



Class-Imbalanced Domain Adaptation: An Empirical Odyssey

Shuhan Tan¹(✉), Xingchao Peng², and Kate Saenko^{2,3}

¹ Sun Yat-Sen University, Guangzhou, China
tanshh@mail2.sysu.edu.cn

² Boston University, Boston, USA
{xpeng, saenko}@bu.edu

³ MIT-IBM Watson AI Lab, Boston, USA

Abstract. Unsupervised domain adaptation is a promising way to generalize deep models to novel domains. However the current literature assumes that the label distribution is domain-invariant and only aligns the feature distributions or *vice versa*. In this work, we explore the more realistic task of *Class-imbalanced Domain Adaptation*: How to align feature distributions across domains *while* the label distributions of the two domains are also different? Taking a practical step towards this problem, we constructed its first benchmark with 22 cross-domain tasks from 6 real-image datasets. We conducted comprehensive experiments on 10 recent domain adaptation methods and find most of them are very fragile in the face of coexisting feature and label distribution shift. Towards a better solution, we further proposed a feature and label distribution CO-ALignment (COAL) model with a novel combination of existing ideas. COAL is empirically shown to outperform most recent domain adaptation methods on our benchmarks. We believe the provided benchmarks, empirical analysis results, and the COAL baseline could stimulate and facilitate future research towards this important problem.

1 Introduction

The success of deep learning models is highly dependent on the assumption that the training and testing data are *i.i.d* and sampled from the same distribution. In reality, they are typically collected from different but related domains, leading to a phenomenon known as *domain shift* [1]. To bridge the domain gap, Unsupervised Domain Adaptation (UDA) transfers the knowledge learned from a labeled source domain to an unlabeled target domain by statistical distribution alignment [2,3] or adversarial alignment [4–6]. Though recent UDA works

S. Tan—Work done while the author was visiting Boston University.

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-66415-2_38) contains supplementary material, which is available to authorized users.

have made great progress, most of them are under the assumption that the prior label distributions of the two domains are identical. Denote the input data as x and output labels as y , and let the source and target domain be characterized by probability distributions p and q , respectively. The majority of UDA methods assume that the conditional label distribution is invariant ($p(y|x) = q(y|x)$), and only the *feature shift* ($p(x) \neq q(x)$) needs to be tackled, neglecting potential *label shift* ($p(y) \neq q(y)$)¹. However, we claim that this assumption makes current UDA methods not applicable in the real world, for the following reasons: **1)** this assumption hardly holds true in real applications, as label shift across domains is commonly seen in the real world. For example, an autonomous driving system should be able to handle constantly changing frequencies of pedestrians and cars when adapting from a rural to a downtown area; or from a rainy to a sunny day. In addition, it is hard to guarantee $p(y) = q(y)$ without any information about $q(y)$ in the real world. **2)** recent theoretical work [9] has demonstrated that if label shift exists, current UDA methods could lead to significant *performance drop*. This is also empirically proved by our experiments. **3)** we cannot check whether label shift exists in real applications. This prevents us from safely applying current UDA methods because we cannot predict the potential risk of performance drop. Therefore, we claim that an applicable UDA method must be able to handle feature shift and label shift at the same time.




Fig. 1. We propose the Class-imbalanced Domain Adaptation setting, where we consider feature shift and label shift simultaneously. We provide the first empirical evaluation of this setting, showing that existing UDA methods are very fragile in the face of label shift. This is because learning marginal domain-invariant features will incorrectly align samples from different categories, leading to negative transfer. We propose an alternate, more robust approach that combines self-training and conditional feature alignment to tackle feature and label shift.

To formulate the above problem, we propose *Class-imbalanced Domain Adaptation* (CDA), a more challenging but practical domain adaptation setting where the conditional feature shift and label shift are required to be tackled *simultaneously*. Specifically, in addition to Covariate Shift assumption ($p(x) \neq q(x)$, $p(y|x) = q(y|x)$), we further assume $p(x|y) \neq q(x|y)$ and $p(y) \neq q(y)$. The main challenges of CDA are: **1)** label shift hampers the effectiveness of

¹ Different from some works [7, 8], we do not assume $p(x|y) = q(x|y)$ for label shift.

mainstream domain adaptation methods that only marginally aligns feature distributions, **2**) aligning the conditional feature distributions ($p(x|y), q(x|y)$) is difficult in the presence of label shift, and **3**) when data in one or both of the domains are unequally distributed across different categories, it is difficult to train an unbiased classifier. An overview of CDA is shown in Fig. 1.

Aligned with our idea, several works [10–12] provide theoretical analyses on domain adaptation with both feature and label shift. However, they do not provide sufficient empirical analysis of current UDA methods under this setting. In addition, no practical algorithm that can solve real-world cross-domain problems has been proposed by these works. Therefore, although this problem has been known for years, most recent UDA methods are still not able to handle it. In this paper, we aim raise concerns and interests towards this important problem by taking one practical step. Firstly, we create CDA benchmarks with 22 cross-domain tasks across 6 real-world image classification datasets. We believe this would facilitate future domain adaptation research towards robustly applicable methods. Secondly, we extensively evaluate 10 state-of-the-art domain adaptation methods to analysis how well CDA is solved currently. We find most of these methods cannot handle CDA well and often lead to negative transfer. Thirdly, towards a better solution, we provide a theoretically-motivated novel combination of existing ideas, which works well as a baseline for future research.

In this work, we visited domain adaptation methods in three categories. Mainstream unsupervised domain adaptation aligns the feature distributions of two domains by methods that include minimizing the Maximum Mean Discrepancy [2, 3], aligning high-order moments [13, 14], or adversarial training [4, 5]. However, these models are limited when applied to the CDA task as they only align the feature distribution, ignoring the issue of label shift [9]. Another line of works [7, 8] assume that only label shift exists ($p(y) \neq q(y)$) between two domains and the conditional feature distribution is invariant ($p(x|y) = q(x|y)$). These methods have achieved good performance when the data in both domains are sampled from the same feature distribution but under different label distributions. However, these models cannot handle the CDA task as the feature distribution is not well aligned. Recently, several works consider the domain adaptation problem where the categories of the source and target domain are not fully overlapped [15–17]. This setting can be seen as a special case of CDA where for some class i we have either $p(y = i) = 0$ or $q(y = i) = 0$. In our experiments, we showed that 8 out of 10 methods we evaluated on CDA tasks frequently lead to negative transfer (produce worse performance than no-adaptation baseline), while the rest methods only leads to limited improvement over the baseline on average. This limited performance showed that current UDA methods are not robust enough to be practically applied, and motivated us to reconsider the solution to the CDA problem.

