

Universal Domain Adaptive Object Detection via Dual Probabilistic Alignment

Yuanfan Zheng^{1,2*}, Jinlin Wu^{1,2*}, Wuyang Li³, Zhen Chen^{1†}

¹CAIR, HKISI-CAS

²MAIS, Institute of Automation, Chinese Academy of Sciences

³The Chinese University of Hong Kong

zyfone.cs@gmail.com, jinlin.wu@nlpr.ia.ac.cn, wymanbest@outlook.com, zchen.francis@gmail.com

Abstract

Domain Adaptive Object Detection (DAOD) transfers knowledge from a labeled source domain to an unannotated target domain under closed-set assumption. Universal DAOD (UniDAOD) extends DAOD to handle open-set, partial-set, and closed-set domain adaptation. In this paper, we first unveil two issues: domain-private category alignment is crucial for global-level features, and the domain probability heterogeneity of features across different levels. To address these issues, we propose a novel Dual Probabilistic Alignment (DPA) framework to model domain probability as Gaussian distribution, enabling the heterogeneity domain distribution sampling and measurement. The DPA consists of three tailored modules: the Global-level Domain Private Alignment (GDPA), the Instance-level Domain Shared Alignment (IDSA), and the Private Class Constraint (PCC). GDPA utilizes the global-level sampling to mine domain-private category samples and calculate alignment weight through a cumulative distribution function to address the global-level private category alignment. IDSA utilizes instance-level sampling to mine domain-shared category samples and calculates alignment weight through Gaussian distribution to conduct the domain-shared category domain alignment to address the feature heterogeneity. The PCC aggregates domain-private category centroids between feature and probability spaces to mitigate negative transfer. Extensive experiments demonstrate that our DPA outperforms state-of-the-art UniDAOD and DAOD methods across various datasets and scenarios, including open, partial, and closed sets.

Code — <https://github.com/zyfone/DPA>

Introduction

Object detection has made significant progress in recent years (Li et al. 2021b; Jia et al. 2023; Zhao et al. 2024). However, the well-trained object detector failed to generalize in novel domain scenarios due to domain shift. Domain Adaptive Object Detection (DAOD) (Krishna, Ohashi, and Sinha 2023; Huang et al. 2024) transferring from the source domain to the unlabelled target domain to overcome domain shift and has been widely applied in medical analysis (Pu

*These authors contributed equally.

†Corresponding author.

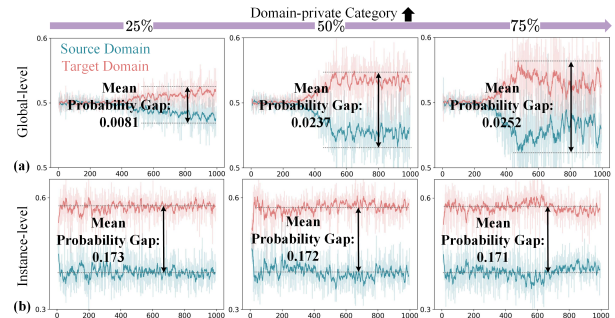


Figure 1: The visualization of domain probability in the domain discriminator¹. (a) As the number of domain-private categories increases, a more distinct gap emerges between the two domains at the global level, suggesting the **alignment of the domain-private category**. (b) With the increases of the domain-private category, the mean probability gap remains approximately constant at the instance level, indicating the **alignment of domain-shared category**.

et al. 2024; Ali et al. 2024; Liu, Li, and Yuan 2023), autonomous driving (Cai et al. 2024; Shi, Zheng, and Chen 2024) and robotic understanding (Chapman et al. 2023; Li et al. 2024). However, DAOD is limited by the closed-set assumption (Ben-David et al. 2010; Li et al. 2023c) failing to generalize to real-world scenarios. To address this, Universal DAOD (UniDAOD) endows DAOD with open-set domain adaption capabilities to overcome label shifts without prior knowledge of categories. The former work, US-DAF (Shi et al. 2022), leverages the threshold filter mechanism and the scale-sensitive domain alignment. CODE (Shi et al. 2024b) adopts virtual domain alignment to avoid aligning domain-private category samples to mitigate negative transfer. Other methods (Lang et al. 2023; Shi et al. 2024a) adopt dynamic weighting for domain-private categories to facilitate positive transfer. Essentially, existing UniDAOD approaches align shared categories at both the global and instance levels while ignoring the alignment of domain-private categories.

Despite significant progress in recent years, the current

¹The horizontal axis is training iteration ($\times 100$), and the vertical axis is the probability of domain discriminator.

UniDAOD paradigm encounters two major issues that result in suboptimal domain alignment. The first issue is that they overlook global feature alignment with domain-private categories. Due to the agnostic prior category knowledge, existing methods (Lang et al. 2023; Shi et al. 2024a) primarily focus on estimating the domain-shared category set to mitigate negative transfer, **wrongly assuming** that both global and instance features fairly contribute to the domain-shared alignment. For the first time, we empirically reveal the issue of this assumption in Fig. 1. We claim *the fact is that global-level features tend to align domain private categories, while instance-level features tend to align domain shared categories*, which can also be thoroughly justified in Fig. 4. This phenomenon motivates us to revisit the UniDAOD domain alignment, focusing on the domain-private category alignment at the global level for UniDAOD.

The second issue is the heterogeneity of features at different feature levels. Since the global feature is a rough representation of entire input images, while the instance feature corresponds to object instances, the feature gap results in significant differences in domain probabilities. Existing approaches address this by employing different thresholds and entropy functions, but these methods require manual parameter tuning. In addition to that, adopting advanced UniDA frameworks such as clustering (Saito et al. 2020; Li et al. 2021a), optimal transport (Chang et al. 2022), and mutual learning (Lu et al. 2024) are complex and challenging to adapt to detection tasks.

