

A Balanced and Uncertainty-aware Approach for Partial Domain Adaptation

Jian Liang¹[0000–0003–3890–1894], Yunbo Wang²[0000–0002–6215–8888], Dapeng Hu¹, Ran He³[0000–0002–3807–991X], and Jiashi Feng¹[0000–0001–6843–0064]

¹ Department of ECE, National University of Singapore (NUS)
liangjian92@gmail.com, dapeng.hu@u.nus.edu, elefjia@nus.edu.sg

² Peking University wangyunbo09@gmail.com

³ Institute of Automation, Chinese Academy of Sciences rhe@nlpr.ia.ac.cn

Abstract. This work addresses the unsupervised domain adaptation problem, especially in the case of class labels in the target domain being only a subset of those in the source domain. Such a partial transfer setting is realistic but challenging and existing methods always suffer from two key problems, negative transfer and uncertainty propagation. In this paper, we build on domain adversarial learning and propose a novel domain adaptation method BA³US with two new techniques termed Balanced Adversarial Alignment (BAA) and Adaptive Uncertainty Suppression (AUS), respectively. On one hand, negative transfer results in misclassification of target samples to the classes only present in the source domain. To address this issue, BAA pursues the balance between label distributions across domains in a fairly simple manner. Specifically, it randomly leverages a few source samples to augment the smaller target domain during domain alignment so that classes in different domains are symmetric. On the other hand, a source sample would be denoted as uncertain if there is an incorrect class that has a relatively high prediction score, and such uncertainty easily propagates to unlabeled target data around it during alignment, which severely deteriorates adaptation performance. Thus we present AUS that emphasizes uncertain samples and exploits an adaptive weighted complement entropy objective to encourage incorrect classes to have uniform and low prediction scores. Experimental results on multiple benchmarks demonstrate our BA³US surpasses state-of-the-arts for partial domain adaptation tasks. Code is available at <https://github.com/tim-learn/BA3US>.

Keywords: Partial Transfer Learning; Domain Adaptation; Adversarial Alignment; Uncertainty Propagation; Object Recognition.

1 Introduction

Over the past two decades, many research efforts have been devoted to unsupervised domain adaptation (UDA), which aims to leverage labeled source domain data to learn to classify unlabeled target domain data. Typically, existing UDA methods minimize the discrepancy between two domains by matching their

statistical distribution moments [41,37,47,18] or by domain adversarial learning [39,10,40,26]. Once the domain shift is mitigated, source classifiers can be easily transferred to the target domain even with no labeled target domain data available. However, both UDA strategies always assume that different domains share the same label space. Such an assumption may not hold in practice and target domain labels may be only a subset of source domain labels. This introduces an unsupervised partial domain adaptation (PDA) problem that receives increasing research attention recently [4,49,5,30].

The PDA problem is challenging since source-only classes may occur in the target domain during distribution alignment, which is well-known as class mismatch that potentially causes *negative transfer*. Several previous PDA approaches [4,5] mitigate negative transfer by jointly filtering out source-only classes and promote positive transfer by matching the data distribution in the shared classes. Samples from source-only classes are expected to have lower weights in the adaptation module such that the marginal distributions of two domains can be aligned well. However, it is rather risky to rule out the source-only classes, especially when the estimation of label distribution in the target domain is inaccurate.

To match two non-identical label spaces, we view the PDA problem from a new perspective and propose to augment the target label space to be the same as the source label space. Specifically, we develop a simple balanced alignment solution, termed Balanced Adversarial Alignment (BAA), that borrows fewer and fewer samples from the source domain to the target domain within an iterative adversarial learning framework. We expect that the augmented target domain looks much more similar to the source domain w.r.t. the label distribution, and the challenging PDA problem can be transformed to a well-studied UDA task. To focus on the originally shared classes, we propose to filter out the source-only classes via a class-level weighting scheme meanwhile, making the large UDA task more compact.

Besides, existing domain adaptation methods always employ a conventional cross-entropy loss to merely promote the prediction score of ground-truth classes but neglect to suppress those of incorrect classes, which may result in a new problem termed *uncertainty propagation*. Intuitively, if incorrect classes have relatively high prediction scores for source data, some wrong classes would possibly have the largest prediction scores for the aligned target data around them. This problem is quite critical but has been always ignored in the domain adaptation field. To circumvent the issue, we develop an uncertainty suppression solution termed Adaptive Uncertainty Suppression (AUS) that exploits complement entropy [7] in the labeled source domain to prohibit possibly high prediction scores from incorrect classes. Specifically, we emphasize more the uncertain samples corresponding to smaller cross-entropy loss (confidence) and propose a confidence-weighted complement entropy objective in addition to the primary cross-entropy objective.

Generally, our baseline is built on the seminal domain adversarial networks [9,10] and exploits conditional entropy minimization and an entropy-aware weight

strategy [26]. In this paper, we equip the baseline model with two proposed techniques mentioned above and finally formulate a unified framework BA³US which well addresses the negative transfer and uncertainty propagation problems in partial domain adaptation. We also empirically discover that the uncertainty suppression technique works well for vanilla closed-set domain adaptation.

To sum up, we make the following contributions. To our best knowledge, this is the first work that tackles partial domain adaptation by augmenting the target domain and transforming it into a UDA-like problem. The proposed balanced augmentation technique is fairly simple and works very well for PDA tasks. Besides, we address an overlooked issue in this field named uncertainty propagation by designing an adaptive weighted complement entropy for the source domain. Extensive results demonstrate that our approach yields new state-of-the-art results on several visual benchmark datasets, including Office31 [34], Office-Home [42], and ImageNet-Caltech [6].

2 Related Work

The past two decades have witnessed remarkable progress in domain adaptation. Interested readers can refer to [8,50,19] for taxonomy and survey.

Unsupervised Domain Adaptation (UDA). Compared with its supervised counterpart, UDA is more practical and challenging since no labeled data in the target domain are available. Recently, deep convolutional neural networks have achieved great success for visual recognition tasks, and we focus on deep UDA methods in this work. They can be categorized into three main groups. The first group aims to minimize the domain discrepancy by matching different statistic moments like maximum mean discrepancy (MMD) [41,25,27,22] and higher-order moment matching [37,47,18]. The second group that is widely used introduces a domain discriminator and exploits the idea of adversarial learning [12] to encourage domain confusion so that the discriminator can not decide which domain the data come from. Some typical examples are [39,10,40,2]. A third group is a reconstruction-based approach, assuming the reconstruction of both source and target domain samples to be important and helpful. Among them, [11,51] utilize encoder-decoder reconstruction and adversarial reconstruction, respectively. Despite their success for vanilla UDA, they are easily stuck by negative transfer for PDA due to the mismatched marginal label distributions.