We postulate that it is essential to align the conditional feature distributions as well as the label distributions to tackle the CDA task. In this work, we address CDA with feature distribution and label distribution **CO-Alignment (COAL)**. Specifically, to deal with feature shift and label shift in an unified

way, we proposed a simple baseline method that combines the ideas of *prototype-based conditional distribution alignment* [18] and *class-balanced self-training* [19]. First, to tackle feature shift in the context of label shift, it is essential to align the conditional rather than marginal feature distributions, to avoid the negative transfer effects caused by matching marginal feature distributions [9] (illustrated in Fig. 1). To this end, we use a prototype-based method to align the conditional feature distributions of the two domains. The *source* prototypes are computed by learning a similarity-based classifier, which are moved towards the *target* domain with a minimax entropy algorithm [18]. Second, we align the label distributions in the context of feature shift by training the classifier with estimated target label distribution through a class-balanced self-training method [19]. We incorporate the above feature distribution and label distribution alignment into an end-to-end deep learning framework, as illustrated in Fig. 2. Comprehensive experiments on standard cross-domain recognition benchmarks demonstrate that COAL achieves significant improvements over the state-of-the-art methods on the task of CDA.

The main contributions of this paper are highlighted as follows: **1)** to the best of our knowledge, we provide the first set of benchmarks and practical solution for domain adaptation under joint feature and label shift in deep learning, which is important for real-world applications; **2)** we deliver extensive experiments to demonstrate that state-of-the-art methods *fail* to align feature distribution in the presence of label distribution, or vice versa; **3)** we propose a simple yet effective feature and label distribution CO-ALignment (COAL) framework, which could be a useful baseline for future research towards practical domain adaptation. We believe the provided benchmarks, empirical analysis and the baseline model could trigger future research works towards more practical domain adaptation.

2 Related Work

Domain Adaptation for Feature Shift. Domain adaptation aims to transfer the knowledge learned from one or more source domains to a target domain. Recently, many unsupervised domain adaptation methods have been proposed. These methods can be taxonomically divided into three categories [20]. The first category is the discrepancy-based approach, which leverages different measures to align the marginal feature distributions between source and target domains. Commonly used measures include Maximum Mean Discrepancy (MMD) [3, 21], \mathcal{H} -divergence [22], Kullback-Leibler (KL) divergence [23], and Wasserstein distance [24, 25]. The second category is the adversarial-based approach [4, 26, 27] which uses a domain discriminator to encourage domain confusion via an adversarial objective. The third category is the reconstruction-based approach. Data are reconstructed in the new domain by an encoder-decoder [28, 29] or a GAN discriminator, such as dual-GAN [30], cycle-GAN [31], disco-GAN [32], and CyCADA [33]. However, these methods mainly consider aligning the marginal distributions to decrease feature shift, neglecting label shift. To the best of our knowledge, we are the first to propose an end-to-end deep model to tackle both of the two domain shifts between the source and target domains.

Domain Adaptation for Label Shift. Despite its wide applicability, learning under label shift remains under-explored. Existing works tackle this challenge by importance reweighting or target distribution estimation. Specifically, [10] exploit importance reweighting to enhance knowledge transfer under label shift. Recently, [34] introduce a test distribution estimator to detect and correct for label shift. These methods assume that the source and target domains share the same feature distributions and only differ in the marginal label distribution. In this work, we explore transfer learning between domains under label shift and label shift simultaneously. As a special case of label shift, some works consider the domain adaptation problem where the categories in the source domain and target domain are not fully overlapped. [35] propose *open set domain adaptation* where the class set in the source domain is a proper subset of that of the target domain. Conversely, [15] introduce *partial domain adaptation* where the class set of the source domain is a proper superset of that of the target domain. In this direction, [36] introduce a theoretical analysis to show that only learning domain-invariant features is not sufficient to solve domain adaptation task when the label priors are not aligned. In a related work, [12] propose asymmetrically-relaxed distribution alignment to overcome the limitations of standard domain adaptation algorithms which aims to extract domain-invariant representations.

Domain Adaptation with Self-training. In domain adaptation, self-training methods are often utilized to compensate for the lack of categorical information in the target domain. The intuition is to assign pseudo-labels to unlabeled samples based on the predictions of one or more classifiers. [37] leverage an asymmetric tri-training strategy to assign pseudo-labels to the unlabeled target domain. [38] propose to assign pseudo-labels to all target samples and use them to achieve semantic alignment across domains.

Recently, [39] propose to progressively label the target samples and align the prototypes of source domain and target domain to achieve domain alignment. However, to the best of our knowledge, self-training has not been applied for DA with label shift.

3 CO-ALIGNMENT OF FEATURE AND LABEL DISTRIBUTION

In Class-imbalanced Domain Adaptation, we are given a *source* domain $\mathcal{D}_S = \{(x_i^s, y_i^s)_{i=1}^{N_s}\}$ with N_s labeled examples, and a *target* domain $\mathcal{D}_T = \{(x_i^t)_{i=1}^{N_t}\}$ with N_t unlabeled examples. We assume that $p(y|x) = q(y|x)$ but $p(x|y) \neq q(x|y)$, $p(x) \neq q(x)$, and $p(y) \neq q(y)$. We aim to construct an end-to-end deep neural network which is able to transfer the knowledge learned from \mathcal{D}_S to \mathcal{D}_T , and train a classifier $y = \theta(x)$ which can minimize task risk in target domain $\epsilon_T(\theta) = \Pr_{(x,y) \sim q}[\theta(x) \neq y]$.

Previous works either focus on aligning the marginal feature distributions [2, 4] or aligning the label distributions [34]. These approaches are not able to fully tackle CDA as they only align one of the two marginal distributions. Motivated by theoretical analysis, in this work we propose to tackle CDA with feature distribution and label distribution CO-ALIGNMENT. To this end, we combine

the ideas of *prototype-based conditional alignment* [18] and *class-balanced self-training* [19] to tackle feature and label shift respectively. An overview of COAL is shown in Fig. 2.




Fig. 2. Overview of the proposed COAL model. Our model is trained iteratively between two steps. In step A, we forward the target samples through our model to generate the pseudo labels and mask. In step B, we train our models by *self-training* with the pseudo-labeled target samples to align the label distributions, and *prototype-based conditional alignment* with the minimax entropy.

