

Large Self-Supervised Models Bridge the Gap in Domain Adaptive Object Detection

Marc-Antoine Lavoie Anas Mahmoud Steven L. Waslander
 University of Toronto Robotics Institute

{marc-antoine.lavoie, anas.mahmoud, steven.waslander}@robotics.utoronto.ca

Abstract

The current state-of-the-art methods in domain adaptive object detection (DAOD) use Mean Teacher self-labelling, where a teacher model, directly derived as an exponential moving average of the student model, is used to generate labels on the target domain which are then used to improve both models in a positive loop. This couples learning and generating labels on the target domain, and other recent works also leverage the generated labels to add additional domain alignment losses. We believe this coupling is brittle and excessively constrained: there is no guarantee that a student trained only on source data can generate accurate target domain labels and initiate the positive feedback loop, and much better target domain labels can likely be generated by using a large pretrained network that has been exposed to much more data. Vision foundational models are exactly such models, and they have shown impressive task generalization capabilities even when frozen. We want to leverage these models for DAOD and introduce DINO Teacher, which consists of two components. First, we train a new labeller on source data only using a large frozen DINOv2 backbone and show it generates more accurate labels than Mean Teacher. Next, we align the student’s source and target image patch features with those from a DINO encoder, driving source and target representations closer to the generalizable DINO representation. We obtain state-of-the-art performance on multiple DAOD datasets. Code available at https://github.com/TRAILab/DINO_Teacher.

1. Introduction

Modern deep learning methods can achieve strong performance on a wide range of benchmark datasets for many computer vision tasks, including classification, object detection, and segmentation. However, this performance often degrades significantly when tested on data outside the training distribution. Domain adaptation addresses this is-

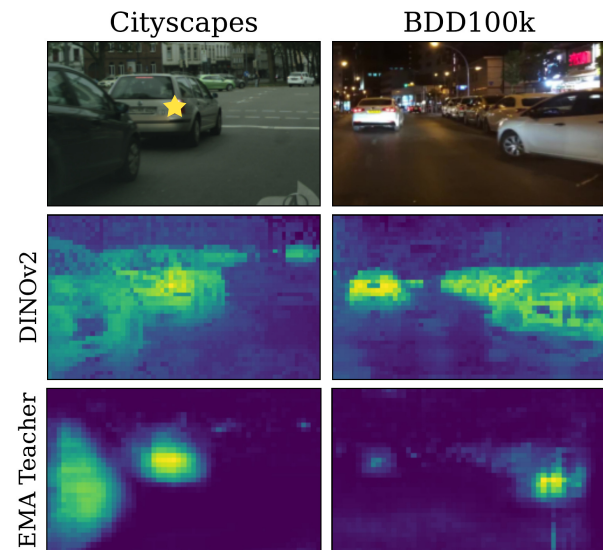


Figure 1. **Cosine similarity of patch feature from Cityscapes to BDD100k.** We evaluate the similarity between the yellow star region in the Cityscapes image and all other regions in both images. Compared to the EMA Teacher, DINOv2 generates semantically stable features across domains, justifying the choice to use it for feature alignment and as a frozen backbone for our labeller.

sue by transferring the performance of a network trained on a labelled source dataset to an unlabelled target dataset [17]. In this work, we specifically consider the domain adaptive object detection (DAOD) task.

Two common techniques for achieving domain adaptation are domain invariance and self-labelling. Domain invariance attempts to minimize the difference between features from the source and target domains, either by using an adversarial strategy with a domain discriminator [6, 17, 22, 47, 54, 56] or by minimizing some distance metric between domains [8, 26, 36]. Self-labelling methods instead attempt to generate their own pseudo-labels on the target dataset. The most common framework for this is Mean Teacher [10, 41, 53], where the labeller is directly

derived from the model being trained. This is a bootstrapping method, where the model selects confident proposals to use as ground truth to train on the target domain, thereby improving its performance and generating more and higher quality pseudo-labels at the next iteration. Because the performance of this self-labelling approach is sensitive to the choice of pseudo-labels to train on, many methods attempt to refine label selection by going beyond simple class confidence thresholding [11, 27, 34].

At the intersection of the two sets of methods are works that use self-generated labels to refine domain invariance approaches [2, 36, 70]. By using labels, it is possible to perform alignment at the proposal [2] or pixel levels [60] and constrained to instances of the same class, improving on the earlier image-level methods. However, these methods are limited by the performance of the teacher network and can perform poorly if some classes are undersampled in the target domain. These issues make it particularly difficult to adapt to small datasets with significant domain gaps, such as severe adverse weather in autonomous driving.

Recently, large-scale vision foundation models (VFMs), whether purely vision-based (DINOv2 [44], SAM [30]) or combining vision and language (CLIP [45]), have become widely available. These models combine large architectures with self-supervised pre-training on large-scale training datasets and have shown impressive performance on multiple computer vision tasks, including as guides in augmentation and generation methods for domain adaptation [40] and generalization [1, 15, 65]. Given this, we believe the current Mean Teacher paradigm for DAOD is outdated. While there are obvious applications of domain adaptation on small architectures for robotics or autonomous driving, and pseudo-labelling remains a powerful strategy, we believe it is no longer reasonable to simply use a duplicate of the student model as the teacher. By decoupling labelling from training the student model, we can use more modern methods to solve domain adaptation while maintaining the size of the trained model. Here, we propose to use vision foundation models with minimal adjustments as the source of the labels, highlighting the usefulness of VFMs in multiple adaptation tasks. Similarly, while class-conditional alignment between domains can be beneficial, we choose to decouple cross-domain alignment from accurate instance labels. To this end, we propose to align student representations with those generated from a VFM backbone. Features from these models have been shown to generalize well to different tasks and new target domains. We illustrate these consistent features for similar objects across different datasets in Fig. 1.

