

Assessing domain gap for continual domain adaptation in object detection

Anh-Dzung Doan ^{a,*}, Bach Long Nguyen ^b, Surabhi Gupta ^c, Ian Reid ^a, Markus Wagner ^b, Tat-Jun Chin ^a

^a Australian Institute for Machine Learning, The University of Adelaide, Adelaide, SA 5000, Australia

^b Department of Data Science and Artificial Intelligence, Monash University, Melbourne, VIC 3800, Australia

^c Safran Electronics & Defense Australasia, Botany, NSW 2019, Australia

ARTICLE INFO

Communicated by Concetto Spampinato

MSC:

41A05

41A10

65D05

65D17

Keywords:

Domain gap

Continual domain adaptation

Object detection

ABSTRACT

To ensure reliable object detection in autonomous systems, the detector must be able to adapt to changes in appearance caused by environmental factors such as time of day, weather, and seasons. Continually adapting the detector to incorporate these changes is a promising solution, but it can be computationally costly. Our proposed approach is to selectively adapt the detector only when necessary, using new data that does not have the same distribution as the current training data. To this end, we investigate three popular metrics for domain gap evaluation and find that there is a correlation between the domain gap and detection accuracy. Therefore, we apply the domain gap as a criterion to decide when to adapt the detector. Our experiments show that our approach has the potential to improve the efficiency of the detector's operation in real-world scenarios, where environmental conditions change in a cyclical manner, without sacrificing the overall performance of the detector. Our code is publicly available <https://github.com/dadung/DGE-CDA>.

1. Introduction

Object detection is the task of localising instances of objects of a certain category within images. It plays a key role in the scene understanding, which is a fundamental problem in various applications, such as, autonomous driving (Arnold et al., 2019), AR/VR (Zhang et al., 2022), and robotics (Sachdeva et al., 2022). A common approach to object detection is to train a deep neural network on a large dataset with ground truth bounding boxes and object categories. Current state-of-the-art methods include Faster-RCNN (Ren et al., 2015), RetinaNet (Lin et al., 2017b), YOLO (Redmon et al., 2016), etc.

In long-term operations, an object detector (OD) must be able to handle variations in appearance caused by natural factors such as time of day, seasons, and weather. These changes can cause the appearance of images to differ from the training data, leading to a decline in detection accuracy. To address this issue, recent works suggest continuously collecting new data and using it to update the perception systems (Doan et al., 2019, 2020; Churchill and Newman, 2013). An ideal approach is to manually label all new data (target data) and combine it with the current training samples (source data) to create a new training database, which is then used to fine-tune the OD (Chen et al., 2022). However, the significant manual annotation effort required for this approach makes it impractical.

To address the impracticality of manual annotation, researchers have focused on two main strategies. The first strategy is to use weakly

supervised learning, which only requires labels indicating the presence of certain object categories in the image (Bilen and Vedaldi, 2016; Kantorov et al., 2016; Song et al., 2014). While this reduces annotation cost, it results in detectors with only about half the accuracy of those trained with fully supervised methods. The second strategy is active learning (Roy et al., 2018; Yuan et al., 2021; Yu et al., 2022; Choi et al., 2021), which selects the most informative samples from the target data for annotation. Typically, target images that cause the greatest uncertainty in the OD are chosen for annotation. However, as object detection involves both regression and classification, accurately characterising the uncertainty of predicted bounding boxes is a complex task.

Despite its potential, active learning still requires human effort in the adaptation process. To achieve a fully autonomous adaptation mechanism, unsupervised domain adaptation (UDA) (Oza et al., 2021) has gained popularity in the research community. In particular, adversarial feature learning (Chen et al., 2018, 2021; Pasqualino et al., 2021) aims to minimise the domain gap between source and target data in the feature space by introducing an adversarial loss with respect to a domain discriminator. A gradient reversal layer (Ganin and Lempitsky, 2015) is used to effectively incorporate adversarial learning into the backpropagation process. However, without proper tuning of the hyperparameters, this approach will likely struggle to achieve satisfactory results on challenging datasets.

* Corresponding author.

E-mail address: dzung.doan@adelaide.edu.au (A.-D. Doan).

Furthermore, image-to-image translation is a widely used approach in UDA for object detection (Arruda et al., 2019; Gao et al., 2021; Schutera et al., 2020). It uses the GAN principle (Goodfellow et al., 2020) to transform the appearance of images from the target domain to resemble the source domain. However, training a GAN-based model is a complex task due to its minimax optimisation problem. This can lead to various challenges such as non-convergence, mode collapse, and diminished gradient.

Another popular technique in UDA for object detection is self-training (RoyChowdhury et al., 2019; Khodabandeh et al., 2019; D’Innocente et al., 2020). It involves using the detector trained on the source data to generate pseudo-labels for the target data. As the pseudo-labels are often noisy, heuristics are proposed to filter out potentially incorrect ones. However, this technique introduces a potential point of failure, as the heuristics may not be able to filter all noisy pseudo-labels, leading to errors that accumulate over time during long-term operations.

In general, existing works primarily focus on efficiently adapting ODs, assuming that the adaptation will always be carried out when the OD encounters a new target domain. However, in practical scenarios, changes in environmental conditions are often cyclical, such as different time of day, weather, and seasons. This means that the appearance of target images will likely resemble that of the source images after a period of time. For example, if the source images were captured in clear conditions from 10am to 3pm, target images captured in cloudy conditions at 1pm will likely be closer to the source images than target images captured in clear conditions at 11pm. This raises the question of whether it is necessary to always adapt the OD during its course of operation. The domain adaptation should only be performed if it can significantly improve the OD’s accuracy or a new domain has a considerable negative impact on the detection performance. Therefore, selectively adapting the OD can save a significant amount of computational cost.

Contributions Prior techniques mainly focus on “how to efficiently adapt ODs?” leaving the question “when to adapt ODs?” unaddressed. Therefore, in this paper, we will investigate the latter. Our hypothesis is that if the distribution of target data is similar to current training data, adapting ODs will not lead to significant improvements in accuracy while resulting in the unnecessary use of resources. To address this, we propose using domain gap as a criterion for deciding when adaptation is necessary in continually changing environmental conditions, based on our finding of a correlation between domain gap and detection accuracy. Our experiments demonstrate that this solution can save adaptation cost without sacrificing the overall performance of the OD.

2. Related work

In order to achieve a robust object detector (OD) for long-term operations, it is vital to continually adapt the OD. There are three main approaches to do this in the literature: weakly supervised object detection (WSOD), active learning for object detection (ALOD), and unsupervised domain adaptation for object detection (UDAOD).

In WSOD, human effort is required to provide image-level labels, indicating which object categories are present in the image. Then, the problem is addressed through Multiple Instance Learning (MIL). The standard pipeline of WSOD consists of two phases (Bilen et al., 2014; Song et al., 2014; Cinbis et al., 2016). Firstly, a set of candidate bounding boxes (instances) likely containing objects are generated by object proposal methods (e.g., selective search (Uijlings et al., 2013) or Edge Boxes (Zitnick and Dollár, 2014)). Secondly, the proposal classification phase alternates between two steps: (i) positive proposals with high confident scores are selected as pseudo instance-level labels and (ii) the OD (instances classifier) is trained under MIL framework. In this pipeline, (Bilen and Vedaldi, 2016) propose an end-to-end weakly supervised deep detection network. The main idea is to estimate the image classification score as a linear combination of weighted proposal

scores. Since then, a number of works have been proposed to improve the pipeline of Bilen and Vedaldi (2016), such as, multitasking with segmentation (Gao et al., 2019), leveraging spatial information (Kantorov et al., 2016), click supervision (Papadopoulos et al., 2017), and improving the proposals quality (Cheng et al., 2020). Recently, Inoue et al. (2018), Hou et al. (2021), Ouyang et al. (2021), Xu et al. (2022) apply WSOD to the domain adaptation, where the performance of the OD in the target domain can be improved by using source data with full annotations.

Apart from WSOD, ALOD is another direction that aims to efficiently leverage human supervision in adapting ODs by exploiting model uncertainty to select the most informative images for human annotation. This helps to minimise the total annotation cost. Specifically, Roy et al. (2018) consider the classification outputs of the detector to find images which the OD is uncertain most. Accordingly, Roy et al. (2018) propose black-box and white-box methods, where the black-box method can be applied to a variety of network architectures while the white-box method is tailored for Single Shot Multibox Detector (Liu et al., 2016). Concurrently, Brust et al. (2018) propose to aggregate confidence scores of bounding boxes to represent the model uncertainty. Then, Yuan et al. (2021) propose to re-weight bounding boxes to eliminate noisy ones before aggregation. Since previous works primarily focus on classification branch of ODs, Choi et al. (2021) employ mixture density networks to estimate the model uncertainty from both regression and classification branches.

