

“©2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.”

Instance-Invariant Domain Adaptive Object Detection via Progressive Disentanglement

Aming Wu^{1*}, Yahong Han¹, Linchao Zhu², Yi Yang²

¹College of Intelligence and Computing, Tianjin University, Tianjin, China

²ReLER, University of Technology Sydney, Australia

{tjwam,yahong}@tju.edu.cn, {Linchao.Zhu,yi.yang}@uts.edu.au

Abstract

Most state-of-the-art methods of object detection suffer from poor generalization ability when the training and test data are from different domains, e.g., with different styles. To address this problem, previous methods mainly use holistic representations to align feature-level and pixel-level distributions of different domains, which may neglect the instance-level characteristics of objects in images. Besides, when transferring detection ability across different domains, it is important to obtain the instance-level features that are domain-invariant, instead of the styles that are domain-specific. Therefore, in order to extract instance-invariant features, we should disentangle the domain-invariant features from the domain-specific features. To this end, a progressive disentangled framework is first proposed to solve domain adaptive object detection. Particularly, based on disentangled learning used for feature decomposition, we devise two disentangled layers to decompose domain-invariant and domain-specific features. And the instance-invariant features are extracted based on the domain-invariant features. Finally, to enhance the disentanglement, a three-stage training mechanism including multiple loss functions is devised to optimize our model. In the experiment, we verify the effectiveness of our method on three domain-shift scenes. Our method is separately 2.3%, 3.6%, and 4.0% higher than the baseline method [36].

1. Introduction

Recently, great efforts have been made on object detection [11, 32, 15, 24, 31]. Though most state-of-the-art methods achieve outstanding detection performance on many benchmarks [9, 25], they suffer from poor generalization ability when the training and test images are from different domains, which is cast into the setting of domain adaptive object detection (DAOD). In the task of DAOD, domain gap always exists between the source/training and target/test im-

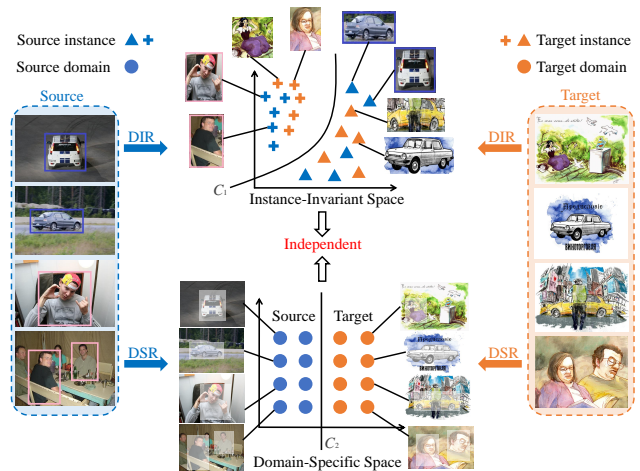


Figure 1. The process of our disentangled method for domain adaptive object detection. We decompose source and target image representations into domain-invariant representations (DIR) and domain-specific representations (DSR). Then, we extract from DIR the instance-invariant representations that lie in an instance-invariant space, in which the instance-invariant features are used to describe the characteristics of objects. In the instance-invariant space, we conduct instance classification (i.e., via C_1) for the adaptive object detection. And different domains could be easily distinguished (i.e., via C_2) in the domain-specific space.

ages, e.g., with different illuminations and different styles etc. Although the performance could be improved via collecting additional images with well-labeled objects from the target domain, it is time-consuming and labor-intensive.

In order to alleviate the impact of domain-shift [10], representative methods [5, 36, 14] towards DAOD employ unsupervised domain adaptation [34, 29, 44] to align distributions of different domains, e.g., via adversarial training [10] or style translation [21]. Distribution alignment is always conducted in a holistic representation (e.g., in feature-level [6, 22] or pixel-level [12, 3, 35]) of source and target images, which may neglect the instance-level characteristics of objects in images, such as object locations or basic shapes of objects etc. When transferring detection ability from

*This work was done when Aming Wu visited ReLER Lab, UTS.

source images to target images, it is the instance-level features that really count, which are always domain-invariant, not the illuminations and painting styles that are domain-specific. Therefore, in order to obtain the instance-invariant features and bridge the domain gap in DAOD, we should try to disentangle the domain-invariant representations (DIR) from the domain-specific representations (DSR).

As a method of feature decomposition, disentangled learning [8, 28] has been demonstrated to be effective in tasks of few-shot learning [33, 38] and image translation [23, 16]. The purpose of disentangled learning is to uncover a set of independent factors that give rise to the current observation [8]. And the major advantage is that disentangled representations could contain all the information presents in the current observation in a compact and interpretable structure while being independent of the current task [28, 2]. In this paper, we propose to employ disentangled learning to disentangle an image representation into a domain-invariant representation (DIR) and a domain-specific representation (DSR) (see Fig. 1), so as to obtain the instance-invariant representation (IIR). Taking the IIR as a bridge, we have great potential to strengthen the transferring ability of a detection model trained on source images.

Particularly, in the proposed detection network, we devise a progressive process to decompose the DIR and DSR with two disentangled layers. The goal of the first layer is to enhance the domain-invariant information in a middle-layer feature map. We utilize a domain classifier to ensure that DSR contains much more domain-specific information. And a mutual information (MI) loss is employed to enlarge the gap between DIR and DSR. Taking the sum of the feature map and DIR as the input, the second layer aims at obtaining the instance-invariant representations (IIR) with a regional proposal networks (RPN) [32, 41]. Moreover, to enhance the disentanglement, we devise a training mechanism including three stages to optimize our model: (i) the stage of feature decomposition aiming at learning disentanglement, (ii) the stage of feature separation aiming at enlarging the gap between DIR and DSR, and (iii) the stage of feature reconstruction aiming at keeping the DIR and DSR contain all the content of the input. For each stage, we use different loss functions to optimize different components of our network, respectively. Experiments on three domain-shift scenes of DAOD demonstrate that our method is effective and achieves a new state-of-the-art performance.

The contributions of this paper are summarized as:

(1) Different from reducing the domain gap with distribution alignment, we propose to enhance the transferring detection ability via a bridge of disentangled instance-invariant representations.

(2) A progressive disentangled network is first proposed to successfully extract instance-invariant features. Meanwhile, a three-stage training mechanism is proposed to fur-

ther enhance the disentangled ability.

(3) On three domain-shift scenes, i.e., Cityscapes [7] → FoggyCityscapes [37], Pascal [9] → Watercolor [17], and Pascal → Clipart [17], our method is separately 2.3%, 3.6%, and 4.0% higher than the baseline method [36].

