

Dynamic Weighted Learning for Unsupervised Domain Adaptation

Ni Xiao, Lei Zhang^(✉)

Learning Intelligence & Vision Essential (LiVE) Group

School of Microelectronics and Communication Engineering, Chongqing University, China

{nixiao, leizhang}@cqu.edu.cn

Abstract

Unsupervised domain adaptation (UDA) aims to improve the classification performance on an unlabeled target domain by leveraging information from a fully labeled source domain. Recent approaches explore domain-invariant and class-discriminant representations to tackle this task. These methods, however, ignore the interaction between domain alignment learning and class discrimination learning. As a result, the missing or inadequate tradeoff between domain alignment and class discrimination are prone to the problem of negative transfer. In this paper, we propose Dynamic Weighted Learning (DWL) to avoid the discriminability vanishing problem caused by excessive alignment learning and domain misalignment problem caused by excessive discriminant learning. Technically, DWL dynamically weights the learning losses of alignment and discriminability by introducing the degree of alignment and discriminability. Besides, the problem of sample imbalance across domains is first considered in our work, and we solve the problem by weighing the samples to guarantee information balance across domains. Extensive experiments demonstrate that DWL has an excellent performance in several benchmark datasets. Code is available at <https://github.com/NiXiao-cqu/TransferLearning-dwl-cvpr2021>.

1. Introduction

Collecting a large number of labeled samples is a time-consuming and laborious task. To tackle these problems, many domain adaptation algorithms have been proposed [18, 42, 28, 30, 17, 36]. Unsupervised/semi-supervised domain adaptation applies the model learned in the labeled source domain to the unlabeled/partially labeled target domain by alleviating the domain divergence.

In recent years, many domain adaptation methods combined with deep neural networks have achieved remarkable performance [23, 8, 2, 26, 19, 37], which explore deep models to reduce the domain shift between the training set

(source domain) and the testing set (target domain). The mainstream approaches tend to learn domain-invariant representations. While class discriminability is also critical to domain adaptation. Many domain adaptation methods that explore the feature discriminability have been proposed recently [5, 16, 3, 43, 28]. Specifically, Li et al. [16] investigates the discriminability by using a discriminability criterion based on Linear Discriminant Analysis (LDA). Deng et al. [5] forces the target structure to be discriminative by introducing a clustering loss. Chen et al. [3] guaranteed the discriminability by restricting the value of the eigenvalue. These methods are helpful to explore the discriminability of feature representations. However, *they independently consider the domain alignment and class discriminability of feature representations*. In fact, excessive alignment will easily lead to the loss of discriminability, while the excessive pursuit of discriminability will lead to the domain misalignment. Moreover, for different cross-domain data scenarios, the degree of alignment and discriminability is different. For data scenarios with a small distribution discrepancy between domains, the model should pay more attention to the learning of discriminability, while for tasks with a large distribution discrepancy, the model should pay more attention to the alignment of distributions. Furthermore, the previous methods did not consider the problem of the imbalanced sample size of the two domains. On the whole, the domain with a large sample size is equivalent to having larger weight in the process of optimizing the model. Such imbalance easily leads to model bias during training and leads to poor alignment and discrimination. As a result, it is easy to cause negative transfer.

To address the above challenges, on the one hand, we propose dynamically weighted learning between domain alignment and class discriminability. This dynamic learning idea is scalable to most previous DA models. On the other hand, we weight the samples according to the imbalance degree of the sample number of the two domains. Since the degree of alignment and discriminability will change with the iterative learning process, the degree of discriminability and alignment are monitored in real-time and fed back to the

model during training. Before that, to avoid the model bias during training, we first measure the degree of imbalance of the sample size of two domains universally and weight the samples according to the measurement results. The contributions are summarized as follows:

- In this paper, we propose a dynamic weighted learning method (DWL) for unsupervised domain adaptation. By monitoring the degree of alignment and discriminability in real-time, our method dynamically adjusts the weight of alignment learning and discriminability learning, so as to avoid excessive alignment or excessive pursuit of discriminability.
- We put forward a sample weighting method to avoid the model bias during training caused by the imbalanced sample size of the two domains.
- Through the dynamic weighting mechanism, our method can be more universal and applicable to cross-domain data scenarios with various feature types and sample sizes.

2. Related Work

Conventional domain adaptation methods [34, 20, 42, 6, 35, 44, 7] have achieved success to minimize the discrepancy across domains. With the development of deep networks, many deep methods show strong competitiveness in recognition accuracy. DDC [33] adopts a pre-trained deep network and introduces an adaptive layer based on the Maximum Mean Discrepancy (MMD). DANN [9] is inspired by the generative adversarial network (GAN) [12] and proposes an adversarial learning method. ADDA [32] proposes a general adversarial adaptation framework, which includes the base model, weights constraint, adversarial objective, and other factors. These methods achieve good performance in domain alignment learning. However, class discriminability is not considered sufficiently.

For discriminability, TPN [24] proposes to align three corresponding score distributions of each class from different domains. BSP [3] guaranteed discriminability by restricting the eigenvalue. ETD [15] exploits attention-aware transport distance and entropy-based regularization for discriminant information. The alignment and discriminability of features are considered in these methods. However, they ignore the interaction between domain alignment and class discriminability, and the weights of alignment learning and discriminability learning are not well controlled. Such an imbalance between alignment learning and discriminability learning tends to generate feature representations with excessive alignment or discriminability, and hard or false-easy feature representations.

3. Imbalance Between Alignment and Discrimination




Figure 1. Imbalance in alignment and discrimination of the two baselines DANN [9] and MCD [28]. The left vertical axis of the red curve represents the MMD distance, which is used to measure the degree of alignment across domains. The smaller value of MMD distance implies better alignment. The right vertical axis corresponding to the blue curve represents the degree of discriminability based on LDA. It is the ratio of the between-class scatter matrix to the within-class scatter matrix. The larger $\max J(\mathbf{W})$ implies the better discriminability.

