

Multi-Source Domain Adaptation for Object Detection

Xingxu Yao^{1,3}, Sicheng Zhao², Pengfei Xu³, Jufeng Yang¹

¹Nankai University, China ²Columbia University, USA ³Didi Chuxing, China

yxx.hbgd@163.com, schzhao@gmail.com, xupengfeipf@didiglobal.com, yangjufeng@nankai.edu.cn

Abstract

To reduce annotation labor associated with object detection, an increasing number of studies focus on transferring the learned knowledge from a labeled source domain to another unlabeled target domain. However, existing methods assume that the labeled data are sampled from a single source domain, which ignores a more generalized scenario, where labeled data are from multiple source domains. For the more challenging task, we propose a unified Faster R-CNN based framework, termed Divide-and-Merge Spindle Network (DMSN), which can simultaneously enhance domain invariance and preserve discriminative power. Specifically, the framework contains multiple source subnets and a pseudo target subnet. First, we propose a hierarchical feature alignment strategy to conduct strong and weak alignments for low- and high-level features, respectively, considering their different effects for object detection. Second, we develop a novel pseudo subnet learning algorithm to approximate optimal parameters of pseudo target subnet by weighted combination of parameters in different source subnets. Finally, a consistency regularization for region proposal network is proposed to facilitate each subnet to learn more abstract invariances. Extensive experiments on different adaptation scenarios demonstrate the effectiveness of the proposed model.

1. Introduction

As a fundamental task in computer vision, object detection has drawn much attention in the past decade [34, 32, 25]. With the development of convolutional neural networks (CNNs), some modern CNN-based detectors like Faster R-CNN [35] have emerged and been successfully applied to many tasks, such as autonomous driving [10, 38], face and pedestrian detection [17, 27], etc. However, the high-quality performance of detectors is based on large-scale training images with annotated bounding boxes. In the real world, variances exist between training and test images in many aspects, including object appearance, background, even taken time. Due to these domain discrepan-



Figure 1. An example of domain shift in the multi-source scenario for object detection. The above images are sampled from different subsets of BDD100k [50]. Source domain S_j and S_k are taken in daytime and night, respectively, while the images in target domain T are taken in dawn/dusk. As shown in the above results, directly combining images from multiple sources and conducting single-source domain adaptation (DA) will cause performance decay compared with only using the best single source domain. Note that mAP denotes mean average precision.

cies, the performance on the test images may decrease dramatically. Although annotating more training data of the new domain is able to alleviate the phenomenon, it is not an optimal strategy due to large time and labor costs.

To mitigate the domain gap, unsupervised domain adaptation (UDA) has been widely used for object detection [6, 20, 22, 19]. Domain adaptive Faster R-CNN [6] is a milestone study developed for tackling the domain shift problem in object detection. In the work, image-level features and instance-level features are respectively aligned through an adversarial manner. Following [6], a series of Faster R-CNN based adaptation models [57, 56] have emerged recently. Considering that image-level alignment transfers massive unnecessary information background for object detection, Zhu *et al.* [56] and Saito *et al.* [37] pay more attention to align informative local regions. To minimize the

domain distribution disparity on each block, He *et al.* [14] propose multi-adversarial Faster R-CNN to conduct layer-wise domain feature alignment. However, existing algorithms for domain adaptive object detection assume that the source data are sampled from a single domain, which limits the generalization of the model. We consider a more practical scenario that the source data are collected from multiple domains with different distributions. As shown in Figure 1, directly combining two sources and performing single-source domain adaptation (DA) cause performance decay compared with the best result of a single domain, *i.e.* 29.9% vs. 31.4%. It is mainly resulted from the mutual interference between different sources that exist serious domain discrepancy, which is also demonstrated in other tasks [54, 29, 36]. Therefore, we need to design a specialized framework for domain adaptive object detection from multiple sources.

Though multi-source domain adaptation has been explored for other tasks, such as image classification [54, 48] and segmentation [53], they all belong to the straightforward classification task regardless of image-level or pixel-level. However, detection model like Faster R-CNN is a complex system for both regression and classification, and it contains multiple components including feature extractor, region proposal network (RPN), *etc.* In this paper, for the new task, we develop a framework termed Divide-and-Merge Spindle Network (DMSN), in which multiple source subnets and a pseudo subnet are included. First, we propose a hierarchical feature alignment strategy. Since low-level features have high resolution that is important to localization [30], we perform strong alignment for them of different domains. For the high-level features that are important for object recognition, we weakly align each source and target in corresponding supervised source subnet. Second, to approximate optimal parameters for the target domain, we develop a pseudo subnet learning (PSL) algorithm, in which the pseudo subnet is updated via exponential moving average (EMA) [42] parameters of different source subnets. Finally, a new consistency regularization on region proposals are conducted between each source subnet and the pseudo target, which enables the model to capture more robust instance-level invariances for object detection. In the testing phase, the predictions of each subnet are merged into final inference.

In summary, our contributions are threefold.

- We propose to conduct domain adaptation for object detection from multiple sources. To the best of our knowledge, this is the first work on multi-source domain adaptation for object detection that contains both classification and regression.
- We propose a unified framework termed DMSN to tackle the new problem. The characteristics of each

source domain are preserved in each independent supervised source subnet. Meanwhile, an optimal pseudo subnet is approximated by aggregating parameters of different source subnets.

- A new consistency regularization is designed to facilitate each source subnet to propose similar region proposals with the pseudo subnet, so that the model is able to learn more abstract invariances. The results of extensive experiments demonstrate the effectiveness of our proposed framework.

2. Related Work

2.1. Object Detection

Object detection [35, 12] is a fundamental task in computer vision, and it has received booming development during the last decade. The methods that are based on the deep convolutional neural network (CNNs) can be classified into two types: one-stage model [25, 32, 33] and two-stage model [35]. Two-stage model employs region proposal-based detectors, in which a mass of region proposals are generated in advance for refining and revision. Originally, Fast R-CNN [11] utilizes selective search to obtain candidate bounding boxes. Later in Faster R-CNN [35], region proposals are learned by the proposed RPN, which obviously saves computing consumptions. Recently, some one-stage models are introduced to accelerate the speed of detection. These methods apply predefined anchor boxes or points near the center of objects to predict bounding boxes. In this paper, we use Faster R-CNN as our base detector following [39, 56, 16], and transfer knowledge from multiple source domains to the new target domain.

