

Bi-Directional Generation for Unsupervised Domain Adaptation

Guanglei Yang,[†] Haifeng Xia,[‡] Mingli Ding,[†] Zhengming Ding^{*,b}

[†]School of Instrument science and engineering, Harbin Institute of Technology

[‡]Department of Electrical & Computer Engineering, Indiana University-Purdue University Indianapolis

^bDepartment of Computer, Information and Technology, Indiana University-Purdue University Indianapolis
{yangguanglei, dingml}@hit.edu.cn, {haifxia, zd2}@iu.edu

Abstract

Unsupervised domain adaptation facilitates the unlabeled target domain relying on well-established source domain information. The conventional methods forcefully reducing the domain discrepancy in the latent space will result in the destruction of intrinsic data structure. To balance the mitigation of domain gap and the preservation of the inherent structure, we propose a Bi-Directional Generation domain adaptation model with consistent classifiers interpolating two intermediate domains to bridge source and target domains. Specifically, two cross-domain generators are employed to synthesize one domain conditioned on the other. The performance of our proposed method can be further enhanced by the consistent classifiers and the cross-domain alignment constraints. We also design two classifiers which are jointly optimized to maximize the consistency on target sample prediction. Extensive experiments verify that our proposed model outperforms the state-of-the-art on standard cross domain visual benchmarks.

Introduction

Deep learning gains huge success in diverse applications across many fields, such as computer vision, data mining, and natural language processing. In practical application, it is usually easy to acquire abundant target data. However, the insufficiency and absence of labels is still a challenge (Ding, Shao, and Fu 2018; Saito et al. 2018) leading to exceedingly time-consuming and expensive manual annotation. To solve this problem, domain adaptation is introduced and has shown excellent performance. It transfers knowledge from external well-labeled source domain to the target domain (Wei, Ke, and Goh 2016; Ding and Fu 2016; Yan et al. 2017). The theory of domain adaptation is to discover the common latent factors across the source and target domains, and thereby reducing domain mismatch across domains. Different domain adaptation methods have been developed, such as feature alignment and domain confusion adaptation (Pan et al. 2019).

A common practice in domain adaption is to align feature distribution between source and target domains by min-

imizing domain shift through metrics such as correlation distances (Yao et al. 2015; Sun, Feng, and Saenko 2016) or maximum mean discrepancy (MMD) (Ding et al. 2018; Ding and Fu 2018). Among the recent efforts on domain adaption, several works (Pan et al. 2019; Pinheiro 2018; Kang et al. 2019) aim to minimize the MMD-based distance of cross-domain in the deep neural network. For example, (Pan et al. 2019) applies the pairwise reproducing kernel Hilbert space (RKHS) between the prototypes of the same class from different domains to reduce the discrepancy of cross-domain. (Kang et al. 2019) divides the domain discrepancy into two kinds of class-aware domain discrepancy. Another branch of domain adaptation is to exploit the domain adversarial training (Tzeng et al. 2017; Zhang et al. 2018; Saito et al. 2018). For example, (Saito et al. 2018) maximizes the discrepancy between two classifiers' outputs to determine target samples without the support of the source.

Despite the success of previous methods, they suffer from following limitations. First of all, almost methods apply a shared generator, which means their domain adaption is monodirectional. For example, DAN (Long et al. 2015) aligns source domain to target domain. The features of the target domain are merely simulated instead of positively taking part in domain adaption. Therefore, such methods cannot fully take advantage of the target data. Meanwhile, the target data has many unknown factors because it may come from complicated situations. The boundary among different classes can be not clear, and thereby it hurts the knowledge transfer. Secondly, most of the previous methods calculate the domain discrepancy at the domain level, but they neglect the class level where the samples come. Class-agnostic adaptation aligning source and target data at the domain-level is possible to cause sub-optimal solutions. These solutions are likely to overfit in the source domain, causing that performance in the target domain is not as expectancy.

To address these above issues, we propose a Bi-Directional Generation (BDG) method for unsupervised domain adaptation with dual consistent classifiers. The dual generators interpolate two intermediate domains and synthesize more effective data so that there are more samples to train the classifiers. Meanwhile, the target samples are

*Corresponding author: Zhengming Ding (zd2@iu.edu). This work is done under the supervision of Dr. Zhengming Ding.
Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

labeled by the pre-trained network. Apart from minimizing the adversarial loss on source branch, we employ domain confusion loss in the target branch so that both domain information can be more sufficiently utilized. Moreover, we exploit class-wise loss to preserve the class-level semantic information when the generators synthesize two intermediate domains reducing domain discrepancy. Furthermore, the consistent loss alternatively reduces the classifier discrepancy by minimizing itself directly while the previous methods implement adversarial training in classifiers by discrepancy loss. In conclusion, the contributions of our paper are highlighted in three folds as follows:

- We propose a novel bi-directional cross-domain generation module, which aims to synthesize the intermediate domains conditioned on each domain. The augmented samples play as a bridge to reduce the domain discrepancy and preserve the class-level structure during domain adaption process.
- We explore dual classifiers to enhance the bi-directional cross-domain generation. Furthermore, a consistent loss is developed to improve the prediction performance on unlabeled target samples.

Related Works

In this section, we will briefly introduce two branches of domain adaptation, then highlight the difference of our model.