DANN is a classic UDA method for domain alignment based on adversarial learning, which can make the distribution of the two domains aligned well. However, from the analysis results on task USPS→MNIST in Digits dataset as shown in Fig. 1 (a), with the training of the model, the domain alignment become better (red curve) while the class discriminability become worse (blue curve). Compared with DANN, as shown in Fig. 1 (b), MCD maintains the discriminability better. However, by observing Fig. 1 (b), we can find that the degree of class discriminability (blue curve) also shows a downward trend from epoch 25 although the MMD distance continues to decrease. This means that with the domain alignment learning, the class discriminability cannot be always promoted and even worse once excessive alignment (negative transfer) happens from epoch 25. However, both the domain alignment and class discrimination can affect the final classification accuracy. Therefore, it is important to propose a reasonable tradeoff between alignment learning and discriminability learning with dynamic control, that is, none of them can cause the other to be worse in their learning process.

4. Theoretical Insight

In the UDA task, n_s labeled samples $\{(x_i^s, y_i^s)\}_{i=1}^{n_s}$ from source domain \mathcal{D}_s and n_t unlabeled samples $\{x_j^t\}_{j=n_s+1}^{n_s+n_t}$ from target domain \mathcal{D}_t are given under the assumption that they are subject to different marginal and conditional distributions. The theory proposed by Ben-David et al. [1] provides an upper bound of the expected error on the target samples with regard to the classifier trained on labeled source data. The upper bound of the expected error on the target samples is mainly relative to three terms: (i) the ex-

pected error on the source domain $\epsilon_S(h)$; (ii) the $\mathcal{H}\Delta\mathcal{H}$ -divergence term between two domains $d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}_s, \mathcal{D}_t)$; (iii) the combined error λ of the ideal joint hypothesis h^* . It is defined as follows,

$$\epsilon_T(h) \leq \epsilon_S(h) + \frac{1}{2}d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}_s, \mathcal{D}_t) + \lambda \quad (1)$$

where the definition of $\mathcal{H}\Delta\mathcal{H}$ -divergence in the hypothesis space \mathcal{H} is:

$$d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}_s, \mathcal{D}_t) = 2 \sup_{h, h' \in \mathcal{H}} |\Pr_{x \sim \mathcal{D}_s}(h(x) \neq h'(x)) - \Pr_{x \sim \mathcal{D}_t}(h(x) \neq h'(x))| \quad (2)$$

and the combined error λ of the ideal joint hypothesis is

$$\lambda = \epsilon_S(h^*) + \epsilon_T(h^*) \quad (3)$$

where $h^* = \arg \min_{h \in \mathcal{H}} \epsilon_S(h) + \epsilon_T(h)$.

In Inequality (1), the first term, $\epsilon_S(h)$, is expected to be small because reliable labels are owned in the source domain. As for the second term, several domain adaptation methods [8, 29, 11, 21, 31] can make the domain discrepancy $d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}_s, \mathcal{D}_t)$ small by alignment learning. The third term, λ , is considered sufficiently small in most adversarial domain adaptation methods.

Potential problem. The goal of the UDA model is to limit the upper bound of the expected error on the target samples to a small value. To achieve this goal, the designed model needs to keep all three terms of the above to a small value. In other words, if one of the three terms goes down and another term goes up, then this model can not effectively reduce the upper bound of $\epsilon_T(h)$. When there is an excessive alignment, $d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}_s, \mathcal{D}_t)$ can be enabled to be smaller, however, the discriminability will be easily reduced and the data will become inseparable, which will make λ to be larger. Therefore, the value of λ is closely related with the class discriminability.

Rationality of dynamic balance. In our work, by the dynamic weighting of domain alignment learning and class discrimination learning, we make $d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}_s, \mathcal{D}_t)$ and λ decrease synchronously. That is, in our method, the $d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}_s, \mathcal{D}_t)$ is further reduced under the premise that this operation does not increase λ , and similarly, λ is further reduced under the premise that this operation does not increase $d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}_s, \mathcal{D}_t)$. For achieving that, we first propose a dynamically mutual interaction idea between $\mathcal{H}\Delta\mathcal{H}$ -divergence and the combined error λ , rather than only independent training in previous models. In this way, the upper bound of $\epsilon_T(h)$ can be effectively reduced in our work.

5. Approach

In this section, we provide the details of the proposed DWL. The overall architecture of DWL is depicted in Fig. 2.

We first weight the samples of the two domains to avoid the model bias during training caused by the imbalanced sample size of the two domains. Finally, the dynamic weighted learning with a balance factor τ between domain alignment and class discrimination is presented. The balance factor is estimated according to the naturally designed degree of alignment and discriminability.

5.1. Sample Weighting

In the problem of UDA, when the sample size of one domain is greatly larger than another domain, on the whole, the domain with large sample size is equivalent to having larger weight during training. Such imbalance easily leads to model bias during training and leads to poor alignment and discrimination, which tends to produce a negative transfer. To avoid the imbalance problem of sample sizes of two domains, we propose to intuitively weight the samples of the two domains in our method. The weights of the samples in each domain are inversely proportional to their proportion in the total sample size of the two domains. Specifically, we weight the samples of each domain as follows:

$$\hat{x}_i^s = a \left(1 + \frac{n_t}{n_s} \right) x_i^s \quad , \quad i = 1, 2, \dots, n_s \quad (4)$$

$$\hat{x}_j^t = a \left(1 + \frac{n_s}{n_t} \right) x_j^t \quad , \quad j = 1, 2, \dots, n_t \quad (5)$$

where $a \in (0, 1]$ is a hyper-parameter that controls the degree of sample weighting.