Partial Domain Adaptation (PDA). In reality, PDA can be considered as a special case of imbalanced domain adaptation, where the target label distribution is quite dissimilar to that of the source domain. Until recent years, [31] first introduces an imbalanced scenario where label numbers of the source and target domains are not the same, which draws the attention of many researchers [38,4,5,49,30,6,17]. Different from shallow methods [31,38], deep methods [4,5,49,30] are mainly based on the domain adversarial learning framework and achieve promising recognition accuracy. Selective adversarial network (SAN) [4] exploits a multi-discriminator domain adversarial network and tries to select source-only classes by imposing different localized weights

on different discriminators. Importance weighted adversarial nets (IWAN) [49] apply only one domain discriminator and weigh each source sample with the probability of being a target sample. Partial adversarial domain adaptation (PADA) [5] estimates the target label distribution and then feeds the class-wise weights to both the source discriminator and the domain discriminator, while two weighted inconsistency-reduced networks (TWINs) [30] leverage two independent networks to estimate the target label distribution and minimize the domain difference measured by the classifiers’ inconsistency on the target samples. Deep Residual Correction Network (DRCN) [21] proposes a weighted class-wise matching strategy to explicitly align target data with the most relevant source subclasses. Recently, Example Transfer Network (ETN) [6] jointly learns domain-invariant representations across domains and a progressive weighting scheme to quantify the transferability of source examples, which achieves state-of-the-art results on several benchmark datasets. Generally, all the PDA methods above attempt to filter out the large source domain to match the small target domain. Comparatively, our method tries to augment the small target domain to match the source domain from a different perspective.

Data Synthesis and Augmentation. Recently, synthesis and augmentation techniques like CycleGAN [51] and *mixup* [48] are favored by UDA and semi-supervised learning methods for improving performance. For example, [16] directly exploits CycleGAN to generate target-like images from source samples to narrow the domain shift for adaptive semantic segmentation. [29,43] extend *mixup* to domain adaptation and generate pseudo training samples via interpolating between certain source samples and uncertain target samples. To some degree, target augmentation in this paper is like a special case of *mixup* where the mixup coefficient is always binary. However, the motivations are totally different. Our method considers neither the interpolated semantic label nor the interpolated domain label for a classification loss.

3 Proposed Method

We elaborate on the proposed framework for partial domain adaptation (PDA) in this section. First, we give definitions and notations. We follow the protocol of unsupervised PDA where we have a labeled source domain dataset $\mathcal{D}_s = \{(x_i^s, y_i^s)\}_{i=1}^{n_s}$, $x_i^s \in \mathbb{R}^d$ and an unlabeled target domain dataset $\mathcal{D}_t = \{(x_i^t)\}_{i=1}^{n_t}$, $x_i^t \in \mathbb{R}^d$ during the training stage. These two domains have different feature distributions: $p_s(x_s) \neq p_t(x_t)$ due to the domain shift. Notably, different from vanilla UDA, the target labels are a subset of the source labels for PDA: $\mathcal{Y}_t \subseteq \mathcal{Y}_s$, and C denotes the total number of classes in \mathcal{Y}_s .

We aim to learn a deep neural network $h : \mathcal{X} \rightarrow \mathcal{Y}$ that consists of two components: $h = g \circ f$. Here $f : \mathcal{X} \rightarrow \mathcal{Z}$ denotes the feature extractor and $g : \mathcal{Z} \rightarrow \mathcal{Y}$ denotes a class predictor. Since we target at learning domain-invariant features, the prediction function is assumed identical, i.e., $g = g_s = g_t$. For simplicity, we also share the feature extractor f for different domains. We introduce an adver-

serial classifier $D : \mathcal{Z} \rightarrow \{0, 1\}$ to mitigate the distribution discrepancy across domains as explained later.

3.1 Domain Adversarial Learning Revisited

Generative adversarial network (GAN) [12] learns two competing components: the discriminator D and the generator F which play a minimax two-player game, where F tries to fool D by generating examples that are as realistic as possible and D tries to classify the generated samples as fake. Inspired by the idea of GAN, domain adversarial neural networks (DANN) [9,10] develops a two-player game for UDA, where one player $D(\cdot; \theta_d)$ (i.e., domain discriminator) tries to distinguish the source domain datum from that of the target domain and the other player $F(\cdot; \theta_f)$ (i.e., feature extractor) is trained to confuse the domain discriminator $D(\cdot; \theta_d)$. Generally, the minimax game of DANN is formulated as

$$\begin{aligned} & \min_{\theta_f, \theta_g} \max_{\theta_d} \mathcal{L}_{cls}(\theta_f, \theta_g) + \lambda \mathcal{L}_{adv}(\theta_f, \theta_d), \\ \mathcal{L}_{adv}(\theta_f, \theta_d) &= \frac{1}{n_s} \sum_{i=1}^{n_s} \log[D(F(x_i^s))] + \frac{1}{n_t} \sum_{j=1}^{n_t} \log[1 - D(F(x_j^t))], \\ \mathcal{L}_{cls}(\theta_f, \theta_g) &= \frac{1}{n_s} \sum_{i=1}^{n_s} l_{ce}(G(F(x_i^s)), y_i^s), \end{aligned} \quad (1)$$

where $l_{ce}(\cdot, \cdot)$ represents the softmax cross-entropy loss, $G(\cdot; \theta_g)$ denotes the source classifier, $F(\cdot; \theta_f)$ represents the domain-shared feature extractor, and λ is a hyper-parameter to trade-off the source risk and domain adversary. Different from a label flipping step in GAN, a gradient reversal layer (GRL) is further defined in [9] to optimize the objective in Eq. (1). Due to its simplicity and effectiveness, the idea of domain adversarial learning has been adopted in many previous domain adaptation works [24,40,26].

As shown in Eq. (1), each sample from both source and target domains is equally involved in the adversarial loss \mathcal{L}_{adv} , which seems not reasonable. If we only have samples distributed in the margin of the classifier (called ‘hard’ samples) from different domains and pursue domain alignment via them, it may perform badly for those ‘easy’ samples across domains. As such, we expect those ‘hard’ samples and ‘easy’ samples to own lower and higher weights during domain adversarial alignment. Specifically, we quantify the difficulty via the entropy criterion $H(h) = -\sum_{c=1}^C h_c \log(h_c)$ and adopt the same weighting strategy as [26] to impose an entropy-aware weight $w(x) = 1 + e^{-H(h(x))}$ on each sample,

$$\mathcal{L}_{adv}^e(\theta_f, \theta_d) = \frac{1}{n_s} \sum_{i=1}^{n_s} w(x_i^s) \log[D(F(x_i^s))] + \frac{1}{n_t} \sum_{j=1}^{n_t} w(x_j^t) \log[1 - D(F(x_j^t))]. \quad (2)$$

Obviously, if there is no domain shift, UDA (including PDA) degenerates to a typical semi-supervised learning problem. On one hand, we aim to mitigate the domain shift; on the other hand, this motivates us to adopt the popular entropy minimization principle [13] in semi-supervised learning to the PDA task. It is