Deep Domain Adaptation

The aim of domain adaptation is to improve the target learning by using the labeled source knowledge, whose distribution is different from the target domain. Spurred by the recent advances in computer vision using deep convolutional neural networks (DCNNs), a number of methods based have been proposed for unsupervised domain adaptation. In particular, one common solution is to guide DCNNs to learn domain feature by minimizing the domain discrepancy with Maximum Mean Discrepancy (MMD) in reproducing kernel Hilbert space (RKHS) (Gretton et al. 2012). MMD is a useful and popular non-parametric metric for the measurement in domain discrepancy between source and target domains. Long et al. learns more transferable features through multi-kernel MMD loss (Long et al. 2015). Then, they further extended this work by adding varieties of MMD loss. For example, RTN (Long et al. 2016) improves DAN by replacing the multi-kernel MMD loss in DAN with a single tensor-based MMD loss and adds the residual module to AlexNet; JAN (Long et al. 2017) exploits joint maximum mean discrepancy criterion to the loss in order to learn a transfer network. Ding et al. explore deep low-rank coding to extract domain-invariant features by adding MMD as the domain alignment loss (Ding and Fu 2018). However, these methods only measuring the domain discrepancy neglect the difference among classes. Domain adaptation without class level domain discrepancy transfers source data to target domain leading to negative transfer. Most recently, SimeNet (Pinheiro 2018) solves this problem by learning domain-invariant features and the categorical prototype representations. In the same way, CAN (Kang et al. 2019) optimizes

the network by considering the discrepancy of intra-class domain and the inter-class domain.

Generative Domain adaptation

Another branch of unsupervised domain adaptation in DCNNs attempts to exploit a domain discriminator to address the domain confusion (Ganin et al. 2016; Tzeng et al. 2017; Zhang et al. 2018). Adversarial adaptation methods minimize the domain discrepancy with a domain discriminator. Due to the adversarial loss in the objective function, more transferable representations can be generated (Tzeng et al. 2017). By adversarial adaptation methods, in theory, generator produces outputs identically distributed as source domain. Meanwhile, with sufficient dataset, a GAN network with encode-decode construction can transfer the same set of samples in source domain to any random distribution of images in the target domain. Moreover, adversarial losses with auxiliary loss can make sure that the learned function can transfer an individual source sample to the desired domain more effectively (Zhu et al. 2017). To further reduce domain discrepancy, Zhu et al. introduces an identity loss to make sure the transfer can preserve semantic feature between the input and output. The traditional methods like cycle loss and identity loss in cycleGAN are too restricted for domain adaption. Meanwhile, the previous process in domain adaption only focuses on the global transform. Although it can reduce the distribution difference as cross domain, it destroys the class semantic feature in each sample. It leads to that the classifier trained in this way cannot get the ideal result in the target domain.

Recently, GTA (Sankaranarayanan et al. 2018) solves this problem by modifying Auxiliary Classifier GAN (AC-GAN) and merging the class label and real/fake label. In addition to GTA, SymNets (Zhang et al. 2019) takes care of this problems by making the asymmetric design of source and target task classifiers sharing with them its layer neurons. These methods apply the MMD in a single branch GAN structure and get an excellent result. However, in their result, we find that when the number of samples in the source domain is insufficient compared with the target domain, the class accuracy can not get the best performance. At the same time, another method to solve this problem employs the multi-branch general structure (Wang et al. 2019). And CADA (Kurmi, Kumar, and Namboodiri 2019) proposes to find adaptable regions using some estimate of the discriminator. It pays most attention to detect whether the part of the picture is transferable or not and does not take the class-level semantic feature into count.

Differently, we propose a dual generative cross-domain generation framework by interpolating two intermediate domains to bridge the domain gap. Our proposed method leverages bi-directional cross-domain generators to make two intermediate domains and use additional target data with pseudo labels for learning two task-specific classifiers. Our work is on the exploitation of building bi-directional generation network with two classifiers, which has not been fully explored in the literature.

The Proposed Algorithm

Preliminary and Motivation

Given the unlabeled target domain $X_t = \{x_t\}_{j=1}^M$ and an auxiliary source domain $X_s = \{x_s\}_{i=1}^N$ with corresponding label y_s , the target of unsupervised domain adaptation is to utilize the classifier trained in source domain to identify sample from target domain. Without loss of generality, the distribution of source data and target data are generally denoted as $x_s \sim p_{data}(x_s)$ and $x_t \sim p_{data}(x_t)$ respectively, and $p_{data}(x_t) \neq p_{data}(x_s)$. In order to solve the discrepancy of cross-domain, adversarial network based methods (Tzeng et al. 2017; Bousmalis et al. 2017; Mancini et al. 2019) include two strategies: align two domains into a latent domain-invariant space (Lee et al. 2019; Chen et al. 2019) and reconstruct one domain in another domain space (Zhang et al. 2019; Roy et al. 2019).

The second strategy transferring source domain into target domain attempts to enable the distribution of source domain to be similar with that of target domain (Long et al. 2015; 2017). Under this circumstance, the performance of the classifier trained in transferring source domain might be further promoted in target domain. However, due to the fact that samples from target domain have none label information, the boundary of class in target domain will be overlapping, which triggers overfitting and mismatching situation. Meanwhile unbalanced domain size may cause insufficient class-semantic feature when we make unsupervised domain adaptation. Furthermore, the domain classifiers might find fewer counterparts to align due to cross-domain translations, rotations, or other transformations. Bi-Directional generative domain adaptation model with dual adversarial classifiers successfully remedy this weakness, enabling BDG to achieve excellent results.