Thus, we propose DINO Teacher (DT) for the DAOD task which exploits large vision foundation models in two ways. First, we show that a simple detector added to the frozen encoder of a large-scale model and trained only on

source data is a better target domain labeller than Mean Teacher, demonstrating the drawbacks of the symmetric architecture and limited training data of the Mean Teacher approach. Next, we propose to use the feature space of a large-scale VFM as a proxy for domain alignment. By aligning the student backbone features with those generated by a VFM independently for source and target images, we drive student source and target features to be closer. Our key contributions are:

- We show that using pseudo-labels from a simple detector trained on a frozen large self-supervised encoder gives substantially more accurate pseudo-labels than a student-derived teacher, and using these new labels greatly improves performance.
- Next, we show that aligning the patch-level output of the small student backbone with the frozen large VFM on both source and target images drives the model to learn features that generalize better without using any labels, decoupling feature alignment from pseudo-label quality.
- We achieve state-of-art performance on multiple DAOD benchmarks, with improvements of +7.6% on BDD100k and +2.3% on Foggy Cityscapes. We also present new results on the ACDC test splits, which cover more extensive adverse weather conditions, and significantly improve our baseline.

2. Related Works

2.1. Domain Adaptive Object Detection

Domain adaptive object detection (DAOD) builds on earlier works in domain adaptation for image classification. The DAOD task is to adapt a detector trained on a labelled source domain to an unlabelled target domain. A common approach for domain adaptation is self-labelling or pseudo-labelling. These methods generate labels for the unsupervised target domain by sampling confident proposals from a model initially trained on the source domain only. This often uses the Mean Teacher framework (MT) [10, 41], where an iteration-averaged version of the model generates more robust labels. This framework forms the baseline of many recent works [2, 11, 29, 35]. Because this self-labelling approach attempts to bootstrap itself into generating more and better labels, it is sensitive to errors or biases in the generated pseudo-labels. Many works try to reduce the bias against rare classes by changing the selection threshold [27, 34, 43] or oversampling instances of rare classes [29, 68]. Other methods improve box localization by estimating box uncertainty and using it to select or reweigh boxes [4, 11, 62]. We believe that the coupling between labelling and the performance of the student model in the Mean Teacher framework limits the quality of generated labels. Instead, we propose to decouple label generation from solving the domain adaptation task on the student model by

training an independent labeller to provide the best possible pseudo-labels to the student. Because we still need to generate the best labels to train the student, some of the recent improvements such as oversampling rare classes [29] could be complimentary to our framework.

Another common strategy in domain adaption is domain invariance. This was first formulated as an adversarial loss [17, 54] that pushes image-level backbone features to contain no domain information. Other methods obtain domain invariant features by instead directly minimizing a distance measure between the feature representations of the source and target domains. Some of the distance measures used include contrastive loss [2, 26, 60], graph alignment [18, 36, 37, 39] or optimal transport [8]. The simplest invariance or alignment methods operate on image-level features, but there is no guarantee that image-level alignment leads to the instance-level alignment needed for detection and segmentation tasks [31]. Alternatively, alignment methods can be combined with the self-labelling approaches discussed above, allowing instance-level and class-conditional alignment [2, 6, 37, 61, 70, 72]. However, coupling the alignment term with self-labelling leads to a bind: good labels are required for robust alignment, but domain alignment is also required to generate good labels. Instead, we propose to indirectly align the domains by encouraging the feature space of the student model to be similar to the feature space of a self-supervised VFM on both the source and the target images without requiring any labels.

Other methods use stronger augmentations, reducing overfitting on the source domain and the limited generated target domain labels. This includes implementations of Mixup [28, 68, 69] and CutMix [42, 67], the weak-strong approach of Unbiased Teacher [41] or various frequency-space transforms [40, 63, 65], which assume that certain frequency bands or that the magnitude spectrum are domain-specific and can be perturbed or transferred across images. Similarly, some methods use generative models to learn a mapping between source and target domain [10, 50, 73] to transform labelled source images into target-like images. A final set of methods uses masked reconstruction [23, 71] in the feature space to learn better target domain features. We believe these methods could be complementary to ours.

2.2. Large Models in Domain Adaptation and Generalization

Recently, many methods have begun leveraging VFMs for domain adaptation. Some methods attempt to refine large pre-trained models by fine-tuning only a few parameters [5, 24] or by modifying the input prompt to the encoder [14, 19, 52, 58]. Other methods are similar to the generative approaches discussed above, with methods such as CLIP-Styler [32], PØDA [15], FAMix [16] and ULDA [64] all showing that given aligned vision-language embeddings, it

is possible to synthesize new target-like images or image features by replicating a domain shift encoded by text in the aligned image feature space, allowing adaptation without requiring any target images. While all these methods are interesting avenues of research, we focus on the standard DAOD task that attempts to maximize performance on an unlabelled target dataset given a relatively small model architecture (*e.g.* VGG16 [51]), which are essential when deploying networks in the real world.

Closer to our work, UFR [40] recently used Segment Anything [30] to refine instance masks for augmentation but still trains and evaluates on smaller architectures. While augmentation methods are useful for domain adaptation, self-labelling methods generally show stronger adaptation performance. Given this, we investigate the use of VFMs as labellers. As mentioned previously, one of the advantages of these large models is their generalizability to other downstream tasks. A recent work, REIN [58], shows that fully frozen backbones are nearly on par with fully fine-tuned ones for domain generalized semantic segmentation, confirming strong feature generalizability for computer vision tasks. We want to leverage this capability to use a large model as a labeller on the unlabelled target domain data and as an alignment target for the student network.

3. Method

In this section, we summarize the domain adaptive object detection problem, present the baseline method and explain our contributions. We build primarily on the Adaptive Teacher [38] framework and frequently contrast our design choices with that method. Fig. 2 summarizes our proposed DINO Teacher framework.

3.1. Problem Definition

The problem we consider is domain adaptive object detection, where we want to transfer performance from a labelled source domain, \mathcal{D}_S , to an unlabelled target domain, \mathcal{D}_T , which shares the same classes. We thus have images, X_S , bounding boxes, B_S , and class labels, Y_S , on the source domain while we only have images, X_T , on the target domain. To train on target images, we first have to generate target pseudo-labels, \tilde{B}_T and \tilde{Y}_T . We use Faster R-CNN [46] as our detector, and the training loss is

$$\begin{aligned} \mathcal{L}^{det} &= \mathcal{L}_S^{det}(X_S, B_S, Y_S) + \mathcal{L}_T^{det}(X_T, \tilde{B}_T, \tilde{Y}_T) \quad (1) \\ &= \mathcal{L}_S^{rpn}(X_S, B_S, Y_S) + \mathcal{L}_S^{roi}(X_S, B_S, Y_S) \\ &\quad + \mathcal{L}_T^{rpn}(X_T, \tilde{B}_T, \tilde{Y}_T) + \mathcal{L}_T^{roi}(X_T, \tilde{B}_T, \tilde{Y}_T), \end{aligned}$$

where \mathcal{L}^{rpn} and \mathcal{L}^{roi} are the standard Faster R-CNN region proposal network (RPN) box proposal and Region of Interest (RoI) box refinement and classification losses.