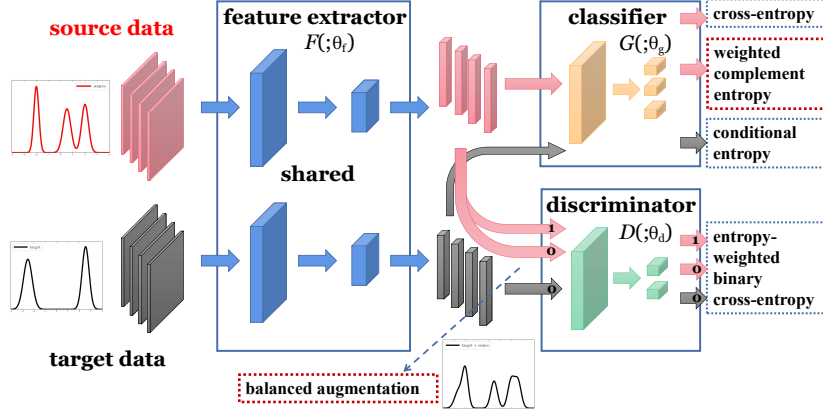


Fig. 1. Architecture of our domain adaptation method. There are three modules: a shared feature extractor F , a classifier G and a domain discriminator D . Different from domain adversarial learning [9], it contains two extra components with marked red border, i.e., balanced augmentation and weighted complement entropy.

desirable that all the unlabeled target samples have highly-confident predictions. This is encouraged via the following conditional entropy term,

$$\mathcal{L}_{ent}(\theta_f, \theta_g) = \frac{1}{n_t} \sum_{j=1}^{n_t} H(G(F(x_j^t))). \quad (3)$$

Inspired by the observation in [4] that redundant information is not beneficial for adaptation, we adopt the following optimization objective of Entropy-regularized DANN (**E-DANN**) as an initial model of our method,

$$\min_{\theta_f, \theta_g} \max_{\theta_d} \mathcal{L}_{cls}^w(\theta_f, \theta_g) + \alpha \mathcal{L}_{ent}(\theta_f, \theta_g) + \lambda \mathcal{L}_{adv}^e(\theta_f, \theta_d), \quad (4)$$

where $\mathcal{L}_{cls}^w(\theta_f, \theta_g) = \frac{1}{n_s} \sum_{i=1}^{n_s} m(y_i^s) l_{ce}(G(F(x_i^s)), y_i^s)$, m denotes the normalized estimated class-level weight vector via the target domain, and α, λ are two empirical trade-off parameters.

3.2 Balanced Adversarial Alignment (BAA)

The joint distribution shift is the actual root to negative transfer [45]. For example, the marginal label distributions are not symmetric in PDA, and thus source-only classes are prone to be matched with target classes, resulting in a negative transfer problem. Generally, a class-level source re-weighting scheme sounds a natural choice for PDA, since it is expected to filter out source-only classes and promote positive transfer between the shared classes across domains. Previous methods [5, 30] resort to target predictions to generate class-level weights and effectively avoid negative transfer to some degree. Ideally, PDA with a weighting scheme behaves like a small UDA problem, but it heavily relies on accurate target predictions to calculate suitable weights.

In contrast, we propose an extremely simple scheme as shown in Fig. 1 for the challenging distribution alignment in PDA. Intuitively, we pursue the balance between different label distributions across domains with quite an opposite solution, i.e., augmenting the target domain using original source samples instead of weighting the source domain. This idea looks weird but is actually reasonable because we readily turn the PDA task into a large UDA-like task and the negative transfer effects caused by source-only classes can be well alleviated. The detailed formulation of balanced augmentation (alignment) is shown below,

$$\begin{aligned} \mathcal{L}_{adv}^{ba}(\theta_f, \theta_d) = & \frac{1}{n_s} \sum_{i=1}^{n_s} w(x_i^s) m(y_i^s) \log[D(F(x_i^s))] + \\ & \frac{1}{n_t} \sum_{j=1}^{n_t} w(x_j^t) \log[1 - D(F(x_j^t))] + \frac{\rho}{n_s} \sum_{i=1}^{n_s} w(x_i^s) m(y_i^s) \log[1 - D(F(x_i^s))]. \end{aligned} \quad (5)$$

Specifically, we adopt a progressive strategy for the target augmentation scheme. We borrow the source samples with different ratios, and the ratio ρ gradually decreases to 0 as the number of iterations increases. This is because the learned feature representations in early iterations are not quite transferable and we need more source samples to avoid class mismatch. When the learned features are desirably discriminative and transferable, the estimation of class-level weights becomes more accurate, making the augmentation trivial.

Note that we borrow original samples from the source domain rather than exploiting a generative model like CycleGAN [51] to synthesize target-like source images. The reason is that obtaining data-dependent translation models between such heterogeneous domains is quite time-consuming and we find translation does not even improve the adaptation results.

3.3 Adaptive Uncertainty Suppression (AUS)

Previous DA methods focus on strengthening the feature transferability by developing various domain alignment strategies, but they mostly ignore the feature discriminability and simply use the conventional cross-entropy loss in the labeled source domain to learn the features. In that case, even though the domain shift is mitigated, the classifier may perform worse on target data. This is because source classes are not equally separated from each other, which may lead to the propagation of the confusion (uncertainty) to target predictions and thus confusing features in the target domain, which is termed as the uncertainty propagation problem. Take a 3-way classification problem as an example. There is no class overlap in the source domain, but class 1 and 2 may be close to each other. A toy example of how class 1




Fig. 2. Mitigating the effects of uncertainty propagation from source. [blue: source, red: target, gray: adversarial alignment.]

and class 2 behave is shown in Fig. 2 where a few samples lie close to the decision boundary between class 1 and class 2. During domain alignment, some unlabeled target data are enforced to match these source data, which would be easily misclassified. However, such a critical problem has always been overlooked in prior literature. Looking back at the cross-entropy loss $l_{ce}(\hat{y}, y) = -\sum_i y_i \log(\hat{y}_i)$, it only exploits the information from the ground-truth class while ignoring that from other incorrect classes. For example, the source output $[0.6, 0.3, 0.1]$ is more uncertain than $[0.6, 0.2, 0.2]$, but both have the same cross-entropy loss.