To address these issues, we propose a novel Dual Probabilistic Alignment (DPA) framework. For the first issue, we conduct a theoretical analysis to unveil that domain-private alignment is crucial for global-level features. Therefore, we propose a Global-level Domain Private Alignment (GDPA) module that includes global-level sampling, alignment weight calculation, and global-level domain alignment. Global-level sampling aims to mine domain-private category samples. Alignment weight calculation involves the cumulative distribution function to refine the distribution distance estimation as the weight, thereby conducting global-level domain alignment to address the domain-private alignment issue. For the second problem, we conduct a tailored domain-shared category alignment at the instance-level features. To effectively obtain the domain-shared category, we propose a novel unsupervised clustering perspective. We set the domain label as the center and map samples to domain probabilities to calculate the gradient norm (distance). We then model the frequency of the gradient norm as a Gaussian distribution using bins. The continuous frequency bins of the samples represent those within a certain radius (the sum of bins). Therefore, we propose an Instance-level Domain Shared Alignment (IDSA) method consisting of instance-level sampling, alignment weight calculation, and instance-level domain alignment. The instance-level sampling utilizes a Gaussian distribution modeling to select domain-shared category samples. Alignment weight calculation involves the Gaussian distribution statistical properties as the weight, thereby conducting instance-level domain alignment to address the heterogeneity of features across different levels. According to the upper bound obtained by theoretical anal-

ysis, the PCC module aggregates domain-private category centroids and conducts cross-space consistency in the private category to mitigate negative transfer. In conclusion, our key contributions are as follows:

- We first reveal that domain-specific alignment is crucial for global-level features. Additionally, we provide a theoretical analysis of the upper bound of UniDAOD to support this observation.
- A novel unsupervised clustering perspective is proposed to sample the instance samples through continuous frequency bins of the gradient norm as the sampling radius on the Gaussian distribution modeling.
- We propose a novel Dual Probabilistic Alignment (DPA) framework. DPA aligns domain-private categories at the global and domain-shared categories at the instance level. In addition, the DPA aggregates domain-private categories centroid between feature and probability spaces to mitigate negative transfer.
- Extensive experiments across open-set, partial-set, and closed-set scenarios demonstrate that the DPA framework achieves state-of-the-art performance, significantly surpassing existing UniDAOD methods.

Related Work

Domain Adaptive Object Detection (DAOD)

DAOD addresses the covariate shift from labeled data in the source domain to the unlabeled target domain under the closed-set categories. Existing DAOD methods can be categorized into adversarial training and mean teacher paradigms. As for adversarial training, DAF (Chen et al. 2018) incorporates global and instance alignment modules based on the Faster-RCNN detector, while incremental variants (Krishna, Ohashi, and Sinha 2023) improve global and instance alignment. ATMT (Li et al. 2023a) explores the potential of self-supervised learning, and EPM (Hsu et al. 2020) introduces a new FCOS detector. Li et al. propose graph-based alignment methods (Li et al. 2022; Li, Liu, and Yuan 2022b) to align class-conditional features, with IGG (Li et al. 2023b) enhancing graph generation by effectively addressing non-informative noise. As for the mean teacher paradigm, existing works (Chen et al. 2022b; Deng et al. 2023; Cao et al. 2023; Li, Guo, and Yuan 2023; Liu et al. 2022) focus on generating pseudo-labels for the target domain. In general, existing DAOD works have limited generalizability in open-world scenarios.

Universal Domain Adaptation (UniDA)

UniDA (You et al. 2019) is a general paradigm for partial-set (Zhang et al. 2018), open-set (Panareda Busto and Gall 2017), and closed-set domain adaptation (Tzeng et al. 2017). The existing UniDA can be categorized into four paradigms, including threshold, clustering, optimal migration, and mutual learning. The threshold methods (You et al. 2019; Fu et al. 2020; Chen et al. 2022a) estimate inter-sample uncertainty to identify shared categories, often relying heavily on manually set thresholds. Clustering-based UniDA




Figure 2: Illustration of the proposed DPA framework. (a) GDPA establishes the global-level embedding feature to sample the outlier in the feature space, then applies a CDF of Gauss distribution to weighting the probability distribution. (b) IDSA obtains the gradient norm of instance probability and models it as a Gauss distribution for sampling and weighting. (c) PCC obtains the domain-private common centroid and constrains distances of samples to the centroid between feature and probability spaces.

methods (Saito et al. 2020; Li et al. 2021a) have been developed to distinguish shared categories. UniOT (Chang et al. 2022) introduces optimal transport to detect shared categories. Despite diversity advancements in classification tasks, UniDAOD remains in its early stages. US-DAF (Shi et al. 2022) conducts scale domain alignment and uses thresholds for sample mining. Recently, UCF (Lang et al. 2023) and W-adapt (Shi et al. 2024a) employ weighting mechanisms to mitigate the negative transfer of private categories. CODE (Shi et al. 2024b) leverages virtual domain labels to avoid domain-private samples alignment. In general, existing UniDAOD methods perform domain-shared category alignment at both the global and instance levels but overlook the alignment of domain-private categories.

Theoretical Motivation

We theoretically analyze the error risk upper of UniDAOD for shared and private categories based on Unsupervised Domain Adaptation (UDA) Theory (Ben-David et al. 2010).

Definition 1. Universal Domain Adaptation (UniDA). We define source and target domains with data $x_{s/t}$ with distribution $\{\mathcal{D}_{s/t} | \mathcal{P}_{x \sim \mathcal{D}_s} \neq \mathcal{P}_{x \sim \mathcal{D}_t}\}$, and a label function $\psi : x \rightarrow \{c, c^*\}$, where c is the domain-shared categories and c^* is the domain-private categories. For simplicity, we omit the s/t notation unless explicitly indicated. The goal of UniDA is to train a model h that can minimize the error risk of shared categories in target domain $\epsilon_c(h) = \min \mathbb{E}_{(x, \psi(x)) \sim \mathcal{D}_t^c} [h(x) \neq \psi(x)]$.

Definition 2. Private and Shared Category Error Risks. Given the input x draw from the distribution \mathcal{D} with the label function ψ , we can get the error risk as follows:

$$\epsilon(h) = \mathbb{E}_{x \sim \mathcal{D}} |h(x) - \psi(x)| = \int_x |h(x) - \psi(x)| \mathcal{P}_{x, \psi} dx. \quad (1)$$

We can decompose Eq. (1) into the domain-shared categories and domain-private categories and let $\mathcal{F}_{h, \psi}$ denote

