarial feature alignment loss  $\mathcal{L}_{adv}$  comprises four parts: DQFA loss on image features  $\mathcal{L}_{enc}^{DQFA}$ , DQFA loss on object features  $\mathcal{L}_{dec}^{DQFA}$ , BGPA loss  $\mathcal{L}_{dec}^{BGPA}$ , and TIFA loss  $\mathcal{L}_{enc}^{TIFA}$ .



$$\mathcal{L}_{adv} = \lambda_1 \mathcal{L}_{enc}^{DQFA}(f) + \lambda_2 \mathcal{L}_{dec}^{DQFA}(g) + \lambda_3 \mathcal{L}_{dec}^{BGPA}(g) + \lambda_4 \mathcal{L}_{enc}^{TIFA}(f) \quad (4)$$

$\mathcal{L}_{enc}^{DQFA}$  is computed in the following way, and the other three parts are computed in the same way:

$$\mathcal{L}_{enc}^{DQFA}(f) = \mathbb{E}_{f \in \mathcal{D}_{src}} \log D(f) + \mathbb{E}_{f \in \mathcal{D}_{tgt}} \log(1 - D(f)) \quad (5)$$

where  $D$  denotes the domain discriminator.

Then, in the transfer training step, both the student and the teacher model are initialized with the model trained in the first step. The student model is trained on the source and the target domain alternatively, one epoch for each domain, to maintain the object detection performance and learn to transfer between domains. The student model is updated with back-propagation, and the teacher model is updated by the student model’s weights after each source domain trained epoch. The loss for target domain training is made up of two components: UDA object detection loss and adversarial feature alignment loss:

$$\mathcal{L}_{unsup} = \mathcal{L}_{det}(I_{tgt}, \hat{y}_{tgt}) - \mathcal{L}_{adv} \quad (6)$$

where  $\mathcal{L}_{adv}$  is denoted in Eq. 4,  $I_{tgt}$  is the unlabeled target domain images,  $\hat{y}_{tgt}$  is the generated pseudo labels, and  $\mathcal{L}_{det}(I_{tgt}, \hat{y}_{tgt})$  is the pseudo-supervised object detection loss on the target domain.

To summarize, let  $T$  denote the object detector and  $D$  denote the domain discriminators, the final training objective of MTTrans is defined as :

$$\min_T \max_D \mathcal{L}_{det}^{src}(T) + \mathcal{L}_{det}^{tgt}(T) - \mathcal{L}_{adv}(T, D) \quad (7)$$

## 4 Evaluation

In this section, we conduct extensive experiments to demonstrate the advantages of our proposed method. In Section 4.1, the details of the experimental setup are given. In Section 4.2, we demonstrate the performance of MTTrans in three challenging domain adaptation scenarios, including Weather, Scene, and Synthetic to Real Adaptation. We also conduct comprehensive ablation studies to investigate the impact of each component in Section 4.3. Finally, we conduct qualitative analysis to provide a better understanding in Section 4.4.

### 4.1 Experimental Setup

**Datasets** We evaluate our method on four public datasets, including Cityscapes [6], Foggy Cityscapes [27], Sim10k [15], and BDD100k [40]. We present the performance of MTTrans in three domain adaptation scenarios according to these datasets:

- **Weather adaptation.** In this scenario, we use Cityscapes collected from urban scenes as the source dataset. It consists of 3,475 images with pixel-level annotation, and 2,975 of them are used for training, the other 500 are for evaluation. Foggy Cityscapes constructed from Cityscapes by a fog synthesis algorithm is used as the target dataset.
- **Scene Adaptation.** In this condition, Cityscapes still serves as the source dataset. We utilize the daytime subset of BDD100k as the target dataset, which consists of 36,728 training images and 5,258 validation images annotated with bounding boxes.
- **Synthetic to Real Adaptation.** In this scenario, we utilize Sim10k created by the Grand Theft Auto game engine as the source domain. It is designed to conclude 10,000 training images with 58,701 bounding box annotations for cars. Car instances in Cityscapes are for target domain training and evaluation.