2. Related Work

Domain Adaptive Object Detection. Though most methods [11, 31, 15, 26] of object detection have achieved outstanding performance, their transferring abilities are limited for the task of DAOD. Recently, many methods [21, 36, 20] have been proposed to solve the domain-shift problem in object detection. These methods mainly focus on feature-level or pixel-level alignment. For example, the method in [5] utilizes adversarial training [10] to align global feature distributions of the source and target domains, whereas the method in [36] aligns distributions of both global and local features. For pixel-level adaptation, the work [21] devises a generative network to increase the diversity of the source domain, which is similar to data augmentation. However, as the alignment is conducted in holistic representations of images, it is not dedicated to the task of adaptive object detection, which focuses on the bridge of domains with instance-level characteristics. Therefore, in this paper, we focus on extracting instance-level features that are domain-invariant, which are helpful for improving the transferring ability of a detection method.

Disentangled Learning. The purpose of disentangled learning [18, 28, 2, 30] is to correctly uncover a set of independent factors that give rise to the current observation. Recently, disentangled learning has been well explored in tasks of few-shot learning [33, 38] and image translation [23, 16]. Particularly, by decomposing the style of an image, the work [23] proposed a disentangled method to make a diverse image-to-image translation. Liu et al. [27] proposed a model of cross-domain representation disentanglement. Based on generative adversarial networks, this method alleviated the impact of domain-shift and improved the classification performance on multiple datasets. As for adaptive object detection, on one hand, we should remove the domain-shift; on the other hand, it is important to transfer the detection ability via the bridge of the instance characteristics. Thus, it is not straightforward to apply the disentangled learning to the task of DAOD.

In this paper, we devise a new network of progressive disentanglement to decompose image representations into domain-specific and domain-invariant representations, and from which we extract the instance-invariant representations to bridge the detection ability between source and target domains. Experiments on three domain-shift scenes of DAOD demonstrate the effectiveness of our method.

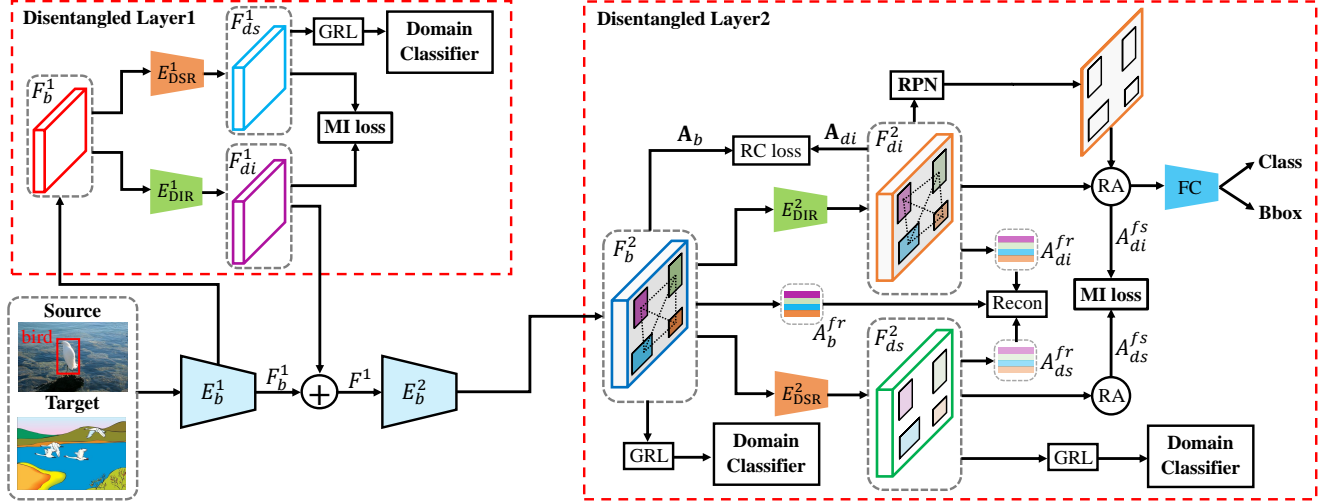


Figure 2. Illustration of the proposed network of progressive disentanglement. ‘Recon’ indicates the reconstruction loss. ‘GRL’ is the gradient reverse layer. ‘RA’ indicates the operation of RoI-Alignment. ‘RC loss’ and ‘MI loss’ separately denote the proposed relation-consistency loss and the mutual information loss. ‘ \oplus ’ is the operation of element-wise sum. And the dot lines indicate the relations existed between the extracted proposals. There are two disentanglement layers in the network. The purpose of the first layer is to enhance the domain-invariant information in a middle-layer feature map. And the goal of the second layer is to obtain the instance-invariant features. During training, in order to enhance the disentanglement, we devise a three-stage optimization mechanism with multiple loss functions. For each stage, we use different loss functions to optimize different components of the network.

3. Instance-Invariant Adaptive Detection

Suppose we have the access to an image x^s including labels y^s and bounding boxes b^s , which are drawn from a set of annotated source images $\{X_s, Y_s, B_s\}$. Here, X_s , Y_s , and B_s separately indicate the set of images, labels, and bounding-box annotations, which are from the source domain. Meanwhile, we could also access to a target image x^t drawn from a set of unlabeled target images $\{X_t\}$.

3.1. The Network of Progressive Disentanglement

As is shown in Fig. 2, we devise two disentanglement layers to extract domain-invariant information progressively.

The First Disentangled Layer. The goal of this layer is to enhance the domain-invariant information in a middle-layer feature map. Concretely, given a source image x^s and target image x^t , we first obtain a feature map F_b^1 that is the output of a middle-layer feature extractor E_b^1 . Then, two different extractors are devised to disentangle the DIR and DSR from F_b^1 . The processes are shown as follows:

$$F_{di}^1 = E_{DIR}^1(F_b^1), F_{ds}^1 = E_{DSR}^1(F_b^1), F^1 = F_b^1 + F_{di}^1. \quad (1)$$

Here, E_{DIR}^1 and E_{DSR}^1 separately indicate the DIR and DSR extractor. The size of F_{di}^1 and F_{ds}^1 is set to the same value as that of F_b^1 . Then, we take the sum F^1 of F_{di}^1 and F_b^1 as the input of the second feature extractor E_b^2 . Since F_{di}^1 contains more domain-invariant information, the sum operation could alleviate the impact of domain-shift on F^1 .

The Second Disentangled Layer. The purpose of this layer is to obtain the instance-invariant features. Particu-

larly, based on the output F_b^2 of the extractor E_b^2 , we devise two extractors, i.e., E_{DIR}^2 and E_{DSR}^2 , to disentangle the DIR and DSR from F_b^2 . The processes are as follows:

$$F_b^2 = E_b^2(F^1), F_{di}^2 = E_{DIR}^2(F_b^2), F_{ds}^2 = E_{DSR}^2(F_b^2). \quad (2)$$

Here, the size of F_{di}^2 and F_{ds}^2 is set to the same value as that of F_b^2 . Next, the RPN is performed on F_{di}^2 to extract a set of instance-invariant proposals. Finally, for an image from the source domain, the detection loss is as follows:

$$L_D = -\frac{1}{n_s} \sum_{j=1}^{n_s} L_{det}(D(A_j), y_j^s, b_j^s), \quad (3)$$

where n_s denotes the number of proposals. A_j indicates the RoI-Alignment [32, 13] result of the j -th proposal. D includes the classification and regression network. L_{det} is assumed to contain all the losses for the detection, e.g., classification and bounding-box regression loss.