In spite of a promising approach, ALOD still requires human effort in annotating data. To address this issue, UDAOD has become an increasingly popular approach in the community. Specifically, Arruda et al. (2019) use CycleGAN (Zhu et al., 2017) to translate image appearances from daytime to night to train an OD. However, CycleGAN can cause object of interests to be distorted, unrealistic, or disappeared, leading to noisy data for ODs. Therefore, detection-guided CycleGAN (Gao et al., 2021) adds a detection branch to guide the image generation process to ensure the quality of objects of interest in translated images. Another strategy is adversarial feature learning. In particular, domain adaptive Faster-RCNN (Chen et al., 2018) minimises the domain discrepancy through aligning image and instance distributions. This is achieved via training domain classifiers in an adversarial manner. The idea of Chen et al. (2018) is then extended to a scale-aware method by Chen et al. (2021), which aligns image distributions in different image scales through employing feature pyramid network (Lin et al., 2017a). A similar idea is also investigated in RetinaNet architecture (Pasqualino et al., 2021). Apart from adversarial feature learning, self-training has recently gained popularity in the community. Specifically, RoyChowdhury et al. (2019) combine detection and tracking to generate pseudo-labels for target datasets. However, since pseudo-labels may be very noisy, RoyChowdhury et al. (2019) propose a label smoothing technique to mitigate this negative effect. Moreover, to more efficiently address the noise present in the pseudo-labels, Khodabandeh et al. (2019) use an additional image classifier to improve the quality of pseudo-labels. Recently, Li et al. (2022) employ the teacher-student model combined with self-training and adversarial feature learning to achieve a state-of-the-art result.

As alluded, continually adapting ODs is indeed an active research topic. However, existing studies primarily focus on developing efficient methods for adaptation and overlook the crucial question of when to adapt. In realistic scenarios, environmental conditions usually change in a cyclical manner. Hence, always adapting ODs is not only computationally expensive but also unlikely to produce significant accuracy improvement.

3. Object detection architecture

RetinaNet (Lin et al., 2017b) is adopted as our OD, see Fig. 1(a). Specifically, the ResNet backbone (He et al., 2016) has three residual blocks C_3 , C_4 , and C_5 , whose outputs are the inputs of blocks P_3 , P_4 , and P_5 of the feature pyramid network (FPN) (Lin et al., 2017a). Then, class and box subnets B_1 , B_2 , and B_3 receive the outputs of FPN to predict the bounding boxes and object categories.

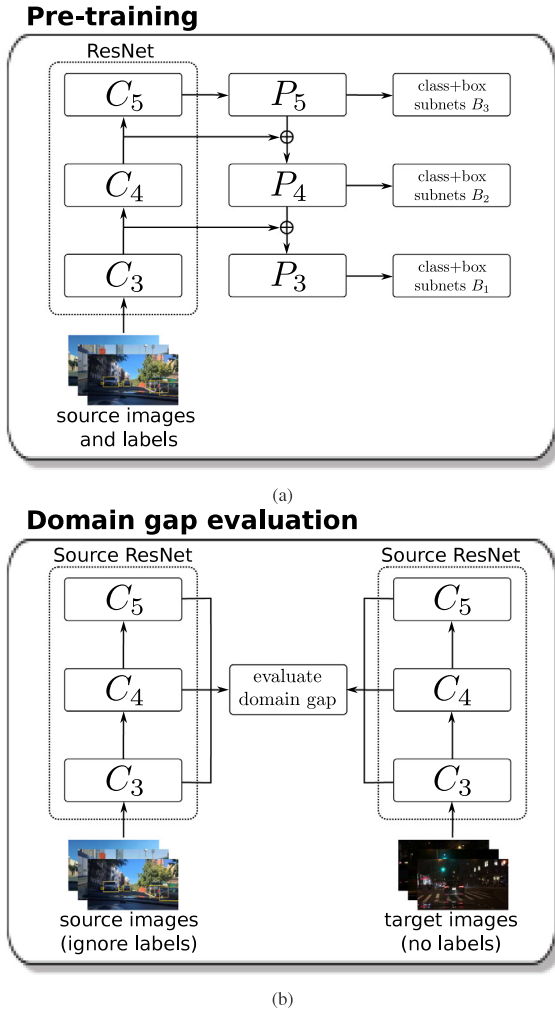


Fig. 1. An overview of our approach. (a) the ResNet is pre-trained on source images with full supervision. (b) The outputs from different blocks C_3 , C_4 , and C_5 are used to evaluate the domain gap between the source and target images. Note that labels are not needed in this step.

4. Domain gap evaluation

Let the source images be $D_s = \{x_i^s\}_{i=1}^{N_s}$ and the target images be $D_t = \{x_j^t\}_{j=1}^{N_t}$. Firstly, we will train ResNet using source images D_s , see Fig. 1(a). Then, to evaluate the domain gap between source domain D_s and target domain D_t , we only use the ResNet backbone, see Fig. 1(b).

For an input image x , we vectorise the outputs of C_3 , C_4 , and C_5 and denote them as $C_k(x)$ where $k = 3, 4$, or 5 . In the following sections, we will present three methods for domain gap evaluation: maximum mean discrepancy (Gretton et al., 2012), sliced Wasserstein distance (Villani, 2009), and distance of second-order statistics (Sun et al., 2016), and show how to apply these methods to the pipeline in Fig. 1(b).

4.1. Maximum mean discrepancy (MMD)

MMD (Gretton et al., 2012) is a non-parametric metric for estimating the distributional gap between two point sets. It has a variety of applications, e.g., domain adaptation (Yan et al., 2019), deep generative models (Dziugaite et al., 2015), anomaly detection (Zhang et al., 2021), etc. Mathematically, MMD is formulated as follows

$$\text{MMD}(D_s, D_t) = \left\| \frac{1}{N_s} \sum_{i=1}^{N_s} \phi(x_i^s) - \frac{1}{N_t} \sum_{j=1}^{N_t} \phi(x_j^t) \right\|^2, \quad (1)$$

where, ϕ is a feature map. To apply Eq. (1) to the pipeline in Fig. 1(b), we consider $\phi(\cdot)$ as $C_k(\cdot)$.

$$\text{MMD}_k(D_s, D_t) = \left\| \frac{1}{N_s} \sum_{i=1}^{N_s} C_k(x_i^s) - \frac{1}{N_t} \sum_{j=1}^{N_t} C_k(x_j^t) \right\|^2, \quad (2)$$

where $k = 3, 4$, or 5 .

To obtain a single value to represent MMD between D_s and D_t , we compute the mean of MMD_k

$$\text{MMD}(D_s, D_t) = \frac{1}{3} \sum_{k=3}^5 \text{MMD}_k(D_s, D_t) \quad (3)$$

4.2. Sliced Wasserstein distance (SWD)

Wasserstein distance (WD) is a popular metric in optimal transport theory (Villani, 2009) and have been receiving an increasing attention from the community in designing loss functions for deep generative models (Arjovsky et al., 2017; Adler and Lunz, 2018; Liu et al., 2019) and domain adaptation (Flamary et al., 2016; Damodaran et al., 2018; Xu et al., 2019). Compared to other measures (e.g., Jensen–Shannon divergence, Kullback–Leibler divergence, and total variation distance), WD takes into account the underlying geometry of probability space. SWD is a variance of WD, which aims to deal with high-dimensional data. Particularly, let empirical distributions of D_s and D_t as

$$\mu_s = \sum_{i=1}^{N_s} p_i^s \delta_{x_i^s}, \quad \mu_t = \sum_{j=1}^{N_t} p_j^t \delta_{x_j^t}, \quad (4)$$

where, $\delta_{x_i^s}$ and $\delta_{x_j^t}$ are the Dirac functions at location x_i^s and x_j^t . p_i^s and p_j^t are the probability masses associated to the i th and j th samples. In practice, one typically sets $p_i^s = \frac{1}{N_s}$ and $p_j^t = \frac{1}{N_t}$.

Next, define the set

$$B = \{ \gamma \in (\mathbb{R}^+)^{N_s \times N_t} \mid \gamma \mathbf{1}_{N_t} = \mu_s, \gamma^T \mathbf{1}_{N_s} = \mu_t \}, \quad (5)$$

where, $\mathbf{1}$ is a vector with all elements equal to 1.

Then, for each $k = 3, 4$, or 5 , we have the WD_k as follows

$$\text{WD}_k(D_s, D_t) = \min_{\gamma \in B} \langle \gamma, \mathbf{D}_k \rangle_F \quad (6)$$

where, $\langle \cdot, \cdot \rangle_F$ is the Frobenius dot product and \mathbf{D}_k is the cost matrix with $\mathbf{D}_k(i, j) = \|C_k(x_i^s) - C_k(x_j^t)\|_2^2$.