$|h(x) - \psi(x)|$ as follows:

$$\epsilon(h) = \int_{x_c} \mathcal{F}_{h, \psi} \mathcal{P}_{x, \psi} dx + \int_{x_{c^*}} \mathcal{F}_{h, \psi} \mathcal{P}_{x, \psi} dx = \epsilon_c(h) + \epsilon_{c^*}(h), \quad (2)$$

where $\epsilon_c(h)$ and $\epsilon_{c^*}(h)$ are the error risk for domain-shared and domain-private categories. Then, we define a symmetric hypothesis space \mathcal{H} based on the error risks upper bound of UDA (Ben-David et al. 2010) and combined with Eq. (2) to obtain the error risks upper bound for UniDAOD as follows:

$$\begin{aligned} \epsilon_c^t(h) \leq & \underbrace{\epsilon^s(h) - \epsilon_{c^*}^t(h)}_{\mathcal{L}_{\text{det}}} + \underbrace{d_{\mathcal{H}}^c(\mathcal{D}_s^c, \mathcal{D}_t^c)}_{\mathcal{L}_{\text{PCC}}} + \underbrace{d_{\mathcal{H}}^{c^*}(\mathcal{D}_s^{c^*}, \mathcal{D}_t^{c^*})}_{\mathcal{L}_{\text{IDSA}}} \\ & + \underbrace{\mathbb{E}_{x \sim \mathcal{D}_s^{c^*}} [\psi_t(x) - \psi_s(x)]}_{\arg \min_{\mathcal{H}} \text{ at global-level}} + \underbrace{\mathbb{E}_{x \sim \mathcal{D}_s^c} [\psi_t(x) - \psi_s(x)]}_{\arg \min_{\mathcal{H}} \text{ at instance-level}}. \end{aligned} \quad (3)$$

Remark 1. Existing UniDAOD methods (Shi et al. 2022; Lang et al. 2023; Shi et al. 2024b,a) employ the domain shared category domain alignment $d_{\mathcal{H}}^c(\mathcal{D}_s^c, \mathcal{D}_t^c)$ for both the global and instance level features, ignore domain-private category alignment $d_{\mathcal{H}}^{c^*}(\mathcal{D}_s^{c^*}, \mathcal{D}_t^{c^*})$ and maximize target domain error risk $\epsilon_{c^*}^t(h)$ of the domain-private categories. This oversight leads to an increase in the upper bound of the domain-shared categories, denoted as $\epsilon_c^t(h)$, in the target domain. In UniDAOD, this issue pertains to the global-level domain-private category alignment and the instance-level domain-shared category alignment. This conclusion is consistent with the observations in Fig. 1.

Dual Probabilistic Alignment Framework

Overview. The proposed DPA framework is depicted in Fig. 2. To minimize the upper bound $\epsilon_c^t(h)$ of the domain-shared category of the target domain, DPA comprises GDPA, IDSA, and PCC to optimal the terms in Eq. (3). GDPA minimizes

the domain-private category c^* domain distribution discrepancy $d_{\mathcal{H}}^{c^*}(\mathcal{D}_s^{c^*}, \mathcal{D}_t^{c^*})$ for global-level features, and IDSA minimizes domain-shared category domain distribution discrepancy $d_{\mathcal{H}}^c(\mathcal{D}_s^c, \mathcal{D}_t^c)$ for instance-level features. Additionally, PCC maximize the domain-private category risk error $\epsilon_{c^*}^t(h)$ of target domain.

Global-level Domain Private Alignment (GDPA)

To align global-level domain-private category features, we sample outliers in feature space and model the batch samples as Gauss distribution for the cumulative distribution function to estimate the domain distribution as the weights.

Global-level Sampling. The global-level sampling process involves constructing the dynamic feature centroid and updating the learnable radius. The embedding feature x_e through the encoder of the domain discriminator computes the dynamic feature centroid and the learnable radius. The dynamic feature centroid, denoted as $\mathbf{C} = \mathbf{M}(y_d)$, is derived from the memory bank $\mathbf{M} \in \mathbb{R}^{2 \times c'}$. The learnable radius applies the softplus activation function to calculate the boundary $d = \log(1 + e^{\nabla(y_d)})$, where $\nabla \in \mathbb{R}^2$ represents the learnable boundary parameters and y_d is the domain label. Subsequently, we can calculate the distance from the sample to the feature centroid and perform sampling with the learnable boundaries as follows:

$$\begin{aligned} \Omega_g^{\text{neg}} &= \{i \mid \|x_i - \mathbf{C}\|_2 > d\}, \\ \Omega_g^{\text{pos}} &= \{i \mid \|x_i - \mathbf{C}\|_2 \leq d\}, \end{aligned} \quad (4)$$

where Ω_g^{neg} represents the negative sample indexes and Ω_g^{pos} represents the positive sample indexes. Finally, we update the feature centroid and the learnable radius. The memory bank $\mathbf{M}(y_d) = \mathbf{M}(y_d) \cdot \pi + \bar{x}_e \cdot (1 - \pi)$ is adjusted using a momentum update as $\pi = \frac{\bar{x}_e \cdot \mathbf{M}(y_d)}{\|\bar{x}_e\|_2 \cdot \|\mathbf{M}(y_d)\|_2}$, where $\bar{x}_e = \frac{1}{n} \sum_i x_{e,i}$ represents the mean of the current batch of embedding features. Additionally, the learnable radius is updated based on the boundary loss $\mathcal{L}_{\text{bound}}$ as follows:

$$\mathcal{L}_{\text{bound}} = \frac{1}{n} \sum_{i=1}^n \epsilon_i (d - \|x_i - \mathbf{C}\|_2) + (1 - \epsilon_i) (\|x_i - \mathbf{C}\|_2 - d), \quad (5)$$

where $\epsilon_i = \mathbb{I}(i \in \Omega_g^{\text{neg}})$ is indicator function. In contrast to the existing UniDAOD threshold methods (Shi et al. 2024b, 2022), GDPA sampling is data-driven adaptive updating.

Calculating Alignment Weight. We obtain the domain probability of embedding feature $p_g = \mathcal{F}_{\mathcal{D}_g}(x_e)$ through the global level domain discriminator \mathcal{D}_g . To model the Gaussian distribution for the probabilities of the current batch, we estimate expectation $\mu_g = \frac{1}{n} \sum_i p_{g,i}$ and variance $\sigma_g^2 = \frac{1}{n} \sum_i (p_{g,i} - \mu_g)^2$. After that, we adopt the cumulative distribution function (CDF), which is calculated as the weight for domain alignment as follows:

$$\Phi(z) = \frac{1}{2} \left[1 + \text{erf}\left(\frac{z - \mu_g}{\sigma_g \sqrt{2}}\right) \right], \quad (6)$$

where $\text{erf}(\cdot)$ is the Gauss error function, and z is the mean of the probability distribution for adversarial training. As illustrated in Fig. 5, we observe that as the shared category

ratio $\beta = \frac{\mathcal{C}_s \cap \mathcal{C}_t}{\mathcal{C}_s \cup \mathcal{C}_t}$ decreases, the weights $\frac{\Phi_s / (1 - \Phi_t)}{\Phi_s + (1 - \Phi_t)}$ exhibit increased scaling to accommodate a substantial domain gap.