3.1 Theoretical Motivations

Conditional Feature Alignment. According to [36], the target error in domain adaptation is bounded by three terms: 1) source error, 2) the discrepancy between the marginal distributions and 3) the distance between the source and target optimal labeling functions. Denote $h \in \mathcal{H}$ as the hypothesis, $\epsilon_S(\cdot)$ and $\epsilon_T(\cdot)$ as the expected error of a labeling function on source and target domain, and f_S and f_T as the optimal labeling functions in the source and target domain. Then, we have:

$$\epsilon_T(h) \leq \epsilon_S(h) + d_{\hat{\mathcal{H}}}(\mathcal{D}_S, \mathcal{D}_T) + \min\{\epsilon_S(f_T), \epsilon_T(f_S)\}, \tag{1}$$

where $d_{\hat{\mathcal{H}}}$ denote the discrepancy of the marginal distributions [36]. The bound demonstrates that the optimal labeling functions f_S and f_T need to generalize well in both domains, such that the term $\min\{\epsilon_S(f_T), \epsilon_T(f_S)\}$ can be bounded. Conventional domain adaptation approaches which only align marginal feature distribution cannot guarantee that $\min\{\epsilon_S(f_T), \epsilon_T(f_S)\}$ is minimized. This motivates us to align the conditional feature distribution, *i.e.* $p(x|y)$ and $q(x|y)$.

Class-Balanced Self-training. Theorem 4.3 in [36] indicates that the target error $\epsilon_T(h)$ can not be minimized if we only align the feature distributions and neglect the shift in label distribution. Denote d_{JS} as the Jensen-Shannon (JS) distance between two distributions, [36] propose:

$$\epsilon_S(h) + \epsilon_T(h) \geq \frac{1}{2}(d_{JS}(p(y), q(y)) - d_{JS}(p(x), q(x)))^2 \tag{2}$$

This theorem demonstrates that when the divergence between label distributions $d_{JS}(p(y), q(y))$ is significant, minimizing the divergence between marginal distributions $d_{JS}(p(x), q(x))$ and the source task error $\epsilon_S(h)$ will enlarge the target task error $\epsilon_T(h)$. Motivated by this, we propose to estimate and align the empirical label distributions with a self-training algorithm.

3.2 Prototype-Based Conditional Alignment for Feature Shift

The mainstream idea in feature-shift oriented methods is to learn domain-invariant features by aligning the marginal feature distributions, which was proved to be inferior in the presence of label shift [9]. Instead, inspired by [18], we align the conditional feature distributions $p(x|y)$ and $q(x|y)$. To this end, we leverage a *similarity-based classifier* to estimate $p(x|y)$, and a minimax entropy algorithm to align it with $q(x|y)$. We achieve conditional feature distribution alignment by aligning the source and target prototypes in an adversarial process.

Similarity-Based Classifier. The architecture of our COAL model contains a feature extractor F and a similarity-based classifier C . Prototype-based classifiers perform well in few-shot learning settings [40], which motivates us to adopt them since in label-shift settings some categories can have low frequencies. Specifically, C is composed of a weight matrix $\mathbf{W} \in \mathbb{R}^{d \times c}$ and a temperature parameter T , where d is the dimension of feature generated by F , and c is the total number of classes. Denote \mathbf{W} as $[\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_c]$, this matrix can be seen as c d -dimension vectors, one for each category. For each input feature $F(x)$, we compute its similarity with the i_{th} weight vector as $s_i = \frac{F(x)\mathbf{w}_i}{T\|F(x)\|}$. Then, we compute the probability of the sample being labeled as class i by $h_i(x) = \sigma\left(\frac{F(x)\mathbf{w}_i}{T\|F(x)\|}\right)$, normalizing over all the classes. Finally, we can compute the prototype-based classification loss for \mathcal{D}_S with standard cross-entropy loss:

$$\mathcal{L}_{SC} = \mathbb{E}_{(x,y) \in \mathcal{D}_S} \mathcal{L}_{ce}(h(x), y) \quad (3)$$

The intuition behind this loss is that the higher the confidence of sample x being classified as class i , the closer the embedding of x is to \mathbf{w}_i . Hence, when optimizing Eq. 3, we are reducing the intra-class variation by pushing the embedding of each sample x closer to its corresponding weight vector in \mathbf{W} . In this way, \mathbf{w}_i can be seen as a representative data point (prototype) for $p(x|y = i)$.

Conditional Alignment by Minimax Entropy. Due to the lack of categorical information in the target domain, it is infeasible to utilize Eq. 3 to obtain target prototypes. Following [18], we tackle this problem by 1) moving each source prototype to be closer to its nearby target samples, and 2) clustering target samples around this moved prototype. We achieve these two objectives jointly by entropy minimax learning. Specifically, for each sample $x^t \in \mathcal{D}_T$ fed into the network, we can compute the mean entropy of the classifier’s output by

$$\mathcal{L}_H = \mathbb{E}_{x \in \mathcal{D}_T} H(x) = -\mathbb{E}_{x \in \mathcal{D}_T} \sum_{i=1}^c h_i(x) \log h_i(x). \quad (4)$$

Larger $H(x)$ indicates that sample x is similar to all the weight vectors (prototypes) of C . We achieve conditional feature distributions alignment by aligning the source and target prototypes in an adversarial process: (1) we train C to *maximize* \mathcal{L}_H , aiming to move the prototypes from the source samples towards the neighboring target samples; (2) we train F to *minimize* \mathcal{L}_H , aiming to make the embedding of target samples closer to their nearby prototypes. By training with these two objectives as a min-max game between C and F , we can align source and target prototypes. Specifically, we add a gradient-reverse layer [5] between C and F to flip the sign of gradient.

3.3 Class-Balanced Self-training for Label Shift

As the source label distribution $p(y)$ is different from that of the target $q(y)$, it is not guaranteed that the classifier C which has low risk on \mathcal{D}_S will have low error on \mathcal{D}_T . Intuitively, if the classifier is trained with imbalanced source data, the decision boundary will be dominated by the most frequent categories in the training data, leading to a classifier biased towards source label distribution. When the classifier is applied to target domain with a different label distribution, its accuracy will degrade as it is highly biased towards the source domain. To tackle this problem, we use the method in [19] to employ *self-training* to estimate the target label distribution and refine the decision boundary. In addition, we leverage *balanced sampling* of the source data to further facilitate this process.