Another strategy aligning the features of two domains in the same latent feature will suffer from the similar problems. They only employ a generator shared by the source domain and target domain (Long et al. 2015; Tsai et al. 2018). However, the classifier is just trained by the part of latent space belonging to the source domain. And the training process neglects the application of target domain. Thus, when the number of samples in the target domain is much more than that in the source domain, the well-trained classifier fails to have great performance in target domain. In order to take full advantage of the information of target domain, the bi-directional generation becomes a promising method. In order to preserve semantic information in the picture, bi-directional generation based methods (Huang et al. 2018) employ the identity loss with L_1 norm making a balance between generating samples and original samples. It can significantly reduce distribution discrepancy. However, these methods obviously lose the class-level semantic structure. Thus, we propose the Bi-Directional Generation for cross-domain learning with dual consistent classifiers to capture class-level semantic information and overcome the problems of overfitting and mismatching.

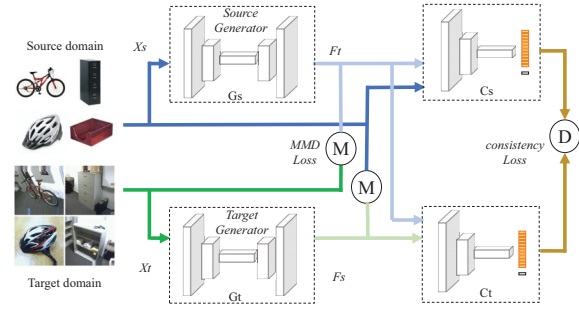


Figure 1: Overview of the proposed bi-directional generation network. It consists of two generators G_s and G_t , two classifiers C_s and C_t . C_s is fed with F_t and X_t ; C_t is fed with F_s and F_t . \mathcal{M} is used to calculate \mathcal{L}_{MMD} . \mathcal{D} calculates discrepancy of the two prediction tensors to product \mathcal{L}_{con} .

Bi-Directional Cross-Domain Generation

As illustrated in Figure 1, X_s, X_t are source and target samples, respectively. We propose the cross-domain generators G_s, G_t to transfer one domain input to the other domain distribution. Specifically, two generators are defined as $G_s : X_s \rightarrow X_t$ and $G_t : X_t \rightarrow X_s$, respectively. Given the source samples X_s , G_s tries to generate F_t that looks similar to target samples X_t . Similarly, With X_t , G_t aims to generate F_s which looks similar to X_s .

We define our bi-directional generation loss in source branch using the following formulation:

$$\begin{aligned} \mathcal{L}_{GAN}^s(X_s) &= \mathcal{L}_{dis_s} + \mathcal{L}_{cls_s} \\ \mathcal{L}_{dis_s} &= \mathbb{E}[\log C_s(X_s)] + \mathbb{E}[\log(1 - C_s(G_s(X_s)))] \\ \mathcal{L}_{cls_s} &= \mathbb{E}[\log C_s(X_s, Y_s)] + \mathbb{E}[\log C_s(G_s(X_s), Y_s)], \end{aligned} \quad (1)$$

where \mathcal{L}_{dis_s} is the discrimination loss and the \mathcal{L}_{cls_s} represents the classification loss. C_s aims to classify the transferred samples F_t and the original samples X_s , which means it can discriminate X_s as real samples and F_t as fake samples. Note that F_t as a newly-augmented intermediate domain across source and target mitigates the domain gap. Meanwhile, F_t generated from G_s not only needs to be identified easily by C_s but also makes C_s have difficult in discriminating which domains it comes from.

Inspired by domain confusion (Tsai et al. 2018), they forward two domain data to a fully-convolutional discriminator by using a cross-entropy loss for the two classes (*i.e.* source and target). In the same way, C_t aims to classify F_s and F_t . We define our bi-directional generation loss in target branch using the following equation as:

$$\begin{aligned} \mathcal{L}_{GAN}^t(X_s, X_t) &= \mathcal{L}_{dis_t} + \mathcal{L}_{cls_t} \\ \mathcal{L}_{dis_t} &= \mathbb{E}[\log C_t(G_s(X_s))] + \mathbb{E}[\log(1 - C_t(G_t(X_t)))] \\ \mathcal{L}_{cls_t} &= \mathbb{E}[\log C_t(F_t, Y_s)] + \mathbb{E}[\log C_t(F_s, \hat{Y}_t)], \end{aligned} \quad (2)$$

In order to align the two domains X_t and F_s on the target distributions, we employ a pre-trained classifier C_0 which is trained on the source domain without any auxiliary

method. It produces the pseudo label of the target domain, $\hat{Y}_t = C_0(X_t)$. It is helpful to transfer data from target domain to source domain because C_0 preserves the class-level semantic information of source domain. The discrimination part of C_t is different from the general idea discriminating the real or fake sample. It distinguishes which domain the sample comes from. Specifically, it regards F_t as samples from source domain and F_s as samples belonging to target domain. Meanwhile, F_t and F_s generated from G_s , G_t respectively enable C_t to easily classify them and make C_t have difficult in recognizing which domain they comes from. Thus, the proposed method can utilize more information from target domain.