Global-level Domain Alignment. To achieve global-level alignment, the gradient reversal layer is employed with focal loss as follows:

$$\begin{aligned} \mathcal{L}_{\text{GDPA}} &= -\frac{1}{n_{\text{neg}}} \sum_{i=1}^{n_{\text{neg}}} \left[\frac{\Phi_s}{\Phi_s + (1 - \Phi_t)} (1 - p_{s,i})^\gamma \log p_{s,i} \right. \\ &\quad \left. + \frac{(1 - \Phi_t)}{\Phi_s + (1 - \Phi_t)} p_{t,i}^\gamma (1 - \log p_{t,i}) \right], \end{aligned} \quad (7)$$

where γ is the gamma parameter, p is the probability of the domain discriminator, and $n_{\text{neg}} \in \Omega_g^{\text{neg}}$ represents the negative samples numbers.

Instance-level Domain Shared Alignment (IDSA)

To efficiently align the domain-shared category at the instance level, we calculate the gradient norm of the instance samples to model a Gaussian distribution to discard outlier samples and estimate the weight of the domain-shared category samples to improve domain alignment.

Instance-level Sampling. The instance-level sampling process involves constructing the probability space and the sampling criteria. First, we build the probability space to model the Gaussian distribution. The number of \hat{n} instance-level features x_v is generated by the domain discriminator $\mathcal{F}_{\mathcal{D}_v}$ to calculate the domain probabilities $p_v = \text{Sigmoid}(\mathcal{F}_{\mathcal{D}_v}(x_v))$. The domain probabilities are used to compute the gradient norm η_v for each instance as $\eta_v = |p_v - y_d|$, where y_d is the domain label. We then construct the gradient norm bins $\hat{\Omega} = \{i \cdot \psi \mid i \in \mathbb{Z}\}$ to calculate the gradient norm frequencies τ , which model the Gaussian distribution with the minimum interval $\psi = \text{argmin}\{(\eta_v^{\text{max}} - \eta_v^{\text{min}}) \cdot \eta_v^{\text{std}}, \delta\}$, where δ is the hyperparameter. For the sampling criteria, we leverage the statistical characteristics of the Gaussian distribution. The first sampling criterion involves filtering out samples in noncontinuous frequency bins. The second criterion is related to the characteristic of the Gaussian distribution, where the frequency of filtered bins is lower than that of continuous bins. The sampling process is as follows:

$$\begin{aligned} \Omega_v^{\text{pos}} &= \{i \mid \tau_i > 0, i_{j+1} - i_j = 1, \forall j\}, \\ \Omega_v^{\text{neg}} &= \{i \mid i \notin \Omega_v^{\text{pos}}, \tau_i < \tau_\omega\}, \end{aligned} \quad (8)$$

where τ_ω denotes the frequency in the first and last continuous bins in positive samples Ω_v^{pos} . The sum of the bins represents the sampling radius from the feature centroid in the feature space, which dynamically adjusts the bins following a Gaussian distribution of the source or target domain data during adversarial training.

Calculating Alignment Weight. These negative instances are excluded from the instance alignment through the instance weight as follows:

$$\mathcal{W}_v = \begin{cases} 0, & \mathcal{W}_v \in \Omega_v^{\text{neg}}, \\ 1 - \frac{|\eta_v^{\text{mean}} - 0.5|}{0.5}, & \mathcal{W}_v \in \Omega_v^{\text{pos}}, \end{cases} \quad (9)$$

where η_v^{mean} is the mean value of the gradient norm η_v .

Methods	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	motor	person	plant	sheep	sofa	mAP
Source Only	31.8	41.2	31.1	34.7	5.1	33.7	23.0	20.7	8.3	43.0	52.7	49.6	40.6	17.0	13.8	29.8
DAF (Chen et al. 2018)	37.2	38.0	26.9	35.9	2.3	35.2	24.0	28.5	4.2	33.8	54.7	59.4	58.4	13.4	17.9	31.3
MAF (He and Zhang 2019)	24.2	42.9	35.1	32.3	11.0	41.7	22.4	32.6	6.7	40.0	59.1	52.7	41.0	24.1	17.9	32.2
HTCN (Chen et al. 2020)	25.9	47.8	36.0	32.8	11.3	39.4	51.7	18.7	10.5	40.9	56.3	57.9	49.4	21.3	20.4	34.7
UAN (You et al. 2019)	26.6	37.7	48.2	31.5	8.6	32.8	23.7	31.6	2.4	36.6	56.6	42.8	44.8	14.7	16.4	30.3
CMU (Fu et al. 2020)	14.7	41.9	52.5	34.7	9.2	36.5	38.1	21.0	7.6	37.0	48.6	55.7	44.5	17.7	21.1	32.1
SFA (Wang et al. 2021)	25.2	30.5	24.3	22.1	1.6	28.3	37.5	23.4	2.6	19.9	43.1	47.0	39.0	16.8	22.0	25.5
US-DAF (Shi et al. 2022)	34.9	40.8	28.9	36.4	17.7	38.4	64.6	28.0	10.3	45.8	64.5	62.5	52.1	25.8	24.8	38.4
UCF (Lang et al. 2023)	36.2	44.3	28.3	37.1	2.2	36.0	61.9	27.7	4.0	39.9	64.7	64.2	52.6	20.9	26.9	36.5
CODE (Shi et al. 2024b)	36.8	45.1	42.0	37.7	18.4	44.5	47.7	33.1	8.4	45.7	69.2	61.7	50.7	25.5	24.6	39.4
DPA (Ours)	32.9	46.0	62.1	41.4	4.2	42.0	64.4	33.3	8.1	40.5	67.1	64.2	57.8	32.1	25.1	41.4

Table 1: Comparison on Pascal VOC to Clipart1k (Open-set: $\beta = 75\%$, $C_s \cap C_t \neq \emptyset$, $C_s \setminus C_t \neq \emptyset$, $C_t \setminus C_s \neq \emptyset$).