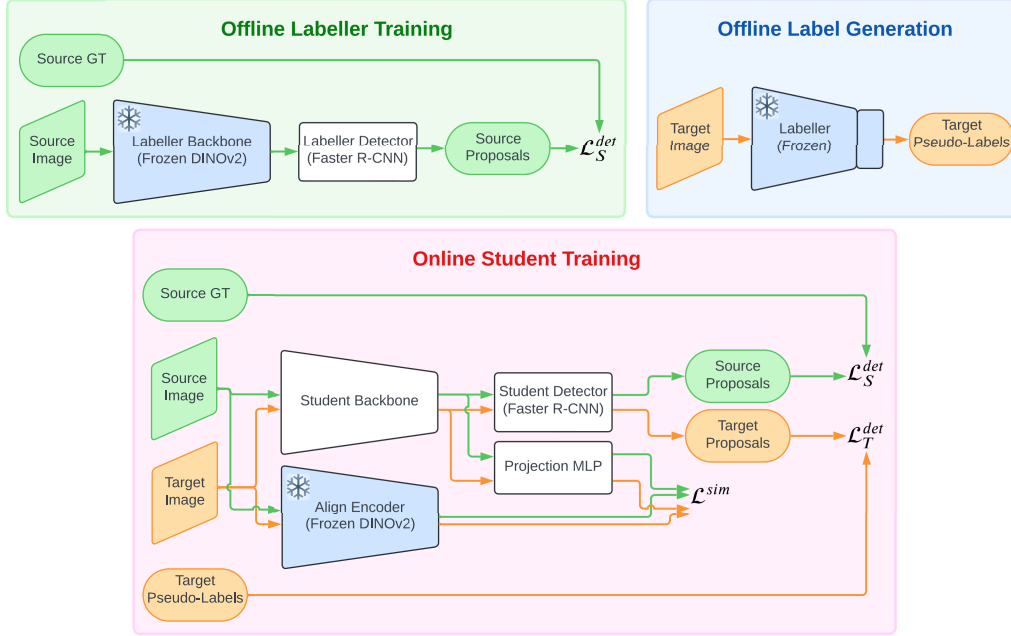


Figure 2. **Diagram of our proposed method.** **Offline Labeller Training:** we add a detector head to a frozen DINOv2 encoder and train it with source images only. **Offline Label Generation:** we combine and freeze the labeller backbone and detector, and generate target pseudo-labels. **Online Student Training:** We train a student network using source ground truth boxes and target pseudo-labels, and align patch features to a frozen DINOv2 encoder. During inference, the alignment encoder and projection MLP are not used.

3.2. Mean Teacher Framework

In the Mean Teacher framework [10], a student model, f_s , trained on the source data is used to instantiate a teacher, f_t , to generate pseudo-labels on the target domain. This teacher is a duplicate of the original network whose parameters, $\bar{\theta}$, are updated as an exponential moving average (EMA) of the student parameters, θ . The teacher is thus essentially an ensemble over iterations of the student training process

$$\bar{\theta} \leftarrow \alpha \bar{\theta} + (1 - \alpha) \theta, \quad (2)$$

where α is the EMA decay factor. To generate pseudo-labels, each target image is augmented using two different augmentations. The first augmented image is fed to the teacher network, which generates pseudo-labels by thresholding box proposals with class probabilities above a given threshold, δ . The second augmented image is sent to the student, and the thresholded teacher proposals are used as ground truth for the student. Pseudo-label training on the target domain starts after some initialization phase of n^{initPL} iterations training exclusively on source images. Unbiased Teacher also proposes to improve learning by using weak-strong augmentations [41]. Simple augmentations (crop, flip) are applied to the teacher images to make labelling easy, while the student images have stronger augmentations (blurring, colour jitter, grayscaleing, Cutout [12]).

3.3. Foundation Models for Pseudo-Labeling

Our method generally follows the same self-labelling strategy, but instead of using the time-averaged student, \bar{f} , as a teacher, we propose to use another model entirely, f^{big} . A key issue with using the student model as a starting point for the teacher is that the model operates outside its training data distribution when labelling target data. Because DINOv2 is trained using self-supervised losses on a large corpus of data [44], we hypothesize that its feature representation will generalize well, including over the considered domain gap from source to target. Thus, we propose to use DINOv2 as the encoder for a new labeller, and we train an object detection head on top of the frozen encoder using only the source data. We confirm that this leads to a transferable detector by evaluating the performance of the DINO labeller on the unseen target data in Tab. 6.

We use DINOv2 ViT-G [13] as our frozen encoder backbone and use a Faster R-CNN detection head. We do not require that the detection head be identical to the student network’s because we only use the output box proposals (see Fig. 2) when training the student, but we use the same detector as the student model to highlight the impact of changing the backbone. We train the labeller detector with the Faster R-CNN loss \mathcal{L}_S^{det} presented in (1).

We generate all pseudo-labels \tilde{B}_T and \tilde{Y}_T for the target domain training examples as a preprocessing step in a

single forward pass over the target training data. To select pseudo-labels, we threshold the bounding box proposals according to class probability similar to Mean Teacher. In our architecture, f^{big} replaces the student-derived teacher \tilde{f} for labelling, but we still generate and evaluate performance on the EMA model \tilde{f} as it is significantly better than the student f on the detection task. Model ensembling [33] improves performance outside of its use as a teacher and is a common training strategy in recent vision models [3, 57].