Self-training. In order to refine the decision boundary, we propose to estimate the target label distribution with self-training. We assign pseudo labels \hat{y} to all the target samples according to the output the classifier C . As we are also aligning the conditional feature distributions ($p(x|y)$ and $q(x|y)$), we assume that the distribution of high-confidence pseudo labels $q(\hat{y})$ can be used as an approximation of the real label distribution $q(y)$ for the target domain. Training C with these pseudo-labeled target samples under approximated target label distribution, we are able to reduce the negative effect of label shift.

To obtain high-confidence pseudo labels, for each category, we select top $k\%$ of the target samples with the highest confidence scores belonging to that category. We utilize the highest probability in $h(x)$ as the classifier’s confidence on sample x . Specifically, for each pseudo-labeled sample (x, \hat{y}) , we set its selection mask $m = 1$ if $h(x)$ is among the top $k\%$ of all the target samples with the same pseudo-label, otherwise $m = 0$. Denote the pseudo-labeled target set as $\hat{\mathcal{D}}_T = \{(x_i^t, \hat{y}_i^t, m_i)_{i=1}^{N_t}\}$, we leverage the input and pseudo labels from $\hat{\mathcal{D}}_T$ to train the classifier C , aiming to refine the decision boundary with target label distribution. The total loss function for classification is:

$$\mathcal{L}_{ST} = \mathcal{L}_{SC} + \mathbb{E}_{(x, \hat{y}, m) \in \hat{\mathcal{D}}_T} \mathcal{L}_{ce}(h(x), \hat{y}) \cdot m \quad (5)$$

where \hat{y} indicates the pseudo labels and m indicates selection masks. In our approach, we choose the top $k\%$ of the highest confidence target samples *within* each category, instead of universally. This is crucial to estimate the real target label distribution, otherwise, the easy-to-transfer categories will dominate $\hat{\mathcal{D}}_T$,

leading to inaccurate estimation of the target label distribution [19]. As training processes, we are able to obtain pseudo labels with higher accuracy. Therefore, we increase k by k_{step} after each epoch until it reaches a threshold k_{max} . Typically, we initialize k with $k_0 = 5$, and set $k_{step} = 5$, $k_{max} = 30$.

Balanced Sampling of Source Data. When coping with label shift, the label distribution of the source domain could be highly imbalanced. A classifier trained on imbalanced categories will make highly-biased predictions for the samples from the target domain [41]. This effect also hinders the self-training process discussed above, as the label distribution estimation will also be biased. To tackle these problems, we apply a balanced mini-batch sampler to generate training data from the source domain and ensure that each source mini-batch contains roughly the same number of samples for each category.

3.4 Training Process

In this section, we combine the above ideas into an end-to-end training pipeline. Denote α as the trade-off between classifier training and feature distribution alignment, we first define the adaptive learning objective as follows:

$$\hat{C} = \arg \min_C \mathcal{L}_{ST} - \alpha \mathcal{L}_H, \quad \hat{F} = \arg \min_F \mathcal{L}_{ST} + \alpha \mathcal{L}_H. \quad (6)$$

Given input samples from source domain \mathcal{D}_S and target domain \mathcal{D}_T , we first pretrain our network F and C with only labeled data \mathcal{D}_S . Then, we iterate between **pseudo-label assignment** (step A) and **adaptive learning** (step B). We update the pseudo labels in each epoch as we obtain better feature representations from adaptive learning, which leads to more accurate pseudo labels. On the other hand, better pseudo labels could also facilitate adaptive learning in the next epoch. This process continues until convergence or reaching the maximum number of iterations. An overview of it is shown in Fig. 2.

4 Experiments

In this section, we first construct the CDA benchmarks with 26 cross-domain adaptation tasks based on 4 **Digits** datasets, **Office-Home** [42] and **Domain-Net** [14]. Then we evaluate and analysis 10 representative state-of-the-art domain adaptation methods as well as our COAL baseline. Finally, we provide additional analysis experiments to further explore the CDA problem.

4.1 Class-Imbalanced Domain Adaptation Benchmark

Domain Shift Protocol. Because the images use are already collected from separate feature domains, we only create label shift for each cross-domain task. To create label shift between source and target domains, we sub-sample the current datasets with **Reversely-unbalanced Source** and **Unbalanced Target (RS-UT)** protocol. In this setting, both the source and target domains have unbalanced label distribution, while the label distribution of the source domain is a

reversed version of that of the target domain. Following [43], the unbalanced label distribution is created by sampling from a Pareto distribution [44]. An illustration of this setting can be found in Fig. 3(b). We refer our reader to supplementary material for detailed data splits and creation process.

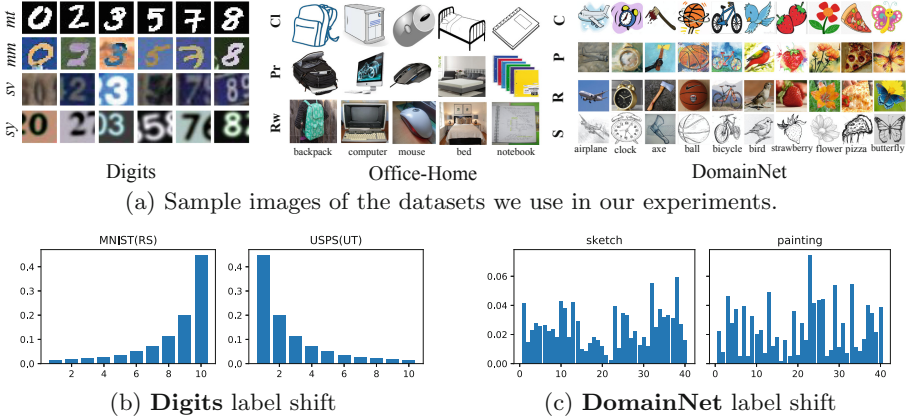


Fig. 3. (a): Image examples from Digits, Office-Home [42], and DomainNet [14]. (b): illustrations of **R**eversely-unbalanced **S**ource (**RS**) and **U**nbalanced **T**arget (**UT**) distribution in MNIST \rightarrow USPS task. (c): Natural label shift of DomainNet.

Digits. We select four digits datasets: MNIST [45], USPS [46], SVHN [47] and Synthetic Digits (SYN) [48] and regard each of them as a separate domain. In this work, we investigate four domain adaptation tasks: **MNIST** \rightarrow **USPS**, **USPS** \rightarrow **MNIST**, **SVHN** \rightarrow **MNIST**, and **SYN** \rightarrow **MNIST**.