2.2. Domain Adaptation for Object Detection

Due to the labor consumption of data annotation and domain shift existing in different domains, domain adaptive object detection [47, 1, 44, 45] has emerged to solve these problems. Domain adaptive Faster R-CNN [6] is a milestone two-stage object detector to diminish the domain gap between different datasets. The work is based on Faster R-CNN [35] and aligns the distribution on image-level and instance-level by domain classifier by adversarial learning. Besides, the consistency between image-level and instance-level alignment is enforced to learn a more robust model. Along this idea, a series of frameworks [4, 20, 18] are developed for domain adaptive object detection. For example, He *et al.* [14] propose to conduct feature alignment in multiple layers, while Saito *et al.* [37] exploit strong and weak adaptation in local and strong features, respectively. Apart from aligning different domains on feature-level, Hsu *et al.* [16] also employ CycleGAN [16] to map source images into an intermediate domain through synthesizing the target distributions on the pixel-level. When directly applying

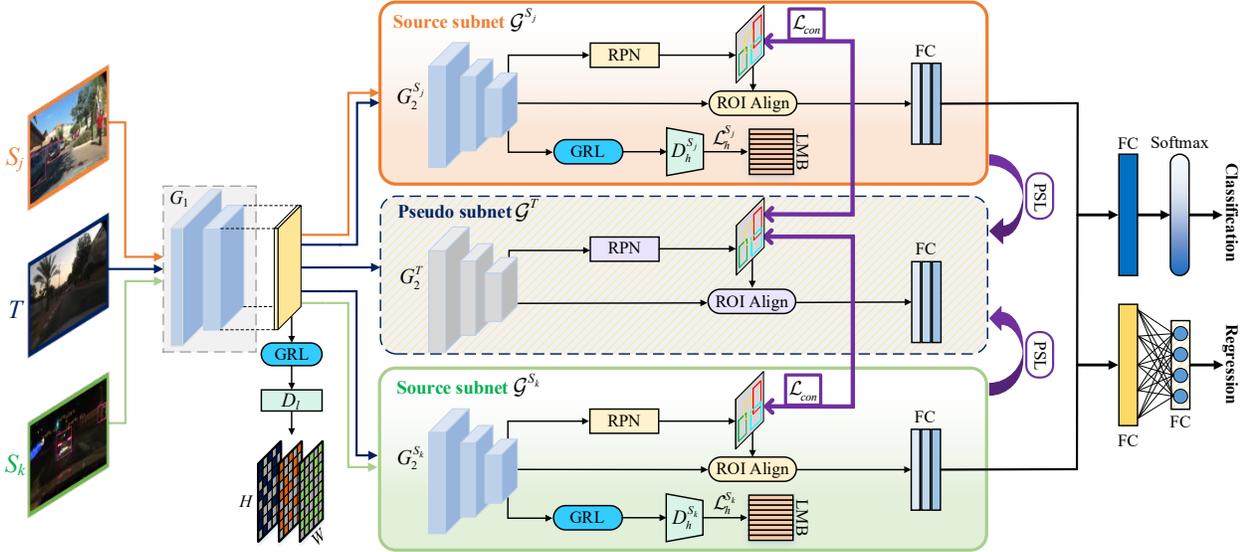


Figure 2. Pipeline of the proposed DMSN. The feature extractor is divided into two parts, including G_1 and G_2 , where $G_2 = \{\{G_2^{S_i}\}_{i=1}^M, G_2^T\}$. In the first training phase, features from G_1 are aligned through GRL and domain discriminator D_l trained by least-squares loss. Then the images of domain S_j and T are fed into the source subnet \mathcal{G}^{S_i} , $i = 1, \dots, M$ to perform the high-level source-target domain alignment and source domain detection learning. In the second training phase, we maintain the learning objective of the first phase and add consistency regularization (\mathcal{L}_{con}) for RPN and pseudo subnet learning (PSL). LMB denotes loss memory bank. Note that the arrows or lines in purple denote the operation in the second training phase.

to adaptive object detection with multiple sources, a more practical scenario, most of these methods suffer from low performances. In this paper, we develop a specialized network to handle source images from multiple domains.

2.3. Multi-Source Domain Adaptation

Multi-source domain adaptation (MSDA) [41, 52, 31, 49, 21] assumes that the training instances are from multiple domains. Recently, several studies follow the routine of single-source DA to utilize a shared network to forcedly align the multiple source distributions. For instance, multi-source domain adversarial network (MDAN) [51] trains a discriminator to distinguish the source and target data, while the feature extractor aims to obtain domain-invariant features by adversarial learning. To better diminish the discrepancy between domains, Peng *et al.* [29] propose to transfer the knowledge from multiple sources to the target by aligning moments of their feature distributions. However, these methods may result in performance decay due to the loss of the discriminative power of each domain.

Besides, there is another routine based on the theoretical analysis [2, 3, 26] for MSDA, which demonstrates that the target distribution can be approximated by a mixture of the multiple source distributions. To calculate more accurate combination weights for source domains, Hoffman *et al.* [15] develop a novel bound using DC-programming. Therefore, weighted combination of multiple source do-

ains has been widely used for MSDA [43]. Xu *et al.* [48] develop a deep cocktail network (DCTN) to train a classifier for each source domain and combine the predictions of different classifiers weighted by perplexity scores. In [54], Zhao *et al.* first fix the specialized feature extractor of each source and adversarially map the target data into the feature space of each source. Then, the prediction is generated by aggregating different source predictions weighted by the discrepancy between each source and the target. Different from them, we introduce a pseudo subnet for the target domain, which is updated by directly weighting parameters of each source subnet. Also, existing MSDA algorithms [54, 23, 48] forcibly align different domains without considering the loss of the feature discriminativity. In this paper, we can not only learn more domain invariance but also preserve discriminative property of features. Besides, we are the first to focus on MSDA for object detection.