3.4. Foundation Models for Feature Alignment

As discussed in Sec. 2, existing domain invariance methods operate in two ways. First, they can align at the image level, which does not guarantee alignment at the instance level and can be dominated by background information, but which does not require labels. Alternatively, they can use pseudo-labels to match semantically related regions or instances for alignment, which generally improves performance [37, 70] at the cost of relying on accurate generated labels. We propose an alternative alignment method that aligns semantically similar regions without using generated labels by using a proxy alignment target, to which we can independently align the features of the student model from both source and target images and without having to match regions across domains.

To this end, we use the feature representation of a VFM as our alignment target. Again, we hypothesize that the DINOv2 features will generalize better to the target domain and that the representation of any particular object class will be better defined across domains than the representation of a small student model trained on source data only. To improve the student, we propose to align patch-level features of the student backbone with those of a frozen DINOv2 encoder on individual images of both the source and target domains. This drives the student representation of both domains towards the large model’s representation, reducing the domain gap. We align patch features of the same image processed by the student and by the DINO encoder, ensuring that we respect local semantic information: the same car patch is encoded by both the student backbone and the DINO alignment teacher. We compare the clustering of the instance-level features of the student network with and without alignment to DINOv2 and the original DINOv2 clustering in Fig. 3 using t-SNE [55]. This shows that DINO features have better rare class separation and that aligning the student to DINOv2 improves separation on both the source domain Cityscapes [7] and the target domain BDD [66].

We take DINOv2 ViT-B [13] as our frozen encoder backbone. Given patch-level DINO teacher features \mathbf{x}^{big} and student features \mathbf{x} , we pass the student features through a 2-layer projection MLP $g(\cdot)$ to match the channel dimensions C and then use bilinear interpolation to match the di-

mensions H and W of the DINO features. The similarity between the normalized features of both backbones is maximized with the alignment loss

$$\mathcal{L}^{\text{sim}} = \frac{1}{NHW} \sum_k \sum_{ij} 1 - \frac{\text{interp}(g(\mathbf{x}))^T \mathbf{x}^{\text{big}}}{\|\text{interp}(g(\mathbf{x}))\|_2 \|\mathbf{x}^{\text{big}}\|_2}, \quad (3)$$

where N , H and W are the batch size, height and width of the final DINO encoder feature map, and where we have suppressed the features indices i, j, k for ease of reading.

In implementation, we align the student and DINO features on the source domain from the start, pushing the initial student representation to be predictive of the DINO representation and to learn the projection MLP. After n^{initSim} , we align with both source and target images. Notably, we start aligning on the target domain before using any pseudo-labels, different from Adaptive Teacher which starts both simultaneously. We begin domain alignment early to reduce issues of unaligned gradients from supervised training in different domains.

In contrast to the labeller described in Sec. 3.3 above, we use the smaller DINOv2 ViT-B as the alignment encoder. While the labeller only runs once on each image of the target training dataset and this is done before training the student, the alignment DINO encoder is used during online training, and we found that inference on the larger ViT models was excessively slow. Our method does not require that the two DINO encoders be the same, and the encoders are never used simultaneously (see Fig. 2).

3.5. Complete Method

Our full loss combines the pseudo-labelling detection loss (1) with the similarity loss (3)

$$\mathcal{L} = \mathcal{L}_S^{\text{det}} + \lambda^{\text{unsup}} \mathcal{L}_T^{\text{det}} + \lambda^{\text{sim}} \mathcal{L}^{\text{sim}}, \quad (4)$$

where λ^{unsup} and λ^{sim} are hyperparameters that tune the respective weights of the losses associated with the pseudo-labels and the similarity term.

4. Experiments

4.1. Datasets

Cityscapes. All of the reported runs use Cityscapes [7] as the source dataset. Cityscapes is an autonomous driving dataset with 2975 training images and 500 validation images, with detection labels for 8 classes.

Foggy Cityscapes. Foggy Cityscapes [48] is a synthesized dataset that simulates foggy weather at three intensities (0.005, 0.01, 0.02) using the original clear images and a depth map. It has 8925 training images and 1500 validation images, one for each image in Cityscapes for each intensity level. We evaluate on the full validation dataset.




Figure 3. **t-SNE of the backbone instance-level embeddings across domains.** Subfigures (a) and (b) are taken from VGG16 features after 20k training iterations of supervised training on source only (Cityscapes), following the protocol defined in Sec. 4. Without alignment (a), there is confusion between all similar classes. With alignment (b), overlap is reduced, particularly between **persons** (•, ×) and **riders** (•, ×). Pretrained DINOv2 (c) has well separated clusters.

BDD100k. The BDD100k [66] dataset is another large autonomous driving dataset. Following previous works [21, 72], we train and evaluate on images of the Day-time split, which corresponds to a domain shift in camera and geographic location, but without significant change in weather. It contains 36,728 training and 5258 validation images. It shares 7 of the 8 classes present in Cityscapes.

ACDC. The Adverse Conditions Dataset with Correspondences (ACDC) [49] is an autonomous driving dataset focused on four adverse weather conditions: fog, night, rain and snow. It shares the same 8 classes as Cityscapes for object detection. Each split has 400 training and 100 validation images (106 for night validation). We train and evaluate on each of the individual splits.

4.2. Experimental Setup

We base our implementation on the Detectron2 framework [59]. For runs on Foggy Cityscapes and BDD100k, we follow past works [2] and use VGG16 [51] pretrained on ImageNet [9] as the backbone. For runs on ACDC, we use the common ResNet-50 [20], also pretrained on ImageNet as the backbone. Single-scale Faster R-CNN is the detector in all experiments. Our models are trained on 2 RTX 6000 and use 8 source and 8 target images in each batch. We generally follow the test setting of Adaptive Teacher [38]. When reimplementing baselines, images are scaled to have a height of 600 pixels for training. When aligning with DINO as per section 3.4, we scale the height to 588 pixels to make it divisible by the ViT patch size of 14. All the results we present, including baselines, are evaluated from an EMA model \bar{f} as discussed in Sec. 3.3.