3.2. Training with the Three-stage Optimization

As is discussed in the section of Introduction, the goal of disentangled learning is to uncover a set of independent factors that give rise to the current observation [8]. And these factors could contain all the information presents in the observation [28]. Therefore, we devise a three-stage training mechanism (see Fig. 3) to enhance the disentanglement.

3.2.1 The Stage of Feature Decomposition

The goal of the first stage is to ensure that our model not only learns the location and classification of the objects but




Figure 3. Illustration of the three-stage training. Here, the red arrow denotes the operation of reconstruction. ‘Stage- f_d ’ is the first stage aiming at learning disentanglement. ‘Stage- f_s ’ is the second stage aiming at keeping the disentangled DIR and DSR independent. And ‘Stage- f_r ’ is the third stage aiming at keeping the DIR and DSR could contain all the content of the input.

also disentangles the image features. Based on F_{di}^2 , we first utilize RPN to obtain a set of object proposals O^{fd} . To ensure that F_b^2 and F_{di}^2 have the same object contents in the same locations, based on the proposals O^{fd} , RoI-Alignment is performed on F_b^2 and F_{di}^2 to obtain A_b^{fd} and A_{di}^{fd} , respectively. Next, we devise two networks D_b and D_{di} to perform the classification and bounding-box regression. Finally, for a source image, the detection loss is defined as:

$$L_D^{fd} = L_D^b(D_b(A_b^{fd})) + L_D^{di}(D_{di}(A_{di}^{fd})), \quad (4)$$

where L_D^b and L_D^{di} indicate the detection loss.

By using the detection loss, F_b^2 and F_{di}^2 are ensured to contain the instance information. Besides, for our method, it is also important to keep the learned F_{ds}^1 and F_{ds}^2 contain more domain-specific information, which could ensure our model owns the ability of feature disentanglement. In this paper, we exploit the method of adversarial domain classification [10] to distinguish the source and target domains. Specifically, we employ four domain classifiers C_b^1 , C_{ds}^1 , C_b^2 , and C_{ds}^2 in our model, which separately take F_b^1 , F_{ds}^1 , F_b^2 , and F_{ds}^2 as the input and output a domain label l_d that indicates the source or target domain: l_d is 0 for the source domain and 1 for the target domain.

Besides, for domain classifiers, during training, we employ Focal Loss (FL) [24, 36] to impose bigger weights on the hard-to-classify examples (i.e., the examples near the classification boundary) than on the easy ones (i.e., the examples far from the classification boundary).

$$FL(p) = -g(p)\log(p), \quad g(p) = \alpha(1-p)^\gamma, \quad (5)$$

where γ controls the weight on the hard-to-classify examples. $p \in [0, 1]$ is the model’s estimated probability for the output domain label l_d . Finally, the loss of the first training stage is denoted as follows:

$$\begin{aligned} L_s^{fd} &= L_D^{fd} + FL_s(C_b^1(F_b^1)) + FL_s(C_{ds}^1(F_{ds}^1)) \\ &\quad + FL_s(C_b^2(F_b^2)) + FL_s(C_{ds}^2(F_{ds}^2)), \\ L_t^{fd} &= FL_t(C_b^1(F_b^1)) + FL_t(C_{ds}^1(F_{ds}^1)) \\ &\quad + FL_t(C_b^2(F_b^2)) + FL_t(C_{ds}^2(F_{ds}^2)), \end{aligned} \quad (6)$$

where L_s^{fd} and L_t^{fd} are the objective functions of the source and target domains. FL_s and FL_t indicate the domain losses. The overall loss L^{fd} is the sum of L_s^{fd} and L_t^{fd} .

With the help of the detection loss L_D^{fd} and domain loss FL , the disentangled DIR and DSR contain instance and domain-specific information, respectively. Next, we will perform the second training stage to keep the disentangled DIR and DSR independent.

3.2.2 The Stage of Feature Separation

In this stage, we first fix the extractor E_b^1 and E_b^2 of the model trained on the first stage. Then, we employ the model to extract F_b^1 , F_{di}^1 , F_{ds}^1 (Eq. (1)), F_b^2 , F_{di}^2 , and F_{ds}^2 (Eq. (2)). RPN is performed on F_{di}^2 to obtain the proposals O^{fs} .

Mutual Information Minimization. In order to enlarge the gap between the DIR and DSR, we minimize the MI loss between A_{di}^{fs} and A_{ds}^{fs} , as well as between F_{di}^1 and F_{ds}^1 , where A_{di}^{fs} and A_{ds}^{fs} indicate the RoI-Alignment results of F_{di}^2 and F_{ds}^2 based on O^{fs} . The process of MI is:

$$I(X; Z) = \int_{X \times Z} \log \frac{d\mathbb{P}_{XZ}}{d\mathbb{P}_X \otimes \mathbb{P}_Z} d\mathbb{P}_{XZ}, \quad (7)$$

where \mathbb{P}_{XZ} indicates the joint probability distribution of $(A_{di}^{fs}, A_{ds}^{fs})$ or (F_{di}^1, F_{ds}^1) . $\mathbb{P}_X = \int_Z d\mathbb{P}_{XZ}$ and $\mathbb{P}_Z = \int_X d\mathbb{P}_{XZ}$ are the marginal distributions. Obviously, by minimizing the MI loss, we could impose independent constraints on the tuples $(A_{di}^{fs}, A_{ds}^{fs})$ and (F_{di}^1, F_{ds}^1) . Besides, since F_{ds}^1 and F_{ds}^2 contain more domain-specific information, MI loss could promote F_{di}^1 and F_{di}^2 to contain more domain-invariant information, which can help strengthen the ability of disentanglement. In this paper, we adopt Mutual Information Neural Estimator (MINE) [1] to compute the MI loss. Concretely, based on Monte-Carlo integration [30], MINE could be computed as follows:

$$I(X, Z) = \frac{1}{n} \sum_{j=1}^n T(x, z, \theta) - \log\left(\frac{1}{n} \sum_{j=1}^n e^{T(x, z', \theta)}\right), \quad (8)$$

where (x, z) is sampled from the joint distribution and z' is sampled from the marginal distribution. Here, we devise a neural network to perform the Monte-Carlo integration.

It is worth noting that, for the second disentangled layer, we use the RoI-Alignment results A_{di}^{fs} and A_{ds}^{fs} , instead of the feature map F_{di}^2 and F_{ds}^2 , to compute MI loss, which could not only reduce the computational costs but also ensure our model pays more attention to regions of objects.

Relation-consistency Loss. To further improve the disentanglement, we devise a relation-consistency loss (Fig. 4). Specifically, since F_{di}^2 and F_b^2 have the same object contents in the same locations, based on the proposals O^{fs} , A_{di}^{fs} and A_b^{fs} should keep similar semantic relations.