However, computing \mathbf{D}_k is expensive as the dimensionality d_k of $C_k(x)$ is very large. Hence, the main idea of SWD is to project $C_k(x)$ to 1-D first, then WD formulation will be applied to the projected 1-D data.

Define a set $\{R_{k,m}\}_{m=1}^M$, where $R_{k,m}$ is the m th one-dimensional linear projection sampled from the uniform on the unit sphere of dimension $d_k - 1$. Then, for each k , SWD is formalised as

$$\text{SWD}_k(D_s, D_t) = \frac{1}{M} \sum_{m=1}^M \min_{\gamma \in B} \langle \gamma, \hat{\mathbf{D}}_k \rangle_F \quad (7)$$

where, $\hat{\mathbf{D}}_k(i, j) = \left\| (R_{k,m})^T C_k(x_i^s) - (R_{k,m})^T C_k(x_j^t) \right\|_2^2$.

Finally, to obtain a scalar to represent the domain gap between D_s and D_t , we compute

$$\text{SWD}(D_s, D_t) = \frac{1}{3} \sum_{k=3}^5 \text{SWD}_k(D_s, D_t) \quad (8)$$

4.3. Distance of second-order statistics (DSS)

DSS is the distance of covariances of source and target features. It is a popular metric in domain adaptation (Sun et al., 2016; Sun and Saenko, 2016; Wang et al., 2017). Specifically, suppose we have

$$\mathbf{C}_k^s = [C_k(x_1^s), \dots, C_k(x_{N_s}^s)]^T, \quad \mathbf{C}_k^t = [C_k(x_1^t), \dots, C_k(x_{N_t}^t)]^T,$$

with $k = 3, 4$, or 5 , i.e., $\mathbf{C}_k^s(i, j)$ (or $\mathbf{C}_k^t(i, j)$) indicates j th feature dimension of source (or target) sample i th. Then, we have the covariance matrices of source and target domains

$$\mathbf{K}_k^s = \frac{1}{N_s - 1} \left((\mathbf{C}_k^s)^T \cdot \mathbf{C}_k^s - \frac{1}{N_s} (\mathbf{1}^T \mathbf{C}_k^s)^T \cdot (\mathbf{1}^T \mathbf{C}_k^s) \right) \quad (9)$$

$$\mathbf{K}_k^t = \frac{1}{N_t - 1} \left((\mathbf{C}_k^t)^T \cdot \mathbf{C}_k^t - \frac{1}{N_t} (\mathbf{1}^T \mathbf{C}_k^t)^T \cdot (\mathbf{1}^T \mathbf{C}_k^t) \right) \quad (10)$$

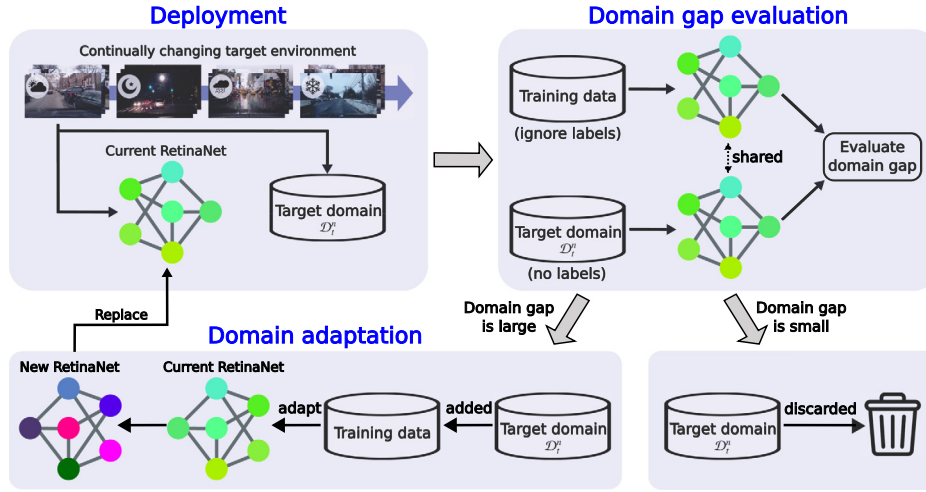


Fig. 2. The application of domain gap evaluation in continually adapting RetinaNet. Specifically, for each target domain D_t^n , the current RetinaNet is used to evaluate the domain gap between the current training data and D_t^n (without the need of labels, as outlined in Section 4). If the domain gap is found to be small, D_t^n is discarded, and no adaptation is needed. However, if the domain gap is found to be large, D_t^n is added to the current training database to refine the current RetinaNet, resulting in a new RetinaNet model which replaces the current RetinaNet in the operating environment.

where, $\mathbf{1}$ is a vector with all elements equal to 1. Finally, we compute the DSS for each k

$$\text{DSS}_k = \frac{1}{4d_k^2} \|\mathbf{K}_k^s - \mathbf{K}_k^t\|_F^2, \quad (11)$$

where, d_k is the dimension of $C_k(x)$.

However, as d_k in the ResNet architecture is very large (i.e., millions of dimensions), computing \mathbf{K}_k^s and \mathbf{K}_k^t is intractable.

To address this issue, we adopt the idea of SWD (see Section 4.2) to project $C_k(x)$ to 1-D first, and then DSS formulation will be applied to the projected 1-D data. Specifically, given a set of one-dimensional linear projections $\{R_{k,m}\}_{m=1}^M$ randomly sampled from the uniform on the unit sphere of dimension $d_k - 1$, for each projection m , let

$$\hat{C}_{k,m}(x) = R_{k,m}^T C_k(x). \quad (12)$$

Next, we form

$$\hat{C}_{k,m}^s = [\hat{C}_{k,m}(x_1^s), \dots, \hat{C}_{k,m}(x_{N_s}^s)]^T, \quad (13)$$

$$\hat{C}_{k,m}^t = [\hat{C}_{k,m}(x_1^t), \dots, \hat{C}_{k,m}(x_{N_t}^t)]^T \quad (14)$$

Covariance matrices of projected features are then computed as

$$\hat{\mathbf{K}}_{k,m}^s = \frac{1}{N_s - 1} \left((\hat{C}_{k,m}^s)^T \hat{C}_{k,m}^s - \frac{1}{N_s} (\mathbf{1}^T \hat{C}_{k,m}^s)^T (\mathbf{1}^T \hat{C}_{k,m}^s) \right) \quad (15)$$

$$\hat{\mathbf{K}}_{k,m}^t = \frac{1}{N_t - 1} \left((\hat{C}_{k,m}^t)^T \hat{C}_{k,m}^t - \frac{1}{N_t} (\mathbf{1}^T \hat{C}_{k,m}^t)^T (\mathbf{1}^T \hat{C}_{k,m}^t) \right) \quad (16)$$

As $\hat{C}_{k,m}(x)$ is 1-D projected feature of $C_k(x)$, computing $\hat{\mathbf{K}}_{k,m}^s$ and $\hat{\mathbf{K}}_{k,m}^t$ is fast.

Finally, the DSS for each k is computed as follows

$$\hat{\text{DSS}}_k = \frac{1}{4M} \sum_{m=1}^M \|\hat{\mathbf{K}}_{k,m}^s - \hat{\mathbf{K}}_{k,m}^t\|_F^2 \quad (17)$$

As we need a scalar to represent the domain gap between D_s and D_t , we calculate the mean of $\hat{\text{DSS}}_k$

$$\text{DSS}(D_s, D_t) = \frac{1}{3} \sum_{k=3}^5 \hat{\text{DSS}}_k \quad (18)$$

5. Application of domain gap evaluation in continual domain adaptation

Continual domain adaptation is defined as follows: Given a fixed task, RetinaNet is required to adapt to a sequence of N target domains

$\{D_t^n\}_{n=1}^N$. Fig. 2 illustrates our pipeline for applying domain gap evaluation in the continual domain adaptation of RetinaNet. Specifically, given a new target domain D_t^n , we will determine if domain adaptation is necessary. This is performed by measuring the domain gap between the current training data and the target domain D_t^n using the current RetinaNet, as described in Section 4. Note that labels are not required for domain gap evaluation.

If the domain gap is found to be smaller than a predefined threshold, the target domain D_t^n is discarded as it is deemed unnecessary. This is because if the domain gap between D_t^n and the training data is small, the two datasets likely share similar distributions. Therefore, adapting the RetinaNet with this D_t^n will not significantly improve its performance and will consume unnecessary resources.

However, if the domain gap is found to be larger than the threshold, D_t^n is added to the training database and used to adapt the current RetinaNet. There are a wide range of methods for adapting object detection models in the literature, as outlined in Section 2. Once adapted, the new RetinaNet model replaces the current model in operation.

6. Experiments

6.1. Datasets

As the primary focus of this paper is on continually adapting ODs to changing environmental conditions, we require object detection datasets that exhibit a variety of conditions. For this reason, we decide to use the DGTA (Kiefer et al., 2022), BDD (Yu et al., 2020), and KITTI (Geiger et al., 2012) datasets in our experiments.