In all of our runs and for all teachers, we use a pseudo-labelling threshold of $\delta = 0.8$ and use weak-strong augmentations [38] as our only augmentation. We set the loss hyperparameters in equation (4) as $\lambda^{unsup} = 1$ and $\lambda^{sim} = 1$. For runs on VGG16, we use a learning rate of 0.04 with-

out decay, an EMA decay factor of $\alpha = 0.9996$, pretrain on source only for $n^{initPL} = 20,000$ iterations and train for a total of 60,000 iterations. For runs on ResNet-50, we use a learning rate of 0.01 without decay, an EMA decay factor of $\alpha = 0.999$, pretrain on source for $n^{initPL} = 20,000$ iterations and train for a total of 40,000 iterations.

We use the same DINOv2-based labeller to generate pseudo-labels for all runs following section 3.3. We train a Faster R-CNN head on top of a ViT-G encoder with frozen DINOv2 weights and use the same training setting and augmentations as the ResNet-50 student model but use only the source Cityscapes images.

We use a frozen DINOv2 ViT-B as the alignment encoder following section 3.4. The projection MLP has a hidden layer dimension of 1024. We initialize the similarity branch on source images until $n^{initSim} = 5000$ iterations. We discuss the impact of the size of the DINO model in Tab. 5.

Because of the small size of the training and validation datasets for the individual ACDC splits, we present the average performance across three runs for both the baseline Adaptive Teacher and our method. When replicating the baseline, we use the same parameters used in [38] ($\lambda_{unsup} = 1$, $\lambda_{dis} = 0.1$) and start both the domain discrimination and the pseudo-labelling after 20,000 iterations.

4.3. Results

We compare the performance of our method with other state-of-the-art methods and report class-averaged precision at 50% IoU (mAP@50). In the tables, we use the following abbreviations: **FR** for Faster R-CNN, **GA** for graph alignment, and **MT** for Mean Teacher. We evaluate across different environments, cameras and adverse weather conditions.

Cityscapes \rightarrow BDD100k. The first test setting considers domain adaptation to BDD100K, corresponding to different environment and camera. We present our results in

Method	Type	Person	Rider	Car	Truck	Bus	Motor	Bicycle	mAP
DA-Faster [6]	FR	28.9	27.4	44.2	19.1	18.0	14.2	22.4	24.9
SIGMA [37]	FCOS+GA	46.9	29.6	64.1	20.2	23.6	17.9	26.3	32.7
TDD [21]	FR+MT	39.6	38.9	53.9	24.1	25.5	24.5	28.8	33.6
PT [4]	FR+MT	40.5	39.9	52.7	25.8	33.8	23.0	28.8	34.9
NSA [72]	FR+MT	-	-	-	-	-	-	-	35.5
REACT [35]	FR+MT	-	-	-	-	-	-	-	35.8
CAT [29]	FR+MT	44.6	41.5	61.2	31.4	34.6	24.4	31.7	38.5
HT [11]	FCOS+MT	53.4	40.4	63.5	27.4	30.6	28.2	38.0	40.2
DT (ours)	FR+MT	51.6	47.0	66.6	44.3	45.9	38.3	40.8	47.8

Table 1. Results for domain adaptive object detection from Cityscapes to BDD100k. Best performing methods are **bolded**, second best are underlined.

Method	Type	Person	Rider	Car	Truck	Bus	Train	Motor	Bicycle	mAP
DA-Faster [6]	FR	29.2	40.4	43.4	19.7	38.3	28.5	23.7	32.7	32.0
DICN [25]	FR	47.3	57.4	64.0	22.7	45.6	29.6	38.6	47.4	44.1
NLTE [39]	FR+GA	43.1	50.7	58.7	33.6	56.7	42.7	33.7	43.3	45.4
PT [4]	FR+MT	43.2	52.4	63.4	33.4	56.6	37.8	41.3	48.7	47.1
MIC [23]	FR+MT	50.9	55.3	67.0	33.9	52.4	33.7	40.6	47.5	47.6
AT [38]	FR+MT	45.5	55.1	64.2	35.0	56.3	54.3	38.5	51.9	50.9
CMT [2]	FR+MT	47.0	55.7	64.5	39.4	63.2	51.9	40.3	53.1	51.9
REACT [35]	FR+MT	51.4	57.9	67.4	37.7	58.4	52.8	44.6	54.6	53.1
DT (ours)	FR+MT	48.5	60.0	65.4	47.2	66.5	52.9	46.2	56.7	55.4

Table 2. Results for domain adaptive object detection from Cityscapes to Foggy Cityscapes for the entire validation set. Best performing methods are **bolded**, second best are underlined.

Tab. 1. We obtain state-of-the-art results, exceeding the performance of Harmonious Teacher [11] by +7.6%, with particularly large improvements on the rarer classes. Given a larger target dataset, we show that improving the quality of labels by using a better labeller improves performance substantially over using the Mean Teacher framework. We discuss the correlation between labeller size and student network performance in Tab. 5a.

Cityscapes \rightarrow **Foggy Cityscapes**. The second test considers domain adaptation across a relatively easy adverse weather condition, the synthetic addition of fog. We present our results in Tab. 2 for the full validation set and obtain state-of-the-art performance on this dataset, improving over the recent REACT [35] method by +2.3% and our baseline architecture AT by +4.5% when using the ViT-G labels. Compared to our results on BDD100k, we show more minor improvements on this dataset. We believe this is due to the small domain gap between Cityscapes and Foggy Cityscapes, which is generated by a consistent procedure on the original Cityscapes images and uses the same labels.

Cityscapes \rightarrow **ACDC splits**. The final test setting considers domain adaptation in the most complex setting: adverse weather, different camera and environment, and with

Method	Fog	Night	Rain	Snow
AT [†] [38]	62.2	29.5	37.7	55.2
DT-G (ours)	68.6	36.4	39.0	56.8

Table 3. Results for domain adaptive object detection from Cityscapes to the 4 ACDC splits. [†]We reimplement AT from publicly available code.

a much smaller target dataset. We present our results in Tab. 3. Here as well, we observe consistent and significant performance improvements compared to the Adaptive Teacher baseline, improving by +6.4% on fog, +6.9% on night, +1.3% on rain and +1.6% on snow. While the DINO pseudo-labels do better than Mean Teacher labels, particularly for the most challenging night split, we believe the quality of pseudo-labels limits our improvements. We discuss this further in section 4.5 below. Another issue is the small size of the target dataset, which leads to overfitting on the pseudo-labels and restricts performance on the test set.