5.2. Domain Alignment Learning and Class Discrimination Learning

Adversarial learning has been successfully introduced to domain alignment by learning domain-invariant feature representations. In order to obtain domain-invariant feature representations, the input weighted samples \hat{x}_s and \hat{x}_t are embedded by feature generator G . The parameters θ_g of G and the parameters θ_d of domain discriminator D are trained by optimizing the following standard minimax domain alignment loss:

$$\min_{\theta_g} \max_{\theta_d} \mathcal{L}_{da}(\theta_g, \theta_d) = \mathbb{E}_{x_i^s \sim \mathcal{D}_s} \log [D(G(\hat{x}_i^s))] + \mathbb{E}_{x_j^t \sim \mathcal{D}_t} \log [1 - D(G(\hat{x}_j^t))] \quad (6)$$

Adversarial learning is effective to achieve domain alignment. However, class discriminability can not be ensured. In order to obtain the feature representations which have good discriminability, we are inspired by MCD [28] which maximizes the discrepancy of two classifiers to obtain strong discriminant features. We propose to optimize

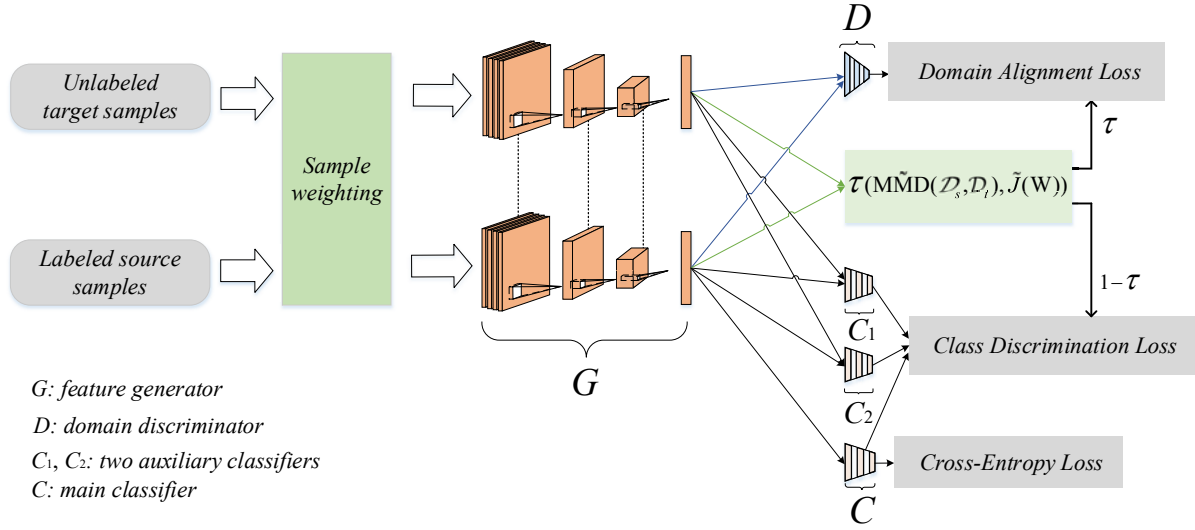


Figure 2. The overall structure of the proposed DWL. We separate the network into three modules: feature generator (G), domain discriminator (D), classifiers (C_1, C_2, C) and with associated parameters $\theta_g, \theta_d, (\theta_{c_1}, \theta_{c_2}, \theta_c)$. τ is a balance factor constructed from two metrics, i.e., $\text{MMD}(\mathcal{D}_s, \mathcal{D}_t)$ and $\max J(\mathbf{W})$. $\text{MMD}(\mathcal{D}_s, \mathcal{D}_t)$ is used to measure the degree of domain alignment and $\max J(\mathbf{W})$ is used to measure the degree of discrimination. $\text{MMD}(\mathcal{D}_s, \mathcal{D}_t)$ and $\tilde{J}(\mathbf{W})$ represents normalized $\text{MMD}(\mathcal{D}_s, \mathcal{D}_t)$ and $\max J(\mathbf{W})$, resp.

the following class discrimination loss :

$$\begin{aligned}
 \min_{\theta_g, \theta_c} \max_{\theta_{c_1}, \theta_{c_2}} \mathcal{L}_{cd}(\theta_g, \theta_c, \theta_{c_1}, \theta_{c_2}) \\
 = \mathbb{E}_{x_j^t \sim \mathcal{D}_t} \|C_1(G(\hat{x}_j^t)) - C_2(G(\hat{x}_j^t))\|_1 \\
 + \|C(G(\hat{x}_j^t)) - C_1(G(\hat{x}_j^t))\|_1 \\
 + \|C(G(\hat{x}_j^t)) - C_2(G(\hat{x}_j^t))\|_1
 \end{aligned} \quad (7)$$

where C, C_1 , and C_2 are three classifiers which are pre-trained by supervised learning in the source domain. By maximizing the discrepancy between C_1 and C_2 in the target domain (fix G and C), C_1 and C_2 can detect the target samples excluded by the support vectors of the source. Then by training the G to minimize the discrepancy (fix C_1 and C_2), the obtained target features can have strong discriminability. Different from [28], we added a main classifier C , whose decision hyperplane is between C_1 and C_2 , to make the distance between the support vectors and the decision boundary larger. $\theta_{c_1}/\theta_{c_2}/\theta_c$ denotes the parameters of the auxiliary classifiers $C_1/C_2/C$, resp.

5.3. Dynamic Weighted Learning

In order to avoid the excessive pursuit of alignment or discriminability independently, and make both improved simultaneously towards a good direction, we propose to measure the degree of alignment and discriminability in each iteration in real-time and then construct a dynamic balance factor to control the weight of domain alignment loss and class discriminability loss. Maximum Mean Discrepancy (MMD) and Linear Discriminant Analysis (LDA) are re-

spectively used to measure the degree of alignment and discriminability of the current feature representations across domains. MMD is a popular estimator to calculate the degree of alignment of data distribution between two domains, and it is defined as follows,

$$\text{MMD}(\mathcal{D}_s, \mathcal{D}_t) = \left\| \mathbb{E}_{x_i^s \sim \mathcal{D}_s} G(\hat{x}_i^s) - \mathbb{E}_{x_j^t \sim \mathcal{D}_t} G(\hat{x}_j^t) \right\|^2 \quad (8)$$

The discriminability estimator $\max J(\mathbf{W})$ based on LDA is defined as follows,

$$\max_{\mathbf{W}} J(\mathbf{W}) = \frac{\text{tr}(\mathbf{W}^\top \mathbf{S}_b \mathbf{W})}{\text{tr}(\mathbf{W}^\top \mathbf{S}_w \mathbf{W})} \quad (9)$$

where \mathbf{S}_b is the between-class scatter matrix and \mathbf{S}_w is the within-class scatter matrix [3]. Clearly, the larger $\max J(\mathbf{W})$ implies better discriminability.