The DGTA dataset (Kiefer et al., 2022) is a synthetic dataset captured from the video game Grand Theft Auto V (Rockstar Games, 2013) and contains five categories, but only the ‘‘boat’’ category is used in this experiment. The source data is made up of images captured during clear conditions between 9am and 3pm (denoted as **clear-9 h-15 h**). To create the target data, the overcast condition is chosen, and images are divided into 12 chunks, each covering a 2-hour time range (0 h–1 h, 2 h–3 h, 4 h–5 h, ..., 22 h–23 h). Each chunk forms a different target dataset, resulting in 12 target datasets in total. These 12 target datasets are denoted as **0 h–1 h**, **2 h–3 h**, **2 h–3 h**, ..., **22 h–23 h**. Table 1 shows the statistics of the obtained datasets, while Fig. 3 illustrates some samples of the source and target datasets.

BDD (Yu et al., 2020) is a real dataset containing 100,000 driving videos under different weather conditions (clear, foggy, overcast,

Table 1

Statistics of DGTA, where “# imgs” and “# anns” are abbreviations of “number of images” and “number of annotations”.

Domain	Training		Testing	
	# imgs	# anns	# imgs	# anns
Clear 9 h–15 h	4,658	48,643	1,165	12,003
Overcast 0 h–1 h	1,826	17,247	457	3,904
Overcast 2 h–3 h	2,038	18,190	510	4,428
Overcast 4 h–5 h	1,595	16,212	399	4,115
Overcast 6 h–7 h	1,549	17,279	388	4,324
Overcast 8 h–9 h	1,545	16,037	387	3,801
Overcast 10 h–11 h	1,678	18,828	420	4,817
Overcast 12 h–13 h	1,436	14,776	359	3,571
Overcast 14 h–15 h	1,308	12,430	327	3,139
Overcast 16 h–17 h	1,368	12,987	343	3,392
Overcast 18 h–19 h	1,708	19,689	428	4,871
Overcast 20 h–21 h	1,639	17,490	410	4,275
Overcast 22 h–23 h	1,784	17,884	447	4,215

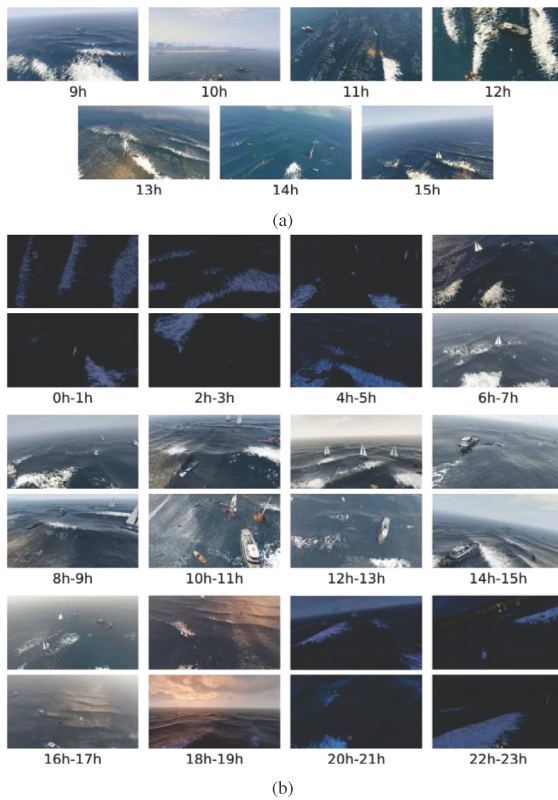


Fig. 3. Samples of DGTA dataset, where (a) source domain in clear-9 h-15 h and (b) target domains in overcast in 0 h–1 h, 2 h–3 h, 4 h–5 h, ..., 22 h–23 h.

cloudy, rainy, and snowy) and time of day (daytime and night). The source domain is selected from the **clear-daytime** condition, while the target domains are selected from **clear-night**, **cloudy-daytime**, **overcast-daytime**, **rainy-daytime**, **rainy-night**, **snowy-daytime**, and **snowy-night** conditions. Remaining conditions are not considered due to the small number of images. Then, two following settings are created for BDD dataset.

- Single-class BDD (denoted as **sBDD**): only the “car” category is used.
- Multi-class BDD (denoted as **mBDD**): four categories (pedestrian, car, traffic light, traffic sign) are used.

Tables 2 and 3 respectively present the statistics of sBDD and mBDD, while Fig. 5 illustrates their samples.



Fig. 4. Samples of KITTI benchmark, where source domain is KITTI images, and target domains are Fog-750 m/375 m/150 m/75 m/30 m.

Table 2

Statistics of sBDD, where “# imgs” and “# anns” are abbreviations of “number of images” and “number of annotations”.

Domain	Training		Testing	
	# imgs	# anns	# imgs	# anns
Clear-daytime	12,347	138,754	1,750	19,745
Clear-night	22,428	198,137	3,222	30,350
Cloudy-daytime	4,242	48,649	632	7,432
Overcast-daytime	7,501	89,463	1,032	12,453
Rainy-daytime	2,497	24,851	393	3,958
Rainy-night	2,181	18,002	281	2,372
Snowy-daytime	2,835	28,040	418	4,477
Snowy-night	2,214	19,092	266	2,247

Table 3

Statistics of mBDD, where “# imgs” and “# anns” are abbreviations of “number of images” and “number of annotations”.

Domain	Training		Testing	
	# imgs	# anns	# imgs	# anns
Clear-daytime	12,450	225,523	1,764	32,331
Clear-night	22,871	354,717	3,274	52,910
Cloudy-daytime	4,261	80,261	638	12,264
Overcast-daytime	7,550	152,106	1,039	20,840
Rainy-daytime	2,521	44,954	396	7,007
Rainy-night	2,208	33,384	286	4,564
Snowy-daytime	2,861	50,992	422	7,726
Snowy-night	2,248	35,588	273	4,383



Fig. 5. Samples of BDD dataset, where source domain is clear-daytime, and target domains are clear-night, cloudy-daytime, overcast-daytime, rainy-daytime, rainy-night, snowy-daytime, and snowy-night.

KITTI (Geiger et al., 2012) is a widely used dataset in autonomous driving, which contains seven object categories, but only the “car” category is used in this experiment. The source domain is composed of KITTI images, while the target domains are generated using the technique of Halder et al. (2019), resulting in five different levels of fog: **fog-750 m**, **fog-375 m**, **fog-150 m**, **fog-75 m**, and **fod-30 m**; see Fig. 4 for sample images. The training and testing sets are randomly split, leading to 3740 training images with 14,655 annotations and 3741 testing images with 14,087 annotations.

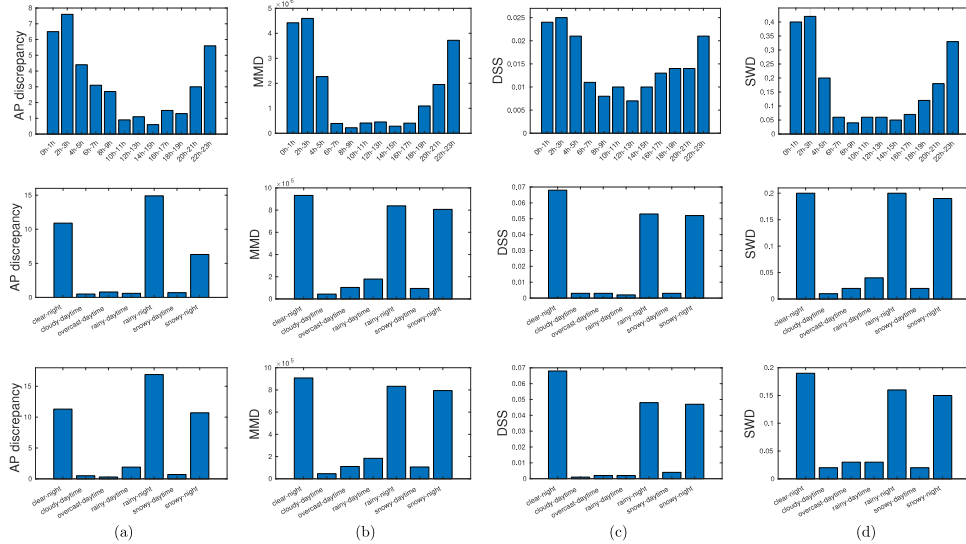


Fig. 6. Histograms of AP discrepancy and domain gap in DGTA (1st row), sBDD (2nd row) and mBDD (3rd row) datasets, where (a) AP discrepancy, (b)–(d) domain gap between source and target domains evaluated by MMD, DSS and SWD. The source model is trained on source data **clear-9 h-15 h** (for DGTA) and **clear-daytime** (for sBDD and mBDD), and target domains are specified in *x*-axis values.