3. Method

3.1. Overview

In unsupervised MSDA, we consider the scenario that there are M labeled source domains S_1, S_2, \dots, S_M and one unlabeled target domain T . In the i -th source domain $S_i = \{(x_i^j, B_i^j)\}_{j=1}^{N_i}$, suppose x_i^j represents the j -th images and B_i^j denotes the corresponding bounding box annotation. Note that N_i is the number of images in the i -th source

domain. In the unlabeled target domain $T = \{x_T^j\}_{j=1}^{N_T}$, the j -th image is represented by x_T^j , and N_T denotes the images of target domain. In this problem, our goal is to learn an object detector that can correctly detect the object in images from the target domain based on multiple labeled source domains $\{S_i\}_{i=1}^M$ and unlabeled training data in the target domain T . We emphasize the following two hypotheses in the paper: (1) homogeneity, *i.e.*, $x_i^j \in \mathbb{R}^d$, $x_T^j \in \mathbb{R}^d$, indicating that the feature of data from all the domains are in identical space with the same dimension but different distributions; (2) closed set, meaning that the data from all the domains share the same category space.

We propose a novel framework termed Divide-and-Merge Spindle (DMSN) and the pipeline is shown in Figure 2. Our learning process is divided into two phases. In the first training phase, we strongly align low-level features for all the domains, while the high-level features are weakly aligned between each source domain and the target domain. Meanwhile, we also train each labeled source subnet \mathcal{G}^{S_i} separately supervised by instance-level annotations, which enables it to effectively learn the characteristics of each source. In the second training phase, the learning objective in the first training phase is maintained. We conduct a consistency learning algorithm on region proposals. Besides, we propose a dynamic weighting (DW) strategy to aggregate the parameters of other source subnets for updating the parameters of the pseudo subnet.

3.2. Hierarchical Feature Alignment

As shown in Figure 2, the backbone is divided into two parts, *i.e.* G_1 and $G_2 = \{\{G_2^{S_i}\}_{i=1}^M, G_2^T\}$. We perform different alignment strategies for the features of different levels. At the same time, the supervised learning for object detection for each labeled source domain is conducted in the corresponding source subnet.

Low-level feature alignment. Considering the fact that low-level features are scarcely associated with high-level semantics, we aim to learn domain-invariant features with shared parameters. Furthermore, with the high resolution, low-level features benefit to improve the localization ability [30, 5]. Therefore, we strongly conduct align local features in lower layers using a least-squares loss inspired by [37, 55] to train the domain discriminator D_l , consisting of 1×1 convolutional layers. A gradient reversal layer (GRL) [9] is placed between the backbone and domain discriminator D_l to implement adversarial learning. The low-level feature maps extracted from G_1 are input into D_l , and then the dimension of the channel is reduced to $M + 1$, representing M source domains and one target domain. The width and height of the feature maps are denoted with W and H , respectively. We use $D_l(G_1(x))_{wh}^k$ to represent each location in k -th channel-wise output for image x , where $k = 1, 2, \dots, M + 1$. Different from [37], for images $\{x_i^j\}_{j=1}^{N_i}$ of i -th source domain, we aim to not only align

the features with the target domain but also bridge the gap between source domains:

$$\mathcal{L}_l^{S_i} = \frac{1}{N_i W H} \sum_{j=1}^{N_i} \sum_{w=1}^W \sum_{h=1}^H \left((1 - D_l(G_1(x_i^j))_{wh}^i)^2 + \sum_{k \neq i}^{M+1} (D_l^k(G_1(x_i^j))_{wh}^k)^2 \right). \quad (1)$$

For target images $\{x_T^j\}_{j=1}^{N_T}$, the least-squares loss is formulated as:

$$\mathcal{L}_l^T = \frac{1}{N_T W H} \sum_{j=1}^{N_T} \sum_{w=1}^W \sum_{h=1}^H (1 - D_l(G_1(x_T^j))_{wh}^{M+1})^2. \quad (2)$$

Since the two losses are conducted on each location of the feature maps, the local features can be well aligned. The overall loss on low-level features is formulated as:

$$\mathcal{L}_l = \sum_{i=1}^M \mathcal{L}_l^{S_i} + \mathcal{L}_l^T. \quad (3)$$

High-level feature alignment. In the higher layers, the network is divided into $M + 1$ branches, including M source subnet $\{\mathcal{G}^{S_i}\}_{i=1}^M$ and a pseudo target subnet \mathcal{G}^T . The source instances are fed into one of the corresponding specialized subnets, while the target instances are fed into all the subnets. We define $G_2^{S_i}$ as the second-stage feature extractor for the i -th source domain, while G_2^T denotes the second-stage feature extractor in the pseudo subnet. In the source subnet \mathcal{G}^{S_i} , we suppose that the high-level holistic feature of j -th image from the i -th source and target are respectively represented as $F_{S_{ij}}$ and $F_{T_{ij}}$, respectively, where $F_{S_{ij}} = G_2^{S_i}(G_1(x_i^j))$ and $F_{T_{ij}} = G_2^{S_i}(G_1(x_T^j))$. To preserve characteristics of each source, we aim to weakly align distributions of holistic features for each $G_2^{S_i}$ via GRL and source-target domain discriminator $D_h^{S_i}$. Inspired by [37], we employ focal loss [24] to train $D_h^{S_i}$ of $G_2^{S_i}$ as follows:

$$\mathcal{L}_h^{S_i} = \frac{1}{N_i} \sum_{j=1}^{N_i} (1 - D_h^{S_i}(F_{S_{ij}}))^\gamma \log(D_h^{S_i}(F_{S_{ij}})) + \frac{1}{N_T} \sum_{j=1}^{N_T} D_h^{S_i}(F_{T_{ij}})^\gamma \log(1 - D_h^{S_i}(F_{T_{ij}})), \quad (4)$$

where $\mathcal{L}_h^{S_i}$ denotes the domain classifier loss on high-level features between the i -th source and the target. γ is used to controls the weight on hard examples. The overall source-target domain discriminator loss is written as $\mathcal{L}_h = \sum_{i=1}^M \mathcal{L}_h^{S_i}$.