Class-wise Cross-Domain Alignment

The Maximum Mean Discrepancy (MMD) (Gretton et al. 2012) is a useful and popular non-parametric metric for measure discrepancy between cross domains. The MMD loss is regarded as the discrepancy measure to compare distinct distributions, which computes the domain discrepancy between the sample means of the source and target data. Without this loss, we experimentally find that the performance of the proposed algorithm drops significantly. Inspired by (Wang et al. 2019; Kang et al. 2019), we apply the class-level MMD loss as well. Our MMD loss includes two terms: *global MMD* reducing the distance between two domains center point; and *class MMD* reducing center point distance in each class between the source domain and target domain. We express the objective of MMD constraint in source and target domains as:

$$\mathcal{L}_{MMD}^{s/t} = \mathcal{L}_{gMMD}^{s/t} + \frac{1}{C} \mathcal{L}_{cMMD}^{s/t}, \quad (3)$$

where C is the class number. The global MMD in source branch is defined as:

$$\mathcal{L}_{gMMD}^s = \left\| \frac{1}{n_{s_{x_s \in X_s}}} \sum G_s(x_s) - \frac{1}{n_{t_{x_t \in X_t}}} \sum x_t \right\|_2, \quad (4)$$

where n_s, n_t are the number of the source domain and target domain. It is easy to write similar equations for the target domain branch according to (4). And the class MMD in source branch is computed as:

$$\mathcal{L}_{cMMD}^s = \sum_c \left\| \frac{1}{n_{s_{x_s \in X_s^c}^c}} \sum G_s(x_s) - \frac{1}{n_{t_{x_t \in X_t^c}^c}} \sum x_t \right\|_2, \quad (5)$$

where n_s^c, n_t^c denote the number of its domain in class c . It is easy to write similar equations for the target domain branch according to (5). The difference between our work and others is that we calculate two MMD losses: the MMD loss between F_t and X_t as well as the MMD loss between F_s and X_s . With pseudo labels of X_t and F_s , we can easily calculate their \mathcal{L}_{cMMD}^t . This strategy aims to take full use of target domain and reduce the domain gap with two directions.

Dual Consistent Classifiers

Dual consistent classifiers refer to that no matter where the classifiers are trained, they should obtain very similar pre-

diction outputs given the target data. In our work, we employ two classifiers C_t and C_s , which are trained by samples from different domains. However, in ideal situation, C_t and C_s should have similar prediction ability on target samples. So we introduce the \mathcal{L}_{con} to regular the two target prediction outputs. The dual consistent classifiers loss is defined as:

$$\mathcal{L}_{con} = \|C_t(F_s) - C_s(F_s)\|_1, \quad (6)$$

where \mathcal{L}_{con} denotes the function measuring L1-norm between two probabilistic outputs computed from C_t and C_s . This term indicates how the two classifiers agree on their predictions. Our goal is to make the two classifiers' prediction consistent so that it can reduce the domain gap further. Meanwhile, by utilizing this loss, the two classifiers can trigger each other. It is helpful to make the whole framework stable during whole training process.

Overall Objective and Optimization

To integrate cross-domain generation, discriminative class-wise alignment and dual consistent classifiers into one unified framework, we have our overall objective function as follows:

$$\mathcal{L} = \mathcal{L}_{GAN}^s + \mathcal{L}_{GAN}^t + \lambda(\mathcal{L}_{MMD}^s + \mathcal{L}_{MMD}^t) + \gamma \mathcal{L}_{con}, \quad (7)$$

where λ, γ control the relative importance of the two terms. The \mathcal{L}_{GAN}^s has been defined in equation (1) and the \mathcal{L}_{GAN}^t can be calculated by equation (2). The $\mathcal{L}_{MMD}^{s/t}$ has been declared in equation (3).

To sum up the previous discussion, we need to train two classifiers which take inputs from the generator and minimize classifier loss, and the generators which try to mimic the the all objective loss. We optimize this objective in three steps.

Step A First, we train classifier C_0 to classify the source samples correctly and produce the pseudo label \hat{Y}_t of the target domain X_t and their alignment domain F_s . The pseudo label \hat{Y}_t is utilized to calculate the $\mathcal{L}_{cMMD}^{s/t}$ and make a clear boundary of class in target domain, which is helpful to take full advantage of information from target domain. They are also beneficial to keep the balance of two branches' domain adaption, because the pseudo labels still maintain abundant domain-invariant representations to guide the G_t . This step is crucial to implement next steps. We train the networks C_0 , whose structure is the same as the classifiers trained by next step, to minimize cross-entropy loss. The objective is as follows:

$$\min_{C_0} \mathcal{L}(C_0, X_s) = \mathbb{E}[\log C_0(X_s)]. \quad (8)$$

Step B In this step, we train the classifiers (C_s, C_t) by fixing generators (G_s, G_t). In order to take advantage of rich domain-invariant representations, we utilize data and their labels from both domains and train the final models for both branches. The C_s is fed by X_s and F_t . Meanwhile C_t is trained by F_t with y_s , and F_s with \hat{y}_t . Under this circumstance, the discrepancy between source domain and target domain will be reduced dramatically. In addition, the classifiers is not only to minimize the GAN losses but also they

can also keep coincident in predicting target domain samples. To realize this goal, we add a consistent loss on the loss of classifier. The objective is as follows:

$$\min_{C_s, C_t} \mathcal{L}_{GAN}^s + \mathcal{L}_{GAN}^t + \gamma \mathcal{L}_{con}, \quad (9)$$

where these two classifiers can be optimized.

Step C We train the generators to minimize the full objective loss by using equation (7) for fixed classifiers. Unlike the general adversarial process in generator training, which minimizes the \mathcal{L}_{dis_t} by regarding the fake samples as real samples, our method switch the both domain label at the same time. In detail, the F_s drawn from target domain is regarded as source domain samples and the F_t is considered as target domain samples.

