

# Contrastive Adaptation Network for Unsupervised Domain Adaptation

Guoliang Kang<sup>1</sup>, Lu Jiang<sup>2</sup>, Yi Yang<sup>1,3\*</sup>, Alexander G. Hauptmann<sup>4</sup>

<sup>1</sup>CAI, University of Technology Sydney, <sup>2</sup>Google AI, <sup>3</sup>Baidu Research, <sup>4</sup>Carnegie Mellon University  
 kgl.prml@gmail.com, lujiang@google.com, Yi.Yang@uts.edu.au, alex@cs.cmu.edu

## Abstract

*Unsupervised Domain Adaptation (UDA) makes predictions for the target domain data while manual annotations are only available in the source domain. Previous methods minimize the domain discrepancy neglecting the class information, which may lead to misalignment and poor generalization performance. To address this issue, this paper proposes Contrastive Adaptation Network (CAN) optimizing a new metric which explicitly models the intra-class domain discrepancy and the inter-class domain discrepancy. We design an alternating update strategy for training CAN in an end-to-end manner. Experiments on two real-world benchmarks Office-31 and VisDA-2017 demonstrate that CAN performs favorably against the state-of-the-art methods and produces more discriminative features.*

## 1. Introduction

Recent advancements in deep neural networks have successfully improved a variety of learning problems [40, 8, 26, 19, 20]. For supervised learning, however, massive labeled training data is still the key to learning an accurate deep model. Although abundant labels may be available for a few pre-specified domains, such as ImageNet [7], manual labels often turn out to be difficult or expensive to obtain for every ad-hoc target domain or task. The absence of in-domain labeled data hinders the application of data-fitting models in many real-world problems.

In the absence of labeled data from the target domain, Unsupervised Domain Adaptation (UDA) methods have emerged to mitigate the domain shift in data distributions [2, 1, 5, 37, 30, 18, 3, 17]. It relates to unsupervised learning as it requires manual labels only from the source domain and zero labels from the target domain. Among the recent work on UDA, a seminal line of work proposed by Long *et al.* [22, 25] aims at minimizing the discrepancy between the source and target domain in the deep neural network, where the domain discrepancy is measured by Maximum

\*Corresponding author. Part of this work was done when Yi Yang was visiting Baidu Research during his Professional Experience Program.

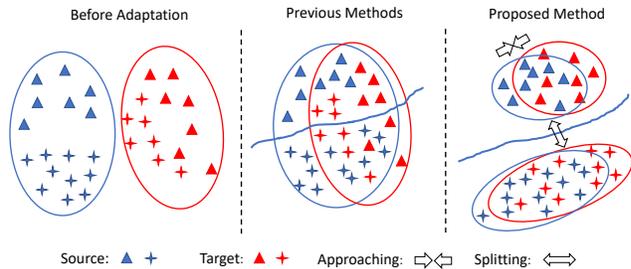


Figure 1. Comparison between previous domain-discrepancy minimization methods and ours. **Left:** The domain shift exists between the source and target data before adaptation. **Middle:** Class-agnostic adaptation aligns source and target data at the domain-level, neglecting the class label of the sample, and hence may lead to sub-optimal solutions. Consequently, the target samples of one label may be misaligned with source samples of a different label. **Right:** Our method performs class-aware alignment across domains. To avoid the misalignment, only the intra-class domain discrepancy is minimized. The inter-class domain discrepancy is maximized to enhance the model’s generalization ability.

Mean Discrepancy (MMD) [22] and Joint MMD (JMMD) [25]. MMD and JMMD have proven effective in many computer vision problems and demonstrated the state-of-the-art results on several UDA benchmarks [22, 25].

Despite the success of previous methods based on MMD and JMMD, most of them measure the domain discrepancy at the *domain* level, neglecting the class from which the samples are drawn. These class-agnostic approaches, hence, do not discriminate whether samples from two domains should be aligned according to their class labels (Fig. 1). This can impair the adaptation performance due to the following reasons. First, samples of different classes may be aligned incorrectly, *e.g.* both MMD and JMMD can be minimized even when the target-domain samples are misaligned with the source-domain samples of a different class. Second, the learned decision boundary may generalize poorly for the target domain. There exist many sub-optimal solutions near the decision boundary. These solutions may overfit the source data well but are less discriminative for the target.