3.3. Pseudo Subnet Learning

After the first training phase, each source subnet has learned characteristics of corresponding source domain. According to the theoretical proof in [2, 3, 26], the target

distribution can be approximated by the weighted combination of the distribution of different sources. Existing studies on MSDA [48, 54] mainly focus on weighting different predictions in the inference phase. In this paper, from a new perspective, we aim to directly learn a pseudo subnet \mathcal{G}^T for target by weighting the parameters of each source subnet \mathcal{G}^{S_i} . Inspired by mean teacher method [4, 42, 8] used in semi-supervised learning, we obtain the pseudo subnet by exponential moving average (EMA) parameters of source subnets. Different from [4], which only averages parameters on the temporal sequence, we also consider how to aggregate the parameters of different source subnets. To reasonably integrate the knowledge of each source, we design a dynamic weighting (DW) strategy based on the discrepancy between each source and target to conduct greater weight for more relevant sources.

As mentioned in Section 3.2, source-target classifier $D_h^{S_i}$ in \mathcal{G}^{S_i} is employed to distinguish target images and source images of the i -th domain. The larger value of $\mathcal{L}_h^{S_i}$ indicates that the distributions of the i -th source and the target are more similar. Based on the important cues, we design a dynamic loss memory bank (LMB) to store the loss values of the past training steps. We use \mathcal{U}_i^t to denote LMB of the i -th source domain in t -th training step, *i.e.* $\mathcal{U}_i^t = [\mathcal{L}_h^{S_i}(t - c + 1), \mathcal{L}_h^{S_i}(t - c + 2), \dots, \mathcal{L}_h^{S_i}(t)]$, where $\mathcal{L}_h^{S_i}(t)$ denotes the loss value of t -th training step and c represents the length of \mathcal{U}_i^t . The mean value of all the elements in \mathcal{U}_i^t is denoted as \mathcal{V}_i^t . The larger the value of \mathcal{V}_i^t is, the more similar the distributions of the i -th source and the target are, and vice versa. Therefore, we measure the discrepancy of S_i and T based on \mathcal{V}_i^t at iteration t . Considering LMB of all the sources, we define the relative similarity of the i -th source and the target at iteration t as $\beta_i^t = \mathcal{V}_i^t / \sum_j^M \mathcal{V}_j^t$. In the second training phase, β_i^t is regarded as the weight of the i -th source subnet \mathcal{G}^{S_i} . Therefore, the pseudo subnet parameters $P_{\mathcal{G}^T}^t$ at the t -th training step are updated as EMA of aggregated parameters of each \mathcal{G}^{S_i} :

$$P_{\mathcal{G}^T}^t = \alpha P_{\mathcal{G}^T}^{t-1} + (1 - \alpha) \cdot \sum_i^M \beta_i^t P_{\mathcal{G}^{S_i}}^t, \quad (5)$$

where $P_{\mathcal{G}^{S_i}}^t$ denotes the parameters of \mathcal{G}^{S_i} at training step t and α is a smoothing coefficient parameter.

3.4. Consistency Learning of RPN

Although we have learned a specialized subset for each source in the first training phase, we employ an implicit strategy to bridge the discrepancy between each source and the target in the second phase. Considering that the performance of RPN is closely related to the detection results, we conduct a consistency regularization on the region proposals from $G_p^{S_i}$ and G_p^T for the same target image, where $G_p^{S_i}$ and G_p^T denote the RPN of \mathcal{G}^{S_i} and \mathcal{G}^T , respectively. For a target image, it will be fed into different subnets after G_1 . Then each RPN will generate top- \mathcal{N} ranked pro-

posed bounding boxes for detection. We expect that the region proposals from different RPNs are consistent through each source subnet which has learned specialized domain knowledge. We use $\mathcal{R}^i = \{r_j^i\}_{j=1}^{\mathcal{N}}$ to represent region proposal set of the i -th source subnet. Similarly, the region proposal set of \mathcal{G}^T is denoted as $\mathcal{R}^T = \{r_j^T\}_{j=1}^{\mathcal{N}}$. Since \mathcal{G}^T is updated through aggregating the knowledge from labeled sources with DW strategy, it will progressively be more friendly for the target domain during training. Therefore, we impose consistency regularization between each \mathcal{R}^i and \mathcal{R}^T . For the n -th region proposal $r_n^T \in \mathcal{R}^T$, the highest Intersection-over-Union (IoU) overlap between $\{r_j^i\}_{j=1}^{\mathcal{N}}$ and r_n^T is denoted as $\mathcal{O}_n^i = \max_{j \in [1, \mathcal{N}]} (\text{IoU}(r_j^i, r_n^T))$, and the corresponding serial number of region proposal in \mathcal{R}^i is $n_i^* = \arg \max_{j \in [1, \mathcal{N}]} (\text{IoU}(r_j^i, r_n^T))$. The overall consistency cost is formalized as:

$$\mathcal{L}_{con} = \frac{1}{\mathcal{N}} \sum_{n=1}^{\mathcal{N}} \sum_{i=1}^M \mathcal{O}_n^i |n_i^* - n|, \quad (6)$$

where $|n_i^* - n|$ represents the distance between n_i^* and n .

3.5. Overall Objective

The supervised learning loss for the detection of labeled source images is denoted as \mathcal{L}_{det} , which is composed of classification and regression error for RPN and RCNN. Combining detection loss \mathcal{L}_{det} and our introduced losses for domain adaptation, the final loss function of DMSN is written as:

$$\mathcal{L} = \mathcal{L}_{det} + \lambda(\mathcal{L}_l + \mathcal{L}_h + \mathcal{L}_{con}), \quad (7)$$

where λ is a trade-off parameter to balance the detection loss and domain adaptation loss. Placing GRL between backbone and domain discriminator, \mathcal{L}_l and \mathcal{L}_h are optimized by an adversarial manner. While domain discriminators aim to minimize \mathcal{L}_l and \mathcal{L}_h , the base network optimizes its parameters to maximize the losses due to reversed gradient. Note that \mathcal{L}_{con} and PSL can automatically participate in the second training phase without manual work, and DMSN is optimized in an end-to-end training manner. No gradients are propagated through pseudo subnet \mathcal{G}^T in the training. The parameters of \mathcal{G}^T is updated via Equation (5).

4. Experiments

In evaluation, we conduct extensive experiments of our framework in different domain adaptation scenarios, including cross camera adaptation and cross time adaptation.

4.1. Datasets

KITTI. The KITTI dataset [10] is collected from an autonomous driving platform, containing street scenarios taken in cities, highways, and rural areas. A total of 7,481 training images are used as source images.