Office-Home [42] is a dataset collected in office and home environment with 65 object classes and four domains: Real World (**Rw**), Clipart (**Cl**), Product (**Pr**), Art (**Ar**). Since the “Art” domain is too small to sample an imbalanced subset, we focus on the remaining domains and explore all the six adaptation tasks.

DomainNet [14] is a large-scale testbed for domain adaptation, which contains six domains with about 0.6 million images distributed among 345 classes. Since some domains and classes contains many mislabeled outliers, we select 40 commonly-seen classes from four domains: Real (**R**), Clipart (**C**), Painting (**P**), Sketch (**S**). Different from the two datasets above, the existed label shift in DomainNet is significant enough, as illustrated in Fig. 3(c). Therefore, we use the original label distributions without sub-sampling for this dataset.

Evaluated Methods. To form a comprehensive empirical analysis, we evaluated recent domain adaptation methods from three categories, including **1**) conventional UDA methods that only aligns feature distribution: **DAN** [2],

JAN [21], **DANN** [5], **MCD** [6] and **BSP** [49]; **2)** method that only aligns label distribution: **BBSE**[34]; **3)** methods that align feature distribution while assuming non-overlapping label spaces: **PADA** [15], **ETN** [16] and **UAN** [17]. We also evaluated **FDANN** [12], which relaxes the feature distribution alignment objective in DANN to deal with potential label shift.

Table 1. Per-class mean accuracy on Digits. Our model achieves **84.33%** average accuracy across four tasks, outperforming other evaluated methods.

Methods	USPS \rightarrow MNIST	MNIST \rightarrow USPS	SVHN \rightarrow MNIST	SYN \rightarrow MNIST	AVG
Source Only	75.31 \pm 0.09	87.92 \pm 0.74	50.25 \pm 0.81	85.74 \pm 0.49	74.81
UAN (2019)	55.72 \pm 2.06	83.23 \pm 0.75	50.20 \pm 1.75	71.26 \pm 2.82	65.10
ETN (2019)	62.85 \pm 2.20	79.27 \pm 1.29	52.82 \pm 1.50	72.42 \pm 6.53	66.84
FDANN (2019)	72.59 \pm 1.61	81.62 \pm 2.38	45.65 \pm 2.93	82.07 \pm 1.65	70.48
JAN (2017)	75.75 \pm 0.75	78.82 \pm 0.93	53.21 \pm 3.94	75.64 \pm 1.42	70.86
BBSE (2018)	75.01 \pm 3.68	78.84 \pm 10.73	49.01 \pm 2.02	85.69 \pm 0.71	72.14
BSP (2019)	71.99 \pm 1.52	89.74 \pm 0.77	50.61 \pm 1.67	77.30 \pm 1.20	72.41
PADA (2018)	73.66 \pm 0.15	78.59 \pm 0.23	54.13 \pm 1.61	85.06 \pm 0.60	72.86
MCD (2018)	77.18 \pm 5.65	85.34 \pm 4.07	53.52 \pm 4.23	76.37 \pm 3.48	73.10
DAN (2015)	<u>79.12</u> \pm 1.34	87.15 \pm 1.71	53.63 \pm 1.80	80.89 \pm 2.00	75.20
DANN (2015)	77.28 \pm 2.13	<u>91.88</u> \pm 0.74	<u>57.16</u> \pm 1.83	77.60 \pm 1.29	<u>75.98</u>
COAL (Ours)	88.12 \pm 0.37	93.04 \pm 1.67	65.67 \pm 1.29	90.60 \pm 0.44	84.33

Table 2. Per-class mean accuracy on Office-Home dataset. Our model achieve **58.87%** average accuracy across six tasks.

Methods	Rw \rightarrow Pr	Rw \rightarrow Cl	Pr \rightarrow Rw	Pr \rightarrow Cl	Cl \rightarrow Rw	Cl \rightarrow Pr	AVG
Source Only	70.75	35.51	65.65	34.99	51.27	51.11	51.55
BSP (2019)	66.15	23.48	65.42	20.81	34.54	31.04	40.24
PADA (2018)	60.77	32.28	57.09	26.76	40.71	38.34	42.66
BBSE (2018)	61.10	33.27	62.66	31.15	39.70	38.08	44.33
MCD (2018)	66.18	32.32	62.66	28.40	41.41	38.59	44.93
DAN (2015)	67.85	38.17	66.86	34.24	52.95	51.64	45.02
UAN (2019)	70.85	41.15	67.26	36.82	56.24	55.77	48.62
ETN (2019)	71.69	34.03	<u>70.45</u>	40.74	60.48	55.19	52.14
FDANN (2019)	68.56	40.57	67.32	37.33	55.84	53.67	53.88
JAN (2017)	71.22	<u>43.12</u>	68.20	37.03	57.97	56.80	55.72
DANN (2015)	71.78	46.08	67.98	39.45	58.40	<u>57.39</u>	<u>56.85</u>
COAL (Ours)	73.65	42.58	74.46	<u>40.61</u>	<u>59.22</u>	62.71	58.87

Implementation Details. We implement all our experiments in Pytorch platform. We used the official implements for all the evaluated methods except for DANN [5], BBSE [34] and FDANN [12], which are reproduced by ourselves. For fair comparison, we use the same backbone networks for all the methods.

Table 3. Per-class mean accuracy on DomainNet dataset with natural label shifts. Our method achieve **75.89%** average accuracy across the 12 experiments. Note that DomainNet contains about 0.6 million images, it is non-trivial to have even one percent performance boost.

Method	R → C	R → P	R → S	C → R	C → P	C → S	P → R	P → C	P → S	S → R	S → C	S → P	AVG
Baseline	58.84	67.89	53.08	76.70	53.55	53.06	84.39	55.55	60.19	74.62	54.60	57.78	62.52
BBSE	55.38	63.62	47.44	64.58	42.18	42.36	81.55	49.04	54.10	68.54	48.19	46.07	55.25
PADA	65.91	67.13	58.43	74.69	53.09	52.86	79.84	59.33	57.87	76.52	66.97	61.08	64.48
MCD	61.97	69.33	56.26	79.78	56.61	53.66	83.38	58.31	60.98	81.74	56.27	66.78	65.42
DAN	64.36	70.65	58.44	79.44	56.78	60.05	84.56	61.62	62.21	79.69	65.01	62.04	67.07
FDANN	66.15	71.80	61.53	81.85	60.06	61.22	84.46	66.81	62.84	81.38	69.62	66.50	69.52
UAN	71.10	68.90	67.10	83.15	63.30	64.66	83.95	65.35	67.06	82.22	70.64	68.09	72.05
JAN	65.57	<u>73.58</u>	67.61	85.02	64.96	67.17	<u>87.06</u>	67.92	66.10	84.54	72.77	67.51	72.48
ETN	69.22	<u>72.14</u>	63.63	<u>86.54</u>	65.33	63.34	85.04	65.69	68.78	84.93	72.17	68.99	73.99
BSP	<u>67.29</u>	73.47	69.31	86.50	<u>67.52</u>	<u>70.90</u>	86.83	70.33	68.75	84.34	72.40	71.47	74.09
DANN	63.37	73.56	72.63	86.47	65.73	70.58	86.94	<u>73.19</u>	<u>70.15</u>	<u>85.73</u>	75.16	70.04	<u>74.46</u>
Ours	73.85	75.37	<u>70.50</u>	89.63	69.98	71.29	89.81	68.01	70.49	87.97	<u>73.21</u>	<u>70.53</u>	75.89