Though several previous methods [20, 35] incorporate virtual adversarial training [32] in the classification term to implicitly increase the margin between different classes, more computation complexity is required and several parameters need to be tuned. Inspired by [7], we exploit a complement entropy that expects uniform and low prediction scores for incorrect classes for labeled source samples. To accurately suppress uncertainty, we further place more emphasis on the uncertain samples that own smaller cross-entropy loss (confidence) and propose a confidence-weighted complement entropy objective below,

$$\mathcal{L}_{wce}^w(\theta_f, \theta_g) = \frac{1}{n_s \log(K-1)} \sum_{i=1}^{n_s} m(y_i^s) l_{wce}(G(F(x_i^s)), y_i^s), \quad (6)$$

where $l_{wce}(\hat{y}, y) = (1 - \hat{y}_a)^\xi \sum_{j \neq a} \frac{\hat{y}_j}{1 - \hat{y}_a} \log\left(\frac{\hat{y}_j}{1 - \hat{y}_a}\right)$,

where ξ is a hyper-parameter and a is the index of ground-truth class in y . Different from the complementary training strategy in [7], we exploit the adaptive weighted complement entropy \mathcal{L}_{wce} as a regularizer, which is more efficient. We also assign a class-level weight for each sample in Eq. (6) like that in \mathcal{L}_{cls}^w .

3.4 Unified Minimax Optimization Problem of BA³US

Finally, we integrate all the terms in Eqs. (4, 5, 6) on both source and target samples to avoid negative transfer and uncertainty prorogation and derive a unified framework for PDA. The overall min-max objective is formulated as

$$\min_{\theta_f, \theta_g} \max_{\theta_d} \mathcal{L}_{cls}^w(\theta_f, \theta_g) + \alpha \mathcal{L}_{ent}(\theta_f, \theta_g) + \beta \mathcal{L}_{wce}^w(\theta_f, \theta_g) + \lambda \mathcal{L}_{adv}^{ba}(\theta_f, \theta_d), \quad (7)$$

where β is another trade-off hyper-parameter to balance the complement entropy and the cross-entropy term. Again, $m \propto \sum_{j=1}^{n_t} G(F(x_i^t))$ is the normalized estimated class-level weight vector. To optimize the objective above, we follow [9] to introduce a gradient reversal layer and adopt the same progressive strategy for parameter λ , i.e., increasing λ from 0 to 1 as the number of iterations grows.

Our method is closely related to DANN [9], sharing similar formalism of the domain adaptation theory [1] that the expected target risk $\epsilon_T(H)$ on the target examples is bounded by the source risk $\epsilon_S(H)$ and other two terms below,

$$\epsilon_T(H) \leq \epsilon_S(H) + |\epsilon_S(H, H^*) - \epsilon_T(H, H^*)| + [\epsilon_S(H^*) + \epsilon_T(H^*)], \quad (8)$$

where $H^* = \arg \min_{x \in \mathcal{H}} [\epsilon_S(x) + \epsilon_T(x)]$ is the ideal joint hypothesis for the combined risk. DANN further discovers the second term $\mathcal{H}\Delta\mathcal{H}$ -distance can be

upper bounded by error of the domain adversarial classifier. As we do not have labels of the target domain, we expect the entropy minimization term on the target domain to help reduce the last term. Besides, the proposed complement entropy objective alleviates uncertainty propagation to make the weight estimation more accurate, and thus our method would turn PDA into a small UDA task. In this way, we can expect that our method minimizes the empirical target risk $\mathbf{E}_{(x_t, y_t) \sim p_t}[f(x_t) \neq y_t]$. The detailed optimization is summarized in Algorithm 1.

4 Experiments

4.1 Setup

Datasets. **Office31** dataset [34] includes images of 31 object classes from three different domains, i.e., Amazon, DSLR, and Webcam. We follow the standard protocol used in [5] and pick up images of 10 categories shared by Office31 and Caltech256 [14] as target domains. **Office-Home** dataset [42] consists of 4 different domains with each containing 65 kinds of everyday objects, i.e., Artistic, Clipart, Product images, and Real World images. Likewise, we follow [5] to select the first 25 categories (in alphabetic order) in each domain as a partial target domain. **ImageNet-Caltech** is a large-scale object recognition dataset that consists of two subsets, ImageNet-1K [33] and Caltech256 [14]. Here we use images from the public validation set of ImageNet-1K for the target domain. In reality, each source domain contains 1,000 and 256 classes, and each target domain contains 84 classes.

Baseline methods. We utilize all the source and target samples and report the average classification accuracy and standard deviation over 3 random trials. $\mathbf{A} \rightarrow \mathbf{B}$ means \mathbf{A} is the source domain and \mathbf{B} is the partial target domain. For comprehensive comparison, we provide the recognition results of our methods including E-DANN, Ours (w/ BAA) and Ours (BA³US) on each dataset, and compare them with some popular UDA methods [40, 26] and existing PDA methods, including SAN [4], IWAN [49], PADA [5], SSPDA [3], MWPDA [17], DRCN [21], ETN [6], and SAFN [46].

Implementation details. If not specified, all the methods adopt **ResNet-50** [15] as backbone. We fine-tune the pre-trained ImageNet model in **PyTorch** using a NVIDIA Titan X (12 GB memory). The adversarial layer and classifier layer are trained through back-propagation, and the learning rate of the classifier layer is 10 times that of lower layers. We adopt mini-batch SGD with momentum of 0.9 and the learning rate annealing strategy as [10, 26]: the learning rate is adjusted by $\eta_p = \eta_0(1 + \hat{\alpha}p)^{-\hat{\beta}}$, where p denotes the training progress changing from 0 to 1, and $\eta_0 = 0.01$, $\hat{\alpha} = 10$, $\hat{\beta} = 0.75$ as suggested by [26]. Based on the number of classes and balancing analysis in [7], we use $\beta = 5$ for Office31 and ImageNet-Caltech, and $\beta = 1$ for Office-Home. Besides, we set the number of intervals N/N_u to 10, and N_u is set to 200, 500, and 4,000 for Office31, Office-Home, and ImageNet-Caltech, respectively. Note that our method does not use the ten-crop technique [5, 26] at the evaluation phase for better performance.

Table 1. Accuracy (%) on **Office-Home** dataset for *partial domain adaptation* via ResNet-50 [15]. The best in **bold red**; the second best in *italic blue*.