Table 4
Kullback–Leibler divergence between domain gap and AP discrepancy.

	DGTA	sBDD	mBDD
MMD	0.15	0.08	0.05
DSS	0.08	0.06	0.08
SWD	0.09	0.07	0.08

6.2. Experimental setup

RetinaNet with a ResNet-50 backbone is trained on the source domain (i.e., **clear-9 h-15 h** for DGTA and **clear-daytime** for sBDD and mBDD) to obtain a source OD. Specifically, the OD is initialised from the OD pretrained on COCO (Lin et al., 2014). We train the OD with a learning rate of 0.0002 for 30 k iterations, then reduce the learning rate to 0.00002 for another 30 k iterations. Unless stated otherwise, the batch size of 4, momentum of 0.9, weight decay of 0.001, and $M = 10$ are used in the experiment.

The performance of the detection is evaluated based on the average precision and average recall metrics, both measured with a threshold of 50%. These metrics are abbreviated as AP50 and AR50, respectively. Each experiment is conducted in an NVIDIA GeForce RTX 2080 Ti and implemented using Detectron2 (Wu et al., 2019).

6.3. Correlation between domain gap and detection accuracy

6.3.1. Setup

Inspired by Ben-David et al. (2010, Theorem 2) stating that the discrepancy of source and target errors are upper-bounded by the domain gap, we firstly define the discrepancy of AP

$$AP \text{ discrepancy} = |AP_{\text{source}} - AP_{\text{target}}| \quad (19)$$

where, AP_{source} and AP_{target} are APs of source OD on source and target testing sets. Then, the domain gap will be evaluated by the three metrics presented in Section 4.

6.3.2. Results

The correlation between the AP discrepancy and domain gap is shown in Fig. 6. Specifically, when the gap between the source and target domains is large, the AP discrepancy will also be large, meaning that the OD trained on the source domain will perform poorly in the

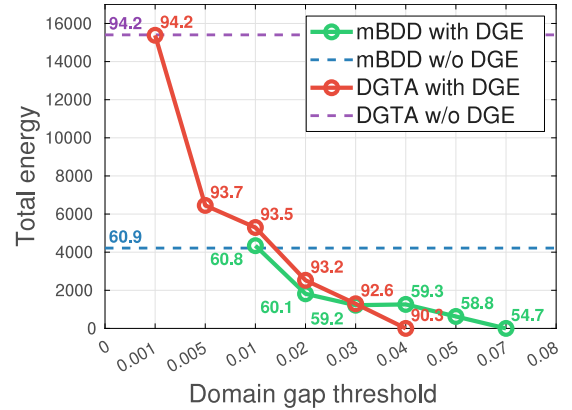


Fig. 7. Correlation between total energy consumption for continual domain adaptation, the domain gap threshold, and AP50, where AP50 is reported adjacent to the markers ○ and ○.

target domain. Conversely, when the gap is small, the source OD will deliver an accuracy on the target domain that is similar to its accuracy on the source domain

To quantitatively compare the performance of MMD, DSS, and SWD, we compute the Kullback–Leibler divergence between the histograms of domain gap and AP discrepancy, as presented in Table 4. The results show that DSS and SWD perform consistently well across the DGTA, sBDD, and mBDD datasets, while MMD performs slightly worse on the DGTA dataset. This is likely because MMD only considers first-order statistics, while DSS considers second-order statistics and SWD takes into account data geometry.

In the following section, we will use DSS to evaluate the benefit of domain gap evaluation in continual domain adaptation of object detectors.¹

6.4. Continual domain adaptation of RetinaNet

6.4.1. Setup

To simulate a continual change of environmental conditions, we consider the following configuration

¹ MMD and SWD are expected to deliver a similar outcome as DSS

Table 5

Comparison between AP50 (top) and AR50 (bottom) in continual domain adaptation on DGTA, where target domain is sequentially changed **0 h–1 h** → **1 h–2 h** → ... → **22 h–23 h**.

Time	t →											Mean	
	0 h–1 h	2 h–3 h	4 h–5 h	6 h–7 h	8 h–9 h	10 h–11 h	12 h–13 h	14 h–15 h	16 h–17 h	18 h–19 h	20 h–21 h		22 h–23 h
No DA	86.1±0.3 91.6±0.4	85.0±0.1 90.5±0.2	88.6±0.1 92.5±0.1	90.2±0.2 93.1±0.2	90.8±0.1 93.6±0.1	94.1±0.1 96.3±0.1	94.5±0.1 96.7±0.1	92.9±0.2 95.0±0.1	91.7±0.2 94.0±0.2	92.1±0.1 94.7±0.1	90.2±0.1 93.6±0.2	87.1±0.2 92.0±0.2	90.3±0.1 93.6±0.1
SDA w/o DGE	94.2±0.1 96.4±0.2	94.5±0.1 96.5±0.1	93.5±0.2 95.9±0.2	93.1±0.2 95.5±0.2	93.3±0.1 95.5±0.1	95.2±0.2 97.0±0.1	95.5±0.2 97.0±0.1	94.7±0.3 96.5±0.3	94.2±0.2 96.1±0.3	94.2±0.3 96.3±0.2	93.5±0.4 95.6±0.3	94.9±0.1 96.6±0.1	94.2±0.1 96.3±0.1
SDA with DGE	94.3±0.1 96.5±0.2	94.6±0.1 96.7±0.1	92.6±0.3 95.1±0.3	90.8±0.2 93.7±0.2	91.3±0.2 94.1±0.1	94.5±0.1 96.5±0.2	94.9±0.1 96.7±0.1	93.3±0.1 95.3±0.2	92.1±0.3 94.3±0.3	92.4±0.2 95.0±0.2	92.7±0.3 95.1±0.2	94.2±0.3 96.2±0.2	93.1±0.2 95.4±0.2
UDA w/o DGE	89.1±0.8 92.9±0.8	89.5±0.5 93.1±0.5	89.8±0.7 92.9±0.6	90.9±0.2 93.9±0.3	91.3±0.8 94.1±0.8	94.3±0.2 96.3±0.2	94.9±0.1 96.8±0.1	93.5±0.1 95.7±0.2	92.5±0.5 94.6±0.4	92.1±0.5 94.8±0.5	90.5±0.4 93.6±0.5	89.0±0.4 92.5±0.3	91.5±0.1 94.3±0.1
UDA with DGE	89.4±0.8 93.2±0.4	89.2±0.8 93.0±0.5	89.8±0.4 93.0±0.3	91.0±0.4 94.1±0.2	91.2±0.3 94.2±0.3	94.1±0.3 96.3±0.3	94.6±0.1 96.6±0.1	93.2±0.3 95.3±0.2	92.3±0.2 94.5±0.2	92.3±0.2 94.9±0.1	92.1±0.2 94.0±0.3	90.0±0.5 93.3±0.4	91.5±0.2 94.4±0.2

Table 6

Comparison between AP50 (top) and AR50 (bottom) in continual domain adaptation on KITTI, where target domain is sequentially changed **fog-750 m** → **fog-375 m** → **fog-150 m** → **fog-75 m** → **fog-30 m**.

Time	t →					Mean
	fog-750 m	fog-375 m	fog-150 m	fog-75 m	fog-30 m	
No DA	89.1±0.1 94.7±0.1	86.6±0.1 92.5±0.1	79.3±0.1 85.5±0.2	66.6±0.5 72.6±0.6	35.0±0.6 37.3±0.6	71.3±0.2 76.5±0.2
SDA w/o DGE	93.0±0.1 97.3±0.1	92.9±0.1 96.9±0.1	92.5±0.1 96.7±0.1	91.6±0.1 96.3±0.3	85.6±0.2 95.5±0.2	91.1±0.1 96.5±0.2
SDA with DGE	89.1±0.1 94.7±0.1	86.6±0.1 92.5±0.1	92.0±0.1 97.2±0.1	91.4±0.1 96.6±0.2	86.3±0.1 96.1±0.4	89.1±0.1 95.4±0.2
UDA w/o DGE	89.0±0.4 93.8±0.4	85.8±0.6 90.0±0.6	79.5±0.8 82.9±1.1	71.3±1.0 74.2±1.2	43.6±1.0 45.6±1.1	73.9±0.8 77.3±0.9
UDA with DGE	89.1±0.1 94.7±0.1	86.6±0.1 92.5±0.1	80.4±1.1 84.9±1.5	69.8±0.9 73.1±0.7	42.1±0.8 43.7±0.9	73.6±0.6 77.8±0.7

- For DGTA, the source domain is still **clear-9 h-15 h**, but the target domain in overcast will be sequentially changed **0 h–1 h** → **1 h–2 h** → **2 h–3 h** → ... → **22 h–23 h**.
- For sBDD and mBDD, the source domain is **clear-daytime**, but the target domain will be continually changed **clear-night** → **cloudy-daytime** → **overcast-daytime** → **rainy-daytime** → **rainy-night** → **snowy-daytime** → **snowy-night**.