Concretely, we first obtain the average-pooling results $P_{di}^{fs} \in \mathbb{R}^{k \times m}$ and $P_b^{fs} \in \mathbb{R}^{k \times m}$ of A_{di}^{fs} and A_b^{fs} ,




Figure 4. Illustration of relation-consistency Loss. ‘P’ indicates the ‘Person’ class. The goal of the loss is to ensure the relations (the red solid lines) between object proposals in F_b^2 and the relations between object proposals in F_{di}^2 are consistent. The purple dot lines denote the consistency between the two red lines.

where k and m indicate the numbers of proposals and channels. Then we separately construct a graph $G_{di} = \{\mathcal{V}_{di}, \mathcal{E}_{di}\}$ and $G_b = \{\mathcal{V}_b, \mathcal{E}_b\}$. Here, we take P_{di}^{fs} and P_b^{fs} as the nodes \mathcal{V}_{di} and \mathcal{V}_b , respectively. \mathcal{E}_{di} and \mathcal{E}_b are used to indicate the edges (relations) between proposals. Next, we define two adjacency matrix for two undirected graphs, i.e., $\mathbf{A}_b = \text{softmax}_r((P_b^{fs})(P_b^{fs})^T)$ and $\mathbf{A}_{di} = \text{softmax}_r((P_{di}^{fs})(P_{di}^{fs})^T)$. And softmax_r indicates we make softmax operation across the row directions. The relation-consistency loss is computed as:

$$L_{rel} = \|\mathbf{A}_{di} - \mathbf{A}_b\|_2^2. \quad (9)$$

Note that the computation of the relation-consistency loss does not need any parameters. Finally, the loss of the second training stage is denoted as follows:

$$\begin{aligned} L_s^{fs} &= L_D^{di}(D_{di}(A_{di}^{fs})) + FL_s(C_{ds}^2(F_{ds}^2)) + I_s^2 \\ &\quad + L_{rel}^s + FL_s(C_{ds}^1(F_{ds}^1)) + I_s^1, \\ L_t^{fs} &= FL_t(C_{ds}^2(F_{ds}^2)) + I_t^2 + L_{rel}^t \\ &\quad + FL_t(C_{ds}^1(F_{ds}^1)) + I_t^1, \end{aligned} \quad (10)$$

where L_D^{di} is the detection loss based on A_{di}^{fs} . L_s^{fs} and L_t^{fs} are the training objectives of the source and target domain, respectively. I^1 and I^2 indicate MI loss computed on the first and second disentangled layer, respectively. The overall loss L^{fs} is the sum of L_s^{fs} and L_t^{fs} . After this stage, the gap between DIR and DSR could be enlarged. Next, we will perform the third training stage aiming at keeping the disentangled DIR and DSR contain all the content of the input used for disentanglement.

3.2.3 The Stage of Feature Reconstruction

We employ a reconstruction loss to attain the purpose of this training stage. Concretely, we first use the model trained on the second stage to extract F_b^2 , F_{di}^2 , and F_{ds}^2 (Eq. (2)). Then, RPN is performed on F_{di}^2 to extract proposals O^{fr} . The reconstruction loss is computed as follows:

$$A_r^{fr} = R(\langle A_{di}^{fr}, A_{ds}^{fr} \rangle), \quad L_{recon} = \|A_r^{fr} - A_b^{fr}\|_2^2, \quad (11)$$

where A_{di}^{fr} , A_{ds}^{fr} , and A_b^{fr} are the RoI-Alignment results of F_{di}^2 , F_{ds}^2 , and F_b^2 based on the proposals O^{fr} . R is the

Algorithm 1 Instance-Invariant Adaptive Object Detection

Input: source images $\{X_s, Y_s, B_s\}$; target images $\{X_t\}$; feature extractors E_b^1 and E_b^2 ; disentangled extractors $E_{DIR}^1, E_{DSR}^1, E_{DIR}^2$, and E_{DSR}^2 ; detection networks D_b and D_{di} ; domain classifiers C_b^1, C_{ds}^1, C_b^2 , and C_{ds}^2 ; MI estimators T^1 and T^2 ; reconstruction network R .

Output: trained $\hat{E}_b^1, \hat{E}_{DIR}^1, \hat{E}_{DSR}^1, \hat{E}_b^2, \hat{E}_{DIR}^2, \hat{E}_{DSR}^2$, and detection networks \hat{D}_b, \hat{D}_{di} .

- 1: **while** not converged **do**
- 2: Sample a mini-batch from $\{X_s, Y_s, B_s\}$ and $\{X_t\}$;
- 3: **Feature Decomposition:**
- 4: Update $E_b^1, E_{DIR}^1, E_b^2, E_{DIR}^2, D_b, D_{di}$ by Eq. (4);
- 5: Update $E_{DSR}^1, C_b^1, C_{ds}^1, E_b^2, E_{DSR}^2, C_b^2$, and C_{ds}^2 by Eq. (5);
- 6: **Feature Separation:**
- 7: Update $E_{DIR}^1, E_{DIR}^2, D_{di}$ by L_D^{di} in Eq. (10);
- 8: Update $E_{DSR}^1, C_{ds}^1, E_{DSR}^2, C_{ds}^2$ by FL in Eq. (10);
- 9: Calculate the MI loss between F_{di}^1 and F_{ds}^1 with T^1 , and between A_{di}^{fs} and A_{ds}^{fs} with T^2 ;
- 10: Update $E_{DIR}^1, E_{DSR}^1, T^1, E_{DIR}^2, E_{DSR}^2, T^2$ by Eq. (8);
- 11: Update E_{DIR}^2 by Eq. (9);
- 12: **Feature Reconstruction:**
- 13: Reconstruct RoI-Alignment result A_b^{fr} by $(A_{di}^{fr}, A_{ds}^{fr})$;
- 14: Update E_{DIR}^2, E_{DSR}^2, R by Eq. (11);
- 15: **end while**
- 16: **return** $\hat{E}_b^1 = E_b^1; \hat{E}_{DIR}^1 = E_{DIR}^1; \hat{E}_{DSR}^1 = E_{DSR}^1; \hat{E}_b^2 = E_b^2; \hat{E}_{DIR}^2 = E_{DIR}^2; \hat{E}_{DSR}^2 = E_{DSR}^2; \hat{D}_b = D_b; \hat{D}_{di} = D_{di}$.

reconstruction network. $\langle a, b \rangle$ indicates the concatenation of a and b . Here, in order to make the model pay more attention to instance content, the reconstruction loss is only computed on the regions of the proposals. Besides, since the output of the first disentangled layer includes the entire F_b^1 , to reduce the computational costs, we do not calculate the reconstruction loss on the first layer.

In this paper, our model is trained in an end-to-end way. The detailed training procedures are presented in Algorithm 1. During each training stage, the parameters that do not appear in the current stage are considered to be fixed.

4. Experiments

We evaluate our approach on three domain-shift scenes, i.e., Cityscapes [7] \rightarrow FoggyCityscapes [37], Pascal VOC [9] \rightarrow Watercolor [17], and Pascal VOC \rightarrow Clipart [17].