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 [15]	46.33	67.51	75.87	59.14	59.94	62.73	58.22	41.79	74.88	67.40	48.18	74.17	61.35
ADDA [40]	45.23	68.79	79.21	64.56	60.01	68.29	57.56	38.89	77.45	70.28	45.23	78.32	62.82
CDAN+E [26]	47.52	65.91	75.65	57.07	54.12	63.42	59.60	44.30	72.39	66.02	49.91	72.80	60.73
IWAN [49]	53.94	54.45	78.12	61.31	47.95	63.32	54.17	52.02	81.28	76.46	56.75	82.90	63.56
SAN [4]	44.42	68.68	74.60	67.49	64.99	77.80	59.78	44.72	80.07	72.18	50.21	78.66	65.30
PADA [5]	51.95	67.00	78.74	52.16	53.78	59.03	52.61	43.22	78.79	73.73	56.60	77.09	62.06
SSPDA [3]	52.02	63.64	77.95	65.66	59.31	73.48	70.49	51.54	84.89	76.25	<i>60.74</i>	80.86	68.07
MWPD [17]	55.39	77.53	81.27	57.08	61.03	62.33	68.74	56.42	86.67	76.70	57.67	80.06	68.41
ETN [6]	<i>59.24</i>	77.03	79.54	62.92	65.73	75.01	68.29	55.37	84.37	75.72	57.66	<i>84.54</i>	70.45
DRCN [21]	54.00	76.40	83.00	62.10	64.50	71.00	70.80	49.80	80.50	77.50	59.10	79.90	69.00
SAFN [46]	58.93	76.25	81.42	70.43	72.97	77.78	72.36	55.34	80.40	75.81	60.42	79.92	71.83
	±0.50	±0.33	±0.27	±0.46	±1.39	±0.52	±0.31	±0.46	±0.78	±0.37	±0.83	±0.20	
E-DANN	54.05	74.12	84.06	67.06	64.95	75.15	71.29	53.09	83.42	76.00	58.17	81.53	70.24
	±0.45	±0.12	±0.19	±0.19	±1.10	±0.59	±0.19	±0.37	±0.22	±0.65	±0.51	±0.26	
Ours (w/ BAA)	56.20	<i>79.55</i>	<i>86.21</i>	<i>70.86</i>	69.94	<i>81.06</i>	<i>72.51</i>	<i>57.91</i>	86.47	<i>77.10</i>	59.34	83.64	<i>73.40</i>
	±0.28	±0.45	±0.32	±0.67	±2.52	±0.30	±0.32	±0.30	±0.32	±0.62	±0.46	±0.39	
Ours (BA ³ US)	60.62	83.16	88.39	71.75	<i>72.79</i>	83.40	75.45	61.59	<i>86.53</i>	79.25	62.80	86.05	75.98
	±0.45	±0.12	±0.19	±0.19	±1.10	±0.59	±0.19	±0.37	±0.22	±0.65	±0.51	±0.26	

4.2 Quantitative Results for Partial Domain Adaptation

The results on three object recognition datasets including **Office-Home**, **Office31** and **ImageNet-Caltech** for PDA are shown in Table 1 and 2, with some baseline results directly reported from ETN [6] with the same protocol. Obviously, BA³US achieves the best or second-best results on 10 out of 12 transfer tasks on the **Office-Home** dataset, and Ours (w/ BAA) and SAFN obtain the second and third best results, respectively. Regarding the average accuracy, BA³US advances the state-of-the-art result on **Office-Home** in SAFN [46] by 5.78%, from 71.83% to 75.98%. For two specific tasks $Cl \rightarrow Ar$ and $Pr \rightarrow Rw$, BA³US merely performs slightly worse than the best method. Besides, a state-of-the-art UDA approach CDAN+E [26] performs worse than ResNet-50, which implies the difficulty of the partial transfer task. Further compared with UDA methods, even PDA methods like IWAN [49] and PADA [5] do not work well, which again indicates the partial transfer is quite challenging.

On the small-scale **Office31** dataset, BA³US again obtains the best average accuracy and performs the best in 3 out of 6 transfer tasks. For transfer tasks from a large source domain A to small target domains (D , W), BA³US remarkably outperforms other PDA methods. BA³US performs slightly worse for the $D \rightarrow A$ task because the target domain D is very small, making the proposed target augmentation and adaptive uncertainty suppression techniques inefficient. On the large-scale **ImageNet-Caltech** dataset, BA³US performs the best for both transfer tasks and still holds the best average accuracy with significant improvements. Moreover, UDA methods do not always perform better than ResNet-50, which implies they may suffer from the negative transfer problem.

Besides the ResNet-50 backbone network, we further investigate the effectiveness of BA³US with another backbone network VGG-16 [36] on **Office31** and compare it with state-of-the-art methods in Table 3. It can be clearly observed that BA³US achieves the best or second-best results for all the tasks,

significantly advancing the average accuracy from 93.88% to 95.84%. Compared with the results in Table 2, we find BA³US (97.81%→95.84%) is also robust than ETN (96.73%→93.88%) w.r.t. the change of the backbone network.

Table 2. Accuracy (%) on **Office31** and **ImageNet-Caltech** for *partial domain adaptation* via ResNet-50 [15]. The best in **bold red**; the second best in *italic blue*.

Method	Office31							ImageNet-Caltech			
	A → D	A → W	D → A	D → W	W → A	W → D	Avg.	I → C	C → I	Avg.	
ResNet-50 [15]	83.44±1.12	75.59±1.09	83.92±0.95	96.27±0.85	84.97±0.86	98.09±0.74	87.05	69.69±0.78	71.29±0.74	70.49	
ADDA [40]	83.41±0.17	75.67±0.17	83.62±0.14	95.38±0.23	84.25±0.13	99.85±0.12	87.03	71.82±0.45	69.32±0.41	70.57	
CDAN+E [26]	77.07±0.90	80.51±1.20	93.58±0.07	98.98±0.00	91.65±0.00	98.09±0.00	89.98	72.45±0.07	72.02±0.13	72.24	
IWAN [49]	90.45±0.36	89.15±0.37	95.62±0.29	99.32±0.32	94.26±0.25	99.36±0.24	94.69	78.06±0.40	73.33±0.46	75.70	
SAN [4]	94.27±0.28	93.90±0.45	94.15±0.36	99.32±0.52	88.73±0.44	99.36±0.12	94.96	77.75±0.36	75.26±0.42	76.51	
PADA [5]	82.17±0.37	86.54±0.31	92.69±0.29	99.32±0.45	95.41±0.33	100.0±0.00	92.69	75.03±0.36	70.48±0.44	72.76	
SSPDA [3]	90.87	91.52	90.61	92.88	94.36	98.94	93.20	-	-	-	
MWPDA [17]	95.12	96.61	95.02	100.0	95.51	100.0	97.05	-	-	-	
DRCN [21]	86.00	88.05	95.60	100.0	95.80	100.0	94.30	75.30	78.90	77.10	
ETN [6]	95.03±0.22	94.52±0.20	96.21±0.27	100.0±0.00	94.64±0.24	100.0±0.00	96.73	83.23±0.24	74.93±0.44	79.08	
E-DANN	92.36±0.00	93.22±0.00	94.61±0.05	100.0±0.00	94.71±0.05	98.73±0.00	95.60	78.31±0.81	77.69±0.25	78.00	
Ours (w/ BAA)	93.63±0.00	93.90±0.00	94.89±0.09	100.0±0.00	94.78±0.00	100.0±0.00	96.20	82.97±0.49	79.34±0.08	81.16	
Ours (BA ³ US)	99.36±0.00	98.98±0.28	94.82±0.05	100.0±0.00	94.99±0.08	98.73±0.00	97.81	84.00±0.15	83.35±0.28	83.68	

Table 3. Accuracy (%) on **Office31** for *partial domain adaptation* via VGG-16 [36]. The best in **bold red**; the second best in *italic blue*.