4.4. Ablations

We run our ablations on the Cityscapes to BDD100k test case, as the large dataset size and significant domain gap best exemplify performance improvements from different methodological choices.

Method	Label.	Align.	n^{initPL}	Best mAP
AT [†] [38]	MT	\mathcal{L}^{dis}	28.5	31.8
case 1	MT	\mathcal{L}^{sim}	32.5	35.3
case 2	DL	\mathcal{L}^{dis}	28.5	46.8
DT	DL	\mathcal{L}^{sim}	33.0	47.8

Table 4. **Ablation of the DT components on Cityscapes to BDD100k.** We use **MT** for Mean Teacher labels and **DL** for DINO labeller. [†]We reimplement AT from publicly available code.

Ablations on Components. Tab. 4 presents the ablation on the contributions of each component of our approach. Starting from the Adaptive Teacher baseline, we substitute the standard EMA teacher labeller **MT** with our DINO labeller **DL** (section 3.3) and substitute the domain adversarial loss \mathcal{L}^{dis} with our patch similarity loss \mathcal{L}^{sim} (section 3.4). We report the best performance and performance at n^{initPL} iterations, just before any pseudo-labels are used, to separate the impacts of changes in alignment and labelling methods. We see that while our alignment loss \mathcal{L}^{sim} improves performance by +4.0% before pseudo-labels are introduced and overall by +3.5% (AT vs case 1), it doesn’t allow the Mean Teacher framework to generate improvements similar to the DINO labeller (case 1 vs DT). Using DINO labels improves performance by over 10% regardless of the type of alignment used (AT vs case 2 or DT).

Teacher Size. Next, we consider the impact of the size of the DINOv2 model used as the backbone to generate pseudo-labels (section 3.3) and as the alignment target (section 3.4). Results are presented in Tab. 5. The ViT-S, B, L and G models have 21, 86, 300 and 1,100 million parameters respectively. Tab. 5a presents the ablation on DINO labeller size, and scaling the backbone from ViT-B to ViT-G improves our results by +1.8%. Interestingly, the final student model is better than the labeller for ViT-B and ViT-L, meaning labeller performance is not a direct upper bound. While the labeller is trained on source only, the student uses a high-confidence subset of target labels and true source labels to learn a representation that generalizes to both domains. From Tab. 5b, we also find that aligning with the larger ViT-B model improves by +0.4% compared to ViT-S, and both do better than the baseline.

4.5. DINO Teacher vs Other Baselines

Finally, we consider how our DINO Teacher compares to other strategies for object detection on the target dataset and present the results in Tab. 6. First, the ViT-G DINO labeller generalizes much better than the source-only student model, validating our choice of using a VFM labeller for pseudo-label generation. Our DT approach significantly outperforms source-only student training and target-only training on the small ACDC splits, but it lags behind target-only

Method	Lab.	Stud.	Method	n^{initPL}	mAP
MT	-	35.3	\mathcal{L}^{dis}	28.5	46.8
DL-B	42.7	46.0	$\mathcal{L}^{sim}\mathbf{S}$	31.9	47.4
DL-L	45.7	46.9	$\mathcal{L}^{sim}\mathbf{B}$	33.0	47.8
DL-G	51.1	47.8			

(a) Labeller size

(b) Alignment teacher size

Table 5. **Ablations of the size of DINOv2 backbone model for the labeller and alignment teacher.** We consider ViT-B, L and G for the DINO labeller (**DL-B**, **DL-L**, **DL-G**), and ViT-S and B for the alignment loss ($\mathcal{L}^{sim}\mathbf{S}$, $\mathcal{L}^{sim}\mathbf{B}$).

Method	BDD	ACDC			
		Fog	Night	Rain	Snow
Source only	29.1	61.2	19.7	37.0	46.3
Target only	54.4	48.6	24.6	30.5	46.0
ViT-G labeller	51.1	65.8	39.4	47.3	49.3
w/ GT labels [†]	51.6	70.3	37.2	44.0	57.5
DT*	47.8	68.6	36.4	39.0	56.8

Table 6. **Performance of the DT model compared to other strategies.** *Reported results for DT on ACDC are the average of 3 runs. [†]Network trained on source and target with all labels.

training on the larger BDD dataset, showing the advantages of domain adaptation for small datasets. An alternative upper bound for domain adaptation is the performance with an ideal labeller, where ground-truth target domain labels are available and training is done on source and target images, denoted w/ GT labels in Tab. 6. This improves performance compared to target-only training on ACDC but not for BDD, meaning training with the source data can reduce performance given sufficient labels on the larger target dataset. Our DT model is close to this bound, but the gap indicates that a better labeller could yield improvements.

5. Conclusion

We present a novel framework for domain adaptive object detection, DINO Teacher, in which we leverage the generalizability of the features of vision foundational models such as DINOv2. Our DINO Teacher approach is composed of two components. First, we create a new labeller by training a detection head on a frozen VFM backbone using only source data. We show that this simple framework outperforms the standard Mean Teacher self-labelling approach, particularly on the BDD100k dataset, demonstrating the limitations of the standard Mean Teacher architecture. Next, we show that aligning image features from the backbone of the student model with those from a frozen VFM encoder independently for source and target images pushes the student to learn features that are more discriminative and generalize better from the source to the target domain, improving performance.