Methods	bus	car	cat	chair	cow	table	dog	horse	motor	person	mAP
Source Only	43.3	33.0	8.4	32.1	24.0	28.7	6.9	34.9	51.8	42.5	30.6
DAF (Chen et al. 2018)	37.5	32.8	10.2	40.3	27.2	31.3	4.1	41.0	55.5	52.0	33.2
MAF (He and Zhang 2019)	37.1	31.1	9.7	38.1	19.9	29.1	2.5	37.3	50.7	50.0	30.6
HTCN (Chen et al. 2020)	29.5	34.4	17.3	33.8	50.6	14.0	3.6	46.9	74.7	58.5	36.3
UAN (You et al. 2019)	48.9	26.4	14.6	36.7	49.9	30.0	3.2	39.9	56.1	52.0	35.8
CMU (Fu et al. 2020)	33.3	32.8	8.1	41.5	55.5	24.6	5.6	43.3	54.9	60.4	36.0
SFA (Wang et al. 2021)	29.4	28.3	14.4	30.5	29.4	13.0	3.1	26.8	61.6	2.9	27.9
US-DAF (Shi et al. 2022)	31.3	41.9	7.3	42.9	64.3	30.0	5.7	44.8	69.5	61.9	40.0
UCF (Lang et al. 2023)	32.4	37.4	4.0	33.5	59.3	40.2	4.4	39.0	58.1	61.7	37.0
CODE (Shi et al. 2024b)	50.0	38.8	19.5	42.0	46.9	34.5	13.5	43.7	64.3	52.9	40.6
DPA (Ours)	45.0	41.3	13.7	38.0	64.9	30.0	13.8	45.8	74.0	60.2	42.7

Table 2: Comparison on Pascal VOC to Clipart1k (Open-set: $\beta = 50\%$, $C_s \cap C_t \neq \emptyset$, $C_s \setminus C_t \neq \emptyset$, $C_t \setminus C_s \neq \emptyset$).

Instance-level Domain Alignment. This processing aims to provide instance-level features into the domain discriminator for adversarial training to achieve domain alignment. Based on the obtained weights \mathcal{W}_v , the loss function of the IDSA module is as follows:

$$\mathcal{L}_{\text{IDSA}} = -\frac{1}{\hat{n}} \sum_i^{\hat{n}} \mathcal{W}_v \cdot (1 - p_v) \log(p_v) + \mathcal{W}_v \cdot p_v (1 - \log(p_v)), \quad (10)$$

where \hat{n} is the number of instance proposals. By optimizing the function $\mathcal{L}_{\text{IDSA}}$ in adversarial training, the IDSA module mitigates negative transfer caused by domain-private feature alignment. It calibrates the domain-shared feature distribution according to a Gaussian distribution to enhance positive transfer between source and target domains.

Private Class Constraint (PCC)

Given the instance-level feature x_v and the domain probabilities p_v , we first perform a classifier head to establish domain-private categories for both the source domain and the target domain: $\{c_i^* \mid \hat{y}_{s,i}^* \cap \hat{y}_{t,i}^* = \emptyset\}_i^{n^*}$. To aggregate the centroids of domain-private categories in feature and probability spaces, we calculate the feature centroid \bar{x}_v and the probability centroid \bar{p}_v . We then conduct the cosine similarity distance $\mathcal{G}_i = \|x_{v,i} - \bar{x}_v\|_2$ and probability samples $\mathbf{g}_i = \|p_{v,i} - \bar{p}_v\|_2$ to centroid to measure intra-domain distances. To measure the intra-domain distance, we use cosine

similarity, defined as $\varepsilon_{s/t} = \frac{1}{n^*} \sum_{i=1}^{n^*} \frac{\mathcal{G}_i \cdot \mathbf{g}_i}{\|\mathcal{G}_i\|_2 \cdot \|\mathbf{g}_i\|_2}$. Finally, we employ the mean squared error (MSE) loss function to minimize the inter-domain distance, as follows:

$$\mathcal{L}_{\text{PCC}} = (\varepsilon_s - \varepsilon_t)^2. \quad (11)$$

The loss function \mathcal{L}_{PCC} optimizes the network by adopting a gradient detach for the source domain.

Optimization

The training loss of the DPA is represented as \mathcal{L}_{DPA} , which consists of the following loss terms:

$$\mathcal{L}_{\text{DPA}} = \mathcal{L}_{\text{det}} + \mathcal{L}_{\text{GDPA}} + \mathcal{L}_{\text{IDSA}} + \alpha \mathcal{L}_{\text{PCC}}, \quad (12)$$

where \mathcal{L}_{det} is the Faster-RCNN detector loss. $\mathcal{L}_{\text{GDPA}}$ and $\mathcal{L}_{\text{IDSA}}$ are the domain alignment losses for the GDPA and IDSA modules at the global and instance levels, respectively. The \mathcal{L}_{DPA} is optimized using the SGD optimizer. The bound loss $\mathcal{L}_{\text{bound}}$ is optimized using the Adam optimizer with a learning rate set to 0.1. The hyperparameter of α is 0 for the initial epoch and 0.1 thereafter.

Experiments

Implementation Details

We conduct extensive experiments following the setting (Shi et al. 2022) for three benchmarks: open-set, partial-set, and

Methods	plane	bicycle	bird	boat	bottle	mAP
Source Only	33.2	55.7	25.4	29.2	41.6	37.0
DAF	31.5	42.5	25.2	34.4	50.8	36.9
MAF	29.3	57.0	27.1	33.9	41.8	37.8
HTCN	32.5	53.0	24.1	27.0	48.4	37.0
UAN	35.6	55.9	27.1	28.2	44.2	38.2
CMU	45.5	52.7	28.8	29.4	40.1	39.3
SFA	28.4	32.4	27.2	34.2	34.2	31.3
US-DAF	44.2	57.5	27.9	32.2	40.5	40.5
UCF	35.8	52.9	28.6	20.8	55.7	38.8
CODE	42.1	61.4	26.2	32.1	44.1	41.2
DPA (Ours)	40.8	58.1	28.2	33.7	52.0	42.5

Table 3: Comparison on Pascal VOC to Clipart1k (Open-set: $\beta = 25\%$, $C_s \cap C_t \neq \emptyset$, $C_s \setminus C_t \neq \emptyset$, $C_t \setminus C_s \neq \emptyset$).

Methods	bicycle	bird	car	cat	dog	person	mAP
Source Only	29.8	50.2	47.1	62.2	51.5	57.8	49.8
DAF	29.5	53.8	50.6	58.1	48.1	56.5	49.4
MAF	28.5	50.0	46.8	59.4	50.2	58.6	48.9
HTCN	26.4	43.0	46.5	50.8	44.0	53.9	44.1
UAN	33.6	52.1	53.8	62.4	52.2	56.1	51.7
CMU	36.9	51.2	53.3	59.3	51.7	59.9	52.0
CODE	39.6	53.1	54.7	57.6	56.1	57.8	53.1
SFA	4.4	17.4	11.1	14.9	20.1	18.4	14.4
US-DAF	35.0	52.4	52.7	63.1	54.3	59.8	52.9
UCF	34.8	52.0	53.8	61.9	54.2	60.5	52.9
DPA (Ours)	31.3	55.2	56.4	61.1	62.1	58.3	54.1

Table 4: Comparison WaterColor to Pascal VOC (Partial-set: $\beta = 30\%$, $C_s \subset C_t$).