Method	A → D	A → W	D → A	D → W	W → A	W → D	Avg.
VGG-16 [36]	76.43±0.48	60.34±0.84	72.96±0.56	97.97±0.63	79.12±0.54	<i>99.36</i> ±0.36	81.03
IWAN [49]	<i>90.95</i> ±0.33	82.90±0.31	89.57±0.24	79.75±0.26	93.36±0.22	88.53±0.16	87.51
SAN [4]	90.70±0.20	83.39±0.36	87.16±0.23	<i>99.32</i> ±0.45	91.85±0.35	100.0 ±0.00	92.07
PADA [5]	81.73±0.34	<i>86.05</i> ±0.36	93.00±0.24	100.0 ±0.00	<i>95.26</i> ±0.27	100.0 ±0.00	92.54
ETN [6]	89.43±0.17	85.66±0.16	95.93 ±0.23	100.0 ±0.00	92.28±0.20	100.0 ±0.00	<i>93.88</i>
Ours (BA ³ US)	95.54 ±0.00	89.83 ±0.00	<i>94.92</i> ±0.05	<i>99.32</i> ±0.00	95.41 ±0.00	100.0 ±0.00	95.84

Ablation study. As shown in Tables 1 and 2, we provide the results of E-DANN and Ours (w/ BAA) along with BA³US on three datasets. Ours (w/ BAA) extends E-DANN by using the proposed balanced adversarial alignment technique in Sec. 3.2 instead, while Ours (BA³US) extends Ours (w/ BAA) by considering the proposed adaptive complement entropy objective in Eq. (6). Firstly, the baseline method E-DANN always performs well for partial domain adaptation since it removes the irrelevant classes in the source classification term like [5]. Secondly, results on all three datasets demonstrate that Ours (BA³US) performs better than Ours (w/ BAA) and Ours (w/ BAA) performs better than E-DANN, which verify the effectiveness of two proposed techniques in Sec. 3.2 and Sec. 3.3.

4.3 Quantitative Results for Closed-set Domain Adaptation

This section further investigates the effectiveness of the proposed uncertainty suppression technique for vanilla closed-set domain adaptation. Here we con-

sider integrating them with DANN and CDAN [26] respectively, and compare our methods with state-of-the-art UDA approaches including [44,46] on the most-favored **Office-Home** dataset. As shown in Table 4, the proposed method BA³US built on CDAN obtains the best average accuracy and ranks the top two in 9 out of 12 different transfer tasks. It is obvious that the adaptive uncertainty suppression technique works well for closed-set domain adaptation, advancing the average accuracy from 67.6% to 68.7% and from 68.0% to 69.2%. In fact, we also study the effectiveness of balanced adversarial alignment but find it hardly improve the performance since the label distributions in UDA have already been symmetric.

Table 4. Accuracy (%) on **Office-Home** dataset for *vanilla unsupervised domain adaptation* via ResNet-50 [15]. Methods* utilize augmentation during evaluation.

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.
CDAN+E [26]*	50.7	70.6	76.0	57.6	70.0	70.0	57.4	50.9	77.3	70.9	56.7	81.6	65.8
MCS [23]	55.9	73.8	79.0	57.5	69.9	71.3	58.4	50.3	78.2	65.9	53.2	82.2	66.3
DRCN [21]	50.6	72.4	76.8	61.9	69.5	71.3	60.4	48.6	76.8	72.9	56.1	81.4	66.6
CDAN+TransNorm [44]*	50.2	71.4	77.4	59.3	72.7	73.1	61.0	53.1	79.5	71.9	59.0	82.9	67.6
SAFN [46]*	54.4	73.3	77.9	65.2	71.5	73.2	63.6	52.6	78.2	72.3	58.0	82.1	68.5
SAFN [46]	52.0	71.7	76.3	64.2	69.9	71.9	63.7	51.4	77.1	70.9	57.1	81.5	67.3
Ours (w/ BAA)	50.9	72.0	77.5	61.2	72.6	72.7	62.8	52.7	79.9	70.8	56.6	82.7	67.6
Ours (BA ³ US)	51.2	73.8	78.1	63.3	73.4	73.6	63.3	54.5	80.4	72.6	56.7	83.7	68.7
Ours (w/ BAA)+CDAN	52.2	73.1	77.9	61.4	72.7	73.2	61.0	51.8	80.0	72.0	57.8	83.3	68.0
Ours (BA ³ US)+CDAN	54.1	74.2	77.7	62.9	73.6	74.6	63.4	54.9	80.4	73.1	58.2	83.6	69.2

4.4 Qualitative Results for Partial Domain Adaptation

We study our methods with different numbers of target classes in Fig. 3(a). The performance decreases when the number is larger than 15, and BA³US always obtains the best results. As expected, the proposed augmentation technique becomes important when the number of target classes is small.

Weight visualization. As stated before, the weighting scheme plays an important role in PDA. Thus we investigate the quality of the estimated class-level weight m in Algorithm 1. As shown in Fig. 3(b-c), we plot the estimated weights of BA³US for the two specific tasks. Since the *Cl* domain and *Ar* domain are quite different, making the estimation challenging. This is also evidenced in other PDA methods in Table 1 since the accuracy is around 55%. For the relatively easy task A→D, the weight estimation seems accurate, resulting in high classification accuracy.

Parameter Sensitivity. We study the sensitivity of parameters ξ and β of BA³US on the **Office-Home** and **Office31** datasets. The mean accuracy is reported in Table 5 and Table 6, respectively. Note that the complement entropy [7] can be considered as a special case of the proposed adaptive one where $\xi = 0$. The results indicate using an adaptive objective is much better, and the performance is relatively stable. Regarding the parameter β , the accuracy around 1.0 and 5.0 is also stable, implying our method is not sensitive.




Fig. 3. (a) Accuracy with varying number of target classes. (b-c) Estimated class-level weights. Yellow bins denote ground-truth classes. Best viewed in color.

Table 5. Sensitivity of parameter ξ .

Avg. (%)	0.0 [7]	0.1	0.3	0.5	0.7	0.9	1.0
Office-Home	75.32	75.88	76.28	76.10	76.10	75.81	75.98
Office31	97.68	97.71	97.65	97.67	97.64	97.84	97.81

Table 6. Sensitivity of parameter β .

Avg. (%)	0.0	0.1	0.5	1.0	5.0	10.0
Office-Home	73.40	73.58	75.09	75.98	75.69	73.25
Office31	96.20	96.13	96.50	96.63	97.81	97.83

Convergence performance. As shown in Fig. 4, we study the convergence performance of the proposed methods for Ar→Cl and A→D. Obviously, the ‘source only’ method works worse without the domain alignment module, and E-DANN performs much better than it and quickly converges after 1,000 iterations. Besides, both BA³US and Ours (w/ BAA) obtain similar promising results, and BA³US performs slightly better since it further considers the adaptive complement entropy to diminish the uncertainty in source predictions.