To address the above issues, we introduce a new *Contrastive Domain Discrepancy (CDD)* objective to enable class-aware UDA. We propose to minimize the intra-class discrepancy, *i.e.* the domain discrepancy within the same class, and maximize the inter-class margin, *i.e.* the domain discrepancy between different classes. Considering the toy example in Fig. 1, CDD will draw closer the source and target samples of the same underlying class (*e.g.* the blue and red triangles), while pushing apart the samples from different classes (*e.g.* the blue triangle and the red star).

Unfortunately, to estimate and optimize with CDD, we may not train a deep network out-of-the-box as we need to overcome the following two technical issues. First, we need labels from both domains to compute CDD, however, target labels are unknown in UDA. A straightforward way, of course, is to estimate the target labels by the network outputs during training. However, because the estimation can be noisy, we find it can harm the adaptation performance (see Section 4.3). Second, during the mini-batch training, for a class  $C$ , the mini-batch may only contain samples from one domain (source or target), rendering it infeasible to estimate the intra-class domain discrepancy of  $C$ . This can result in a less efficient adaptation. The above issues require special design of the network and the training paradigm.

In this paper, we propose *Contrastive Adaptation Network (CAN)* to facilitate the optimization with CDD. During training, in addition to minimizing the cross-entropy loss on labeled source data, CAN alternatively estimates the underlying label hypothesis of target samples through clustering, and adapts the feature representations according to the CDD metric. After clustering, the ambiguous target data (*i.e.* far from the cluster centers) and ambiguous classes (*i.e.* containing few target samples around the cluster centers) are zeroed out in estimating the CDD. Empirically we find that during training, an increasing amount of samples will be taken into account. Such progressive learning can help CAN capture more accurate statistics of data distributions. Moreover, to facilitate the mini-batch training of CAN, we employ the class-aware sampling for both source and target domains, *i.e.* at each iteration, we sample data from both domains for each class within a randomly sampled class subset. Class-aware sampling can improve the training efficiency and the adaptation performance.

We validate our method on two public UDA benchmarks: Office-31 [30] and VisDA-2017 [29]. The experimental results show that our method performs favorably against the state-of-the-art UDA approaches, *i.e.* we achieve the best-published result on the Office-31 benchmark and very competitive result on the challenging VisDA-2017 benchmark. Ablation studies are presented to verify the contribution of each key component in our framework.

In a nutshell, our contributions are as follows,

- We introduce a new discrepancy metric *Contrastive*

*Domain Discrepancy (CDD)* to perform class-aware alignment for unsupervised domain adaptation.

- We propose a network *Contrastive Adaptation Network* to facilitate the end-to-end training with CDD.
- Our method achieves the best-published result on the Office-31 benchmark [30] and competitive performance compared to the state-of-the-art on the challenging VisDA-2017 benchmark [29].

## 2. Related Work

**Class-agnostic domain alignment.** A common practice for UDA is to minimize the discrepancy between domains to obtain domain-invariant features [10, 4, 25, 22, 24, 36, 21]. For example, Tzeng *et al.* [38] proposed a kind of domain confusion loss to encourage the network to learn both semantically meaningful and domain invariant representations. Long *et al.* proposed DAN [22] and JAN [25] to minimize the MMD and Joint MMD distance across domains respectively, over the domain-specific layers. Ganin *et al.* [10] enabled the network to learn domain invariant representations in adversarial way by back-propagating the reverse gradients of the domain classifier. Unlike these domain-discrepancy minimization methods, our method performs *class-aware* domain alignment.