Remarks: MCD (Saito et al. 2018) employs two classifiers trained by source samples jointly with a feature generator. However, through many experiments, we find that the total loss becomes diverge with epoch prolonging. In our method, we adapt the discrepancy loss without classifier adversarial process. We train the two classifiers C_s processing X_s as well as F_t and C_t with the input F_t and F_s .

Obviously, the C_t, C_s are different classifiers. Compared with the (Saito et al. 2018) which shares generator between both domain, we employ two independent generators G_t, G_s to implement bi-direction domain adaption. We add the MMD loss taking the class-level structure information into count, into the total loss to replace the discrepancy loss in (Saito et al. 2018). Meanwhile we adapt discrepancy loss as consistent loss, which is optimized by classifiers and generators. Thus, we do not maximize the discrepancy loss when we train the classifiers. In other words, reducing MMD loss aims to guide the generators on how to map the source domain to the target domain.

Experiments and Results

We evaluate BDG method by the standard benchmarks including Office-31 and Office-Home, compared with state of the art domain adaption methods.

Datasets and Experimental Setup

Office-31 (Saenko et al. 2010), a standard benchmark for visual domain adaptation, contains 4,652 images and 31 categories from three distinct domains, *i.e.*, images collected from the 1) Amazon website (Amazon domain), 2) web camera (Webcam domain), and 3) digital SLR camera (DSLR domain) under different settings, respectively. The dataset is imbalanced across domains, with 2,817 images in Amazon domain, 795 images in Webcam domain, and 498 images in DSLR domain. We follow the standard evaluation protocols for unsupervised domain adaptation (Ganin et al. 2016; Long et al. 2015).

Office-Home (Venkateswara et al. 2017) is a more challenging dataset for domain adaptation evaluation. It consists of around 15,500 images in total from 65 categories of everyday objects in office and home scenes. There are four significantly different domains: Artistic images (Ar), Clip Art (Cl), Product images (Pr), and Real-World images (Rw).

The images of these domains have substantially different appearances and backgrounds. The the number of categories is much larger than that of Office-31, making it more difficult to transfer across domains. We evaluate all methods on all 12 adaptation tasks.

Implementation Details

The hyper-parameters λ, γ in the equation(7) are selected as 1 throughout all experiments. We use ResNet-50(He et al. 2016) models pre-trained on the ImageNet dataset (Russakovsky et al. 2015) as the backbone and we remove its last FC layer. And We fine-tune all convolutional and pooling layers and apply back-propagation to train the classifiers and generators. There optimizer of the classifiers C_s, C_t is mini-batch stochastic gradient descent (SGD) with the momentum of 0.9. At the same time, we use adaptive moment estimation (Adam) to train the generators G_s, G_t , as in (Salimans et al. 2016). And the learning rate is set to 5.0×10^{-4} in all experiments. We report the accuracy result after 20,000 iterations.

To empirically verify the merit of our proposed model, We compare with both conventional and the state of the art transfer learning methods including: **RTN** (Long et al. 2016), **ADDA** (Tzeng et al. 2017), **JAN** (Long et al. 2017), **SimeNet** (Pinheiro 2018), **GTA** (Sankaranarayanan et al. 2018), **TADA** (Wang et al. 2019), **STA** (Liu et al. 2019), **SymNets** (Zhang et al. 2019), **SAFN** (Xu et al. 2019). All of them have been introduced in related works section. we also compare with other domain adaption methods: **ResNet-50** (He et al. 2016) directly exploits the classification model trained on the source domain to classify target samples; **DANN** (Ganin et al. 2016) designs a domain regularizer to calculated the \mathcal{H} -divergence; **MADA** (Pei et al. 2018) takes multi-mode structures to get the suitable alignment of different data distributions based on multiple domain discriminators; **DSR** (Cai et al. 2019) reconstructs the semantic latent variables and domain latent variables by employing a variational auto-encoder.

Comparison Results

The classification accuracy on the Office-31 dataset for unsupervised domain adaptation based on ResNet-50 are shown in Table 1. For a fair comparison, the results of all baselines are directly reported from their original papers wherever available. To be honest, the results on D→A is little lower than existing methods. However, we get reasonable results in almost task, *i.e.* A→W and A→D. Moreover, the performance in task W→D and D→W is over all methods. Taking all into count, the BDG model outperforms all compared methods on mean accuracy.

According to Table 2, the BDG approach overpasses the compared methods on all transfer tasks on Office-Home. Moreover, it improves their accuracy significantly in many tasks, even though this dataset has abundant categories. Compared with TADA(Wang et al. 2019), the BDG improves more than 6 percents in some transfer learning tasks, such as Cl→Ar, Pr→Ar. It illustrates that BDG yields larger improvements on such difficult transfer learning tasks. And it also suggests that BDG can maintain more class-semantic

Table 1: Accuracy (%) on Office-31 for unsupervised domain adaption (ResNet)