4.1. Dataset and Implementation Details

Dataset. For Cityscapes \rightarrow FoggyCityscapes, we use Cityscapes as the source domain. FoggyCityscapes is used as the target domain, which is rendered from Cityscapes and simulates the change of weather condition. Both of them contain 2,975 images in the training set and 500 images in the validation set. And this adaptation scene involves 8 categories. We utilize the training set during training and evaluate on the validation set.

Method	backbone	person	rider	car	truck	bus	train	motorcycle	bicycle	mAP
Source Only	VGG16	24.7	31.9	33.1	11.0	26.4	9.2	18.0	27.9	22.8
DAF [5]	VGG16	25.0	31.0	40.5	22.1	35.3	20.2	20.0	27.1	27.6
DT [17]	VGG16	25.4	39.3	42.4	24.9	40.4	23.1	25.9	30.4	31.5
SC-DA(Type3) [45]	VGG16	33.5	38.0	48.5	26.5	39.0	23.3	28.0	33.6	33.8
DMRL [21]	VGG16	30.8	40.5	44.3	27.2	38.4	34.5	28.4	32.2	34.6
MTOR [4]	ResNet50	30.6	41.4	44.0	21.9	38.6	40.6	28.3	35.6	35.1
MLDA [43]	VGG16	33.2	44.2	44.8	28.2	41.8	28.7	30.5	36.5	36.0
FSDA [42]	VGG16	29.1	39.7	42.9	20.8	37.4	24.1	26.5	29.9	31.3
MAF [14]	VGG16	28.2	39.5	43.9	23.8	39.9	33.3	29.2	33.9	34.0
RLDA [19]	IncepV2 [40]	35.10	42.15	49.17	30.07	45.25	26.97	26.85	36.03	36.45
SW (B) [36]	VGG16	29.9	42.3	43.5	24.5	36.2	32.6	30.0	35.3	34.3
Ours	VGG16	33.12	43.41	49.63	21.98	45.75	32.04	29.59	37.08	36.57
Ours	ResNet101	32.82	44.37	49.57	33.02	46.10	37.97	29.90	35.26	38.63

Table 1. Results (%) on adaptation from Cityscapes to FoggyCityscapes. ‘B’ indicates the baseline method. ‘Source Only’ indicates the model is only trained based on the data from the source domain and does not use the target data.

For Pascal \rightarrow Watercolor and Pascal \rightarrow Clipart, Pascal VOC dataset is used as the real source domain. The images of this dataset include rich bounding box annotations. And the number of object classes is 20. Following a prevalent setting [21, 36], we use Pascal VOC 2007 and 2012 training and validation set for training, which results in about 15K images. Watercolor and Clipart datasets are taken as the target domain. Watercolor contains 6 categories in common with VOC and 2k images in total. Clipart contains 1k images in total, which has the same 20 categories as VOC. For these two target datasets, the splits of training and test set are the same as the work [36].

Implementation Details. Our method is based on Faster-RCNN [32] with RoI-Alignment [13]. For Focal Loss (Eq. (5)), α and γ are set to 1.0 and 2.0. Besides, we separately employ a network including three convolutional layers as the disentangled extractors E_{DIR}^1 , E_{DSR}^1 , E_{DIR}^2 , and E_{DSR}^2 . For the domain classifiers C_b^1 , C_{ds}^1 , C_b^2 , and C_{ds}^2 , we respectively employ a network which includes three fully-connected layers. Meanwhile, for the MI estimators T^1 and T^2 , we separately utilize a network consisting of three fully-connected layers. Finally, one convolutional layer is used as the reconstruction network R . During training, we employ the SGD optimizer with momentum [39]. We first train the model with a learning rate of 0.001 for 50K iterations, then with a learning rate of 0.0001 for 30K more iterations. In the test, we use mean average precisions (mAP) as the evaluation metric.

4.2. Experimental Results

Results on FoggyCityscapes. Table 1 shows the performance of our method on the FoggyCityscapes dataset. Here, we use VGG16 and ResNet101 as the backbone of Faster-RCNN, respectively. We can see that our method outperforms all the methods in Table 1. Particularly, based on the

Method	bike	bird	car	cat	dog	person	mAP
Source Only	68.8	46.8	37.2	32.7	21.3	60.7	44.6
BDC-Faster [36]	68.6	48.3	47.2	26.5	21.7	60.5	45.5
DAF [5]	75.2	40.6	48.0	31.5	20.6	60.0	46.0
SW (B) [36]	82.3	55.9	46.5	32.7	35.5	66.7	53.3
Ours	95.8	54.3	48.3	42.4	35.1	65.8	56.9

Table 2. Results (%) on adaptation from Pascal to Watercolor.

VGG16 backbone and mAP metric, our method is around 2.3% higher than the SW baseline method [36]. Compared with RLDA [19] using InceptionV2 [40] as the strong backbone, our method still outperforms it. These all show our method is effective. Moreover, employing the backbone of ResNet101 could improve the performance of our method significantly. This shows our method is more effective with a better backbone. Fig. 5 shows two detection examples. Compared with the raw images, for object detection, the foggy scene is much more challenging. Meanwhile, compared with the SW method, our method could locate and recognize objects existing in the two images accurately. Particularly, regardless of distance, our method could locate and discriminate the truck accurately. These further demonstrate the effectiveness of our method.

Results on Watercolor and Clipart. Table 2 and 3 separately show the performance of our method on Watercolor and Clipart dataset. Here, we all use ResNet101 as the backbone of Faster-RCNN. For Watercolor scene, our method is 3.6% higher than the SW method. Particularly, for the class of bike, our method outperforms SW by around 13%. This shows our method is effective for the task of DAOD. Fig. 6 shows two examples of Watercolor. We can see that our method could locate and recognize the classes of person and bird accurately. This further shows that our disentangled method indeed alleviates the problem of domain-shift and improves the detection performance.



Figure 5. Detection results on the “Cityscapes \rightarrow FoggyCityscapes” scene. ‘GT’ indicates the groundtruth result. ‘One Disentangled layer’ denotes we only use the second disentangled layer in the model. We can see that our method, i.e., using two disentangled layers, could locate and recognize objects existing in the two foggy images accurately, e.g., the **truck**, **car**, and **bicycle**.