**Discriminative domain-invariant feature learning.** Some previous works pay efforts to learn more discriminative features while performing domain alignment [35, 13, 31, 32, 28, 39]. Adversarial Dropout Regularization (ADR) [31] and Maximum Classifier Discrepancy (MCD) [32] were proposed to train a deep neural network in adversarial way to avoid generating non-discriminative features lying in the region near the decision boundary. Similar to us, Long *et al.* [23] and Pei *et al.* [28] take the class information into account while measuring the domain discrepancy. However, our method differs from theirs mainly in two aspects. Firstly, we explicitly model two types of domain discrepancy, *i.e.* the intra-class domain discrepancy and the inter-class domain discrepancy. The inter-class domain discrepancy, which has been ignored by most previous methods, is proved to be beneficial for enhancing the model adaptation performance. Secondly, in the context of deep neural networks, we treat the training process as an alternative optimization over target label hypothesis and features.

**Intra-class compactness and inter-class separability modeling.** This paper is also related to the work that explicitly models the intra-class compactness and the inter-class separability, *e.g.* the contrastive loss [12] and the triplet loss [33]. These methods have been used in various applications, *e.g.* face recognition [6], person re-identification [16], etc. Different from these methods designed for a single domain, our work focuses on adaptation across domains.

### 3. Methodology

Unsupervised Domain Adaptation (UDA) aims at improving the model's generalization performance on target domain by mitigating the domain shift in data distribution of the source and target domain. Formally, given a set of source domain samples  $\mathcal{S} = \{(\mathbf{x}_1^s, y_1^s), \dots, (\mathbf{x}_{N_s}^s, y_{N_s}^s)\}$ , and target domain samples  $\mathcal{T} = \{\mathbf{x}_1^t, \dots, \mathbf{x}_{N_t}^t\}$ ,  $\mathbf{x}^s, \mathbf{x}^t$  represent the input data, and  $y^s \in \{0, 1, \dots, M-1\}$  denote the source data label of  $M$  classes. The target data label  $y^t \in \{0, 1, \dots, M-1\}$  is unknown. Thus, in UDA, we are interested in training a network using labeled source domain data  $\mathcal{S}$  and unlabeled target domain data  $\mathcal{T}$  to make accurate predictions  $\{\hat{y}^t\}$  on  $\mathcal{T}$ .

We discuss our method in the context of deep neural networks. In deep neural networks, a sample owns hierarchical features/representations denoted by the activations of each layer  $l \in \mathcal{L}$ . In the following, we use  $\phi_l(\mathbf{x})$  to denote the outputs of layer  $l$  in a deep neural network  $\Phi_\theta$  for the input  $\mathbf{x}$ , where  $\phi(\cdot)$  denotes the mapping defined by the deep neural network from the input to a specific layer.

In the rest of this section, we start our discussions by briefly reviewing the relevant concepts in MMD in Section 3.1. Section 3.2 introduces a new domain discrepancy metric. Finally, Section 3.3 and Section 3.4 discuss the objective and the training procedure of proposed deep network.

#### 3.1. Maximum Mean Discrepancy Revisit

In Maximum Mean Discrepancy (MMD),  $\{\mathbf{x}_i^s\}$  and  $\{\mathbf{x}_i^t\}$  are *i.i.d.* sampled from the marginal distributions  $P(\mathbf{X}^s)$  and  $Q(\mathbf{X}^t)$  respectively. Based on the observed samples, MMD [34] performs a kernel two-sample test to determine whether to accept the null hypothesis  $P = Q$  or not. MMD is motivated by the fact that if two distributions are identical, all of their statistics should be the same. Formally, MMD defines the difference between two distributions with their mean embeddings in the reproducing kernel Hilbert space (RKHS), *i.e.*

$$\mathcal{D}_{\mathcal{H}}(P, Q) \triangleq \sup_{f \sim \mathcal{H}} (\mathbb{E}_{\mathbf{X}^s}[f(\mathbf{X}^s)] - \mathbb{E}_{\mathbf{X}^t}[f(\mathbf{X}^t)])_{\mathcal{H}}, \quad (1)$$

where  $\mathcal{H}$  is class of functions.