Since the estimated values of the two evaluation criteria are usually not in the same order of magnitude, we further normalize the estimated values. To normalize the two evaluation values more reasonably and apply them to construct the dynamic balancing factor, we apply min-max scaling to linearly transform the evaluation values and map the results to the range of [0,1]. We define $\tilde{\text{MMD}}(\mathcal{D}_s, \mathcal{D}_t)$ and $\tilde{J}(\mathbf{W})$ to represent the normalized $\text{MMD}(\mathcal{D}_s, \mathcal{D}_t)$ and $\max J(\mathbf{W})$, resp.

$$\begin{aligned}
 \tilde{\text{MMD}}(\mathcal{D}_s, \mathcal{D}_t) \\
 = \frac{\text{MMD}(\mathcal{D}_s, \mathcal{D}_t) - \text{MMD}(\mathcal{D}_s, \mathcal{D}_t)_{\min}}{\text{MMD}(\mathcal{D}_s, \mathcal{D}_t)_{\max} - \text{MMD}(\mathcal{D}_s, \mathcal{D}_t)_{\min}} \quad (10)
 \end{aligned}$$

$$\tilde{J}(\mathbf{W}) = \frac{J(\mathbf{W}) - J(\mathbf{W})_{\min}}{J(\mathbf{W})_{\max} - J(\mathbf{W})_{\min}} \quad (11)$$

According to the normalized estimators $\tilde{\text{MMD}}(\mathcal{D}_s, \mathcal{D}_t) \in [0, 1]$ and $\tilde{J}(\mathbf{W}) \in [0, 1]$ on domain alignment and class discriminability in Eq. (10) and Eq. (11), we construct a dynamic balance factor as follows:

$$\tau = \frac{\tilde{\text{MMD}}(\mathcal{D}_s, \mathcal{D}_t)}{\tilde{\text{MMD}}(\mathcal{D}_s, \mathcal{D}_t) + (1 - \tilde{J}(\mathbf{W}))} \quad (12)$$

In Eq. (12), smaller $\tilde{\text{MMD}}(\mathcal{D}_s, \mathcal{D}_t)$ indicates better domain alignment and smaller $1 - \tilde{J}(\mathbf{W})$ indicates better class discriminability. When the degree of domain alignment is far better than the class discriminability, the $\tilde{\text{MMD}}(\mathcal{D}_s, \mathcal{D}_t)$ approaches 0, the $1 - \tilde{J}(\mathbf{W})$ approaches 1, and the τ approaches 0. When the degree of domain alignment is far worse than the degree of class discrimination, the $\tilde{\text{MMD}}(\mathcal{D}_s, \mathcal{D}_t)$ approaches 1, the $1 - \tilde{J}(\mathbf{W})$ approaches 0, and the τ approaches 1. Notably, the τ approximates to 0.5 when the degree of domain alignment is equal to the degree of class discriminability. Based on the above good properties of τ , we adopt τ as the weight of the domain alignment loss and $1 - \tau$ as the weight of the class discrimination loss. Therefore, the dynamic weighting model of domain alignment and class discrimination is obtained as follows:

$$\min_{\theta_g, \theta_c} \max_{\theta_d, \theta_{c_1}, \theta_{c_2}} \tau \cdot \mathcal{L}_{da}(\theta_g, \theta_d) + (1 - \tau) \cdot \mathcal{L}_{cd}(\theta_g, \theta_c, \theta_{c_1}, \theta_{c_2}) \quad (13)$$

Physically, based on the good properties of τ , when the effectiveness of domain alignment learning is far worse than that of class discrimination learning, the model increases the weight of domain alignment learning. On the contrary, when the learning effect of class discrimination learning is far worse than that of domain alignment learning, the model increases the weight of class discrimination learning. Under this dynamic weighted learning mechanism, our model can maintain the consistency between domain alignment learning and class discrimination learning, thus avoiding excessive domain alignment or class discriminability.

5.4. Overall Training Objective

The overall training objective of our DWL integrates sample weighting, domain alignment learning, class discrimination learning, and the dynamic weighting learning. Besides, we also need to minimize the expected source error $\epsilon_S(h)$ for the labeled source samples. Hence, we propose to solve the following ultimate minimax objective:

$$\min_{\theta_g, \theta_c} \max_{\theta_d, \theta_{c_1}, \theta_{c_2}} \sum_{i=1}^{n_s} \mathcal{L}_{ce}(C(G(x_i^s; \theta_g); \theta_c), y_i^s) + \tau \cdot \mathcal{L}_{da}(\theta_g, \theta_d) + (1 - \tau) \cdot \mathcal{L}_{cd}(\theta_g, \theta_c, \theta_{c_1}, \theta_{c_2}) \quad (14)$$

where \mathcal{L}_{ce} is the standard cross-entropy loss. Through the dynamic weighted learning in Eq. (14), it is natural and effective to avoid discriminability vanishing caused by excessive alignment or domain misalignment caused by excessive discriminability pursuit. Additionally, the sample weighting can alleviate the model bias towards the domain of large sample size.

6. Experiments

6.1. Datasets

VisDA-2017 [25] is a synthetic-to-real image dataset with two domains: synthetic images and real images. We take the synthetic images as the source domain and the real images as the target domain.

Digits Datasets. We construct the adaptation tasks of digits among three datasets MNIST [39], USPS, and SVHN [41]. The MNIST (M) and USPS (U) image datasets are both handwritten digits datasets containing 10 classes of digits. The SVHN (S) is a real-world digits dataset of house numbers in google street view images. We conduct experiments in 3 directions: $M \rightarrow U$, $U \rightarrow M$, and $S \rightarrow M$.