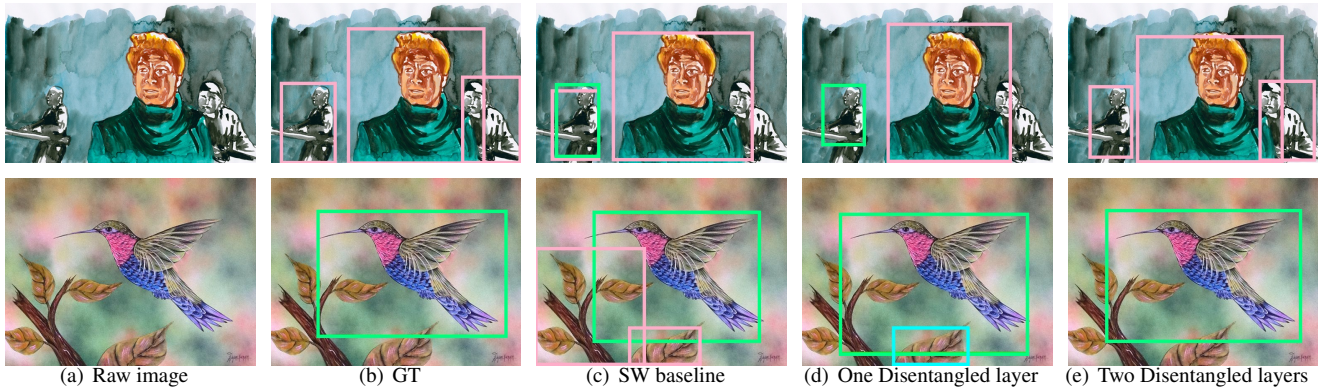


Figure 6. Detection results on the “Pascal VOC \rightarrow Watercolor” scene. We can see that our method, i.e., using two disentangled layers, could locate and recognize objects existing in the two watercolor images accurately, e.g., the **person**, **bird**, and **cat**.

As for Clipart scene which involves more classes than the other two datasets, our method outperforms SW by 4.0%, in terms of the mAP metric. Meanwhile, in Table 3, we can see that our method outperforms the baseline method in multiple categories significantly. For example, for the aeroplane and dog class, our method is around 15% and 16% higher than the SW method. These all demonstrate the good performance of our method.

4.3. Ablation Analysis

In this section, we will make some ablation analysis on our method. Table 4 shows the ablation results. Here, ‘C \rightarrow F’ and ‘V \rightarrow W’ separately indicate the adaptation from Cityscapes to FoggyCityscapes and the adaptation from Pascal VOC to Watercolor. And for the ‘C \rightarrow F’ case, we use VGG16 as the backbone. For the ‘V \rightarrow W’ case, we use ResNet101 as the backbone. ‘OW’ indicates we integrate all loss functions existing in our method and use one training stage. ‘1st’, ‘2nd’, and ‘3rd’ indicate we use the first training stage of Algorithm 1, the first two training stages of Algorithm 1, and the three training stages to optimize our model, respectively. For our progressive method (Two layers), we can see that the three-stage training mechanism is effective. For example, for the ‘C \rightarrow F’ case, the

performance is improved from 33.6% to 36.6%. Meanwhile, we can see that from the first training stage to the third stage, the performance is improved continuously. This shows that for the disentangled learning, the stage of feature separation and feature reconstruction is necessary. Using these two stages does enhance the disentanglement and improve the detection performance. Besides, we can also see that the relation-consistency loss (RC) improves the performance of our method significantly. For example, for the ‘V \rightarrow W’ scene, the performance is improved from 55.2% to 56.9%. This demonstrates the relation-consistency loss helps strengthen the ability of disentanglement.

To further verify the effectiveness of the progressive method, we make a comparison with the method of only using the second disentangled layer (One layer). We can see from Table 4 that our progressive method improves the detection performance significantly, e.g., for the ‘C \rightarrow F’ case, the performance is improved from 34.1% to 36.6%. This shows that using the progressive mechanism is indeed helpful for obtaining a better disentangled representation. Besides, in Fig. 5 and 6, we can see that compared with One layer method, employing two disentangled layers does improve the accuracy of location and recognition. Particularly, taking the first image in Fig. 6 as an example, our




Figure 7. Visualization of feature maps of the second disentangled layer. Here, ‘O-DIR’ (F_{di}^2) and ‘O-DSR’ (F_{ds}^2) indicate we only use the second disentangled layer to extract DIR and DSR based on ‘O-Base’ (F_b^2) and do not use the first disentangled layer. ‘P-DIR’ (F_{di}^2) and ‘P-DSR’ (F_{ds}^2) indicate we use the progressive method to extract DIR and DSR based on ‘P-Base’ (F_b^2). For each feature map, the channels corresponding to the maximum value are selected for visualization. For ‘O-DIR’ and ‘P-DIR’, the bright regions indicate the presence of object-relevant content. For ‘O-DSR’ and ‘P-DSR’, the bright regions indicate the presence of domain-specific information.

Method	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mAP
Source Only	35.6	52.5	24.3	23.0	20.0	43.9	32.8	10.7	30.6	11.7	13.8	6.0	36.8	45.9	48.7	41.9	16.5	7.3	22.9	32.0	27.8
BDC-Faster [36]	20.2	46.4	20.4	19.3	18.7	41.3	26.5	6.4	33.2	11.7	26.0	1.7	36.6	41.5	37.7	44.5	10.6	20.4	33.3	15.5	25.6
DAF [5]	15.0	34.6	12.4	11.9	19.8	21.1	23.2	3.1	22.1	26.3	10.0	10.0	19.6	39.4	34.6	29.3	1.0	17.1	19.7	24.8	19.8
SW (B) [36]	26.2	48.5	32.6	33.7	38.5	54.3	37.1	18.6	34.8	58.3	17.0	12.5	33.8	65.5	61.6	52.0	9.3	24.9	54.1	49.1	38.1
Ours	41.5	52.7	34.5	28.1	43.7	58.5	41.8	15.3	40.1	54.4	26.7	28.5	37.7	75.4	63.7	48.7	16.5	30.8	54.5	48.7	42.1

Table 3. Results (%) on adaptation from Pascal VOC to Clipart. Here, we use ResNet101 as the backbone of Faster-RCNN.

Method	OW	1st	2nd	3rd	RC	C → F	V → W
Two layers	✓				✓	34.1%	52.9%
Two layers		✓				33.6%	53.5%
Two layers			✓		✓	35.3%	55.3%
Two layers				✓	✓	35.5%	55.2%
Two layers				✓	✓	36.6%	56.9%
One layer				✓	✓	34.1%	54.6%
Two layers				✓	✓	36.6%	56.9%

Table 4. Ablation analysis of the proposed progressive disentanglement. Here, we use mAP as the metric.

method accurately locates and classifies the three persons existing in the watercolor image. This further demonstrates the good performance of our method.

4.4. Visualization Analysis

In Fig. 7, taking two watercolor images as examples, a visualization analysis is made to show the learned disentangled representations. We can see both the method of only using the second disentangled layer and the progressive method could learn good disentangled representations. Particularly, compared with the ‘O-Base’ and ‘P-Base’ used for disentanglement, the learned DIR and DSR separately contain much stronger object-relevant information and domain-specific information. These results demonstrate that our method can successfully learn disentangled representations. Besides, compared with ‘O-Base’, ‘P-Base’ contains much

less domain-specific information, e.g., the background information in the first image and the color wall in the second image. This shows the first disentangled layer indeed enhances the domain-invariant information. Meanwhile, compared with ‘O-DIR’, our progressive method can extract a better DIR. Particularly, for these two images, ‘P-DIR’ is much smoother and contains much less domain-specific information. For example, the leaf and background information in the first image, and the flowers in the second image are much less in ‘P-DIR’, which is helpful for the location and recognition of objects. These all show our progressive method really owns the disentanglement ability and learns better instance-invariant features that lead to a better detection performance. More visualization examples can be seen in Fig. 8.