Cityscapes. The Cityscapes dataset [7] contains driving scenarios taken in cities. There are 2,975 images in the training set and 500 images in the validation set with pixel-level labels. All of them are transformed into bounding boxes annotations.

BDD100k. BDD100k dataset [50] contains 100k images, including 70k training images and 10k validation images with bounding boxes annotations. The images are taken in the different times of day including *daytime*, *night*, and *dawn/dusk*. The images taken in *daytime* as the target domain in cross camera adaptation, including 36,728 training and 5,258 validation images. Besides, we use the images taken in all the times in cross time adaptation.

4.2. Compared Baselines

In this paper, we compare DMSN with the following baselines: (1) **Source-only**, *i.e.*, train on the source images and test on the target images directly, indicating the lower bound of domain adaptation. The original Faster R-CNN is treated as the detection model. (2) **Single-source DA**, conduct MSDA via single-source DA, including Strong-Weak [37], SCL [39], DA-ICR-CCR [46], SW-ICR-CCR [46], GPA [47]. (3) **Multi-source DA**, extend existing MSDA methods for classification, *i.e.* MDAN [51] and M³SDA [29], to perform MSDA for object detection by utilizing these algorithms on feature extractor. Note that all the DA methods in the paper are unsupervised, where the labels of target training set are unavailable during training. Besides, we also report oracle results by directly training Faster R-CNN using training target data. For the source-only and single-source domain adaptation, we apply two strategies: (1) single-source, *i.e.* conduct adaptation on each single source domain; (2) source-combined, *i.e.* directly combining all the sources into a unified domain.

4.3. Implementation Details

In our experiments, Faster R-CNN [35] is employed as based object detector in which the pretrained VGG16 [40] is adopted as the backbone, following [6, 46, 37]. Unless otherwise stated, the shorter side of the image has a length of 600 pixels following the implementation of [35, 37] with ROI-alignment [13]. The learning rate is initialized as 0.001, and all the models are trained for 20 epochs with cosine learning rate decay. At the 10-th epoch, PSL and consistency regularization begin to be performed until the end of the training. We set γ of Equation (4) as 5.0 and λ of Equation (7) as 1.0 following [37], while α of Equation (5) and \mathcal{N} of Equation (6) are set to 0.99 and 256, respectively. In the testing phase, the test images of the target domain will be fed into all subnets after G_1 . The final predictions obtained by integrating results from different subnets. We implemented all the experiments with Pytorch [28].

Table 1. Experimental results on adaptation from Cityscapes and KITTI to BDD100k (*daytime*). Average precision (AP, %) on *car* category in target domain is evaluated. The best result is highlight with **bold**.

Standards	Methods	AP on car
Source-only	Cityscapes	44.6
	KITTI	28.6
	Cityscapes+KITTI	43.2
Cityscapes-only DA	Strong-Weak [37]	45.5
	SCL [39]	46.3
	DA-ICR-CCR [46]	45.3
	SW-ICR-CCR [46]	46.5
KITTI-only DA	Strong-Weak [37]	29.6
	SCL [39]	31.1
	DA-ICR-CCR [46]	29.2
	SW-ICR-CCR [46]	30.8
Source-combined DA	Strong-Weak [37]	41.9
	SCL [39]	43.0
	DA-ICR-CCR [46]	41.3
	SW-ICR-CCR [46]	43.6
Multi-source DA	MDAN [51]	43.2
	M ³ SDA [29]	44.1
	DMSN (Ours)	49.2
Oracle	Faster R-CNN [35]	60.2

4.4. Cross Camera Adaptation

Different datasets are captured by different devices or setups, which exhibit different scenes and viewpoints, incurring a strong domain shift. In this experiment, we show the adaptation between datasets collected by different cameras. Specifically, KITTI and Cityscapes are used as two source domains, while the subset taken in the daytime of BDD100k is treated as target domain. The experiment is an adaptation from small-scale datasets to a large-scale dataset.

In Table 1, we report experimental results evaluated on the *car* category (common category) in terms of the average precision (AP). In source-only, the results of KITTI and Cityscapes+KITTI are worse than that of Cityscapes, where “Cityscapes+KITTI” means directly combining Cityscapes and KITTI. The performance drop after adding KITTI can mainly attribute to two reasons: (1) the domain discrepancy between two sources is not well bridged. (2) compared with Cityscapes, the domain distribution of KITTI is further with that of BDD100k, and simply combining KITTI with Cityscapes will decrease the performance. Meanwhile, regardless of source-only or single-source DA methods, simply combining images of multiple sources to train a model results in the worse result compared with the best result trained by single domain. Particularly, the Strong-Weak model transferred from cityscapes achieves 45.5% AP, while the source-combined drops to 41.1% AP. It is mainly because that the domain discrepancy also exists in

Table 2. Adaptation between different subsets with in BDD100k dataset. *daytime* and *night* are the sources, while *dawn/dusk* is the target domain. The mean average precision (mAP, %) of 10 categories is evaluated. The best result of each category is highlight with **bold**.