Office-31 [27] contains 4,652 images across 31 classes from three domains: Amazon (A), DSLR (D), and Webcam (W). We conduct experiments in all 6 transfer tasks.

ImageCLEF-DA [21] contains 12 common classes shared by 3 domains: Caltech256 (C), ImageNet ILSVR-C 2012 (I), and Pascal VOC 2012 (P). We conduct experiments in all 6 transfer tasks.

6.2. Implementation Details

We compare our method with the following state-of-the-art domain adaptation methods: **DAN** (Deep Adaptation Networks) [18], **DANN** (Domain-adversarial Neural Networks) [9], **DRCN** (Deep Reconstruction Classification Networks) [10], **GoGAN** (Coupled Generative Adversarial Networks) [22], **ADDA** (Adversarial Discriminative Domain Adaptation) [32], **CyCADA** (Cycle-Consistent Adversarial Domain Adaptation) [14], **CDAN** (Conditional Adversarial Adaptation Networks) [19], **MCD** (Maximum Classifier Discrepancy) [28], **CAT** (Cluster Alignment with a Teacher) [5], **HAFN** (Hard Adaptive Feature Norm) [38], **TPN** (Transferrable Prototypical Networks) [24], **JAN** (Joint Adaptation Networks) [21], **BSP** (Batch Spectral Penalization) [3], **LWC** (Light-weight Calibrator) [40], and **ETD** (Enhanced Transport Distance) [15].

Following the standard protocols for unsupervised domain adaptation, all labeled source samples and unlabeled target samples participate in the training stage. For the Digits tasks, we follow the protocol in [28]. We use 2k images from MNIST and 1.8k images from USPS in the transfer between MNIST and USPS and use the whole training sets for the adaptation from SVHN to MNIST. In this experiment,

Table 1. Accuracy (%) on VisDA-2017 for unsupervised domain adaptation (ResNet-101).

Method	plane	bycl	bus	car	horse	knife	mcycl	person	plant	sktbrd	train	truck	mean
ResNet101	55.1	53.3	61.9	59.1	80.6	17.9	79.7	31.2	81.0	26.5	73.5	8.5	52.4
DAN [18]	87.1	63.0	76.5	42.0	<u>90.3</u>	42.9	85.9	53.1	49.7	36.3	85.8	20.7	61.1
DANN [9]	81.9	<u>77.7</u>	82.8	44.3	81.2	29.5	65.1	28.6	51.9	54.6	82.8	7.8	57.4
MCD [28]	87.0	60.9	<u>83.7</u>	<u>64.0</u>	88.9	79.6	84.7	76.9	<u>88.6</u>	40.3	83.0	25.8	71.9
BSP [3]	92.4	61.0	81.0	57.5	89.0	<u>80.6</u>	90.1	<u>77.0</u>	84.2	77.9	82.1	38.4	<u>75.9</u>
Ours	<u>90.7</u>	80.2	86.1	67.6	92.4	81.5	<u>86.8</u>	78.0	90.6	<u>57.1</u>	<u>85.6</u>	<u>28.7</u>	77.1

the CNN architecture is a modified version of [28]. For optimization, the network weights are trained by ADAM with 0.0005 weight decay. The learning rate is set as 0.0002 and the mini-batch size is 128. We stop training after 200 epochs and adopt the classification accuracy on the target domain as an evaluation metric. For other image datasets, we use PyTorch to implement our method. we adopt ResNet [13] pre-trained on ImageNet [4] to generate the original feature representations with parameters fine-tuned by our method. The classifiers we adopted in the experiment are two-layer network ($2048 \times 1024 \times \#classes$), and the domain discriminator consists of two layers with ReLU and Dropout (0.5) in all the layers. Specifically, we use the mini-batch SGD with momentum 0.9 and the mini-batch size is 8. The learning rate is set as 0.001.

6.3. Experimental Results

In this section, we conduct extensive experiments to evaluate the DWL. All baseline results are taken from related literature. Our algorithm outperforms many SOTA methods on different datasets, which shows the efficiency and universality of our method.

Results on VisDA-2017 are reported in Table 1. The average accuracy of our model achieves 77.1%, which is higher than 75.9% of the BSP. Notably, BSP guaranteed discriminability by restricting the value of the eigenvalue. However, BSP treats transferability and discriminability to be equally important, which is not always right. In our method, we estimate the transferability and discriminability of the current cross-domain tasks in the iterative process in real-time and dynamically weight the two learning functions to keep them within the effective learning state (i.e., both are promoted towards a good direction). Therefore, our method can be integrated into BSP to enhance its universality and improve recognition accuracy.

Results on Digits are reported in Table 2. Our model achieves 97.3% and 97.4% on tasks of MNIST→USPS and USPS→MNIST, resp., which outperforms the state-of-the-arts. Our proposed method focuses on the balance of domain alignment learning and class discrimination learning, which outperforms CAT and MCD paying more attention to discriminability.

Results on Office-31 are reported in Table 3. On the d-

Table 2. Acc (%) on Digits for unsupervised domain adaptation.

Method	M→U	U→M	S→M	Average
DAN [18]	80.3	77.8	73.5	77.2
DRCN [10]	91.8	73.7	82.0	82.5
CoGAN [22]	91.2	89.1	-	-
ADDA [32]	89.4	90.1	76.0	85.2
CyCADA [14]	95.6	96.5	90.4	94.2
CDAN [19]	93.9	96.9	88.5	93.1
MCD [28]	94.2	94.1	96.2	94.8
CAT [5]	90.6	80.9	98.1	89.9
TPN [24]	92.1	94.1	93.0	93.1
LWC [40]	95.6	<u>97.1</u>	97.1	96.6
ETD [15]	<u>96.4</u>	96.3	<u>97.9</u>	<u>96.9</u>
Ours	97.3	97.4	98.1	97.6

ifficult tasks D→A and W→A where the source domain is small and the domain shift is large, resp., our model can still achieve 73.1% and 69.8% which are higher than 71.0% and 67.8% of the optimal transport model ETD. Notably, ETD exploits attention-aware transport distance and entropy-based regularization for domain alignment and discriminant information, resp. The final classification accuracy of ETD depends more on the initial distribution state of two domains. Compared with ETD, by using sample weighting and real-time state estimation of sample distribution, our DWL can be applied to datasets of different initial states, which means the better generalization of our model.