5. Conclusion

In this paper, we focus on obtaining the instance-invariant features for solving domain adaptive object detection. A progressive disentangled framework is first proposed to decompose domain-invariant and domain-specific features. Then, the instance-invariant features are extracted based on the domain-invariant features, which could alleviate the problem of domain-shift. Finally, we propose a three-stage training mechanism to enhance the disentanglement. In the experiment, our method achieves a new state-

of-the-art performance on three domain-shift scenes.

References

- [1] Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeswar, Sherjil Ozair, Yoshua Bengio, Aaron Courville, and R Devon Hjelm. Mine: mutual information neural estimation. *ICML*, 2018.
- [2] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.
- [3] Konstantinos Bousmalis, Nathan Silberman, David Dohan, Dumitru Erhan, and Dilip Krishnan. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3722–3731, 2017.
- [4] Qi Cai, Yingwei Pan, Chong-Wah Ngo, Xinmei Tian, Lingyu Duan, and Ting Yao. Exploring object relation in mean teacher for cross-domain detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11457–11466, 2019.
- [5] Yuhua Chen, Wen Li, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Domain adaptive faster r-cnn for object detection in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3339–3348, 2018.
- [6] Safa Cicek and Stefano Soatto. Unsupervised domain adaptation via regularized conditional alignment. *arXiv preprint arXiv:1905.10885*, 2019.
- [7] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016.
- [8] Kien Do and Truyen Tran. Theory and evaluation metrics for learning disentangled representations. *arXiv preprint arXiv:1908.09961*, 2019.
- [9] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010.
- [10] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. *arXiv preprint arXiv:1409.7495*, 2014.
- [11] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015.
- [12] Rui Gong, Wen Li, Yuhua Chen, and Luc Van Gool. Dlow: Domain flow for adaptation and generalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2477–2486, 2019.
- [13] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.
- [14] Zhenwei He and Lei Zhang. Multi-adversarial faster-rcnn for unrestricted object detection. *arXiv preprint arXiv:1907.10343*, 2019.
- [15] Han Hu, Jiayuan Gu, Zheng Zhang, Jifeng Dai, and Yichen Wei. Relation networks for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3588–3597, 2018.
- [16] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 172–189, 2018.
- [17] Naoto Inoue, Ryosuke Furuta, Toshihiko Yamasaki, and Kiyoharu Aizawa. Cross-domain weakly-supervised object detection through progressive domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5001–5009, 2018.
- [18] Bo Jiang, Ziyang Zhang, Doudou Lin, and Jin Tang. Graph learning-convolutional networks. *ICML*, 2019.
- [19] Mehran Khodabandeh, Arash Vahdat, Mani Ranjbar, and William G Macready. A robust learning approach to domain adaptive object detection. *ICCV*, 2019.
- [20] Seunghyeon Kim, Jaehoon Choi, Taekyung Kim, and Changick Kim. Self-training and adversarial background regularization for unsupervised domain adaptive one-stage object detection. *ArXiv*, abs/1909.00597, 2019.
- [21] Taekyung Kim, Minki Jeong, Seunghyeon Kim, Seokeon Choi, and Changick Kim. Diversify and match: A domain adaptive representation learning paradigm for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12456–12465, 2019.
- [22] Chen-Yu Lee, Tanmay Batra, Mohammad Haris Baig, and Daniel Ulbricht. Sliced wasserstein discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10285–10295, 2019.
- [23] Hsin-Ying Lee, Hung-Yu Tseng, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Diverse image-to-image translation via disentangled representations. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 35–51, 2018.
- [24] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [25] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [26] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016.
- [27] Yen-Cheng Liu, Yu-Ying Yeh, Tzu-Chien Fu, Sheng-De Wang, Wei-Chen Chiu, and Yu-Chiang Frank Wang. Detach and adapt: Learning cross-domain disentangled deep representation. In *Proceedings of the IEEE Conference on Com-*