In practice, for a layer  $l$ , the squared value of MMD is estimated with the empirical kernel mean embeddings

$$\begin{aligned} \hat{\mathcal{D}}_l^{mmd} &= \frac{1}{n_s^2} \sum_{i=1}^{n_s} \sum_{j=1}^{n_s} k_l(\phi_l(\mathbf{x}_i^s), \phi_l(\mathbf{x}_j^s)) \\ &+ \frac{1}{n_t^2} \sum_{i=1}^{n_t} \sum_{j=1}^{n_t} k_l(\phi_l(\mathbf{x}_i^t), \phi_l(\mathbf{x}_j^t)) \\ &- \frac{2}{n_s n_t} \sum_{i=1}^{n_s} \sum_{j=1}^{n_t} k_l(\phi_l(\mathbf{x}_i^s), \phi_l(\mathbf{x}_j^t)), \quad (2) \end{aligned}$$

where  $\mathbf{x}^s \in \mathcal{S}' \subset \mathcal{S}$ ,  $\mathbf{x}^t \in \mathcal{T}' \subset \mathcal{T}$ ,  $n_s = |\mathcal{S}'|$ ,  $n_t = |\mathcal{T}'|$ . The  $\mathcal{S}'$  and  $\mathcal{T}'$  represent the mini-batch source and target data sampled from  $\mathcal{S}$  and  $\mathcal{T}$  respectively. And  $k_l$  denotes the kernel selected for the  $l$ -th layer of deep neural network.

#### 3.2. Contrastive Domain Discrepancy

We propose to explicitly take the class information into account and measure the *intra-class* and *inter-class* discrepancy across domains. The intra-class domain discrepancy is minimized to compact the feature representations of samples within a class, whereas the inter-class domain discrepancy is maximized to push the representations of each other further away from the decision boundary. The intra-class and inter-class discrepancies are jointly optimized to improve the adaptation performance.

The proposed Contrastive Domain Discrepancy (CDD) is established on the difference between *conditional* data distributions across domains. Without any constraint on the type (*e.g.* marginal or conditional) of data distributions, MMD is convenient to measure such difference between  $P(\phi(\mathbf{X}^s)|Y^s)$  and  $Q(\phi(\mathbf{X}^t)|Y^t)$ , *i.e.*  $\mathcal{D}_{\mathcal{H}}(P, Q) \triangleq \sup_{f \sim \mathcal{H}} (\mathbb{E}_{\mathbf{X}^s}[f(\phi(\mathbf{X}^s)|Y^s)] - \mathbb{E}_{\mathbf{X}^t}[f(\phi(\mathbf{X}^t)|Y^t)])_{\mathcal{H}}$ .

Supposing  $\mu_{cc'}(y, y') = \begin{cases} 1 & \text{if } y = c, y' = c'; \\ 0 & \text{otherwise.} \end{cases}$ , for

two classes  $c_1, c_2$  (which can be same or different), the kernel mean embedding estimation for squared  $\mathcal{D}_{\mathcal{H}}(P, Q)$  is

$$\hat{\mathcal{D}}^{c_1 c_2}(\hat{y}_1^t, \hat{y}_2^t, \dots, \hat{y}_{n_t}^t, \phi) = e_1 + e_2 - 2e_3 \quad (3)$$

where

$$\begin{aligned} e_1 &= \sum_{i=1}^{n_s} \sum_{j=1}^{n_s} \frac{\mu_{c_1 c_1}(y_i^s, y_j^s) k(\phi(\mathbf{x}_i^s), \phi(\mathbf{x}_j^s))}{\sum_{i=1}^{n_s} \sum_{j=1}^{n_s} \mu_{c_1 c_1}(y_i^s, y_j^s)} \\ e_2 &= \sum_{i=1}^{n_t} \sum_{j=1}^{n_t} \frac{\mu_{c_2 c_2}(\hat{y}_i^t, \hat{y}_j^t) k(\phi(\mathbf{x}_i^t), \phi(\mathbf{x}_j^t))}{\sum_{i=1}^{n_t} \sum_{j=1}^{n_t} \mu_{c_2 c_2}(\hat{y}_i^t, \hat{y}_j^t)} \\ e_3 &= \sum_{i=1}^{n_s} \sum_{j=1}^{n_t} \frac{\mu_{c_1 c_2}(y_i^s, \hat{y}_j^t) k(\phi(\mathbf{x}_i^s), \phi(\mathbf{x}_j^t))}{\sum_{i=1}^{n_s} \sum_{j=1}^{n_t} \mu_{c_1 c_2}(y_i^s, \hat{y}_j^t)}. \quad (4) \end{aligned}$$