Method	ResNet-50	RTN	DANN	ADDA	JAN	MADA	SAFN	SimNet	GTA	SymNets	TADA	BDG
A→W	68.4±0.2	84.5±0.2	82.0±0.4	86.2±0.5	85.4±0.3	90.0±0.1	88.8±0.4	88.6±0.5	89.5±0.5	90.8±0.1	94.3±0.3	93.6±0.4
D→W	96.7±0.1	96.8±0.1	96.9±0.2	96.2±0.3	97.4±0.2	97.4±0.1	98.4±0.0	98.2±0.2	97.9±0.3	98.8±0.3	98.7±0.1	99.0±0.1
W→D	99.3±0.1	99.4±0.1	99.1±0.1	98.4±0.3	99.8±0.2	99.6±0.1	99.8±0.0	99.7±0.2	99.8±0.4	100.0±0.0	99.8±0.2	100±0.0
A→D	68.9±0.2	77.5±0.3	79.7±0.4	77.8±0.3	84.7±0.3	87.8±0.2	87.7±1.3	85.3±0.3	87.7±0.5	93.9±0.5	91.6±0.3	93.6±0.3
D→A	62.5±0.3	66.2±0.2	68.2±0.4	69.5±0.4	68.6±0.3	70.3±0.3	69.8±0.4	73.4±0.8	72.8±0.3	74.6±0.6	72.9±0.2	73.2±0.2
W→A	60.7±0.3	64.8±0.3	67.4±0.5	68.9±0.5	70.0±0.4	66.4±0.3	69.7±0.2	71.6±0.6	71.4±0.4	72.5±0.5	73.0±0.3	72.0±0.1
Avg	76.1	81.6	82.2	82.9	84.3	85.2	85.7	86.2	86.5	88.4	88.4	88.5

Table 2: Accuracy (%) on Office-Home for unsupervised domain adaption (ResNet)

Method	Ar→Cl	Ar→Pr	Ar→Rw	Cl→Ar	Cl→Pr	Cl→Rw	Pr→Ar	Pr→Cl	Pr→Rw	Rw→Ar	Rw→Cl	Rw→Pr	Avg
ResNet-50	34.9	50.0	58.0	37.4	41.9	46.2	38.5	31.2	60.4	53.9	41.2	59.9	46.1
DANN	45.6	59.3	70.1	47.0	58.5	60.9	46.1	43.7	68.5	63.2	51.8	76.8	57.6
JAN	45.9	61.2	68.9	50.4	59.7	61.0	45.8	43.4	70.3	63.9	52.4	76.8	58.3
DSR	53.4	71.6	77.4	57.1	66.8	69.3	56.7	49.2	75.7	68.0	54.0	79.5	64.9
SymNets	47.7	72.9	78.5	64.2	71.3	74.2	64.2	48.8	79.5	74.5	52.6	82.7	67.6
TADA	53.1	72.3	77.2	59.1	71.2	72.1	59.7	53.1	78.4	72.4	60.0	82.9	67.6
BDG	51.5	73.4	78.7	65.3	71.5	73.7	65.1	49.7	81.1	74.6	55.1	84.8	68.7

feature during effective domain adaptation. Moreover, in the office-home dataset existing unbalanced samples number problem, the classifier can implement sufficient adversarial process, which makes the two classifiers be more consistent. Meanwhile, the large dataset is helpful to calculate $\mathcal{L}_{MMD}^{s/t}$, especially $\mathcal{L}_{cMMD}^{s/t}$, because mismatching problem can be avoided, *i.e.*, some class in F_t or F_s is blank.

Empirical Analysis

t-SNE visualization To further understand the alignment of distribution, we visualize the features in 2D-space using the network activations of the FC layer from task $Amazon \rightarrow webcam$ (31 classes) learned by ResNet, DAN, RevGrad and BDG, respectively using t-SNE embeddings (Donahue et al. 2014). The representations generated by BDG (Figure 2) form exactly 31 clusters with clear boundaries. Compared to ResNet, DAN and RevGrad, as expected, our t-SNE figure demonstrates the closer distance between the same classes in different domain. This shows that our BDG generate more discriminative features for both domain and confirms our improvement in Tables 1 & 2.

Parameter analysis We conduct experiments to investigate the sensitivity of our method to the balance parameters λ , γ . We use control variations method to test two super-parameters, whose value are selected at the range of [0.1, 0.2, 0.5, 1, 1.2, 1.5, 2]. For example, when we change the value of λ , γ is fixed to 1. From Figure 3, as λ go up, the average accuracy of office-31 dataset increases while accuracy of office-home decreases gradually. At the same time, as γ increase, the average accuracy of every dataset raise and then go down gradually. In conclusion, hyper-parameters is

quite effective, but the optimal value in each case is different. We fix $\lambda = 1$ and $\gamma = 1$ in the other experiments because its result is more stable than others.

Table 3: The effect of MMD loss and discrepancy loss. The mean accuracy of Office-home and office-31 are reported.

dataset	Variant1	Variant2	Variant3	Variant4	Variant5	BDG
Office-31	81.1	87.4	82.9	82.5	88.3	88.5
Office-Home	61.0	67.8	63.2	63.5	68.2	68.7

Ablation studies We compare our method (“BDG”) with other variants, to verify the effect of $\mathcal{L}_{MMD}^{s/t}$ and $\mathcal{L}_{con}^{s/t}$. We adopt the single directional cross-domain generation method translating target domain to the source domain without $\mathcal{L}_{MMD}^{s/t}$ as **variant 1**. Compared to **variant 1**, we design another method named **variant 2** which includes the loss function of MMD. **variant 3** is the bi-directional Cross-Domain Generation method with dual classifiers. In addition, we also design the **variant 4** consisting of $\mathcal{L}_{GAN}^{s/t}$ and \mathcal{L}_{con} . Moreover, we also propose the **variant 5** including $\mathcal{L}_{GAN}^{s/t}$ and $\mathcal{L}_{MMD}^{s/t}$ to do comparison with other variants. It can be seen that combining $\mathcal{L}_{MMD}^{s/t}$ and \mathcal{L}_{con} improves the adaption performance. In detail, the mean accuracy of Variant 2(87.4% in the office-31 dataset) increase significantly, compared with the result(81.1%) of variant 1. The main reason is that the $\mathcal{L}_{MMD}^{s/t}$ reduces the domain discrepancy and maintains the class-level information. The similar situation also happens in the office-home dataset. When we compare the variant 3,4,5 and BDG in Table 3, it can be seen that