If the OD encounters a new condition, we will adapt the OD with the following domain adaptation methods

- **No domain adaptation (NoDA)**: The model is only trained on source domain and will never be adapted.
- **Supervised domain adaptation (SDA)**: Annotations are assumed to be fully available, thus the OD can be fine-tuned in a supervised manner. This provides an upper-bound performance for domain adaptation.
- **Unsupervised domain adaptation (UDA)**: Labels are only available in the source data, and the new data is added to the unlabelled target set. Specifically, we use adversarial RetinaNet (Pasqualino et al., 2021) for the adaptation.

In all adaptation methods, the current model is used as the initialisation, then the model is trained for 20k iterations with a fixed learning rate 0.0002. Also, we consider following settings

- **Without domain gap evaluation (w/o DGE)**: The model is always adapted in every condition.
- **With domain gap evaluation (with DGE)**: Using DSS, we will verify if the domain adaptation is necessary for the new data.

To quantitatively evaluate the adaptation cost, we measure the total GPU energy consumption for domain adaptation. The GPU energy is measured using the pyJoules (Spirals, 2019).

6.4.2. Results

Effects of different domain gap thresholds The experiment is conducted on DGTA and mBDD. SDA is used for domain adaptation. AP50 is reported as the detection accuracy.

The result is shown in Fig. 7. When the threshold is set at a relatively low value (i.e., 0.001 for DGTA and 0.01 for mBDD), domain adaptation is carried out under all conditions. This results in a total energy usage for continual domain adaptation and AP50 that are comparable to the results of w/o DGE.

As the thresholds increase (i.e., 0.001 to 0.005 for DGTA and 0.01 to 0.02 for mBDD), there is a slight drop in AP50 (i.e., approximately 0.5% for DGTA and 0.7% for mBDD). However, with DGE can significantly reduce the total energy usage required for continual domain adaptation, by approximately 58% for both DGTA and mBDD.

When the threshold is set at a substantially higher value (i.e., 0.04 for DGTA and 0.07 for mBDD), domain adaptation is not applied, resulting in zero total energy usage for continual domain adaptation. However, this also causes a considerable decrease in AP50, with approximately 4% for DGTA and 6% for mBDD, compared to the results of w/o DGE.

Comparison between with DGE and w/o DGE We repeat each experiment 5 times, then report mean and standard deviation of AP50 and AR50 over 5 runs. The domain gap threshold is fixed to 0.02 for all datasets.

It is unsurprising that SDA and UDA consistently outperform NoDA by a significant margin in datasets DGTA, KITTI, sBDD, and mBDD (see Tables 5, 6, 7 and 8).

According to the results presented in Table 5 of the DGTA dataset, using SDA as the domain adaptation method, SDA w/o DGE yields better overall performance than SDA with DGE, with an improvement of 1.1% in AP50's mean and 0.9% in AR50's mean. However, it is worth noting that the total energy consumption of SDA w/o DGE is approximately 81.6% higher than that of SDA with DGE, as shown in Fig. 8. Similarly, if UDA is used as the domain adaptation method, Table 5 shows that the performance of UDA w/o DGE is comparable to

Table 7

Comparison between AP50 (top) and AR50 (bottom) in continual domain adaptation on sBDD, where target domain is sequentially changed **clear-night** → **cloudy-daytime** → **overcast-daytime** → **rainy-daytime** → **rainy-night** → **snowy-daytime** → **snowy-night**.

Time	t →							Mean
	clear-night	cloudy-daytime	overcast-daytime	rainy-daytime	rainy-night	snowy-daytime	snowy-night	
No DA	66.1±0.1	77.3±0.1	77.6±0.1	77.4±0.1	62.0±0.1	77.5±0.2	70.5±0.2	72.6±0.1
	88.0±0.1	89.8±0.1	90.2±0.1	91.2±0.1	85.2±0.2	90.5±0.2	90.3±0.1	89.3±0.1
SDA w/o DGE	74.2±0.1	77.5±0.2	78.3±0.1	79.3±0.2	74.3±0.3	78.9±0.2	78.1±0.5	77.2±0.1
	91.5±0.1	90.0±0.1	90.3±0.2	92.2±0.2	91.1±0.5	91.4±0.1	93.4±0.4	91.4±0.1
SDA with DGE	74.2±0.3	77.6±0.1	77.7±0.2	77.9±0.5	71.7±0.5	78.3±0.2	77.1±0.1	76.4±0.2
	91.6±0.3	90.0±0.2	90.0±0.2	91.3±0.3	89.9±0.7	91.0±0.3	93.2±0.4	91.0±0.1
UDA w/o DGE	66.5±0.3	77.8±0.3	77.8±0.3	78.8±0.4	62.3±0.4	78.4±0.3	70.8±0.4	73.2±0.1
	88.4±0.4	90.2±0.2	90.3±0.3	91.5±0.3	85.8±0.4	90.8±0.3	90.8±0.3	89.7±0.1
UDA with DGE	66.5±0.2	77.7±0.1	77.9±0.1	77.7±0.2	62.0±0.3	77.5±0.3	70.6±0.3	72.8±0.1
	88.3±0.3	90.2±0.1	90.3±0.1	91.4±0.1	85.4±0.3	90.8±0.3	90.6±0.3	89.6±0.1

Table 8

Comparison between AP50 (top) and AR50 (bottom) in continual domain adaptation on mBDD, where target domain is sequentially changed **clear-night** → **cloudy-daytime** → **overcast-daytime** → **rainy-daytime** → **rainy-night** → **snowy-daytime** → **snowy-night**.

Time	t →							Mean
	clear-night	cloudy-daytime	overcast-daytime	rainy-daytime	rainy-night	snowy-daytime	snowy-night	
No DA	49.3±0.1	60.0±0.2	60.1±0.1	58.7±0.1	43.7±0.2	60.9±0.1	49.7±0.3	54.6±0.1
	73.4±0.1	76.8±0.1	77.3±0.1	78.4±0.1	69.1±0.1	78.2±0.1	73.9±0.3	75.3±0.1
SDA w/o DGE	58.3±0.2	60.6±0.4	61.8±0.2	61.7±0.2	57.7±0.2	63.6±0.2	62.5±0.2	60.9±0.1
	79.1±0.2	76.8±0.5	78.4±0.4	80.1±0.5	79.6±0.5	80.1±0.3	82.3±0.2	79.5±0.1
SDA with DGE	58.3±0.4	60.4±0.5	61.1±0.1	60.5±0.3	55.4±0.1	63.0±0.1	61.0±0.5	60.0±0.1
	79.3±0.3	76.9±0.4	77.7±0.2	79.7±0.4	78.4±0.2	79.8±0.3	81.4±0.5	79.0±0.2
UDA w/o DGE	49.9±0.3	61.5±0.1	61.4±0.1	59.6±0.4	43.9±0.3	61.1±0.2	49.7±0.4	55.3±0.1
	73.9±0.3	78.1±0.1	78.5±0.1	79.1±0.4	69.2±0.3	79.6±0.2	73.5±0.5	75.9±0.1
UDA with DGE	49.8±0.6	60.7±0.3	60.8±0.2	59.3±0.4	44.4±0.5	62.0±0.2	50.9±0.7	55.4±0.4
	73.8±0.6	77.6±0.2	78.1±0.2	79.1±0.4	69.9±0.8	79.1±0.3	74.8±0.3	76.1±0.3

that of UDA with DGE in terms of AP50 and AR50. However, Fig. 8 reveals that the energy consumption of UDA with DGE is 91.6% lower than that of UDA w/o DGE.

Regarding the KITTI dataset, Table 6 reveals that SDA w/o DGE outperforms SDA with DGE by about 2% in AP50’s mean and 1.1% in AR50’s mean. Nonetheless, it is important to note that, as reported in Fig. 8, the total energy consumption of SDA w/o DGE is significantly higher, approximately 45% more than SDA with DGE. Likewise, as indicated in Table 6, both UDA w/o DGE and UDA with DGE exhibit a similar level of accuracy. However, UDA with DGE achieves considerable energy savings, reducing total energy consumption by about 39% in comparison to UDA w/o DGE.

When SDA is used as the domain adaptation method in the sBDD dataset, the performance of SDA w/o DGE is better than that of SDA with DGE, with an average improvement of 0.8% in AP50 and 0.4% in AR50, as reported in the Table 7. However, it is important to note that the total energy consumption of SDA w/o DGE is 59.4% higher than that of SDA with DGE (see Fig. 8). Also, Table 7 demonstrates that the AP50 and AR50 of UDA with DGE are similar to those of UDA w/o DGE. However, the comparison of the total energy consumption in Fig. 8 reveals that UDA with DGE consumes significantly less energy than UDA w/o DGE.