References

- [1] Yasser Benigmim, Subhankar Roy, Slim Essid, Vicky Kalogeiton, and Stéphane Lathuilière. Collaborating foundation models for domain generalized semantic segmentation. In *CVPR*, pages 3108–3119, 2024. 2
- [2] Shengcao Cao, Dhiraj Joshi, Liang-Yan Gui, and Yu-Xiong Wang. Contrastive Mean Teacher for domain adaptive object detectors. In *CVPR*, pages 23839–23848, 2023. 2, 3, 6, 7
- [3] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, pages 9650–9660, 2021. 5
- [4] Meilin Chen, Weijie Chen, Shicai Yang, Jie Song, Xinchao Wang, Lei Zhang, Yunfeng Yan, Donglian Qi, Yueting Zhuang, Di Xie, et al. Learning domain adaptive object detection with Probabilistic Teacher. In *ICML*, pages 3040–3055. PMLR, 2022. 2, 7
- [5] Shoufa Chen, Chongjian Ge, Zhan Tong, Jiangliu Wang, Yibing Song, Jue Wang, and Ping Luo. AdaptFormer: Adapting vision transformers for scalable visual recognition. *NeurIPS*, 35:16664–16678, 2022. 3
- [6] 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*, pages 3339–3348, 2018. 1, 3, 7
- [7] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The Cityscapes dataset for semantic urban scene understanding. In *CVPR*, pages 3213–3223, 2016. 5
- [8] Nicolas Courty, Remi Flamary, Alain Rakotomamonjy, and Devis Tuia. Optimal transport for domain adaptation. In *NIPS 2014, Workshop on Optimal Transport and Machine Learning*, 2014. 1, 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 *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 6
- [10] Jinhong Deng, Wen Li, Yuhua Chen, and Lixin Duan. Unbiased Mean Teacher for cross-domain object detection. In *CVPR*, pages 4091–4101, 2021. 1, 2, 3, 4
- [11] Jinhong Deng, Dongli Xu, Wen Li, and Lixin Duan. Harmonious Teacher for cross-domain object detection. In *CVPR*, pages 23829–23838, 2023. 2, 7
- [12] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017. 4
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2020. 4, 5
- [14] Zhekai Du, Xinyao Li, Fengling Li, Ke Lu, Lei Zhu, and Jingjing Li. Domain-agnostic mutual prompting for unsupervised domain adaptation. In *CVPR*, pages 23375–23384, 2024. 3
- [15] Mohammad Fahes, Tuan-Hung Vu, Andrei Bursuc, Patrick Pérez, and Raoul De Charette. PODA: Prompt-driven zero-shot domain adaptation. In *ICCV*, pages 18623–18633, 2023. 2, 3
- [16] Mohammad Fahes, Tuan-Hung Vu, Andrei Bursuc, Patrick Pérez, and Raoul de Charette. A simple recipe for language-guided domain generalized segmentation. In *CVPR*, pages 23428–23437, 2024. 3
- [17] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, pages 1180–1189. PMLR, 2015. 1, 3
- [18] Changlong Gao, Chengxu Liu, Yujie Dun, and Xueming Qian. CSDA: Learning category-scale joint feature for domain adaptive object detection. In *ICCV*, pages 11421–11430, 2023. 3
- [19] Chunjiang Ge, Rui Huang, Mixue Xie, Zihang Lai, Shiji Song, Shuang Li, and Gao Huang. Domain adaptation via prompt learning. *IEEE Transactions on Neural Networks and Learning Systems*, 2023. 3
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 6
- [21] 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 *CVPR*, pages 9570–9580, 2022. 6, 7
- [22] Judy Hoffman, Dequan Wang, Fisher Yu, and Trevor Darrell. Fcns in the wild: Pixel-level adversarial and constraint-based adaptation. *arXiv preprint arXiv:1612.02649*, 2016. 1
- [23] Lukas Hoyer, Dengxin Dai, Haoran Wang, and Luc Van Gool. MIC: Masked image consistency for context-enhanced domain adaptation. In *CVPR*, pages 11721–11732, 2023. 3, 7
- [24] Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. LoRA: Low-rank adaptation of large language models. In *ICLR*, 2022. 3
- [25] Yifan Jiao, Hantao Yao, and Changsheng Xu. Dual instance-consistent network for cross-domain object detection. *IEEE TPAMI*, 45(6):7338–7352, 2022. 7
- [26] Guoliang Kang, Lu Jiang, Yi Yang, and Alexander G Hauptmann. Contrastive adaptation network for unsupervised domain adaptation. In *CVPR*, pages 4893–4902, 2019. 1, 3
- [27] Purbayan Kar, Vishal Chudasama, Naoyuki Onoe, and Pankaj Wasnik. Revisiting class imbalance for end-to-end semi-supervised object detection. In *CVPR*, pages 4570–4579, 2023. 2
- [28] Mikhail Kennerley, Jian-Gang Wang, Bharadwaj Veeravalli, and Robby T Tan. 2PCNet: Two-phase consistency training for day-to-night unsupervised domain adaptive object detection. In *CVPR*, pages 11484–11493, 2023. 3
- [29] Mikhail Kennerley, Jian-Gang Wang, Bharadwaj Veeravalli, and Robby T Tan. CAT: Exploiting inter-class dynamics for domain adaptive object detection. In *CVPR*, pages 16541–16550, 2024. 2, 3, 7
- [30] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer White-