Specifically, for the Digits dataset, we adopt the network architecture proposed by [6]. For the other two datasets, we utilize ResNet-50 [50] as our backbone network, and replace the last fully-connected layer with a randomly initialized N-way classifier layer (for N categories). For all the compared methods, we select their hyper-parameters on the validation set of P → C task of DomainNet. We refer our reader to supplementary material for code and parameters of COAL.

Evaluation Metric. When the target domain is highly unbalanced, conventional overall average accuracy that treats every class uniformly is not an appropriate performance metric [51]. Therefore, we follow [52] to use the *per-class* mean accuracy in our main results. Formally, we denote $S_i = \frac{n_{(i,i)}}{n_i}$ as the accuracy for class i , where $n_{(i,j)}$ represents the number of class i samples labeled as class j , and $n_i = \sum_{j=1}^c n_{(i,j)}$ represents the number of samples in class i . Then, the per-class mean accuracy is computed as $S = \frac{1}{c} \sum_{i=1}^c S_i$.

4.2 Result Analysis

We first show the experimental results on Digits datasets in Table 1. From the results, we can make the following observations: **(1)** Most current domain adaptation methods cannot solve CDA well. On average, 8 of the 10 evaluated domain adaptation methods perform *worse* than the source-only baselines, leading to negative transfer. This result confirmed the theoretical analysis that only aligning marginal feature distribution leads to performance drop under CDA [9]. **(2)** Method that achieve better results on conventional UDA benchmarks does not lead to better results on CDA problem. For example, although MCD is shown to significantly outperform DAN and DANN on several conventional domain adaptation benchmarks [6], its performance is inferior to these older methods in our experiment. We argue that this is because these newer methods achieve better marginal feature distribution alignment, which yet leads to worse performance

under label shift. **(3)** Our COAL baseline achieves **84.33%** average accuracy across four experimental setting, outperforming the best-performing method by **8.4%**. This result demonstrate that aligning only the feature distributions or only the label distributions can not fully tackle CDA task. In contrast, our framework co-aligns the conditional feature distributions and label distributions.

Next, we show the experimental results on more challenging real-object datasets, *i.e.*, Office-Home and DomainNet, in Table 2 and Table 3, respectively. In Office-Home experiments, we can also have the above observations. For example, we observe that 7 out of 10 methods lead to negative transfer, which is consistent with the results on Digits dataset. Our COAL framework achieves **58.87%** average accuracy across the six CDA tasks, outperforming other evaluated methods, and has **7.32%** improvement from the source-only result.

In DomainNet experiments, due to smaller degree of label shift, most evaluated methods could outperform the source-only baseline. However, we still observe the negative influence of label shift. First, we observe inferior performance of newer methods to older methods. For example, DANN outperformed MCD by 9.04%, due to the negative effect of stronger marginal alignment in MCD. Moreover, our model get **75.89%** average accuracy across the 12 tasks, outperforming all the compared baselines. This shows the effectiveness of feature and label distribution co-alignment in this dataset. Furthermore, we carefully tuned the hyper-parameters for the evaluated domain adaptation methods to have weaker feature distribution alignment². If we directly apply the parameters set by the authors, many of these models have much worse performance.

Table 4. The performance of five models *w.* or *w/o.* source balanced sampler. We observe a significant performance boost when the source balanced sampler is applied, both for our model and the compared baselines, demonstrating the effectiveness of source balanced sampler to CDA task.

Methods	U → M		S → M		Pr → Cl		Cl → Rw		R → S	
	w/o	with	w/o	with	w/o	with	w/o	with	w/o	with
Source Only	71.35	75.31	50.35	50.25	34.99	34.99	50.64	51.11	50.16	53.08
DAN	64.81	79.12	22.05	53.63	32.93	34.24	45.18	51.64	64.78	58.44
DANN	42.77	77.28	27.60	57.16	35.17	39.45	47.19	58.40	68.92	72.63
MCD	20.15	77.18	44.83	53.52	33.06	28.40	49.57	41.41	58.50	56.26
COAL	87.50	88.12	60.12	65.67	34.03	40.61	57.67	59.22	59.23	70.50

4.3 Analysis

Effect of Source Balanced Sampler. Source balanced samplers described in Sect. 3.3 can help us tackle the biased-classifier problem caused by the imbalanced data distribution of source domain. A significant performance boost can

² Please refer to supplementary material for details.

be observed after applying the balanced sampler for our COAL model, as well as the compared baselines. In this section, we specifically show the effect of using source balanced samplers. We show in Table 4 the performance of several methods with and without source balanced samplers on 5 adaptation tasks from multiple datasets. We observe that for 20 of the total 25 tasks (5 models on 5 adaptation tasks), using source balanced samplers will significantly improve the domain adaptation performance. These results show the effectiveness of having a source balanced sampler when tackling CDA task.

Ablation Study. Our COAL method has mainly two objectives: 1) alignment of conditional feature distribution \mathcal{L}_{ST} and 2) alignment of label distribution \mathcal{L}_H . To show the importance of these two objectives in CDA, we show the performance of our method without each of these objectives respectively on multiple tasks. The results in Table 5 showed the importance of both objectives. For example, for USPS \rightarrow MNIST, if we remove the conditional feature distribution alignment objective, the accuracy of our model will drop by 2.6%. Similarly, if we remove the label distribution alignment objective, the accuracy will drop by 2.9%. These results demonstrate that both the alignment of conditional feature distribution and label distribution are important to CDA task.