The mBDD dataset analysis in Table 8 reveals that SDA w/o DGE marginally outperforms SDA with DGE by only 0.9% in AP50’s mean and 0.5% in AR50’s mean, which is statistically insignificant. However, as shown in Fig. 8, SDA with DGE offers significant energy savings, reducing total energy consumption by approximately 56.8% compared

to SDA w/o DGE. In addition, Table 8 shows that there is little difference in the performance of UDA with DGE and UDA w/o DGE, as measured by the AP50 and AR50 metrics. However, Fig. 8 highlights the substantial energy savings achieved by UDA with DGE, reducing total energy usage by 66.5% compared to UDA w/o DGE.

It is important to note that in the sBDD and mBDD datasets, the domain gap challenge presents a significant obstacle to achieving accuracy improvements through UDA. However, our paper’s main focus is on determining “when to adapt” rather than “how to adapt”. We believe that identifying the most appropriate times for adaptation can help overcome the domain gap challenge in future research.

7. Discussion and conclusion

7.1. Discussion

Selection of domain gap threshold Effective application of DGE in continual domain adaptation involves a manual process of threshold selection, which is primarily due to two underlying factors: (i) The range of values for the metrics MMD, DSS, and SWD differs and is dependent on the dataset. Specifically, as shown in Fig. 6, for dataset DGTA, $MMD \in [0, 5.10^6]$, $DSS \in [0, 0.025]$, and $SWD \in [0, 0.4]$. However, for dataset sBDD or mBDD, $MMD \in [0, 10^6]$, $DSS \in [0, 0.07]$, and $SWD \in [0, 0.2]$. (ii) The impact of distinct thresholds on varying datasets is inconsistent. This is demonstrated in Fig. 7, where the relationship between total energy consumption for continual domain adaptation, the

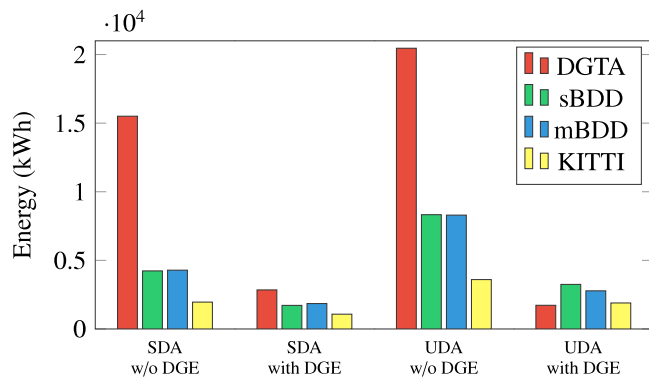


Fig. 8. Total energy usage for continual domain adaptation. Note that we only plot the mean of total energy usage as its standard deviation is insignificant.

threshold for domain gap, and detection accuracy varies between two datasets, DGTA and mBD.

Catastrophic forgetting Domain adaptation can lead to catastrophic forgetting, a phenomenon where the model’s parameters are fine-tuned to fit the target domain, but at the cost of losing its knowledge of the source domain, resulting in poor performance on the source domain [McCloskey and Cohen \(1989\)](#). To mitigate catastrophic forgetting, our paper examines the ideal scenario where all training data can be stored, although this is impractical in real-world applications. To address this issue, popular practical approaches such as replay-based, regularisation-based, and parameter isolation-based methods have been proposed (see [De Lange et al. \(2021\)](#) for a detailed survey). In light of these approaches, developing a DGE method that can be effectively combined with these methods is a promising direction for future research.

Test-time adaptation Test-time adaptation (TTA) aims to adapt the model during the testing phase, when source data is inaccessible due to privacy concerns ([Wang et al., 2021, 2022; Lim et al., 2023](#)). As our work assumes the presence of source data, making it nontrivial to apply to TTA, developing a “when to adapt” method for TTA represents a promising avenue for further investigation.

7.2. Conclusion

This paper examines the use of maximum mean discrepancy, distance of second-order statistics, and sliced Wasserstein distance to evaluate the domain gap in object detection. The findings from experiments on synthetic and real-world datasets indicate a correlation between domain gap and detection accuracy. The paper then applies domain gap to the continual domain adaptation of RetinaNet, resulting in a significant reduction in the overall cost of adaptation.

CRedit authorship contribution statement

Anh-Dzung Doan: Conceptualisation, Methodology, Software, Validation, Formal analysis, Investigation, Writing – original draft, Writing – review & editing. **Bach Long Nguyen:** Conceptualisation, Methodology, Writing – review & editing. **Surabhi Gupta:** Writing – review & editing. **Ian Reid:** Conceptualisation, Writing – review & editing, Funding acquisition. **Markus Wagner:** Conceptualisation, Writing – review & editing, Project administration, Funding acquisition. **Tat-Jun Chin:** Conceptualisation, Methodology, Writing – review & editing, Supervision, Funding acquisition.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Anh-Dzung Doan reports financial support was provided by Australian Research Council. Bach Long Nguyen reports financial support was provided by Australian Research Council. Surabhi Gupta reports financial support was provided by Australian Research Council. Ian Reid reports financial support was provided by Australian Research Council. Markus Wagner reports financial support was provided by Australian Research Council. Tat-Jun Chin reports financial support was provided by Australian Research Council.

Data availability

I have shared my code on GitHub. The data is public benchmark datasets.

Acknowledgements

This research was supported by the Australian Research Council grant LP200200881.

References

- Adler, J., Lunz, S., 2018. Banach wasserstein gan. *Adv. Neural Inf. Process. Syst.*
- Arjovsky, M., Chintala, S., Bottou, L., 2017. Wasserstein generative adversarial networks. In: *International Conference on Machine Learning*. PMLR, pp. 214–223.
- Arnold, E., Al-Jarrah, O.Y., Dianati, M., Fallah, S., Oxtoby, D., Mouzakitis, A., 2019. A survey on 3d object detection methods for autonomous driving applications. *IEEE Trans. Intell. Transp. Syst.*
- Arruda, V.F., Paixão, T.M., Berriel, R.F., De Souza, A.F., Badue, C., Sebe, N., Oliveira-Santos, T., 2019. Cross-domain car detection using unsupervised image-to-image translation: From day to night. In: *International Joint Conference on Neural Networks*.
- Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., Vaughan, J.W., 2010. A theory of learning from different domains. *Mach. Learn.*
- Bilen, H., Pedersoli, M., Tuytelaars, T., 2014. Weakly supervised object detection with posterior regularization. In: *The British Machine Vision Conference*.
- Bilen, H., Vedaldi, A., 2016. Weakly supervised deep detection networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Brust, C.-A., Käding, C., Denzler, J., 2018. Active learning for deep object detection. *arXiv preprint arXiv:1809.09875*.
- Chen, B., Bakhshi, A., Batista, G., Ng, B., Chin, T.-J., 2022. Update compression for deep neural networks on the edge. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Chen, Y., Li, W., Sakaridis, C., Dai, D., Van Gool, L., 2018. Domain adaptive faster r-cnn for object detection in the wild. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Chen, Y., Wang, H., Li, W., Sakaridis, C., Dai, D., Van Gool, L., 2021. Scale-aware domain adaptive faster r-cnn. *Int. J. Comput. Vis.*
- Cheng, G., Yang, J., Gao, D., Guo, L., Han, J., 2020. High-quality proposals for weakly supervised object detection. *IEEE Trans. Image Process.*
- Choi, J., Elezi, I., Lee, H.J., Farabet, C., Alvarez, J.M., 2021. Active learning for deep object detection via probabilistic modeling. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- Churchill, W., Newman, P., 2013. Experience-based navigation for long-term localisation. *Int. J. Robot. Res.*
- Cinbis, R.G., Verbeek, J., Schmid, C., 2016. Weakly supervised object localization with multi-fold multiple instance learning. *IEEE Trans. Pattern Anal. Mach. Intell.*
- Damodaran, B.B., Kellenberger, B., Flamary, R., Tuia, D., Courty, N., 2018. DeepJDOT: Deep joint distribution optimal transport for unsupervised domain adaptation. In: *Proceedings of the European Conference on Computer Vision*.
- De Lange, M., Aljundi, R., Masana, M., Parisot, S., Jia, X., Leonardis, A., Slabaugh, G., Tuytelaars, T., 2021. A continual learning survey: Defying forgetting in classification tasks. *IEEE Trans. Pattern Anal. Mach. Intell.*
- D’Innocente, A., Borlino, F.C., Bucci, S., Caputo, B., Tommasi, T., 2020. One-shot unsupervised cross-domain detection. In: *European Conference on Computer Vision*.
- Doan, A.D., Latif, Y., Chin, T.-J., Liu, Y., Do, T.T., Reid, I., 2019. Scalable place recognition under appearance change for autonomous driving. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- Doan, A.D., Latif, Y., Chin, T.J., Reid, I., 2020. HM⁴: Hidden Markov model with memory management for visual place recognition. *IEEE Robot. Autom. Lett.*
- Dziugaite, G.K., Roy, D.M., Ghahramani, Z., 2015. Training generative neural networks via maximum mean discrepancy optimization. *arXiv preprint arXiv:1505.03906*.