**Implementation Details** Our method is built on the basis of Deformable DETR [47]. We set ImageNet [7] pre-trained ResNet-50 [13] as CNN backbone in all experiments. In the burn-in step, we adopt Adam optimizer [17] during training for 50 epochs. We initialize the learning rate as  $2 \times 10^{-4}$  which decayed by 0.1 after 40 epochs. The batch size is set to 4 for all domain adaptation scenarios. In the second cross-domain training step, the model is trained for 40 epochs, and the initial learning rate is  $2 \times 10^{-6}$ . The learning rate is decayed by 0.1 after 20 epochs. The batch size in this step is set to 2. In addition, we adopt Mean Average Precision (mAP) with a threshold of 0.5 as the evaluation metric. The filtering threshold for pseudo label generation is set to 0.5. All experiments are conducted on NVIDIA Tesla A100 GPUs.

## 4.2 Comparisons with Other Methods

The proposed MTTrans framework is assessed in three different domain adaptation scenarios, and the results obtained are compared with prior works in this section. SOTA results are achieved in all three scenarios.

**Weather Adaptation** Variations in weather conditions are common, yet challenging, and object detectors must be reliable under all conditions. As a result, we evaluate the robustness of models under weather changes by transferring from Cityscapes to Foggy Cityscapes. As shown in Table 1, MTTrans outperforms other end-to-end methods by a significant margin (43.4% vs. 41.3% by the closest and SOTA end-to-end model, SFA). In addition, it achieves competitive results when compared with two-stage approaches.

**Scene Adaptation** Scene layouts are not static and frequently change in real-world applications, especially in autonomous driving scenarios. As a result, model performance under scene adaptation is critical. As shown in Table 2, MTTrans obtains SOTA results (32.6%) with significant improvements over previous

**Table 1.** Results of different methods for weather adaptation, that is, from cityscapes to fog cityscapes. FRCNN and DefDETR are abbreviations for Faster R-CNN and Deformable DETR, respectively.

Method	Detector	person	rider	car	truck	bus	train	mcycle	bicycle	mAP
FasterRCNN [24](Source)	FRCNN	26.9	38.2	35.6	18.3	32.4	9.6	25.8	28.6	26.9
DivMatch [16]	FRCNN	31.8	40.5	51.0	20.9	41.8	34.3	26.6	32.4	34.9
SWDA [26]	FRCNN	31.8	44.3	48.9	21.0	43.8	28	28.9	35.8	35.3
SCDA [46]	FRCNN	33.8	42.1	52.1	26.8	42.5	26.5	29.2	34.5	35.9
MTOR [2]	FRCNN	30.6	41.4	44.0	21.9	38.6	40.6	28.3	35.6	35.1
CR-DA [34]	FRCNN	30.0	41.2	46.1	22.5	43.2	27.9	27.8	34.7	34.2
CR-SW [34]	FRCNN	34.1	44.3	53.5	24.4	44.8	38.1	26.8	34.9	37.6
GPA [36]	FRCNN	32.9	46.7	54.1	24.7	45.7	41.1	32.4	38.7	39.5
ViSGA [25]	FRCNN	38.8	45.9	57.2	<b>29.9</b>	<b>50.2</b>	<b>51.9</b>	31.9	40.9	43.3
FCOS [31] (Source)	FCOS	36.9	36.3	44.1	18.6	29.3	8.4	20.3	31.9	28.2
EPM[14]	FCOS	44.2	46.6	58.5	24.8	45.2	29.1	28.6	34.6	39.0
Def DETR [47] (Source)	DefDETR	37.7	39.1	44.2	17.2	26.8	5.8	21.6	35.5	28.5
SFA [33]	DefDETR	46.5	48.6	62.6	25.1	46.2	29.4	28.3	44.0	41.3
MTTrans(Ours)	DefDETR	<b>47.7</b>	<b>49.9</b>	<b>65.2</b>	25.8	45.9	33.8	<b>32.6</b>	<b>46.5</b>	<b>43.4</b>