Fig. 4. t-SNE visualization for features Source Only (baseline), DAN, DANN and COAL on DomainNet task Real \rightarrow Clipart. Blue and red points represents features from the source domain and target domain, respectively. (Color figure online)

Table 5. Ablation study of different objectives in our method. We randomly select 8 sets of experiments to perform the ablation study.

Methods	U \rightarrow M	M \rightarrow U	S \rightarrow M	Cl \rightarrow Rw	Pr \rightarrow Rw	R \rightarrow C	R \rightarrow P	P \rightarrow R	AVG
w/o \mathcal{L}_{ST}	85.22	85.94	55.17	58.38	69.39	71.92	74.39	77.45	72.23
w/o \mathcal{L}_H	85.57	92.28	63.34	58.17	72.11	71.34	69.92	87.14	74.98
COAL	88.12	93.04	65.67	59.22	74.46	73.85	75.37	89.81	77.44

Feature Visualization. In this section, we plot the learned features with t-SNE [53] in Fig. 4. We investigate the Real to Clipart task in DomainNet experiment with ResNet-50 backbones. From (a)–(d), we observe that our method can better align source and target features in each category, while other methods either

leave the feature distributions unaligned, or incorrectly aligned samples in different categories. These results further show the importance of *prototype-based conditional feature alignment* for CDA task.

Different Degrees of Label Shift. We also investigate the effect of different degrees of label shift. Specifically, we create 4 interval degrees of label shift between the **BS-BT** (**B**lanced **S**ource and **B**lanced **T**arget) and RS-UT setting. With these label set settings, we evaluated the performance of different methods. The results show that the performance of previous domain adaptation methods will be significantly affected by label shift, while our method is much more robust. Please refer to supplementary material for detailed setting and results.

5 Conclusion

In this paper, we first propose the Class-imbalanced Domain Adaptation(**CDA**) setting and demonstrate its importance in practical scenarios. Then we provide the first benchmark of this problem, and conduct a comprehensive empirical analysis on recent domain adaptation methods. The result shows that most existing methods are fragile in the face of CDA, which prevents them from being practically applied. Based on theoretical motivations, we propose a feature distribution and label distribution co-alignment framework, which empirically works well as a baseline for future research.

We believe this work takes an important step towards applicable domain adaptation. We hope the provided benchmarks, empirical results and baseline model would stimulate and facilitate future works to design robust algorithms that can handle more realistic problems. An interesting research direction would be better detecting and correcting label shift under feature shift.

References

1. Quionero-Candela, J., Sugiyama, M., Schwaighofer, A., Lawrence, N.D.: Dataset Shift in Machine Learning. The MIT Press, Cambridge (2009)
2. Long, M., Cao, Y., Wang, J., Jordan, M.: Learning transferable features with deep adaptation networks. In: Bach, F., Blei, D. (eds.) Proceedings of the 32nd International Conference on Machine Learning, Proceedings of Machine Learning Research, Lille, France, PMLR, 07–09 July 2015, vol. 37, pp. 97–105 (2015)
3. Tzeng, E., Hoffman, J., Zhang, N., Saenko, K., Darrell, T.: Deep domain confusion: maximizing for domain invariance. arXiv preprint [arXiv:1412.3474](https://arxiv.org/abs/1412.3474) (2014)
4. Tzeng, E., Hoffman, J., Saenko, K., Darrell, T.: Adversarial discriminative domain adaptation. In: Computer Vision and Pattern Recognition (CVPR), vol. 1, p. 4 (2017)
5. Ganin, Y., Lempitsky, V.: Unsupervised domain adaptation by backpropagation. In: Bach, F., Blei, D. (eds.) Proceedings of the 32nd International Conference on Machine Learning, Proceedings of Machine Learning Research, Lille, France, PMLR, 07–09 July 2015, vol. 37, pp. 1180–1189 (2015)

6. Saito, K., Watanabe, K., Ushiku, Y., Harada, T.: Maximum classifier discrepancy for unsupervised domain adaptation. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), June 2018
7. Lipton, Z., Wang, Y.X., Smola, A.: Detecting and correcting for label shift with black box predictors. In: Dy, J., Krause, A. (eds.) Proceedings of the 35th International Conference on Machine Learning, Proceedings of Machine Learning Research, Stockholmsmässan, Stockholm Sweden, PMLR, 10–15 July 2018, vol. 80, pp. 3122–3130 (2018)
8. Azizzadenesheli, K., Liu, A., Yang, F., Anandkumar, A.: Regularized learning for domain adaptation under label shifts. In: International Conference on Learning Representations (2019)
9. Zhao, H., Combes, R.T.D., Zhang, K., Gordon, G.: On learning invariant representations for domain adaptation. In: Chaudhuri, K., Salakhutdinov, R. (eds.) Proceedings of the 36th International Conference on Machine Learning, Proceedings of Machine Learning Research., Long Beach, California, USA, PMLR, 09–15 June 2019, vol. 97, pp. 7523–7532 (2019)
10. Zhang, K., Schölkopf, B., Muandet, K., Wang, Z.: Domain adaptation under target and conditional shift. In: International Conference on Machine Learning, pp. 819–827 (2013)
11. Gong, M., Zhang, K., Liu, T., Tao, D., Glymour, C., Schölkopf, B.: Domain adaptation with conditional transferable components. In: International Conference on Machine Learning, pp. 2839–2848 (2016)
12. Wu, Y., Winston, E., Kaushik, D., Lipton, Z.: Domain adaptation with asymmetrically-relaxed distribution alignment. arXiv preprint [arXiv:1903.01689](https://arxiv.org/abs/1903.01689) (2019)
13. Zellinger, W., Grubinger, T., Lughofer, E., Natschläger, T., Saminger-Platz, S.: Central moment discrepancy (CMD) for domain-invariant representation learning. CoRR abs/1702.08811 (2017)
14. Peng, X., Bai, Q., Xia, X., Huang, Z., Saenko, K., Wang, B.: Moment matching for multi-source domain adaptation. In: The IEEE International Conference on Computer Vision (ICCV) (2019)
15. Cao, Z., Ma, L., Long, M., Wang, J.: Partial adversarial domain adaptation. In: Proceedings of the European Conference on Computer Vision (ECCV), pp. 135–150 (2018)
16. Cao, Z., You, K., Long, M., Wang, J., Yang, Q.: Learning to transfer examples for partial domain adaptation. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), June 2019
17. You, K., Long, M., Cao, Z., Wang, J., Jordan, M.I.: Universal domain adaptation. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), June 2019
18. Saito, K., Kim, D., Sclaroff, S., Darrell, T., Saenko, K.: Semi-supervised domain adaptation via minimax entropy. In: ICCV (2019)
19. Zou, Y., Yu, Z., Kumar, B.V., Wang, J.: Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In: Proceedings of the European Conference on Computer Vision (ECCV), pp. 289–305 (2018)
20. Wang, M., Deng, W.: Deep visual domain adaptation: a survey. *Neurocomputing* **312**, 135–153 (2018)
21. Long, M., Zhu, H., Wang, J., Jordan, M.I.: Deep transfer learning with joint adaptation networks. In: Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6–11 August 2017, pp. 2208–2217 (2017)

22. Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., Vaughan, J.W.: A theory of learning from different domains. *Mach. Learn.* **79**(1–2), 151–175 (2010)
23. Zhuang, F., Cheng, X., Luo, P., Pan, S.J., He, Q.: Supervised representation learning: transfer learning with deep autoencoders. *IJCA I*, 4119–4125 (2015)
24. Lee, J., Raginsky, M.: Minimax statistical learning with Wasserstein distances. arXiv preprint [arXiv:1705.07815](https://arxiv.org/abs/1705.07815) (2017)
25. Shen, J., Qu, Y., Zhang, W., Yu, Y.: Wasserstein distance guided representation learning for domain adaptation. arXiv preprint [arXiv:1707.01217](https://arxiv.org/abs/1707.01217) (2017)
26. Liu, M.Y., Tuzel, O.: Coupled generative adversarial networks. In: *Advances in Neural Information Processing Systems*, pp. 469–477 (2016)
27. Peng, X., Huang, Z., Sun, X., Saenko, K.: Domain agnostic learning with disentangled representations. In: *ICML* (2019)
28. Bousmalis, K., Trigeorgis, G., Silberman, N., Krishnan, D., Erhan, D.: Domain separation networks. In: *Advances in Neural Information Processing Systems*, pp. 343–351 (2016)
29. Ghifary, M., Kleijn, W.B., Zhang, M., Balduzzi, D., Li, W.: Deep reconstruction-classification networks for unsupervised domain adaptation. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *ECCV 2016*. LNCS, vol. 9908, pp. 597–613. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46493-0_36
30. Yi, Z., Zhang, H.R., Tan, P., Gong, M.: DualGAN: unsupervised dual learning for image-to-image translation. In: *ICCV*, pp. 2868–2876 (2017)
31. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: *2017 IEEE International Conference on Computer Vision (ICCV)* (2017)
32. Kim, T., Cha, M., Kim, H., Lee, J.K., Kim, J.: Learning to discover cross-domain relations with generative adversarial networks. In: Precup, D., Teh, Y.W. (eds.) *Proceedings of the 34th International Conference on Machine Learning, Proceedings of Machine Learning Research, International Convention Centre, Sydney, Australia, PMLR, 06–11 August 2017, vol. 70, pp. 1857–1865* (2017)
33. Hoffman, J., et al.: CyCADA: cycle-consistent adversarial domain adaptation. In: Dy, J., Krause, A. (eds.) *Proceedings of the 35th International Conference on Machine Learning, Proceedings of Machine Learning Research, Stockholm, Sweden, PMLR, 10–15 July 2018, vol. 80, pp. 1989–1998* (2018)
34. Lipton, Z.C., Wang, Y.X., Smola, A.: Detecting and correcting for label shift with black box predictors. arXiv preprint [arXiv:1802.03916](https://arxiv.org/abs/1802.03916) (2018)
35. Panareda Busto, P., Gall, J.: Open set domain adaptation. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 754–763 (2017)
36. Zhao, H., Combes, R.T.d., Zhang, K., Gordon, G.J.: On learning invariant representation for domain adaptation. arXiv preprint [arXiv:1901.09453](https://arxiv.org/abs/1901.09453) (2019)
37. Saito, K., Ushiku, Y., Harada, T.: Asymmetric tri-training for unsupervised domain adaptation. In: *Proceedings of the 34th International Conference on Machine Learning*, vol. 70, pp. 2988–2997 (2017). [JMLR.org](https://jmlr.org)
38. Xie, S., Zheng, Z., Chen, L., Chen, C.: Learning semantic representations for unsupervised domain adaptation. In: *International Conference on Machine Learning*, pp. 5419–5428 (2018)
39. Chen, C., et al.: Progressive feature alignment for unsupervised domain adaptation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 627–636 (2019)
40. Chen, W.Y., Liu, Y.C., Kira, Z., Wang, Y.C., Huang, J.B.: A closer look at few-shot classification. In: *International Conference on Learning Representations* (2019)

41. He, H., Garcia, E.A.: Learning from imbalanced data. *IEEE Trans. Knowl. Data Eng.* **21**(9), 1263–1284 (2009)
42. Venkateswara, H., Eusebio, J., Chakraborty, S., Panchanathan, S.: Deep hashing network for unsupervised domain adaptation. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017)
43. Liu, Z., Miao, Z., Zhan, X., Wang, J., Gong, B., Yu, S.X.: Large-scale long-tailed recognition in an open world. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2019)
44. Reed, W.: The Pareto, Zipf and other power laws. *Econ. Lett.* **74**, 15–19 (2001)
45. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proc. IEEE* **86**, 2278–2324 (1998)
46. Hull, J.J.: A database for handwritten text recognition research. *IEEE Trans. Pattern Anal. Mach. Intell.* **16**(5), 550–554 (1994)
47. Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., Ng, A.Y.: Reading digits in natural images with unsupervised feature learning (2011)
48. Ganin, Y., Lempitsky, V.S.: Unsupervised domain adaptation by backpropagation. In: *ICML* (2015)
49. Chen, X., Wang, S., Long, M., Wang, J.: Transferability vs. discriminability: batch spectral penalization for adversarial domain adaptation. In: Chaudhuri, K., Salakhutdinov, R. (eds.) *Proceedings of the 36th International Conference on Machine Learning, Proceedings of Machine Learning Research., Long Beach, California, USA, PMLR, 09–15 June 2019, vol. 97, pp. 1081–1090* (2019)
50. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778 (2015)
51. He, H., Garcia, E.A.: Learning from imbalanced data. *IEEE Trans. Knowl. Data Eng.* **21**(9), 1263–1284 (2008)
52. Dong, Q., Gong, S., Zhu, X.: Imbalanced deep learning by minority class incremental rectification. *IEEE Trans. Pattern Anal. Mach. Intell.* **41**(6), 1367–1381 (2019)
53. van der Maaten, L., Hinton, G.: Visualizing data using t-SNE (2008)