- Flamary, R., Courty, N., Tuia, D., Rakotomamonjy, A., 2016. Optimal transport for domain adaptation. *IEEE Trans. Pattern Anal. Mach. Intell.*
- Ganin, Y., Lempitsky, V., 2015. Unsupervised domain adaptation by backpropagation. In: *International Conference on Machine Learning*.
- Gao, Y., Liu, B., Guo, N., Ye, X., Wan, F., You, H., Fan, D., 2019. C-midn: Coupled multiple instance detection network with segmentation guidance for weakly supervised object detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- Gao, P., Tian, T., Li, L., Ma, J., Tian, J., 2021. DE-CycleGAN: An object enhancement network for weak vehicle detection in satellite images. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.*
- Geiger, A., Lenz, P., Urtasun, R., 2012. Are we ready for autonomous driving? The KITTI vision benchmark suite. In: *Conference on Computer Vision and Pattern Recognition*.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y., 2020. Generative adversarial networks. *Commun. ACM*
- Gretton, A., Borgwardt, K.M., Rasch, M.J., Schölkopf, B., Smola, A., 2012. A kernel two-sample test. *J. Mach. Learn. Res.*
- Halder, S.S., Lalonde, J.F., de Charette, R., 2019. Physics-based rendering for improving robustness to rain. In: *ICCV*.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Hou, L., Zhang, Y., Fu, K., Li, J., 2021. Informative and consistent correspondence mining for cross-domain weakly supervised object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Inoue, N., Furuta, R., Yamasaki, T., Aizawa, K., 2018. Cross-domain weakly-supervised object detection through progressive domain adaptation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Kantorov, V., Oquab, M., Cho, M., Laptev, I., 2016. Contextlocnet: Context-aware deep network models for weakly supervised localization. In: *European Conference on Computer Vision*.
- Khodabandeh, M., Vahdat, A., Ranjbar, M., Macready, W.G., 2019. A robust learning approach to domain adaptive object detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- Kiefer, B., Ott, D., Zell, A., 2022. Leveraging synthetic data in object detection on unmanned aerial vehicles. In: *International Conference on Pattern Recognition*.
- Li, Y.J., Dai, X., Ma, C.Y., Liu, Y.C., Chen, K., Wu, B., He, Z., Kitani, K., Vajda, P., 2022. Cross-domain adaptive teacher for object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Lim, H., Kim, B., Choo, J., Choi, S., 2023. TTN: A domain-shift aware batch normalization in test-time adaptation. In: *International Conference on Learning Representations*.
- Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S., 2017a. Feature pyramid networks for object detection. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P., 2017b. Focal loss for dense object detection. In: *Proceedings of the IEEE International Conference on Computer Vision*.
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L., 2014. Microsoft COCO: Common objects in context. In: *European Conference on Computer Vision*.
- Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.-Y., Berg, A.C., 2016. Ssd: Single shot multibox detector. In: *European Conference on Computer Vision*.
- Liu, H., Gu, X., Samaras, D., 2019. Wasserstein GAN with quadratic transport cost. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- McCloskey, M., Cohen, N.J., 1989. Catastrophic interference in connectionist networks: The sequential learning problem. In: *Psychology of Learning and Motivation*. Elsevier.
- Ouyang, S., Wang, X., Lyu, K., Li, Y., 2021. Pseudo-label generation-evaluation framework for cross domain weakly supervised object detection. In: *IEEE International Conference on Image Processing*.
- Oza, P., Sindagi, V.A., VS, V., Patel, V.M., 2021. Unsupervised domain adaptation of object detectors: A survey. *arXiv preprint arXiv:2105.13502*.
- Papadopoulos, D.P., Uijlings, J.R., Keller, F., Ferrari, V., 2017. Training object class detectors with click supervision. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Pasqualino, G., Furnari, A., Signorello, G., Farinella, G.M., 2021. An unsupervised domain adaptation scheme for single-stage artwork recognition in cultural sites. *Image Vis. Comput.*
- Redmon, J., Divvala, S., Girshick, R., Farhadi, A., 2016. You only look once: Unified, real-time object detection. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Ren, S., He, K., Girshick, R., Sun, J., 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. *Adv. Neural Inf. Process. Syst.*
- Rockstar Games, 2013. *Grand Theft Auto V*. <https://www.rockstargames.com/gta-v>.
- Roy, S., Unmesh, A., Nambodiri, V.P., 2018. Deep active learning for object detection.. In: *The British Machine Vision Conference*.
- RoyChowdhury, A., Chakrabarty, P., Singh, A., Jin, S., Jiang, H., Cao, L., Learned-Miller, E., 2019. Automatic adaptation of object detectors to new domains using self-training. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Sachdeva, R., Hammond, R., Bockman, J., Arthur, A., Smart, B., Craggs, D., Doan, A.D., Rowntree, T., Schutz, E., Orenstein, A., et al., 2022. Autonomy and perception for space mining. In: *International Conference on Robotics and Automation*. pp. 4087–4093.
- Schutera, M., Hussein, M., Abhau, J., Mikut, R., Reischl, M., 2020. Night-to-day: Online image-to-image translation for object detection within autonomous driving by night. *IEEE Trans. Intell. Veh.*
- Song, H.O., Girshick, R., Jegelka, S., Mairal, J., Harchaoui, Z., Darrell, T., 2014. On learning to localize objects with minimal supervision. In: *International Conference on Machine Learning*.
- Spirals, 2019. *pyJoules*. <https://pyjoules.readthedocs.io/en/latest/>.
- Sun, B., Feng, J., Saenko, K., 2016. Return of frustratingly easy domain adaptation. In: *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Sun, B., Saenko, K., 2016. Deep CORAL: Correlation alignment for deep domain adaptation. In: *European Conference on Computer Vision*.
- Uijlings, J.R., Van De Sande, K.E., Gevers, T., Smeulders, A.W., 2013. Selective search for object recognition. *Int. J. Comput. Vis.*
- Villani, C., 2009. *Optimal Transport: Old and New*, Vol. 338. Springer.
- Wang, Q., Fink, O., Van Gool, L., Dai, D., 2022. Continual test-time domain adaptation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Wang, Y., Li, W., Dai, D., Van Gool, L., 2017. Deep domain adaptation by geodesic distance minimization. In: *Proceedings of the IEEE International Conference on Computer Vision Workshops*.
- Wang, D., Shelhamer, E., Liu, S., Olshausen, B., Darrell, T., 2021. Tent: Fully test-time adaptation by entropy minimization. In: *International Conference on Learning Representations*.
- Wu, Y., Kirillov, A., Massa, F., Lo, W.Y., Girshick, R., 2019. Detectron2. <https://github.com/facebookresearch/detectron2>.
- Xu, P., Gurram, P., Whipps, G., Chellappa, R., 2019. Wasserstein distance based domain adaptation for object detection. *arXiv preprint arXiv:1909.08675*.
- Xu, Y., Sun, Y., Yang, Z., Miao, J., Yang, Y., 2022. H2FA R-CNN: Holistic and hierarchical feature alignment for cross-domain weakly supervised object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Yan, H., Li, Z., Wang, Q., Li, P., Xu, Y., Zuo, W., 2019. Weighted and class-specific maximum mean discrepancy for unsupervised domain adaptation. *IEEE Trans. Multimed.*
- Yu, F., Chen, H., Wang, X., Xian, W., Chen, Y., Liu, F., Madhavan, V., Darrell, T., 2020. BDD100k: A diverse driving dataset for heterogeneous multitask learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Yu, W., Zhu, S., Yang, T., Chen, C., 2022. Consistency-based active learning for object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Yuan, T., Wan, F., Fu, M., Liu, J., Xu, S., Ji, X., Ye, Q., 2021. Multiple instance active learning for object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Zhang, Y., Chen, Y., Wang, J., Pan, Z., 2021. Unsupervised deep anomaly detection for multi-sensor time-series signals. *IEEE Trans. Knowl. Data Eng.*
- Zhang, W., Han, B., Hui, P., 2022. SEAR: Scaling experiences in multi-user augmented reality. *IEEE Trans. Vis. Comput. Graphics*.
- Zhu, J.Y., Park, T., Isola, P., Efros, A.A., 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In: *Proceedings of the IEEE International Conference on Computer Vision*.
- Zitnick, C.L., Dollár, P., 2014. Edge boxes: Locating object proposals from edges. In: *European Conference on Computer Vision*.