Standards	Methods	bike	bus	car	motor	person	rider	light	sign	train	truck	mAP
Source-only	Daytime	35.1	51.7	52.6	9.9	31.9	17.8	21.6	36.3	0	47.1	30.4
	Night	27.9	32.5	49.4	15.0	28.7	21.8	14.0	30.5	0	30.7	25.0
	Daytime+Night	31.5	46.9	52.9	8.4	29.5	21.6	21.7	34.3	0	42.2	28.9
Daytime-only DA	Strong-Weak [37]	34.9	51.2	52.7	15.1	32.8	23.6	21.6	35.6	0	47.1	31.4
	SCL [39]	29.1	51.3	52.8	17.2	32.0	19.1	21.8	36.3	0	47.2	30.7
	GPA [47]	36.6	52.1	53.1	15.6	33.0	23.0	21.7	35.4	0	48.0	31.8
	DA-ICR-CCR [46]	35.6	47.5	52.7	13.9	32.2	22.7	22.8	35.5	0	45.7	30.9
	SW-ICR-CCR [46]	32.8	51.4	53.0	15.4	32.5	22.3	21.2	35.4	0	47.9	31.2
Night-only DA	Strong-Weak [37]	31.4	38.2	51.0	9.9	29.5	22.2	18.7	32.5	0	35.7	26.9
	SCL [39]	25.3	31.7	49.3	8.9	25.8	21.2	15.0	28.6	0	26.2	23.2
	GPA [47]	32.7	38.3	51.8	14.1	29.0	21.5	17.1	31.1	0	40.0	27.6
	DA-ICR-CCR [46]	30.0	32.4	50.1	14.4	29.1	22.8	17.4	32.2	0	29.7	25.8
	SW-ICR-CCR [46]	32.3	45.1	51.6	7.2	29.2	24.9	19.9	33.0	0	41.1	28.4
Source-combined DA	Strong-Weak [37]	29.7	50.0	52.9	11.0	31.4	21.1	23.3	35.1	0	44.9	29.9
	SCL [39]	33.9	47.8	52.5	14.0	31.4	23.8	22.3	35.4	0	45.1	30.9
	GPA [47]	31.7	48.8	53.9	20.8	32.0	21.6	20.5	33.7	0	43.1	30.6
	DA-ICR-CCR [46]	28.2	47.6	51.6	17.6	28.8	21.9	17.4	33.2	0	45.8	29.2
	SW-ICR-CCR [46]	25.3	51.3	52.1	17.0	33.4	18.9	20.7	34.8	0	47.9	30.2
Multi-source DA	MDAN [51]	37.1	29.9	52.8	15.8	35.1	21.6	24.7	38.8	0	20.1	27.6
	M ³ SDA [29]	36.9	25.9	51.9	15.1	35.7	20.5	24.7	38.1	0	15.9	26.5
	DMSN (Ours)	36.5	54.3	55.5	20.4	36.9	27.7	26.4	41.6	0	50.8	35.0
Oracle	Faster R-CNN [35]	27.2	39.6	51.9	12.7	29.0	15.2	20.0	33.1	0	37.5	26.6

different sources, and thus the source-combined images may interfere with each other during training [36]. Compared with existing multi-source DA methods [51, 29] for other tasks, our DMSN for object detection obtains obviously better performance.

4.5. Cross Time Adaptation

In this scenario, we conduct our domain adaptation between different time conditions. Specifically, we conduct experiments using three different time-related subsets of BDD100k, *i.e.* *daytime*, *night*, *dawn/dusk*. 36,728 training images in *daytime* and 27,971 training images in *night* are treated as source data, while 5,027 training images in *dawn/dusk* are treated as target data. The evaluation is conducted on a total of 778 validation dawn/dusk images. Different from cross camera adaptation, this adaption scenario is from the large-scale datasets to a small-scale dataset.

Table 2 shows that our model effectively reduces the domain discrepancy across time conditions and performs favorably against the state-of-the-art methods [37, 39, 46]. Intuitively, when *night* is treated as the source domain, the adaptation performance is worse than *daytime* as the source. It matches human cognition that it is hard to recognize and localize the object at night. Due to the domain gap between *daytime* and *night*, the performance of source-combined DA drops compared with daytime-only DA apart from SCL [39]. Our adaptation model achieves considerable improvement over the best results of source-combined

DA and multi-source DA by 4.1% mAP and 7.4% mAP, respectively. Though multi-source DA (*i.e.*, MDAN and M³SDA) employ different methods to align different domains, the alignment is limited on the image level. However, instance-level information plays an essential role in object detection. Only using forcible image-level alignment may reduce the discriminative ability of local features for objects. It is observed that the results of MDAN and M³SDA are extremely poor on *bus* and *truck* that have similar appearance, and it directly results in performance drop on mAP. It is worth emphasizing that our results surpasses that of oracle by about 10% mAP, which reveals a phenomenon that the oracle results may be not the upper bound of domain adaptation in all the scenarios. We conclude that the above phenomenon is closely related to the number of training data. In our experiments, a total of over 60,000 images of *daytime* and *night* participate in training model, which greatly exceeds 5,027 *dawn/dusk* training images in the oracle experiment.

4.6. Analysis

Ablation study. We perform a detailed ablation study for cross time adaptation, and its results are shown in Table 3. Based on the skeleton of DMSN, the baseline is achieved by only conducting supervised learning for object detection in all the source subnets (SS), without any other operations. When consecutively adding domain discriminator loss on the low-level and high-level features, the performance is im-

Table 3. Ablation study for adaptation from *daytime* and *night* to *dawn/dusk* within BDD100k dataset. SS denotes the source subnets and PS denotes the pseudo subnet. DW represents dynamic weighting strategy. The best result of each category is highlight with **bold**.

SS	\mathcal{L}_l	\mathcal{L}_h	PS	DW	\mathcal{L}_{con}	bike	bus	car	motor	person	rider	light	sign	train	truck	mAP
✓						31.4	23.0	52.3	18.4	35.3	21.4	24.5	39.2	0	35.0	28.0
✓	✓					32.9	33.4	53.2	18.2	36.0	23.8	27.1	40.2	0	34.5	29.9
✓	✓	✓				30.4	52.7	54.7	12.2	34.5	22.6	24.9	38.0	0	46.9	31.7
✓	✓	✓	✓			33.0	47.4	54.8	14.5	36.4	26.0	27.6	41.4	0	47.8	32.9
✓	✓	✓	✓		✓	33.1	54.3	55.6	18.5	36.2	24.8	27.1	41.1	0	50.4	34.2
✓	✓	✓	✓	✓	✓	36.5	54.3	55.5	20.4	36.9	27.7	26.4	41.6	0	50.8	35.0