**Table 2.** Results of different methods for the scene adaptation, i.e., Cityscapes to BDD100k daytime subset.

Methods	Detector	person	rider	car	truck	bus	mcycle	bicycle	mAP
FasterR-CNN [24](Source)	FRCNN	28.8	25.4	44.1	17.9	16.1	13.9	22.4	24.1
DAF [4]	FRCNN	28.9	27.4	44.2	19.1	18.0	14.2	22.4	24.9
SWDA [26]	FRCNN	29.5	29.9	44.8	20.2	20.7	15.2	23.1	26.2
SCDA [46]	FRCNN	29.3	29.2	44.4	20.3	19.6	14.8	23.2	25.8
CR-DA [34]	FRCNN	30.8	29.0	44.8	20.5	19.8	14.1	22.8	26.0
CR-SW [34]	FRCNN	32.8	29.3	45.8	22.7	20.6	14.9	<b>25.5</b>	27.4
FCOS [31](Source)	FCOS	38.6	24.8	54.5	17.2	16.3	15.0	18.3	26.4
EPM [14]	FCOS	39.6	26.8	55.8	18.8	19.1	14.5	20.1	27.8
Def DETR [47](Source)	DefDETR	38.9	26.7	55.2	15.7	19.7	10.8	16.2	26.2
SFA [33]	DefDETR	40.2	27.6	57.5	19.1	23.4	15.4	19.2	28.9
MTTrans(Ours)	DefDETR	<b>44.1</b>	<b>30.1</b>	<b>61.5</b>	<b>25.1</b>	<b>26.9</b>	<b>17.7</b>	23.0	<b>32.6</b>

works. In addition, the performance of six out of seven categories in the target domain dataset has been improved. (As the SOTA method in the weather adaptation scenario ViSGA [25] did not open-source their code and did not report performance in the scene adaptation scenario, it is not included here.)

**Synthetic to Real Adaptation** Images and their corresponding annotations can be created by video games such as GTA, which can considerably minimize the manual cost of data collection and annotating. As a result, it is worthwhile to enable the object detector to learn from synthetic images and then adapt to generic real-world images. Therefore, as shown in Table 3, we test MTTrans in the synthetic to real adaptation scenario, obtaining an accuracy of 57.9%, outperforming the previous state-of-the-art by 5.3%.

### 4.3 Ablation Studies

To better analyze each component in our proposed MTTrans framework, we conduct ablation studies by removing parts of the components in MTTrans. From Table 4, we can observe: (1) Adding the mean teacher framework directly to Deformable DETR (MTTrans-AS0, MT-DefDETR)) can improve its performance

**Table 3.** Results of different methods for the synthetic to real adaptation, i.e. Sim10k to Cityscapes.

Methods	Detector	carAP
FasterRCNN [24](Source)	FRCNN	34.6
DAF [4]	FRCNN	41.9
DivMatch [16]	FRCNN	43.9
SWDA [26]	FRCNN	44.6
SCDA [46]	FRCNN	45.1
MTOR [2]	FRCNN	46.6
CR-DA [34]	FRCNN	43.1
CR-SW [34]	FRCNN	46.2
GPA [36]	FRCNN	47.6
ViSGA [25]	FRCNN	49.3
FCOS [31](Source)	FCOS	42.5
EPM [14]	FCOS	47.3
DefDETR [47](Source)	DefDETR	47.4
SFA [33]	DefDETR	52.6
MTTrans(Ours)	DefDETR	<b>57.9</b>


**Table 4.** Ablation studies on the weather adaptation scenario, with Cityscapes to Foggy Cityscapes. MT stands for the mean teacher framework and SharedQE denotes the shared object queries of decoder inputs. Components of other models that differ from MTTrans are marked in red.