Figure 2: Visualization with t-SNE for different adaptation methods (bested viewed in color).**a:** t-SNE of ResNet. **b:** DAN.**c:** RevGrad.**d:** BDG. The input activations of the last FC layer are used for the computation of t-SNE. The results are on Office-31 task *Amazon* \rightarrow *Webcam*.




Figure 3: Parameter study of our algorithm: averaged accuracy with parameter λ, γ learned by (a) webcam \rightarrow amazon (Office-31), (b) Art \rightarrow Realworld (Office-Home), (c) Art \rightarrow Realworld (Office-Home)(black: λ , red: γ)

the result of variant 4 in two datasets is close to the result of variant 3. In detail, the result of variant 4 in the office-home dataset is over only 0.3% than variant 3. However, when the loss combines the $\mathcal{L}_{MMD}^{s/t}$ and \mathcal{L}_{con} , the performance improves. So we can get some conclusions according to Table 3. First of all, it can be proved that $\mathcal{L}_{MMD}^{s/t}$ takes an essential role in our method. Second, it also can demonstrate \mathcal{L}_{con} is an applicable term in our method. The last but not the least, bi-directional cross-domain generation method does work better than single branch structure method.




Figure 4: The curve of \mathcal{L}_{MMD} (red), $\mathcal{L}_{discrepancy}$ (black) and accuracy(blue) during training on task Webcam \rightarrow Amazon of the office-home dataset.

Convergence Analysis The Figure 4 shows the whole training process of our BDG method on task Rw \rightarrow Pr from office-home dataset. We notice that the total loss becomes diverge with epoch prolonging if the full loss include adversarial loss. Therefore, they can get an excellent result in the middle process but their final result may be much weaker than the best. We solve these problems by adding the $\mathcal{L}_{MMD}^{s/t}$ and \mathcal{L}_{con} . The Figure4 also reports that accuracy gradually increases when the $\mathcal{L}_{MMD}^{s/t}$ and \mathcal{L}_{con} continuously decrease. It approves that the training process avoids becoming diverge with epoch prolonging due to the effect of two terms.

Conclusion

In this paper, we propose a bi-directional generative cross-domain generation framework to explore domain adaptation in an unsupervised manner. Previous adversarial domain adaptation approaches with shared generator neglect class features by interpolating two intermediate domains to bridge the domain gap. We add the MMD loss and consistent loss into loss function to implement bi-directional Cross-Domain Generation method. The proposed method includes MMD loss which takes full use of target domain and reduces the domain gap with two directions, and the consistent loss which utilizes the task-specific classifiers to align the source and target features. Experiments on two standard domain adaptation benchmarks verifies the advanced effectiveness of our algorithm which has better performance than other state-

of-the-art domain adaptation models.