Methods	bicycle	bird	car	cat	dog	person	mAP
Source Only	82.4	51.7	48.4	39.9	30.7	59.2	52.0
DAF	73.4	51.9	43.1	35.6	28.8	63.1	49.3
MAF	70.4	50.3	44.3	36.7	30.6	62.9	49.2
HTCN	74.1	49.8	51.9	35.3	35.3	66.0	52.1
UAN	78.0	53.6	50.4	36.4	35.8	65.6	53.3
CMU	82.0	53.9	48.6	39.6	33.1	66.0	53.9
SFA	37.1	39.3	32.3	52.7	9.9	34.1	34.2
US-DAF	86.5	54.1	50.0	43.0	34.0	63.2	55.2
UCF	84.8	52.1	49.8	40.6	33.8	63.2	54.1
CODE	87.9	55.3	50.7	38.9	34.7	67.5	55.8
DPA (Ours)	86.1	53.5	50.3	44.1	38.1	65.6	56.3

Table 5: Comparison on Pascal VOC to WaterColor (Partial-set: $\beta = 30\%$, $C_s \supset C_t$).

closed-set. The baseline approach (Chen et al. 2018) adopts Faster-RCNN as the base detector with the focal loss, and the backbone is ResNet-101 (He et al. 2016) or VGG-16 in (Simonyan and Zisserman 2014) pre-trained on ImageNet (Deng et al. 2009), which adopt VGG-16 in Cityscape to Foggy Cityscape and other benchmarks adopt the ResNet-

Methods	DA Settings	mAP
SFA (Wang et al. 2021)	DAOD	41.3
SCAN++ (Li, Liu, and Yuan 2022a)	DAOD	42.8
SIGMA++ (Li, Liu, and Yuan 2023)	DAOD	44.5
US-DAF (Shi et al. 2022)	UniDAOD	37.8
UCF (Lang et al. 2023)	UniDAOD	34.2
CODE (Shi et al. 2024b)	UniDAOD	42.1
DPA (Ours)	UniDAOD	46.3

Table 6: Comparison on Cityscape to Foggy Cityscape (Closed-set: $\beta = 100\%$, $C_s = C_t$).

101. The DPA model optimized training iterations are 100k, with an initial learning rate of 1e-3 and a subsequent decay of the learning rate to 1e-4 following 50k iterations. The detection performance is evaluated with the mean Average Precision (mAP) metric, and the threshold of mAP follows the setting (Shi et al. 2022) to set 0.5.

Datasets and Domain Adaptation Settings

We evaluate our DPA framework on five datasets across three domain adaptation scenarios (open-set, partial-set, and closed-set): Foggy Cityscapes (Sakaridis, Dai, and Van Gool 2018), Cityscapes (Cordts et al. 2016), Pascal VOC (Everingham et al. 2010), Clipart1k (Inoue et al. 2018), and Watercolor (Inoue et al. 2018). In the open-set scenario, there are shared and private categories in both the source and target domains. We introduce mutil ratios $\beta = \{0.25, 0.5, 0.75\}$ to construct different shared category ratios $\beta = \frac{C_s \cap C_t}{C_s \cup C_t}$ benchmarks. In the partial-set scenario, the category set of the source domain is the subset for the target domain, and vice versa. In the closed-set scenario, the categories in the target and source domains are identical.

Comparisons with the State-of-the-Arts

Open-set scenario. Tables 1, 2, and 3 present the open-set domain adaptive object detection performance from Pascal VOC to Clipart1k under different category overlap ratios β . The proposed framework consistently achieves state-of-the-art performance across various category settings compared to other methods. Compared to related DAOD methods (HTCN, MAF, DAF, SFA), the proposed DPA framework demonstrates significant performance advantages. Additionally, in comparison with UniDAOD methods (CODE, US-DAF, CMU, UAN, and UCF), the proposed DPA framework also exhibits superior performance.

Partial-set scenario. The results of the partial-set domain adaptation are presented in 4 and 5. In the partial-set scenario, the private categories are exclusively present in the source or target domain, leading to negative transfer. The proposed DPA method effectively addresses this issue and outperforms other approaches by a significant margin, achieving 54.1% and 56.3% mAP. For the domain-private category in the target domain ($C_s \subset C_t$), DPA enhances performance by 1.2% compared to the UniDAOD method (US-DAF), as shown in Table 4. For the domain-private cat-

Ratio β	75%	50%	25%
Baseline	33.9	37.9	38.2
DPA w/o GDPA	38.3	39.7	40.1
DPA w/o IDSA	36.3	40.1	39.9
DPA w/o PCC	40.7	42.4	39.4
DPA (Ours)	41.4	42.7	42.5

Table 7: Ablation study on Pascal VOC to Clipart1k (Open-set: $\beta = 75\%, 50\%, 25\%$).

egory in the source domain ($C_s \supset C_t$), DPA improves performance by 4.2% compared to the DAOD method (HTCN). **Closed-set scenario.** The closed-set scenario results are shown in Table 6 show that although existing UniDAOD methods (Shi et al. 2024b,a) integrated with advanced DAOD methods achieve notable performance improvements on closed-set, the DAOD methods significantly outperform the UniDAOD methods. This advantage comes from DAOD under the closed-set assumption, while UniDAOD prioritizes dealing with label shifts in open environments. The proposed DPA framework exhibits satisfactory performance in closed-set scenarios through probability modeling.

Ablation Study

We conduct ablation experiments on each submodule, with the corresponding results presented in Table 7. In these experiments, each module of the proposed DPA framework improves performance. The GDPA and IDSA modules provide significant gains when the domain-shared categories ratio is high ($\beta = 50\%, 75\%$), while the PCC module leads to more substantial improvements when the ratio is low ($\beta = 25\%$).

Category-wise Performance Analysis

To compare the performance of the proposed method with existing DAOD and UniDAOD methods in terms of positive and negative transfer, we present the performance gains of DAOD and UniDAOD relative to the source-only model in Fig. 3. The DAOD methods exhibit significant negative transfer, where DAF, MAF, and HTCN drop by approximately 2%, 4%, and 1% AP in class 0, respectively. In contrast, the UniDAOD methods mitigate negative transfer, with CODE and DPA achieving positive transfer of around 3% and 10% in class 4, respectively. This category-wise performance analysis proves that the proposed method effectively combats negative transfer and strengthens positive transfer.