- head, Alexander C Berg, Wan-Yen Lo, et al. Segment Anything. In *ICCV*, pages 4015–4026, 2023. 2, 3
- [31] Abhishek Kumar, Prasanna Sattigeri, Kahini Wadhawan, Leonid Karlinsky, Rogerio Feris, Bill Freeman, and Gregory Wornell. Co-regularized alignment for unsupervised domain adaptation. *NeurIPS*, 31, 2018. 3
- [32] Gihyun Kwon and Jong Chul Ye. CLIPstyler: Image style transfer with a single text condition. In *CVPR*, pages 18062–18071, 2022. 3
- [33] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. In *ICLR*, 2017. 5
- [34] Hengduo Li, Zuxuan Wu, Abhinav Shrivastava, and Larry S Davis. Rethinking pseudo labels for semi-supervised object detection. In *AAAI*, pages 1314–1322, 2022. 2
- [35] Haochen Li, Rui Zhang, Hantao Yao, Xin Zhang, Yifan Hao, Xinkai Song, and Ling Li. REACT: Remainder adaptive compensation for domain adaptive object detection. *IEEE TIP*, 2024. 2, 7
- [36] Wuyang Li, Xinyu Liu, Xiwen Yao, and Yixuan Yuan. SCAN: Cross domain object detection with semantic conditioned adaptation. In *AAAI*, pages 1421–1428, 2022. 1, 2, 3
- [37] Wuyang Li, Xinyu Liu, and Yixuan Yuan. SIGMA: Semantic-complete graph matching for domain adaptive object detection. In *CVPR*, pages 5291–5300, 2022. 3, 5, 7
- [38] 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 *CVPR*, pages 7581–7590, 2022. 3, 6, 7, 8
- [39] Xinyu Liu, Wuyang Li, Qiushi Yang, Baopu Li, and Yixuan Yuan. Towards robust adaptive object detection under noisy annotations. In *CVPR*, pages 14207–14216, 2022. 3, 7
- [40] Yajing Liu, Shijun Zhou, Xiyao Liu, Chunhui Hao, Baojie Fan, and Jiandong Tian. Unbiased Faster R-CNN for single-source domain generalized object detection. In *CVPR*, pages 28838–28847, 2024. 2, 3
- [41] Yen-Cheng Liu, Chih-Yao Ma, Zijian He, Chia-Wen Kuo, Kan Chen, Peizhao Zhang, Bichen Wu, Zsolt Kira, and Peter Vajda. Unbiased Teacher for semi-supervised object detection. In *ICLR*, 2021. 1, 2, 3, 4
- [42] Giulio Mattolin, Luca Zanella, Elisa Ricci, and Yiming Wang. ConfMix: Unsupervised domain adaptation for object detection via confidence-based mixing. In *Winter Conf. on Applications of Comput. Vis. (WACV)*, pages 423–433, 2023. 3
- [43] Yuxiang Nie, Chaowei Fang, Lechao Cheng, Liang Lin, and Guanbin Li. Adapting object size variance and class imbalance for semi-supervised object detection. In *AAAI*, pages 1966–1974, 2023. 2
- [44] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. DINOv2: Learning robust visual features without supervision. *Transactions on Machine Learning Research*, pages 1–31, 2024. 2, 4
- [45] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763. PMLR, 2021. 2
- [46] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. *NeurIPS*, 28, 2015. 3
- [47] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. Strong-weak distribution alignment for adaptive object detection. In *CVPR*, pages 6956–6965, 2019. 1
- [48] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Semantic foggy scene understanding with synthetic data. *IJCV*, 126:973–992, 2018. 5
- [49] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. ACDC: The adverse conditions dataset with correspondences for semantic driving scene understanding. In *ICCV*, pages 10765–10775, 2021. 6
- [50] Yuhu Shan, Wen Feng Lu, and Chee Meng Chew. Pixel and feature level based domain adaptation for object detection in autonomous driving. *Neurocomputing*, 367:31–38, 2019. 3
- [51] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 3, 6
- [52] Mainak Singha, Harsh Pal, Ankit Jha, and Biplab Banerjee. AD-CLIP: Adapting domains in prompt space using clip. In *ICCV*, pages 4355–4364, 2023. 3
- [53] Antti Tarvainen and Harri Valpola. Mean Teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *NeurIPS*, 30, 2017. 1
- [54] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *CVPR*, pages 7167–7176, 2017. 1, 3
- [55] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9 (11), 2008. 5
- [56] 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*, pages 4516–4526, 2021. 1
- [57] Chien-Yao Wang, Alexey Bochkovskiy, and Hong-Yuan Mark Liao. Yolov7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. In *CVPR*, pages 7464–7475, 2023. 5
- [58] Zhixiang Wei, Lin Chen, Yi Jin, Xiaoxiao Ma, Tianle Liu, Pengyang Ling, Ben Wang, Huaian Chen, and Jinjin Zheng. Stronger fewer & superior: Harnessing vision foundation models for domain generalized semantic segmentation. In *CVPR*, pages 28619–28630, 2024. 3
- [59] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019. 6
- [60] Binhui Xie, Shuang Li, Mingjia Li, Chi Harold Liu, Gao Huang, and Guoren Wang. SePiCo: Semantic-guided pixel contrast for domain adaptive semantic segmentation. *IEEE TPAMI*, 45(7):9004–9021, 2023. 2, 3
- [61] Minghao Xu, Hang Wang, Bingbing Ni, Qi Tian, and Wenjun Zhang. Cross-domain detection via graph-induced prototype alignment. In *CVPR*, pages 12355–12364, 2020. 3

- [62] Mengde Xu, Zheng Zhang, Han Hu, Jianfeng Wang, Lijuan Wang, Fangyun Wei, Xiang Bai, and Zicheng Liu. End-to-end semi-supervised object detection with soft teacher. In *CVPR*, pages 3060–3069, 2021. [2](#)
- [63] Mingjun Xu, Lingyun Qin, Weijie Chen, Shiliang Pu, and Lei Zhang. Multi-view adversarial discriminator: Mine the non-causal factors for object detection in unseen domains. In *CVPR*, pages 8103–8112, 2023. [3](#)
- [64] Senqiao Yang, Zhuotao Tian, Li Jiang, and Jiaya Jia. Unified language-driven zero-shot domain adaptation. In *CVPR*, pages 23407–23415, 2024. [3](#)
- [65] Yanchao Yang and Stefano Soatto. FDA: Fourier domain adaptation for semantic segmentation. In *CVPR*, pages 4085–4095, 2020. [2, 3](#)
- [66] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. BDD100k: A diverse driving dataset for heterogeneous multitask learning. In *CVPR*, pages 2636–2645, 2020. [5, 6](#)
- [67] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. CutMix: Regularization strategy to train strong classifiers with localizable features. In *ICCV*, pages 6023–6032, 2019. [3](#)
- [68] Fangyuan Zhang, Tianxiang Pan, and Bin Wang. Semi-supervised object detection with adaptive class-rebalancing self-training. In *AAAI*, pages 3252–3261, 2022. [2, 3](#)
- [69] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *ICLR*, 2018. [3](#)
- [70] Liang Zhao and Limin Wang. Task-specific inconsistency alignment for domain adaptive object detection. In *CVPR*, pages 14217–14226, 2022. [2, 3, 5](#)
- [71] Zijing Zhao, Sitong Wei, Qingchao Chen, Dehui Li, Yifan Yang, Yuxin Peng, and Yang Liu. Masked retraining teacher-student framework for domain adaptive object detection. In *ICCV*, pages 19039–19049, 2023. [3](#)
- [72] Wenzhang Zhou, Heng Fan, Tiejian Luo, and Libo Zhang. Unsupervised domain adaptive detection with network stability analysis. In *ICCV*, pages 6986–6995, 2023. [3, 6, 7](#)
- [73] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, pages 2223–2232, 2017. [3](#)