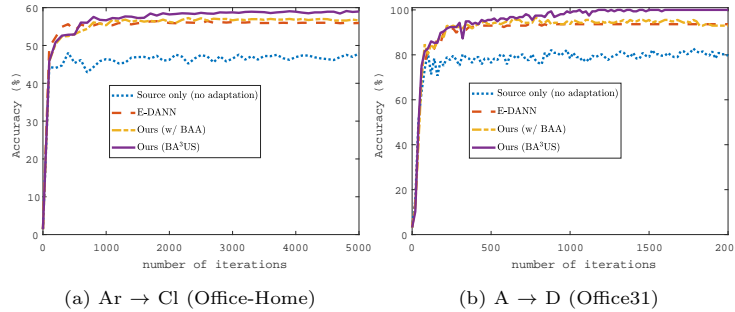


Fig. 4. Convergence analysis of proposed methods on two different transfer tasks. Accuracy (%) is given w.r.t. number of iterations. Best viewed in color.

Feature visualization. We plot in Fig. 5 the t-SNE embeddings [28] of the features learned by ‘source only’, E-DANN and BA³US for the transfer task Ar→Cl. It is easy to discover that the features of target data are rather confusing in ‘no adaptation’ while the balanced alignment module helps mitigate the domain gap and the adaptive uncertainty suppression module helps increase the discrimination of the features.




Fig. 5. t-SNE visualizations for the transfer task Ar→Cl. Blue: source; red: target.

5 Conclusion

We develop a novel adversarial learning-based method BA³US for partial domain adaptation, which well addresses two key problems, negative transfer and uncertainty propagation. To tackle the asymmetric label distributions, BA³US offers a very simple solution by augmenting the target domain with samples from the source domain. Then, it uncovers an overlooked issue in the field termed uncertainty propagation and designs an adaptive complement entropy objective to well suppress the uncertainty in source predictions. Empirical results show it also works well for vanilla closed-set domain adaptation. Further experiments have validated that BA³US improves existing methods with substantial gains, establishing new state-of-the-art.

Acknowledgment

J. Feng was partially supported by NUS ECRA FY17 P08, AISG-100E2019-035, and MOE Tier 2 MOE2017-T2-2-151. The authors also thank Quanhong Fu for her help to improve the technical writing aspect of this paper.

Algorithm 1: Pseudo code of our method termed BA³US.

Input: Labeled source domain \mathcal{D}_s , unlabeled target domain \mathcal{D}_t ;
Parameters: Total training iterations N , updating interval N_u , batch size $B_s = 36$, $\rho_0 = 1/4$, $\xi = 1$, $\alpha = 0.1$, $\beta \in \{1, 5\}$, λ ;
Initialize the model parameters $\theta_f, \theta_g, \theta_d$;
Initialize the class-level weight vector m , $m_i = 1/C$;
for $i = 1$ **to** N **do**
 Obtain B_s samples from \mathcal{D}_s and \mathcal{D}_t , respectively;
 Obtain ρB_s random samples from \mathcal{D}_s ;
 Update $\theta_f, \theta_g, \theta_d$ by optimizing Eq. (7) and gradient reversal layer;
 if $i \% N_u == 0$ **then**
 update the class-level weight vector m ;
 % note that this step is ignored in closed-set domain adaptation
 calculate \mathcal{L}_{ent} in Eq. (3) for model selection;
 $\rho \leftarrow \rho_0 (1 - N_u/N)$;
 end
end
Output: Target outputs corresponding to the minimal value of \mathcal{L}_{ent} .

References

1. Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., Vaughan, J.W.: A theory of learning from different domains. *Machine Learning* **79**(1-2), 151–175 (2010) [8](#)
2. Bousmalis, K., Silberman, N., Dohan, D., Erhan, D., Krishnan, D.: Unsupervised pixel-level domain adaptation with generative adversarial networks. In: *Proc. CVPR*. pp. 3722–3731 (2017) [3](#)
3. Bucci, S., D’Innocente, A., Tommasi, T.: Tackling partial domain adaptation with self-supervision. *arXiv preprint arXiv:1906.05199* (2019) [9](#), [10](#), [11](#)
4. Cao, Z., Long, M., Wang, J., Jordan, M.I.: Partial transfer learning with selective adversarial networks. In: *Proc. CVPR*. pp. 2724–2732 (2018) [2](#), [3](#), [6](#), [9](#), [10](#), [11](#)
5. Cao, Z., Ma, L., Long, M., Wang, J.: Partial adversarial domain adaptation. In: *Proc. ECCV*. pp. 135–150 (2018) [2](#), [3](#), [4](#), [6](#), [9](#), [10](#), [11](#)
6. Cao, Z., You, K., Long, M., Wang, J., Yang, Q.: Learning to transfer examples for partial domain adaptation. In: *Proc. CVPR*. pp. 2985–2994 (2019) [3](#), [4](#), [9](#), [10](#), [11](#)
7. Chen, H.Y., Wang, P.H., Liu, C.H., Chang, S.C., Pan, J.Y., Chen, Y.T., Wei, W., Juan, D.C.: Complement objective training. In: *Proc. ICLR* (2019) [2](#), [8](#), [9](#), [12](#), [13](#)
8. Csurka, G.: Domain adaptation for visual applications: A comprehensive survey. *arXiv preprint arXiv:1702.05374* (2017) [3](#)
9. Ganin, Y., Lempitsky, V.: Unsupervised domain adaptation by backpropagation. In: *Proc. ICML*. pp. 1180–1189 (2015) [2](#), [5](#), [6](#), [8](#)
10. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., Lempitsky, V.: Domain-adversarial training of neural networks. *Journal of Machine Learning Research* **17**(1), 2096–2030 (2016) [2](#), [3](#), [5](#), [9](#)
11. Ghifary, M., Kleijn, W.B., Zhang, M., Balduzzi, D., Li, W.: Deep reconstruction-classification networks for unsupervised domain adaptation. In: *Proc. ECCV*. pp. 597–613 (2016) [3](#)
12. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: *Proc. NeurIPS*. pp. 2672–2680 (2014) [3](#), [5](#)
13. Grandvalet, Y., Bengio, Y.: Semi-supervised learning by entropy minimization. In: *Proc. NeurIPS*. pp. 529–536 (2005) [5](#)
14. Griffin, G., Holub, A., Perona, P.: Caltech-256 object category dataset (2007) [9](#)
15. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proc. CVPR*. pp. 770–778 (2016) [9](#), [10](#), [11](#), [12](#)
16. Hoffman, J., Tzeng, E., Park, T., Zhu, J.Y., Isola, P., Saenko, K., Efros, A., Darrell, T.: Cycada: Cycle-consistent adversarial domain adaptation. In: *Proc. ICML*. pp. 1989–1998 (2018) [4](#)
17. Hu, J., Wang, C., Qiao, L., Zhong, H., Jing, Z.: Multi-weight partial domain adaptation. In: *Proc. BMVC* (2019) [3](#), [9](#), [10](#), [11](#)
18. Koniusz, P., Tas, Y., Porikli, F.: Domain adaptation by mixture of alignments of second-or higher-order scatter tensors. In: *Proc. CVPR*. pp. 7139–7148 (2017) [2](#), [3](#)
19. Kouw, W.M., Loog, M.: A review of single-source unsupervised domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (In press, 2019) [3](#)
20. Kumar, A., Sattigeri, P., Wadhawan, K., Karlinsky, L., Feris, R., Freeman, B., Wornell, G.: Co-regularized alignment for unsupervised domain adaptation. In: *Proc. NeurIPS*. pp. 9345–9356 (2018) [8](#)