Methods	MT	SharedQE	DQFA		TIFA		BGPA		mAP50
			enc	dec	enc	dec	enc	dec	
Deformable DETR (Source)	✗	✗	✗	✗	✗	✗	✗	✗	28.5
MTTrans-AS0(MT-DefDETR)	✓	✓	✗	✗	✗	✗	✗	✗	35.843
MTTrans-AS11	✓	✓	✗	✓	✓	✗	✗	✓	43.183
MTTrans-AS12	✓	✓	✓	✗	✓	✗	✗	✓	42.962
MTTrans-AS13	✓	✓	✓	✓	✗	✗	✗	✓	43.013
MTTrans-AS14	✓	✓	✓	✓	✓	✗	✗	✗	43.003
MTTrans-AS15	✓	✗	✓	✓	✓	✗	✗	✓	42.797
MTTrans-AS21	✓	✓	✓	✓	✓	✓	✗	✗	42.953
MTTrans-AS22	✓	✓	✓	✓	✗	✗	✓	✓	43.068
MTTrans	✓	✓	✓	✓	✓	✗	✗	✓	43.413

on the target domain (+7.34 mAP). We can notice that it is critical to introduce the mean teacher framework in cross-domain adaptation, but there is still much space for improvement due to the poor quality of pseudo labels. (2) Removing any aspect of MTTrans (MTTrans-AS11 to MTTrans-AS14) will result in performance degradation. Removing DQFA for the encoder (MTTrans-AS11) results in the slightest performance drop (-0.230mAP), while removing DQFA for the decoder (MTTrans-AS12) results in the highest drop (-0.451 mAP); (3) Altering the alignment technique for the decoder from BGPA to TIFA (MTTrans-AS21, -0.460 mAP), or replacing TIFA for the encoder with BGPA (MTTrans-AS22, -0.345 mAP) both result in performance drop; (4) Removing the shared object queries between the teacher and student models (MTTrans-AS15) will also decrease MTTrans’s performance (-0.616 mAP).

#### 4.4 Visualization and Analysis

**Pseudo Label Visualization** We show some visualization results of the generated pseudo labels and student model’s prediction, produced by MTTrans and



**Fig. 5.** Visualization results of the generated pseudo labels, ground truth annotations, and student model predictions. As seen in the visualization result, MTTrans can generate pseudo labels of higher quality compared with MT-DefDETR; and the teacher model performs better than the student model. MT-DefDETR stands for directly applying the mean teacher framework to Deformable DETR.

the mean teacher version of Deformable DETR, as shown in Fig. 5. It should be noted that input images are augmented randomly during training, and the student model’s performance can not be compared between the two model’s predictions in the figure. However, the teacher model’s inputs are almost the same, and the pseudo labels generated by MT-DefDETR are of lower quality than MTTrans’s. Furthermore, the pseudo labels generated by the teacher model are better than the student model’s predictions, ensuring that the student model can learn to improve via pseudo labels.

**Detection Results** We show some visualization results of MTTrans on the Foggy Cityscapes dataset, accompanied by ground truth, baseline, and previous state-of-the-art (SOTA) methods. As shown in Fig. 6, MTTrans can better detect complex objects covered by fog in the distance compared with Deformable DETR and SFA. The results are consistent with the numerical assessment results, indicating that MTTrans manages to mitigate the domain shift problem in the UDA Transformer detector. Furthermore, as detection confidence scores are shown in the left-top corner of the bounding box (best zoom in  $\times 8$ ), MTTrans can slightly boost the scores of true positive predictions, further indicating its strength. As shown in the third row of Fig. 6, MTTrans manages to retrieve objects ignored by SFA but discovered by Deformable DETR. From the last row in Fig. 6, even though our proposed MTTrans fails to discover the blurry and indistinguishable person, it succeeds in detecting cars that are not labeled in the ground truth. More visualization of other datasets is presented in the supplement materials.



**Fig. 6.** Visualization of detection results on the Foggy Cityscapes dataset, from left to right, are ground truth, results obtained by Deformable DETR, SFA, and MTTrans. The predicted category and prediction confidence can be seen in the bounding box labels. Recommend to read with computers, and the original image files are attached with supplement materials.

## 5 Conclusions

This paper introduces an end-to-end cross-domain object detection Transformer named MTTrans based on the mean teacher framework and the comprehensive multi-level feature alignment. The mean teacher framework is adopted in MTTrans to fully utilize the unlabeled target domain dataset, making it pseudo-labeled and available in the training of the object detection task. To compensate for the lack of explicit cross-domain ability of the mean teacher framework and to further improve the quality of the generated pseudo labels under domain shift, we propose the multi-level feature alignment which aligns image and object features at local, instance, and global levels based on the cross-scale self-attention mechanism in Deformable DETR [47], including domain query-based feature alignment for global-level image and object features, bi-level graph-based prototype alignment for instance-level object features, and token-wise image feature alignment for local-level image features, respectively. Experimental results demonstrate the effectiveness of our MTTrans. We hope that this paper can inspire future work on cross-domain object detection.