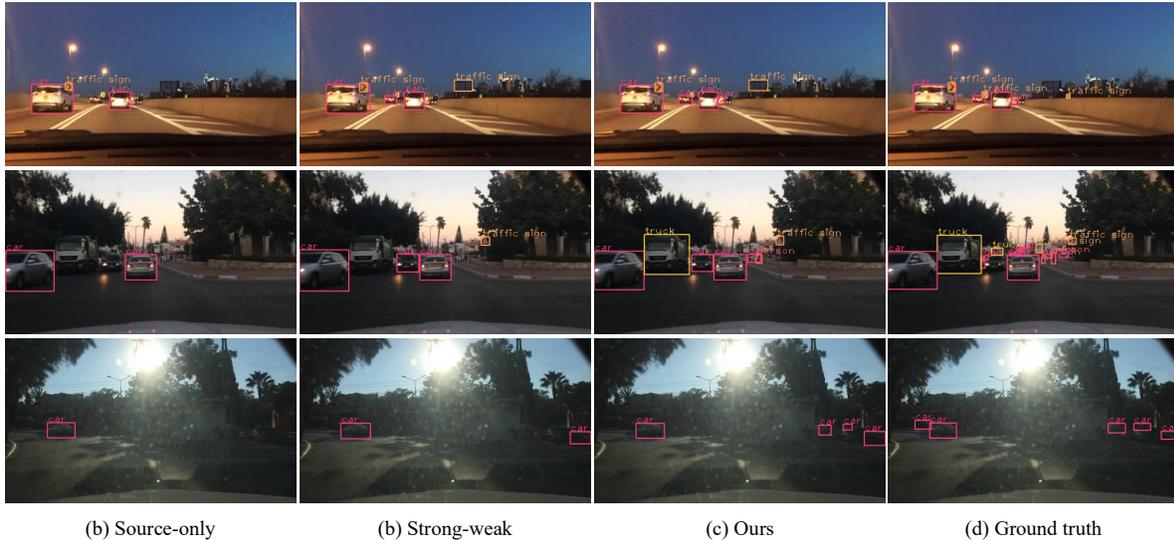


Figure 3. Examples of the qualitative results for adaptation from *daytime* and *night* subsets to *dawn/dusk*. From left to right column, the results of source-only, source-combined DA via Strong-Weak [37], our framework DMSN, and the ground truth.

proved by 1.9% mAP and 1.8%, respectively. In the fourth row, our DW strategy cannot be applied for PSL without consistency regularization, so we treat each source domain equally when updating parameters of the pseudo subnet. The results demonstrate the effectiveness of the pseudo subnet (PS). It is worth highlighting that the newly designed instance-level consistency regularization for region proposals has obvious effectiveness, *i.e.* a 1.3% mAP gain. Finally, the DW strategy for weighting parameters of different source subnets also contributes a slight performance gain.

Visualization. In Figure 3, we present three examples from source-only, compared method Strong-Weak and ours in the cross time adaptation scenario. In general, the proposed algorithm has merits in three aspects compared with existing methods. First, the proposed algorithm accurately detects the small object, *i.e.* *car* in the distance in the first example. For the second, our model correctly localizes the object, *i.e.* *truck*, that has less sample in training data. Finally, in the situation of a strong backlight, *i.e.* the third example, the results demonstrate that our model can also well tackle the hard scenario. Therefore, it is illustrated that the algorithm has a robust generalization ability for these hard cases.

5. Conclusion

In this paper, we develop a framework, *i.e.* DMSN, to address a new task that is adaptation from multiple source domains for object detection. We propose hierarchical feature alignment for features of different levels. Meanwhile, we intend to preserve the characteristics of each source via supervised learning in an independent subnet. We also introduce a pseudo subnet that is updated by exponential moving average parameters in different source subnets. The results on different adaptation scenarios demonstrate the effectiveness of our method. We also reveal that as long as enough source images are collected, it is possible that adaptation results surpass the oracle ones.

Acknowledgment

This work was supported by the National Key Research and Development Program of China Grant (NO. 2018AAA0100403), NSFC (NO.61876094, U1933114, 61701273), Natural Science Foundation of Tianjin, China (NO.20JCJQC00020, 18JCYBJC15400, 18ZXXZNGX00110).

References

- [1] Vinicius F Arruda, Thiago M Paixão, Rodrigo F Berriel, Alberto F De Souza, Claudine Badue, Nicu Sebe, and Thiago Oliveira-Santos. Cross-domain car detection using unsupervised image-to-image translation: From day to night. In *IJCNN*, 2019. 2
- [2] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine Learning*, 79(1-2):151–175, 2010. 3, 4
- [3] John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman. Learning bounds for domain adaptation. In *NeurIPS*, 2008. 3, 4
- [4] Qi Cai, Yingwei Pan, Chong-Wah Ngo, Xinmei Tian, Lingyu Duan, and Ting Yao. Exploring object relation in mean teacher for cross-domain detection. In *CVPR*, 2019. 2, 5
- [5] Chen Chen, Mengyuan Liu, Xiandong Meng, Wanpeng Xiao, and Qi Ju. Refinedelite: A lightweight one-stage object detection framework for cpu-only devices. In *CVPRW*, 2020. 4
- [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*, 2018. 1, 2, 6
- [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*, 2016. 6
- [8] Geoffrey French, Michal Mackiewicz, and Mark Fisher. Self-ensembling for visual domain adaptation. In *ICLR*, 2018. 5
- [9] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, 2015. 4
- [10] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *CVPR*, 2012. 1, 5
- [11] Ross Girshick. Fast r-cnn. In *ICCV*, 2015. 2
- [12] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014. 2
- [13] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017. 6
- [14] Zhenwei He and Lei Zhang. Multi-adversarial faster-rcnn for unrestricted object detection. In *ICCV*, 2019. 2
- [15] Judy Hoffman, Mehryar Mohri, and Ningshan Zhang. Algorithms and theory for multiple-source adaptation. In *NeurIPS*, 2018. 3
- [16] Han-Kai Hsu, Chun-Han Yao, Yi-Hsuan Tsai, Wei-Chih Hung, Hung-Yu Tseng, Maneesh Singh, and Ming-Hsuan Yang. Progressive domain adaptation for object detection. In *WACV*, 2020. 2
- [17] Huaizu Jiang and Erik Learned-Miller. Face detection with the faster r-cnn. In *FG*, 2017. 1
- [18] Mehran Khodabandeh, Arash Vahdat, Mani Ranjbar, and William G Macready. A robust learning approach to domain adaptive object detection. In *CVPR*, 2019. 2
- [19] Seunghyeon Kim, Jaehoon Choi, Taekyung Kim, and Changick Kim. Self-training and adversarial background regularization for unsupervised domain adaptive one-stage object detection. In *ICCV*, 2019. 1
- [20] Taekyung Kim, Minki Jeong, Seunghyeon Kim, Seokeon Choi, and Changick Kim. Diversify and match: A domain adaptive representation learning paradigm for object detection. In *CVPR*, 2019. 1, 2
- [21] Da Li and Timothy M. Hospedales. Online meta-learning for multi-source and semi-supervised domain adaptation. In *ECCV*, 2020. 3
- [22] Wanyi Li, Fuyu Li, Yongkang Luo, Peng Wang, et al. Deep domain adaptive object detection: A survey. In *IEEE Symposium Series on Computational Intelligence*, pages 1808–1813, 2020. 1
- [23] Yitong Li, Michael Murias, Samantha Major, Geraldine Dawson, and David E Carlson. Extracting relationships by multi-domain matching. In *NeurIPS*, pages 6799–6810, 2018. 3
- [24] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, 2017. 4
- [25] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *ECCV*, 2016. 1, 2
- [26] Yishay Mansour, Mehryar Mohri, and Afshin Rostamizadeh. Domain adaptation with multiple sources. In *NeurIPS*, 2009. 3, 4
- [27] Jiayuan Mao, Tete Xiao, Yuning Jiang, and Zhimin Cao. What can help pedestrian detection? In *CVPR*, 2017. 1
- [28] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. 6
- [29] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *ICCV*, 2019. 2, 3, 6, 7
- [30] Zheng Qin, Zeming Li, Zhaoning Zhang, Yiping Bao, Gang Yu, Yuxing Peng, and Jian Sun. Thudernet: Towards real-time generic object detection on mobile devices. In *ICCV*, 2019. 2, 4
- [31] Sayan Rakshit, Dipesh Tamboli, Pragati Shuddhodhan Meshram, Biplab Banerjee, Gemma Roig, and Subhasis Chaudhuri. Multi-source open-set deep adversarial domain adaptation. In *ECCV*, 2020. 3
- [32] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, 2016. 1, 2
- [33] Joseph Redmon and Ali Farhadi. Yolo9000: better, faster, stronger. In *CVPR*, 2017. 2
- [34] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018. 1
- [35] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, 2015. 1, 2, 6, 7