References

- Bousmalis, K.; Silberman, N.; Dohan, D.; Erhan, D.; and Krishnan, D. 2017. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *CVPR*, 3722–3731.
- Cai, R.; Li, Z.; Wei, P.; Qiao, J.; Zhang, K.; and Hao, Z. 2019. Learning disentangled semantic representation for domain adaptation. In *IJCAI*, 2060–2066.
- Chen, C.; Xie, W.; Huang, W.; Rong, Y.; Ding, X.; Huang, Y.; Xu, T.; and Huang, J. 2019. Progressive feature alignment for unsupervised domain adaptation. In *CVPR*, 627–636.
- Ding, Z., and Fu, Y. 2016. Robust transfer metric learning for image classification. *TIP* 26(2):660–670.
- Ding, Z., and Fu, Y. 2018. Deep transfer low-rank coding for cross-domain learning. *IEEE TNNLS* 30(6):1768–1779.
- Ding, Z.; Li, S.; Shao, M.; and Fu, Y. 2018. Graph adaptive knowledge transfer for unsupervised domain adaptation. In *ECCV*, 37–52.
- Ding, Z.; Shao, M.; and Fu, Y. 2018. Robust multi-view representation: A unified perspective from multi-view learning to domain adaptation. In *IJCAI*, 5434–5440.
- Donahue, J.; Jia, Y.; Vinyals, O.; Hoffman, J.; Zhang, N.; Tzeng, E.; and Darrell, T. 2014. Decaf: A deep convolutional activation feature for generic visual recognition. In *ICML*, 647–655.
- Ganin, Y.; Ustinova, E.; Ajakan, H.; Germain, P.; Larochelle, H.; Laviolette, F.; Marchand, M.; and Lempitsky, V. 2016. Domain-adversarial training of neural networks. *JMLR* 17(1):2096–2030.
- Gretton, A.; Borgwardt, K. M.; Rasch, M. J.; Schölkopf, B.; and Smola, A. 2012. A kernel two-sample test. *JMLR* 13:723–773.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*, 770–778.
- Huang, S.-W.; Lin, C.-T.; Chen, S.-P.; Wu, Y.-Y.; Hsu, P.-H.; and Lai, S.-H. 2018. Auggan: Cross domain adaptation with gan-based data augmentation. In *ECCV*, 718–731.
- Kang, G.; Jiang, L.; Yang, Y.; and Hauptmann, A. G. 2019. Contrastive adaptation network for unsupervised domain adaptation. In *CVPR*.
- Kurmi, V. K.; Kumar, S.; and Nambodiri, V. P. 2019. Attending to discriminative certainty for domain adaptation. In *CVPR*, 491–500.
- Lee, C.-Y.; Batra, T.; Baig, M. H.; and Ulbricht, D. 2019. Sliced wasserstein discrepancy for unsupervised domain adaptation. In *CVPR*, 10285–10295.
- Liu, H.; Cao, Z.; Long, M.; Wang, J.; and Yang, Q. 2019. Separate to adapt: Open set domain adaptation via progressive separation. In *CVPR*, 2927–2936.
- Long, M.; Cao, Y.; Wang, J.; and Jordan, M. I. 2015. Learning transferable features with deep adaptation networks. *arXiv preprint arXiv:1502.02791*.
- Long, M.; Zhu, H.; Wang, J.; and Jordan, M. I. 2016. Unsupervised domain adaptation with residual transfer networks. In *Advances in Neural Information Processing Systems*, 136–144.
- Long, M.; Zhu, H.; Wang, J.; and Jordan, M. I. 2017. Deep transfer learning with joint adaptation networks. In *JMLR*, 2208–2217. [JMLR.org](http://jmlr.org).
- Mancini, M.; Bulò, S. R.; Caputo, B.; and Ricci, E. 2019. Adagraph: Unifying predictive and continuous domain adaptation through graphs. In *CVPR*, 6568–6577.
- Pan, Y.; Yao, T.; Li, Y.; Wang, Y.; Ngo, C.-W.; and Mei, T. 2019. Transferrable prototypical networks for unsupervised domain adaptation. In *CVPR*, 2239–2247.
- Pei, Z.; Cao, Z.; Long, M.; and Wang, J. 2018. Multi-adversarial domain adaptation. In *AAAI*.
- Pinheiro, P. O. 2018. Unsupervised domain adaptation with similarity learning. In *CVPR*, 8004–8013.
- Roy, S.; Siarohin, A.; Sangineto, E.; Bulò, S. R.; Sebe, N.; and Ricci, E. 2019. Unsupervised domain adaptation using feature-whitening and consensus loss. In *CVPR*, 9471–9480.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. 2015. ImageNet large scale visual recognition challenge. *IJCV* 115(3):211–252.
- Saenko, K.; Kulis, B.; Fritz, M.; and Darrell, T. 2010. Adapting visual category models to new domains. In *ECCV*, 213–226. Springer.
- Saito, K.; Watanabe, K.; Ushiku, Y.; and Harada, T. 2018. Maximum classifier discrepancy for unsupervised domain adaptation. In *CVPR*, 3723–3732.
- Salimans, T.; Goodfellow, I.; Zaremba, W.; Cheung, V.; Radford, A.; and Chen, X. 2016. Improved techniques for training gans. In *Advances in neural information processing systems*, 2234–2242.
- Sankaranarayanan, S.; Balaji, Y.; Castillo, C. D.; and Chellappa, R. 2018. Generate to adapt: Aligning domains using generative adversarial networks. In *CVPR*, 8503–8512.
- Sun, B.; Feng, J.; and Saenko, K. 2016. Return of frustratingly easy domain adaptation. In *AAAI*.
- Tsai, Y.-H.; Hung, W.-C.; Schuler, S.; Sohn, K.; Yang, M.-H.; and Chandraker, M. 2018. Learning to adapt structured output space for semantic segmentation. In *CVPR*, 7472–7481.
- Tzeng, E.; Hoffman, J.; Saenko, K.; and Darrell, T. 2017. Adversarial discriminative domain adaptation. In *CVPR*, 7167–7176.
- Venkateswara, H.; Eusebio, J.; Chakraborty, S.; and Panchanathan, S. 2017. Deep hashing network for unsupervised domain adaptation. In *CVPR*, 5018–5027.
- Wang, X.; Li, L.; Ye, W.; Long, M.; and Wang, J. 2019. Transferable attention for domain adaptation. In *AAAI*.
- Wei, P.; Ke, Y.; and Goh, C. K. 2016. Deep nonlinear feature coding for unsupervised domain adaptation. In *IJCAI*, 2189–2195.
- Xu, R.; Li, G.; Yang, J.; and Lin, L. 2019. Larger norm more transferable: An adaptive feature norm approach for unsupervised domain adaptation. In *ICCV*, 1426–1435.
- Yan, H.; Ding, Y.; Li, P.; Wang, Q.; Xu, Y.; and Zuo, W. 2017. Mind the class weight bias: Weighted maximum mean discrepancy for unsupervised domain adaptation. In *CVPR*, 2272–2281.
- Yao, T.; Pan, Y.; Ngo, C.-W.; Li, H.; and Mei, T. 2015. Semi-supervised domain adaptation with subspace learning for visual recognition. In *CVPR*, 2142–2150.
- Zhang, Y.; Qiu, Z.; Yao, T.; Liu, D.; and Mei, T. 2018. Fully convolutional adaptation networks for semantic segmentation. In *CVPR*, 6810–6818.
- Zhang, Y.; Tang, H.; Jia, K.; and Tan, M. 2019. Domain-symmetric networks for adversarial domain adaptation. In *CVPR*, 5031–5040.
- Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2223–2232.