**Acknowledgment.** We thank Xiaoqi Li and Zhaoqing Wang for their help with the manuscript revision and the anonymous reviewers for their helpful comments to improve the paper. The authors of this paper are supported by the NSFC through grant No.U20B2053, and we also thanks the support from Beijing Advanced Innovation Center for Future Blockchain and Privacy Computing.

## References

1. Bai, Y., Mei, J., Yuille, A.L., Xie, C.: Are transformers more robust than cnns? *Advances in Neural Information Processing Systems* **34**, 26831–26843 (2021)
2. Cai, Q., Pan, Y., Ngo, C.W., Tian, X., Duan, L., Yao, T.: Exploring object relation in mean teacher for cross-domain detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 11457–11466 (2019)
3. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: *European conference on computer vision*. pp. 213–229. Springer (2020)
4. Chen, Y., Li, W., Sakaridis, C., Dai, D., Van Gool, L.: Domain adaptive faster r-cnn for object detection in the wild. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 3339–3348 (2018)
5. Child, R., Gray, S., Radford, A., Sutskever, I.: Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509* (2019)
6. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 3213–3223 (2016)
7. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *2009 IEEE conference on computer vision and pattern recognition*. pp. 248–255. Ieee (2009)
8. Deng, J., Li, W., Chen, Y., Duan, L.: Unbiased mean teacher for cross-domain object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4091–4101 (2021)
9. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018)
10. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., Lempitsky, V.: Domain-adversarial training of neural networks. *The journal of machine learning research* **17**(1), 2096–2030 (2016)
11. Guo, Q., Qiu, X., Liu, P., Shao, Y., Xue, X., Zhang, Z.: Star-transformer. *arXiv preprint arXiv:1902.09113* (2019)
12. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: *Proceedings of the IEEE international conference on computer vision*. pp. 2961–2969 (2017)
13. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
14. Hsu, C.C., Tsai, Y.H., Lin, Y.Y., Yang, M.H.: Every pixel matters: Center-aware feature alignment for domain adaptive object detector. In: *European Conference on Computer Vision*. pp. 733–748. Springer (2020)
15. Johnson-Roberson, M., Barto, C., Mehta, R., Sridhar, S.N., Rosaen, K., Vasudevan, R.: Driving in the matrix: Can virtual worlds replace human-generated annotations for real world tasks? *arXiv preprint arXiv:1610.01983* (2016)
16. Kim, T., Jeong, M., Kim, S., Choi, S., Kim, C.: Diversify and match: A domain adaptive representation learning paradigm for object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 12456–12465 (2019)
17. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)

18. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. arXiv preprint arXiv:1609.02907 (2016)
19. Kirillov, A., He, K., Girshick, R., Rother, C., Dollár, P.: Panoptic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9404–9413 (2019)
20. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2117–2125 (2017)
21. Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y., Berg, A.C.: Ssd: Single shot multibox detector. In: European conference on computer vision. pp. 21–37. Springer (2016)
22. Luo, Z., Cai, Z., Zhou, C., Zhang, G., Zhao, H., Yi, S., Lu, S., Li, H., Zhang, S., Liu, Z.: Unsupervised domain adaptive 3d detection with multi-level consistency. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8866–8875 (2021)
23. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 779–788 (2016)
24. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems* **28** (2015)
25. Rezaeianaran, F., Shetty, R., Aljundi, R., Reino, D.O., Zhang, S., Schiele, B.: Seeking similarities over differences: Similarity-based domain alignment for adaptive object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9204–9213 (2021)
26. Saito, K., Ushiku, Y., Harada, T., Saenko, K.: Strong-weak distribution alignment for adaptive object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6956–6965 (2019)
27. Sakaridis, C., Dai, D., Van Gool, L.: Semantic foggy scene understanding with synthetic data. *International Journal of Computer Vision* **126**(9), 973–992 (2018)
28. Sohn, K., Berthelot, D., Carlini, N., Zhang, Z., Zhang, H., Raffel, C.A., Cubuk, E.D., Kurakin, A., Li, C.L.: Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in Neural Information Processing Systems* **33**, 596–608 (2020)
29. Sohn, K., Zhang, Z., Li, C.L., Zhang, H., Lee, C.Y., Pfister, T.: A simple semi-supervised learning framework for object detection. arXiv preprint arXiv:2005.04757 (2020)
30. Sun, Q., Li, J., Peng, H., Wu, J., Ning, Y., Yu, P.S., He, L.: Sugar: Subgraph neural network with reinforcement pooling and self-supervised mutual information mechanism. In: Proceedings of the Web Conference 2021. pp. 2081–2091 (2021)
31. Tian, Z., Shen, C., Chen, H., He, T.: Fcos: Fully convolutional one-stage object detection. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 9627–9636 (2019)
32. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
33. Wang, W., Cao, Y., Zhang, J., He, F., Zha, Z.J., Wen, Y., Tao, D.: Exploring sequence feature alignment for domain adaptive detection transformers. In: Proceedings of the 29th ACM International Conference on Multimedia. pp. 1730–1738 (2021)



34. Xu, C.D., Zhao, X.R., Jin, X., Wei, X.S.: Exploring categorical regularization for domain adaptive object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11724–11733 (2020)
35. Xu, M., Zhang, Z., Hu, H., Wang, J., Wang, L., Wei, F., Bai, X., Liu, Z.: End-to-end semi-supervised object detection with soft teacher. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3060–3069 (2021)
36. Xu, M., Wang, H., Ni, B., Tian, Q., Zhang, W.: Cross-domain detection via graph-induced prototype alignment. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12355–12364 (2020)
37. Xu, T., Chen, W., Wang, P., Wang, F., Li, H., Jin, R.: Cdtrans: Cross-domain transformer for unsupervised domain adaptation. arXiv preprint arXiv:2109.06165 (2021)
38. Yang, J., Liu, J., Xu, N., Huang, J.: Tvt: Transferable vision transformer for unsupervised domain adaptation. arXiv preprint arXiv:2108.05988 (2021)
39. Yang, Z., Liu, S., Hu, H., Wang, L., Lin, S.: Reppoints: Point set representation for object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9657–9666 (2019)
40. Yu, F., Xian, W., Chen, Y., Liu, F., Liao, M., Madhavan, V., Darrell, T.: Bdd100k: A diverse driving video database with scalable annotation tooling. arXiv preprint arXiv:1805.04687 2(5), 6 (2018)
41. Zhang, C., Zhang, M., Zhang, S., Jin, D., Zhou, Q., Cai, Z., Zhao, H., Liu, X., Liu, Z.: Delving deep into the generalization of vision transformers under distribution shifts. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7277–7286 (2022)
42. Zhang, J., Huang, J., Luo, Z., Zhang, G., Lu, S.: Da-detr: Domain adaptive detection transformer by hybrid attention. arXiv preprint arXiv:2103.17084 (2021)
43. Zheng, L., Shen, L., Tian, L., Wang, S., Wang, J., Tian, Q.: Scalable person re-identification: A benchmark. In: Proceedings of the IEEE international conference on computer vision. pp. 1116–1124 (2015)
44. Zhong, C., Wang, J., Feng, C., Zhang, Y., Sun, J., Yokota, Y.: Pica: Point-wise instance and centroid alignment based few-shot domain adaptive object detection with loose annotations. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 2329–2338 (2022)
45. Zhou, H., Zhang, S., Peng, J., Zhang, S., Li, J., Xiong, H., Zhang, W.: Informer: Beyond efficient transformer for long sequence time-series forecasting. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 11106–11115 (2021)
46. Zhu, X., Pang, J., Yang, C., Shi, J., Lin, D.: Adapting object detectors via selective cross-domain alignment. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 687–696 (2019)
47. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. arXiv preprint arXiv:2010.04159 (2020)