Qualitative Open-set Alignment Analysis

We further analyze the probability gap in our DPA framework for open-set alignment. As shown in Fig. 4(a), the global-level mean probability gap is more pronounced in our DPA, highlighting its effectiveness in distinguishing domain-private categories. In contrast, Fig. 4(b) shows a smaller mean probability gap at the instance level, demonstrating that our DPA better aligns domain-shared categories. Additionally, we perform a weight quantitative analysis of global-level domain-private alignment, as illustrated in Fig. 5. As the ratio of domain-private categories increases,

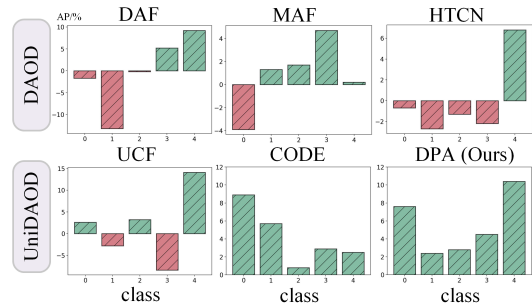


Figure 3: Category-wise performance gain over the source-only model (classes are plane, bicycle, bird, boat, and bottle). Positive transfer is green, and negative transfer is red.




Figure 4: Qualitative analysis of category alignment in terms of the mean probability gap: (a) global-level features and (b) instance-level features. The horizontal axis represents training iterations ($\times 100$), and the vertical axis shows the probability of the domain discriminator. The benchmark is Pascal VOC to Clipart1k. ($\beta = 25\%$).




Figure 5: The weight quantitative analysis of global-level domain-private feature. The horizontal axis is training iteration ($\times 100$), and the vertical axis is weight values $\frac{\Phi_s/(1-\Phi_t)}{\Phi_s+(1-\Phi_t)}$ in source and target domains.

the mean weight gap also increases, indicating that adversarial training adaptively penalizes features associated with domain-private categories through weight adjustments.

Conclusion

We propose a DPA framework for universal domain adaptive object detection with two kinds of probabilistic alignment. Inspired by a theoretical perspective, we propose a GDPA module for aligning global-level private samples and an IDSA module for aligning instance-level domain-shared samples. To combat negative transfer, we propose a PCC module to confuse the discriminability of private categories. Extensive experiments are conducted on open, partial, and closed set scenarios and demonstrate our DPA outperforms state-of-the-art UniDAOD methods by a remarkable margin.

Acknowledgments

This work is supported by the InnoHK program, and the National Natural Science Foundation of China (Grant No.#62306313).

References

- Ali, S.; Ghatwary, N.; Jha, D.; Isik-Polat, E.; Polat, G.; Yang, C.; Li, W.; Galdran, A.; Ballester, M.-Á. G.; Thambawita, V.; et al. 2024. Assessing generalisability of deep learning-based polyp detection and segmentation methods through a computer vision challenge. *Sci. Rep.*
- Ben-David, S.; Blitzer, J.; Crammer, K.; Kulesza, A.; Pereira, F.; and Vaughan, J. W. 2010. A theory of learning from different domains. *Machine learning*, 79: 151–175.
- Cai, M.; Kezierbieke, J.; Zhong, X.; and Chen, H. 2024. Uncertainty-Aware and Class-Balanced Domain Adaptation for Object Detection in Driving Scenes. *IEEE Transactions on Intelligent Transportation Systems*.
- Cao, S.; Joshi, D.; Gui, L.-Y.; and Wang, Y.-X. 2023. Contrastive Mean Teacher for Domain Adaptive Object Detectors. In *CVPR*, 23839–23848.
- Chang, W.; Shi, Y.; Tuan, H.; and Wang, J. 2022. Unified optimal transport framework for universal domain adaptation. *NeurIPS*, 35: 29512–29524.
- Chapman, N. H.; Dayoub, F.; Browne, W.; and Lehnert, C. 2023. Predicting class distribution shift for reliable domain adaptive object detection. *IEEE Robotics and Automation Letters*.
- Chen, C.; Zheng, Z.; Ding, X.; Huang, Y.; and Dou, Q. 2020. Harmonizing transferability and discriminability for adapting object detectors. In *CVPR*, 8869–8878.
- Chen, L.; Lou, Y.; He, J.; Bai, T.; and Deng, M. 2022a. Identical neighborhood contrastive learning for universal domain adaptation. In *AAAI*, volume 36, 6258–6267.
- Chen, M.; Chen, W.; Yang, S.; Song, J.; Wang, X.; Zhang, L.; Yan, Y.; Qi, D.; Zhuang, Y.; Xie, D.; et al. 2022b. Learning Domain Adaptive Object Detection with Probabilistic Teacher. In *ICML*, 3040–3055. PMLR.
- Chen, Y.; Li, W.; Sakaridis, C.; Dai, D.; and Van Gool, L. 2018. Domain adaptive faster r-cnn for object detection in the wild. In *CVPR*, 3339–3348.
- Cordts, M.; Omran, M.; Ramos, S.; Rehfeld, T.; Enzweiler, M.; Benenson, R.; Franke, U.; Roth, S.; and Schiele, B. 2016. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 3213–3223.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *CVPR*, 248–255. Ieee.
- Deng, J.; Xu, D.; Li, W.; and Duan, L. 2023. Harmonious Teacher for Cross-Domain Object Detection. In *CVPR*, 23829–23838.
- Everingham, M.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2010. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88: 303–338.
- Fu, B.; Cao, Z.; Long, M.; and Wang, J. 2020. Learning to detect open classes for universal domain adaptation. In *ECCV*, 567–583. Springer.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*, 770–778.
- He, Z.; and Zhang, L. 2019. Multi-adversarial faster-rcnn for unrestricted object detection. In *ICCV*, 6668–6677.
- Hsu, C.-C.; Tsai, Y.-H.; Lin, Y.-Y.; and Yang, M.-H. 2020. Every pixel matters: Center-aware feature alignment for domain adaptive object detector. In *ECCV*, 733–748. Springer.
- Huang, T.; Huang, C.-C.; Ku, C.-H.; and Chen, J.-C. 2024. Blenda: Domain Adaptive Object Detection Through Diffusion-Based Blending. In *ICASSP*, 4075–4079. IEEE.
- Inoue, N.; Furuta, R.; Yamasaki, T.; and Aizawa, K. 2018. Cross-domain weakly-supervised object detection through progressive domain adaptation. In *CVPR*, 5001–5009.
- Jia, D.; Yuan, Y.; He, H.; Wu, X.; Yu, H.; Lin, W.; Sun, L.; Zhang, C.; and Hu, H. 2023. Detsr with hybrid matching. In *CVPR*, 19702–19712.
- Krishna, O.; Ohashi, H.; and Sinha, S. 2023. MILA: memory-based instance-level adaptation for cross-domain object detection. *arXiv preprint arXiv:2309.01086*.
- Lang, Q.; He, Z.; Fu, X.; and Zhang, L. 2023. Class-aware Memory Guided Unbiased Weighting for Universal Domain Adaptive Object Detection. In *ICCV*, 4345–4354.
- Li, G.; Kang, G.; Zhu, Y.; Wei, Y.; and Yang, Y. 2021a. Domain consensus clustering for universal domain adaptation. In *CVPR*, 9757–9766.
- Li, K.; Wington, C.; Tensmeyer, C.; Morariu, V. I.; Zhao, H.; Manjunatha, V.; Barmpalios, N.; and Fu, Y. 2023a. Improving Cross-Domain Detection With Self-Supervised Learning. In *CVPR*, 4745–4754.
- Li, P.; He, Y.; Yu, F. R.; Song, P.; Yin, D.; and Zhou, G. 2023b. IGG: Improved Graph Generation for Domain Adaptive Object Detection. In *ACM MM*, 1314–1324.
- Li, W.; Chen, Z.; Li, B.; Zhang, D.; and Yuan, Y. 2021b. Htd: Heterogeneous task decoupling for two-stage object detection. *IEEE Transactions on Image Processing*, 30: 9456–9469.
- Li, W.; Guo, X.; and Yuan, Y. 2023. Novel Scenes & Classes: Towards Adaptive Open-set Object Detection. In *ICCV*, 15780–15790.
- Li, W.; Liu, J.; Han, B.; and Yuan, Y. 2023c. Adjustment and alignment for unbiased open set domain adaptation. In *CVPR*, 24110–24119.
- Li, W.; Liu, X.; Ma, J.; and Yuan, Y. 2024. CLIFF: Continual Latent Diffusion for Open-Vocabulary Object Detection. In *ECCV*.
- Li, W.; Liu, X.; Yao, X.; and Yuan, Y. 2022. Scan: Cross domain object detection with semantic conditioned adaptation. In *AAAI*, volume 36, 1421–1428.
- Li, W.; Liu, X.; and Yuan, Y. 2022a. Scan++: Enhanced semantic conditioned adaptation for domain adaptive object detection. *IEEE Transactions on Multimedia*, 25: 7051–7061.