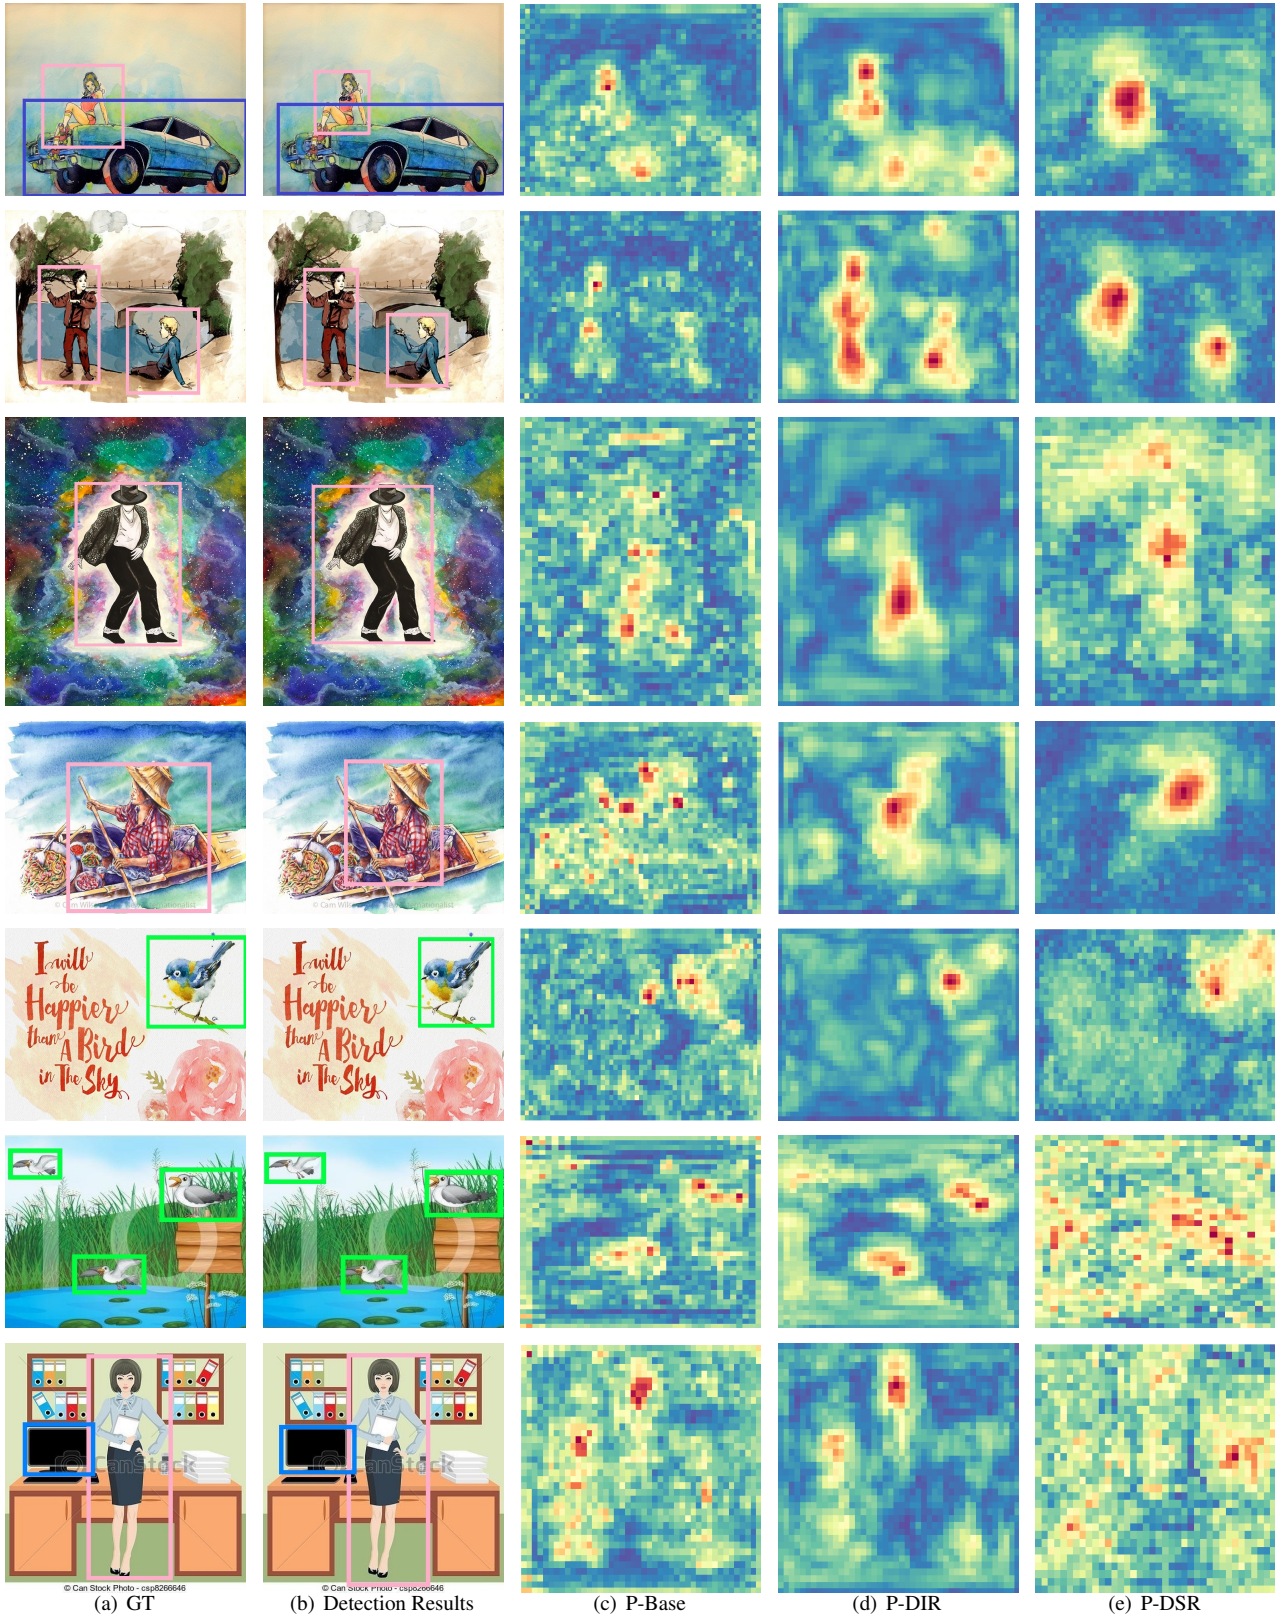


Figure 8. Visualization of feature maps of the second disentangled layer. Here, we use the progressive disentangled method to extract DIR and DSR. ‘Base’ indicates the feature map used for disentanglement. The examples of the first five rows are from the ‘Pascal VOC → Watercolor’ scene. The examples of the last two rows are from the ‘Pascal VOC → Clipart’ scene.

- puter Vision and Pattern Recognition, pages 8867–8876, 2018.
- [28] Francesco Locatello, Stefan Bauer, Mario Lucic, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. Challenging common assumptions in the unsupervised learning of disentangled representations. 2019.
- [29] Yingwei Pan, Ting Yao, Yehao Li, Yu Wang, Chong-Wah Ngo, and Tao Mei. Transferrable prototypical networks for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2239–2247, 2019.
- [30] Xingchao Peng, Zijun Huang, Ximeng Sun, and Kate Saenko. Domain agnostic learning with disentangled representations. *ICML*, 2019.
- [31] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.
- [32] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [33] Karl Ridgeway and Michael C Mozer. Learning deep disentangled embeddings with the f-statistic loss. In *Advances in Neural Information Processing Systems*, pages 185–194, 2018.
- [34] Subhankar Roy, Aliaksandr Siarohin, Enver Sangineto, Samuel Rota Buló, Nicu Sebe, and Elisa Ricci. Unsupervised domain adaptation using feature-whitening and consensus loss. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9471–9480, 2019.
- [35] Paolo Russo, Fabio M Carlucci, Tatiana Tommasi, and Barbara Caputo. From source to target and back: symmetric bi-directional adaptive gan. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8099–8108, 2018.
- [36] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. Strong-weak distribution alignment for adaptive object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6956–6965, 2019.
- [37] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Semantic foggy scene understanding with synthetic data. *International Journal of Computer Vision*, 126(9):973–992, 2018.
- [38] Tyler Scott, Karl Ridgeway, and Michael C Mozer. Adapted deep embeddings: A synthesis of methods for k-shot inductive transfer learning. In *Advances in Neural Information Processing Systems*, pages 76–85, 2018.
- [39] I. Sutskever, J. Martens, G. Dahl, and G. Hinton. On the importance of initialization and momentum in deep learning. In *International Conference on International Conference on Machine Learning*, 2013.
- [40] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.
- [41] Thang Vu, Hyunjun Jang, Trung X Pham, and Chang D Yoo. Cascade rpn: Delving into high-quality region proposal network with adaptive convolution. *arXiv preprint arXiv:1909.06720*, 2019.
- [42] Tao Wang, Xiaopeng Zhang, Li Yuan, and Jiashi Feng. Few-shot adaptive faster r-cnn. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7173–7182, 2019.
- [43] Rongchang Xie, Fei Yu, Jiachao Wang, Yizhou Wang, and Li Zhang. Multi-level domain adaptive learning for cross-domain detection. *arXiv preprint arXiv:1907.11484*, 2019.
- [44] Yabin Zhang, Hui Tang, Kui Jia, and Mingkui Tan. Domain-symmetric networks for adversarial domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5031–5040, 2019.
- [45] Xinge Zhu, Jiangmiao Pang, Ceyuan Yang, Jianping Shi, and Dahua Lin. Adapting object detectors via selective cross-domain alignment. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 687–696, 2019.