Results on ImageCLEF-DA are reported in Table 4. Our method is better than the CAT, which explores the class-conditional structure for the feature space. This indicates that class-conditional structure can be maintained in our method. It is worth noting that our method can be easily plugged and played in the existing UDA methods to enhance their universality.

6.4. Experimental Analysis

Visualization of feature space. Fig. 3 depicts the t-SNE visualizations of features learned in our method on the MNIST to USPS. With the increase of training epochs from (a) to (e), we can observe that feature distribution at epoch 150 has shown good clustering effect. The learned features

Table 3. Accuracy (%) on Office-31 for unsupervised domain adaptation (ResNet-50).

Method	A→W	D→W	W→D	A→D	D→A	W→A	Average
ResNet50 [13]	68.4	96.7	99.3	68.9	62.5	60.7	76.1
DAN [18]	80.5	97.1	99.6	78.6	63.6	62.8	80.4
DANN [9]	82.6	96.9	99.3	81.5	68.4	67.5	82.7
ADDA [32]	86.2	96.2	98.4	77.8	69.5	68.9	82.9
CAT [5]	91.1	98.6	99.6	90.6	70.4	66.5	86.1
ETD[15]	92.1	100.0	100.0	88.0	71.0	67.8	86.2
Ours	89.2	99.2	100.0	91.2	73.1	69.8	87.1

Table 4. Accuracy (%) on ImageCLEF-DA for unsupervised domain adaptation (ResNet-50).

Method	I→P	P→I	I→C	C→I	C→P	P→C	Average
ResNet50 [13]	74.8	83.9	91.5	78.0	65.5	91.2	80.7
DAN [18]	74.5	82.2	92.8	86.3	69.2	89.8	82.5
DANN [9]	75.0	86.0	96.2	87.0	74.3	91.5	85.0
JAN [21]	76.8	88.0	94.7	89.5	74.2	91.7	85.8
HAFN [38]	76.9	89.0	94.4	89.6	74.9	92.9	86.3
CAT [5]	76.7	89.0	94.5	89.8	74.0	93.7	86.3
ETD[15]	81.0	91.7	97.9	93.3	79.5	95.0	89.7
Ours	82.3	94.8	98.1	92.8	77.9	97.2	90.5




Figure 3. The t-SNE visualizations of features generated by the proposed DWL with the increase of the epoch on MNIST → USPS. Red and blue points indicate the source and target samples, resp.

align the source and target domain samples with 10 clusters with clear boundaries.

Convergence analysis. The classification error and the value of balance factor τ in our DWL on task USPS→MNIST (left) and MNIST→USPS (right) are shown in Fig. 4. For each subfigure, the left axis of the red curve represents the classification error and the right axis of the blue curve represents the value of balancing factor τ . It can be found that both of them converge to a flat value gradually with iteration. This means that with the decrease of τ , the class discriminability is emphasised such that the classification error decreases too. In the iteration process, when the change of τ is relatively obvious, the improvement of recognition accuracy is also relatively obvious. We set the initial value of τ as 0.5, and it can be found that τ drops sharply to less than 0.5 in the first epoch, indicating that the model has relatively good alignment but relatively poor discriminability. Indeed, we know that the cross-domain digit

dataset has small distribution difference but relatively poor discriminability, and therefore the results comply with an empirical observation.




Figure 4. Convergence analysis of DWL on the classification error (%) and the dynamic balancing factor τ .

Confusion matrix visualization. Fig. 5 (a) to (d) show the visualizations of confusion matrix for the classifier trained by Source-only and our DWL on tasks $c \rightarrow p$ and $p \rightarrow i$ from ImageCLEF-DA. The Source-only means the classifier is trained only on labeled source data. Hence,

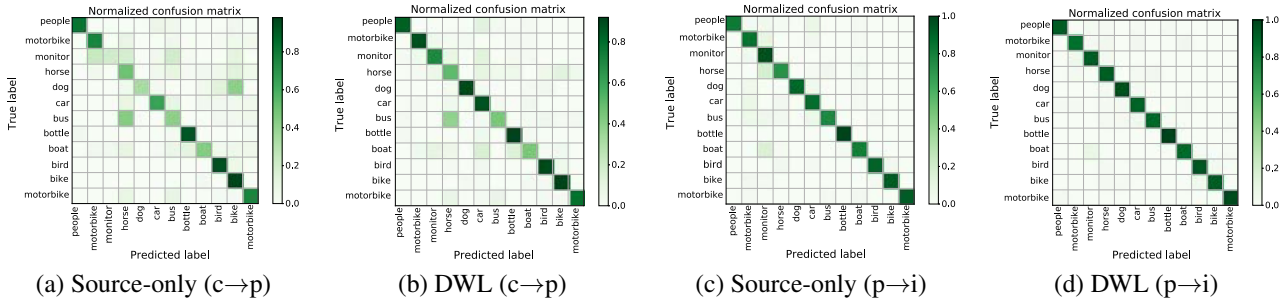


Figure 5. The Confusion Matrix (CM) visualization for proposed DWL.




Figure 6. Effectiveness analysis of DWL on alignment learning and discrimination learning.

the experimental results of Source-only can reveal the impact of domain discrepancy. As can be seen in Fig. 5, the confusion is reduced for most classes through domain adaptation by our method. Comparing with Fig. 5 (a) and (b), the dog is misclassified as the bike for the Source-only and corrected in DWL. In $p \rightarrow i$, the confusing horse is correctly recognized in DWL. The confusion matrices further reveal the discriminability of our method on the target domain.