- Li, W.; Liu, X.; and Yuan, Y. 2022b. Sigma: Semantic-complete graph matching for domain adaptive object detection. In *CVPR*, 5291–5300.
- Li, W.; Liu, X.; and Yuan, Y. 2023. Sigma++: Improved semantic-complete graph matching for domain adaptive object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(7): 9022–9040.
- Liu, X.; Li, W.; Yang, Q.; Li, B.; and Yuan, Y. 2022. Towards robust adaptive object detection under noisy annotations. In *CVPR*, 14207–14216.
- Liu, X.; Li, W.; and Yuan, Y. 2023. Decoupled Unbiased Teacher for Source-Free Domain Adaptive Medical Object Detection. *TNNLS*.
- Lu, Y.; Shen, M.; Ma, A. J.; Xie, X.; and Lai, J.-H. 2024. MLNet: Mutual Learning Network with Neighborhood Invariance for Universal Domain Adaptation. In *AAAI*, volume 38, 3900–3908.
- Panareda Busto, P.; and Gall, J. 2017. Open set domain adaptation. In *ICCV*, 754–763.
- Pu, B.; Wang, L.; Yang, J.; He, G.; Dong, X.; Li, S.; Tan, Y.; Chen, M.; Jin, Z.; Li, K.; et al. 2024. M3-UDA: A New Benchmark for Unsupervised Domain Adaptive Fetal Cardiac Structure Detection. In *CVPR*, 11621–11630.
- Saito, K.; Kim, D.; Sclaroff, S.; and Saenko, K. 2020. Universal domain adaptation through self supervision. *NeurIPS*, 33: 16282–16292.
- Sakaridis, C.; Dai, D.; and Van Gool, L. 2018. Semantic foggy scene understanding with synthetic data. *International Journal of Computer Vision*, 126: 973–992.
- Shi, C.; Zheng, Y.; and Chen, Z. 2024. Domain Adaptive Thermal Object Detection with Unbiased Granularity Alignment. *ACM Transactions on Multimedia Computing, Communications and Applications*, 20(9): 1–23.
- Shi, W.; Liu, D.; Tan, D.; and Zheng, B. 2024a. A dynamically class-wise weighting mechanism for unsupervised cross-domain object detection under universal scenarios. *Knowledge-Based Systems*, 111987.
- Shi, W.; Liu, D.; Wu, Z.; and Zheng, B. 2024b. Confused and disentangled distribution alignment for unsupervised universal adaptive object detection. *Knowledge-Based Systems*, 112085.
- Shi, W.; Zhang, L.; Chen, W.; and Pu, S. 2022. Universal domain adaptive object detector. In *ACM MM*, 2258–2266.
- Simonyan, K.; and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Tzeng, E.; Hoffman, J.; Saenko, K.; and Darrell, T. 2017. Adversarial discriminative domain adaptation. In *CVPR*, 7167–7176.
- Wang, W.; Cao, Y.; Zhang, J.; He, F.; Zha, Z.-J.; Wen, Y.; and Tao, D. 2021. Exploring sequence feature alignment for domain adaptive detection transformers. In *ACM MM*, 1730–1738.
- You, K.; Long, M.; Cao, Z.; Wang, J.; and Jordan, M. I. 2019. Universal domain adaptation. In *CVPR*, 2720–2729.
- Zhang, J.; Ding, Z.; Li, W.; and Ogunbona, P. 2018. Importance weighted adversarial nets for partial domain adaptation. In *CVPR*, 8156–8164.
- Zhao, Y.; Lv, W.; Xu, S.; Wei, J.; Wang, G.; Dang, Q.; Liu, Y.; and Chen, J. 2024. Detsr beat yolos on real-time object detection. In *CVPR*, 16965–16974.