- [36] Matthew Riemer, Ignacio Cases, Robert Ajemian, Miao Liu, Irina Rish, Yuhai Tu, and Gerald Tesauro. Learning to learn without forgetting by maximizing transfer and minimizing interference. In *ICLR*, 2019. 2, 7
- [37] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. Strong-weak distribution alignment for adaptive object detection. In *CVPR*, 2019. 1, 2, 4, 6, 7, 8
- [38] 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. 1
- [39] Zhiqiang Shen, Harsh Maheshwari, Weichen Yao, and Marios Savvides. Scl: Towards accurate domain adaptive object detection via gradient detach based stacked complementary losses. *arXiv preprint arXiv:1911.02559*, 2019. 2, 6, 7
- [40] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 6
- [41] Shiliang Sun, Honglei Shi, and Yuanbin Wu. A survey of multi-source domain adaptation. *Information Fusion*, 24:84–92, 2015. 3
- [42] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NeurIPS*, 2017. 2, 5
- [43] Hang Wang, Minghao Xu, Bingbing Ni, and Wenjun Zhang. Learning to combine: Knowledge aggregation for multi-source domain adaptation. In *ECCV*, 2020. 3
- [44] Tao Wang, Xiaopeng Zhang, Li Yuan, and Jiashi Feng. Few-shot adaptive faster r-cnn. In *CVPR*, 2019. 2
- [45] Xudong Wang, Zhaowei Cai, Dashan Gao, and Nuno Vasconcelos. Towards universal object detection by domain attention. In *CVPR*, 2019. 2
- [46] Chang-Dong Xu, Xing-Ran Zhao, Xin Jin, and Xiu-Shen Wei. Exploring categorical regularization for domain adaptive object detection. In *CVPR*, 2020. 6, 7
- [47] Minghao Xu, Hang Wang, Bingbing Ni, Qi Tian, and Wenjun Zhang. Cross-domain detection via graph-induced prototype alignment. In *CVPR*, 2020. 2, 6, 7
- [48] Ruijia Xu, Ziliang Chen, Wangmeng Zuo, Junjie Yan, and Liang Lin. Deep cocktail network: Multi-source unsupervised domain adaptation with category shift. In *CVPR*, 2018. 2, 3, 5
- [49] Luyu Yang, Yogesh Balaji, Ser-Nam Lim, and Abhinav Srivastava. Curriculum manager for source selection in multi-source domain adaptation. In *ECCV*, 2020. 3
- [50] Fisher Yu, Wenqi Xian, Yingying Chen, Fangchen Liu, Mike Liao, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving video database with scalable annotation tooling. *arXiv preprint arXiv:1805.04687*, 2018. 1, 6
- [51] Han Zhao, Shanghang Zhang, Guanhang Wu, José MF Moura, Joao P Costeira, and Geoffrey J Gordon. Adversarial multiple source domain adaptation. In *NeurIPS*, 2018. 3, 6, 7
- [52] Sicheng Zhao, Bo Li, Pengfei Xu, and Kurt Keutzer. Multi-source domain adaptation in the deep learning era: A systematic survey. *arXiv preprint arXiv:2002.12169*, 2020. 3
- [53] Sicheng Zhao, Bo Li, Xiangyu Yue, Yang Gu, Pengfei Xu, Runbo Hu, Hua Chai, and Kurt Keutzer. Multi-source domain adaptation for semantic segmentation. In *NeurIPS*, 2019. 2
- [54] Sicheng Zhao, Guangzhi Wang, Shanghang Zhang, Yang Gu, Yaxian Li, Zhichao Song, Pengfei Xu, Runbo Hu, Hua Chai, and Kurt Keutzer. Multi-source distilling domain adaptation. In *AAAI*, 2020. 2, 3, 5
- [55] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017. 4
- [56] Xinge Zhu, Jiangmiao Pang, Ceyuan Yang, Jianping Shi, and Dahua Lin. Adapting object detectors via selective cross-domain alignment. In *CVPR*, 2019. 1, 2
- [57] Chenfan Zhuang, Xintong Han, Weilin Huang, and Matthew R Scott. ifan: Image-instance full alignment networks for adaptive object detection. In *AAAI*, 2020. 1