Analysis of alignment and discriminability degrees.

The MMD distance and the $\max J(W)$ of the feature representation learned in our DWL on task USPS→MNIST are shown in Fig. 6. It can be found that by dynamically weighting the alignment learning and discrimination learning, we make a better balance between them, which makes the two kinds of learning in an effective learning state and both of them can promote each other towards a good direction (i.e., the domain distance is smaller and the discriminability degree is larger).

Ablation study. In this part, we analyze the influence of sample weighting and balancing factor τ on improving performance. The results are shown in Table 5 and Table 6. To analyze the role of sample weighting, the cross-domain task $W \rightarrow A$ with large differences across domains in sample size is chosen for the experiment (The domain W contains 795 samples, and domain A contains 2817 samples). The experimental results are shown in Table 5. It can be found that the accuracy of the model can be significantly improved by introducing the sample weighting and dynamic balancing factor τ . Furthermore, we test the results of static weighting

Table 5. The accuracy (%) of ablation experiment.

Sample weighting	Balancing factor τ	Accuracy (%)
x	x	67.9
✓	x	68.3
x	✓	69.1
✓	✓	69.8

Table 6. Accuracy (%) of different weight between \mathcal{L}_{da} and \mathcal{L}_{cd}

Task	1:1	1:9	9:1	7:3	3:7	$\tau : 1 - \tau$
$M \rightarrow U$	93.8	96.2	94.6	94.5	94.9	97.3
$U \rightarrow M$	96.8	95.6	96.8	95.6	95.6	97.4

mechanism for domain alignment learning \mathcal{L}_{da} and class discrimination learning \mathcal{L}_{cd} on Digits dataset. As shown in Table 6, it can be found that the proposed dynamic weighting factor τ achieves more excellent performance than the static weighting with different ratios. Our original intention for promoting the dynamic balance between domain alignment and class discrimination is confirmed.

7. Conclusions

In this paper, we analyze the interaction between alignment learning and discrimination learning and then propose that the weight of alignment learning and discrimination learning should be better dynamically balanced. For a more generalized model to accommodate different datasets, we propose the Dynamic Weighted Learning (DWL) method to dynamically adjust the weights of domain alignment learning and class discrimination learning to adapt to different cross-domain classification scenarios. Besides, the problem of model bias during training which is caused by the imbalanced sample size of the two domains is also considered in this paper, and we propose a simple but effective sample weighting mechanism to solve it. Extensive experiments demonstrate that the dynamic learning idea between the $\mathcal{H}\Delta\mathcal{H}$ -divergence and the combined error λ is useful and scalable in DA.

References

- [1] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. In *Machine learning*, volume 79, pages 151–175, 2010. [2](#)
- [2] Konstantinos Bousmalis, George Trigeorgis, Nathan Silberman, Dilip Krishnan, and Dumitru Erhan. Domain separation networks. In *International Conference on Neural Information Processing Systems (NIPS)*, pages 343–351, 2016. [1](#)
- [3] Xinyang Chen, Sinan Wang, Mingsheng Long, and Jianmin Wang. Transferability vs. discriminability: Batch spectral penalization for adversarial domain adaptation. In *International Conference on Machine Learning (ICML)*, pages 1081–1090, 2019. [1](#), [2](#), [4](#), [5](#), [6](#)
- [4] J. Deng, W. Dong, R. Socher, L. Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255, 2009. [6](#)
- [5] Zhijie Deng, Yucen Luo, and Jun Zhu. Cluster alignment with a teacher for unsupervised domain adaptation. In *IEEE International Conference on Computer Vision (ICCV)*, pages 9943–9952, 2019. [1](#), [5](#), [6](#), [7](#)
- [6] Z. Ding and Y. Fu. Robust transfer metric learning for image classification. *IEEE Transactions on Image Processing*, 26(2):660–670, 2017. [2](#)
- [7] Basura Fernando, Amaury Habrard, Marc Sebban, and Tinne Tuytelaars. Unsupervised visual domain adaptation using subspace alignment. In *IEEE International Conference on Computer Vision (ICCV)*, pages 2960–2967, 2013. [2](#)
- [8] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International Conference on Machine Learning (ICML)*, pages 1180–1189, 2015. [1](#), [3](#)
- [9] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, Francois Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *Journal of Machine Learning Research*, 17(1):2096–2030, 2017. [2](#), [5](#), [6](#), [7](#)
- [10] Muhammad Ghifary, W. Bastiaan Kleijn, Mengjie Zhang, David Balduzzi, and Wen Li. Deep reconstruction-classification networks for unsupervised domain adaptation. In *European Conference on Computer Vision (ECCV)*, pages 597–613. Springer, 2016. [5](#), [6](#)
- [11] M. Gong, K. Zhang, T. Liu, D. Tao, C Glymour, and B Scholkopf. Domain adaptation with conditional transferable components. In *International Conference on Machine Learning (ICML)*, pages 2839–2848, 2016. [3](#)
- [12] I.J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *International Conference on Neural Information Processing Systems (NIPS)*, pages 2672–2680. [2](#)
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Sun Jian. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. [6](#), [7](#)
- [14] Judy Hoffman, Eric Tzeng, Taesung Park, Jun Yan Zhu, Phillip Isola, Kate Saenko, Alexei A. Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *International Conference on Machine Learning (ICML)*, pages 1994–2003, 2018. [5](#), [6](#)
- [15] Mengxue Li, Yi Ming Zhai, You Wei Luo, Peng Fei Ge, and Chuan Xian Ren. Enhanced transport distance for unsupervised domain adaptation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13933–13941, 2020. [2](#), [5](#), [6](#), [7](#)
- [16] Shuang Li, Shiji Song, and Gao Huang. Domain invariant and class discriminative feature learning for visual domain adaptation. *IEEE transactions on image processing*, 27(9):4260–4273, 2018. [1](#)
- [17] Hong Liu, Mingsheng Long, Jianmin Wang, and Michael I Jordan. Transferable adversarial training: A general approach to adapting deep classifiers. In *International Conference on Machine Learning (ICML)*, June 2019. [1](#)
- [18] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael I. Jordan. Learning transferable features with deep adaptation networks. In *International Conference on Machine Learning (ICML)*, pages 97–105, 2015. [1](#), [5](#), [6](#), [7](#)
- [19] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I. Jordan. Conditional adversarial domain adaptation. In *International Conference on Neural Information Processing Systems (NIPS)*, pages 1647–1657, 2018. [1](#), [5](#), [6](#)
- [20] Mingsheng Long, Jianmin Wang, Guiguang Ding, Jianguang Sun, and Philip S. Yu. Transfer feature learning with joint distribution adaptation. In *IEEE International Conference on Computer Vision (ICCV)*, pages 2200–2207, 2014. [2](#)
- [21] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Deep transfer learning with joint adaptation networks. In *International Conference on Machine Learning (ICML)*, pages 2208–2217, 2017. [3](#), [5](#), [7](#)
- [22] Liu Ming-Yu and Tuzel Oncel. Coupled generative adversarial networks. In *International Conference on Neural Information Processing Systems (NIPS)*, pages 469–477, 2016. [5](#), [6](#)
- [23] Pietro Morerio, Riccardo Volpi, Ruggero Ragonesi, and Vittorio Murino. Generative pseudo-label refinement for unsupervised domain adaptation. In *Winter Conference on Applications of Computer Vision (WACV)*, pages 3119–3128, 2020. [1](#)
- [24] Yingwei pan, Ting Yao, Yehao Li, Wang Yu, Chong-Wah Ngo, and Tao Mei. Transferrable prototypical networks for unsupervised domain adaptation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2234–2242, 2019. [2](#), [5](#), [6](#)
- [25] Xingchao Peng, Ben Usman, Neela Kaushik, Judy Hoffman, Dequan Wang, and Kate Saenko. Visda: The visual domain adaptation challenge. [5](#)
- [26] Sanjay Purushotham, Wilka Carvalho, Tanachat Nilanon, and Yan Liu. Variational recurrent adversarial deep domain adaptation. In *International Conference on Learning Representations (ICLR)*, 2017. [1](#)
- [27] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *European Conference on Computer Vision (ECCV)*, pages 213–226. Springer, 2010. [5](#)