21. Li, S., Liu, C.H., Lin, Q., Wen, Q., Su, L., Huang, G., Ding, Z.: Deep residual correction network for partial domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (In press, 2020) [4](#), [9](#), [10](#), [11](#), [12](#)
22. Liang, J., He, R., Sun, Z., Tan, T.: Aggregating randomized clustering-promoting invariant projections for domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **41**(5), 1027–1042 (2019) [3](#)
23. Liang, J., He, R., Sun, Z., Tan, T.: Distant supervised centroid shift: A simple and efficient approach to visual domain adaptation. In: *Proc. CVPR*. pp. 2975–2984 (2019) [12](#)
24. Liu, M.Y., Tuzel, O.: Coupled generative adversarial networks. In: *Proc. NeurIPS*. pp. 469–477 (2016) [5](#)
25. Long, M., Cao, Y., Wang, J., Jordan, M.: Learning transferable features with deep adaptation networks. In: *Proc. ICML*. pp. 97–105 (2015) [3](#)
26. Long, M., Cao, Z., Wang, J., Jordan, M.I.: Conditional adversarial domain adaptation. In: *Proc. NeurIPS*. pp. 1647–1657 (2018) [2](#), [3](#), [5](#), [9](#), [10](#), [11](#), [12](#)
27. Long, M., Zhu, H., Wang, J., Jordan, M.I.: Deep transfer learning with joint adaptation networks. In: *Proc. ICML*. pp. 2208–2217 (2017) [3](#)
28. Maaten, L.v.d., Hinton, G.: Visualizing data using t-sne. *Journal of Machine Learning Research* **9**(Nov), 2579–2605 (2008) [13](#)
29. Mao, X., Ma, Y., Yang, Z., Chen, Y., Li, Q.: Virtual mixup training for unsupervised domain adaptation. *arXiv preprint arXiv:1905.04215* (2019) [4](#)
30. Matsuura, T., Saito, K., Harada, T.: Twins: Two weighted inconsistency-reduced networks for partial domain adaptation. *arXiv preprint arXiv:1812.07405* (2018) [2](#), [3](#), [4](#), [6](#)
31. Ming Harry Hsu, T., Yu Chen, W., Hou, C.A., Hubert Tsai, Y.H., Yeh, Y.R., Frank Wang, Y.C.: Unsupervised domain adaptation with imbalanced cross-domain data. In: *Proc. ICCV*. pp. 4121–4129 (2015) [3](#)
32. Miyato, T., Maeda, S.i., Koyama, M., Ishii, S.: Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **41**(8), 1979–1993 (2018) [8](#)
33. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *International Journal of Computer Vision* **115**(3), 211–252 (2015) [9](#)
34. Saenko, K., Kulis, B., Fritz, M., Darrell, T.: Adapting visual category models to new domains. In: *Proc. ECCV*. pp. 213–226 (2010) [3](#), [9](#)
35. Shu, R., Bui, H.H., Narui, H., Ermon, S.: A dirt-t approach to unsupervised domain adaptation. In: *Proc. ICLR* (2018) [8](#)
36. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014) [10](#), [11](#)
37. Sun, B., Saenko, K.: Deep coral: Correlation alignment for deep domain adaptation. In: *Proc. ECCV Workshops*. pp. 443–450 (2016) [2](#), [3](#)
38. Tsai, Y.H.H., Hou, C.A., Chen, W.Y., Yeh, Y.R., Wang, Y.C.F.: Domain-constraint transfer coding for imbalanced unsupervised domain adaptation. In: *Proc. AAAI*. pp. 3597–3603 (2016) [3](#)
39. Tzeng, E., Hoffman, J., Darrell, T., Saenko, K.: Simultaneous deep transfer across domains and tasks. In: *Proc. ICCV*. pp. 4068–4076 (2015) [2](#), [3](#)
40. Tzeng, E., Hoffman, J., Saenko, K., Darrell, T.: Adversarial discriminative domain adaptation. In: *Proc. CVPR*. pp. 2962–2971 (2017) [2](#), [3](#), [5](#), [9](#), [10](#), [11](#)
41. Tzeng, E., Hoffman, J., Zhang, N., Saenko, K., Darrell, T.: Deep domain confusion: Maximizing for domain invariance. *arXiv preprint arXiv:1412.3474* (2014) [2](#), [3](#)

42. Venkateswara, H., Eusebio, J., Chakraborty, S., Panchanathan, S.: Deep hashing network for unsupervised domain adaptation. In: Proc. CVPR. pp. 5018–5027 (2017) [3](#), [9](#)
43. Wang, Q., Li, W., Van Gool, L.: Semi-supervised learning by augmented distribution alignment. In: Proc. ICCV. pp. 1466–1475 (2019) [4](#)
44. Wang, X., Jin, Y., Long, M., Wang, J., Jordan, M.I.: Transferable normalization: Towards improving transferability of deep neural networks. In: Proc. NeurIPS. pp. 1951–1961 (2019) [12](#)
45. Wang, Z., Dai, Z., Póczos, B., Carbonell, J.: Characterizing and avoiding negative transfer. In: Proc. CVPR. pp. 11293–11302 (2019) [6](#)
46. Xu, R., Li, G., Yang, J., Lin, L.: Larger norm more transferable: An adaptive feature norm approach for unsupervised domain adaptation. In: Proc. ICCV. pp. 1426–1435 (2019) [9](#), [10](#), [12](#)
47. Zellinger, W., Grubinger, T., Lughofer, E., Natschläger, T., Saminger-Platz, S.: Central moment discrepancy (cmd) for domain-invariant representation learning. In: Proc. ICLR (2016) [2](#), [3](#)
48. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. In: Proc. ICLR (2018) [4](#)
49. Zhang, J., Ding, Z., Li, W., Ogunbona, P.: Importance weighted adversarial nets for partial domain adaptation. In: Proc. CVPR. pp. 8156–8164 (2018) [2](#), [3](#), [4](#), [9](#), [10](#), [11](#)
50. Zhang, L.: Transfer adaptation learning: A decade survey. arXiv preprint arXiv:1903.04687 (2019) [3](#)
51. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proc. ICCV. pp. 2223–2232 (2017) [3](#), [4](#), [7](#)