- [28] Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3723–3732, 2018. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#)
- [29] Swami Sankaranarayanan, Yogesh Balaji, Carlos D. Castillo, and Rama Chellappa. Generate to adapt: Aligning domains using generative adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8503–8512, 2018. [3](#)
- [30] Jong-Chyi Su, Yi-Hsuan Tsai, Kihyuk Sohn, Buyu Liu, Subhansu Maji, and Manmohan Chandraker. Active adversarial domain adaptation. In *Winter Conference on Applications of Computer Vision (WACV)*, pages 728–737, 2020. [1](#)
- [31] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *ECCV Workshops*, pages 443–450, 2016. [3](#)
- [32] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell. Adversarial discriminative domain adaptation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2962–2971, 2017. [2](#), [5](#), [6](#), [7](#)
- [33] Eric Tzeng, Judy Hoffman, Ning Zhang, Kate Saenko, and Trevor Darrell. Deep domain confusion: Maximizing for domain invariance. *Computer Science*, 2014. [2](#)
- [34] H. Wang, W. Wang, C. Zhang, and F. Xu. Cross-domain metric learning based on information theory. In *Twenty-eighth Aaai Conference on Artificial Intelligence (AAAI)*, pages 2099–2105, 2014. [2](#)
- [35] Shanshan Wang and Lei Zhang. Lstn: Latent subspace transfer network for unsupervised domain adaptation. In *Chinese Conference on Pattern Recognition and Computer Vision (PRCV)*, pages 273–284. Springer, 2018. [2](#)
- [36] Shaoan Xie, Zibin Zheng, Liang Chen, and Chuan Chen. Learning semantic representations for unsupervised domain adaptation. In *International Conference on Machine Learning (ICML)*, pages 5419–5428, 2018. [1](#)
- [37] Minghao Xu, Jian Zhang, Bingbing Ni, Teng Li, Chengjie Wang, Qi Tian, and Wenjun Zhang. Adversarial domain adaptation with domain mixup. In *Thirty-fourth Aaai Conference on Artificial Intelligence (AAAI)*, 2020. [1](#)
- [38] Rujia Xu, Guanbin Li, Jihan Yang, and Liang Lin. Larger norm more transferable: An adaptive feature norm approach for unsupervised domain adaptation. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1426–1435, 2019. [5](#), [7](#)
- [39] Leon Bottou Yann Lecun, Bengio Y, and Haffner Patrick. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*, pages 2278–2324, 1998. [5](#)
- [40] Shaokai Ye, Kailu Wu, Mu Zhou, Yunfei Yang, Sia Huat Tan, Kaidi Xu, Jiebo Song, Chenglong Bao, and Kaisheng Ma. Light-weight calibrator: a separable component for unsupervised domain adaptation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13733–13742, 2020. [5](#), [6](#)
- [41] Netzer Yuval, Wang Tao, Coates Adam, Bissaco Alessandro, Wu Bo, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. In *NIPS workshop on deep learning and unsupervised feature learning*, pages 12–17, 2011. [5](#)
- [42] Lei Zhang, Wangmeng Zuo, and David Zhang. Lsd: Latent sparse domain transfer learning for visual adaptation. *IEEE Trans. IP*, 25(3):1177–1191, 2016. [1](#), [2](#)
- [43] Weichen Zhang, Wanli Ouyang, Li Wen, and Xu Dong. Collaborative and adversarial network for unsupervised domain adaptation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3801–3809, 2018. [1](#)
- [44] Zhen Zhang, Mianzhi Wang, Yan Huang, and Arye Nehorai. Aligning infinite-dimensional covariance matrices in reproducing kernel hilbert spaces for domain adaptation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3437–3445, 2018. [2](#)