Note that Eq. (3) defines two kinds of class-aware domain discrepancy, 1) when  $c_1 = c_2 = c$ , it measures *intra-class* domain discrepancy; 2) when  $c_1 \neq c_2$ , it becomes the *inter-class* domain discrepancy. To compute the mask  $\mu_{c_2 c_2}(\hat{y}_i^t, \hat{y}_j^t)$  and  $\mu_{c_1 c_2}(y_i^s, \hat{y}_j^t)$ , we need to estimate target labels  $\{\hat{y}_i^t\}$ , which will be discussed in Section 3.4.

Based on the above definitions, the CDD is calculated as (The  $\hat{y}_1^t, \hat{y}_2^t, \dots, \hat{y}_{n_t}^t$  is abbreviated as  $\hat{y}_{1:n_t}^t$ )

$$\hat{\mathcal{D}}^{cdd} = \underbrace{\frac{1}{M} \sum_{c=1}^M \hat{\mathcal{D}}^{cc}(\hat{y}_{1:n_t}^t, \phi)}_{intra} - \underbrace{\frac{1}{M(M-1)} \sum_{c=1}^M \sum_{\substack{c'=1 \\ c' \neq c}}^M \hat{\mathcal{D}}^{cc'}(\hat{y}_{1:n_t}^t, \phi)}_{inter}, \quad (5)$$

where the intra- and inter-class domain discrepancies will be optimized in the opposite direction.

Note although the estimation of the labels  $\{\hat{y}_i^t\}$  can be noisy, the CDD (which is established on MMD) in itself is robust to the noise to an extent. Because MMD is determined by the mean embeddings of distributions in the RKHS, the sufficient statistics is less likely to be severely affected by the label noise, especially when the amount of data is large. We will discuss and verify this in Section 4.3.

### 3.3. Contrastive Adaptation Network

Deep convolutional neural networks (CNNs) is able to learn more transferable features than shallow methods. However, the discrepancy still exists for domain-specific layers. Specifically, the convolutional layers extracting general features are more transferable, while the fully-connected (FC) layers which exhibit abstract and domain-specific features should be adapted [22, 25].

In this paper, we start from ImageNet [7] pretrained networks, *e.g.* ResNet [14, 15], and replace the last FC layer with task-specific ones. We follow the general practice that minimizes the domain discrepancy of last FC layers and fine-tunes the convolutional layers through back-propagation. Then our proposed CDD can be readily incorporated into the objective as an adaptation module over the activations of FC layers. We name our network Contrastive Adaptation Network (CAN).

**The overall objective.** In a deep CNN, we need to minimize CDD over multiple FC layers, *i.e.* minimizing

$$\hat{\mathcal{D}}_{\mathcal{L}}^{cdd} = \sum_{l=1}^L \hat{\mathcal{D}}_l^{cdd}. \quad (6)$$

Besides, we train the network with labeled source data through minimizing the cross-entropy loss,

$$\ell^{ce} = -\frac{1}{n'_s} \sum_{i'=1}^{n'_s} \log P_{\theta}(y_{i'}^s | \mathbf{x}_{i'}^s) \quad (7)$$

where  $y^s \in \{0, 1, \dots, M-1\}$  is the ground-truth label of sample  $\mathbf{x}^s$ .  $P_{\theta}(y|\mathbf{x})$  denotes the predicted probability of label  $y$  with the network parameterized by  $\theta$ , given input  $\mathbf{x}$ .

Therefore, the overall objective can be formulated as

$$\min_{\theta} \ell = \ell^{ce} + \beta \hat{\mathcal{D}}_{\mathcal{L}}^{cdd} \quad (8)$$

where  $\beta$  is the weight of the discrepancy penalty term. Through minimizing  $\hat{\mathcal{D}}_{\mathcal{L}}^{cdd}$ , the intra-class domain discrepancy is minimized and the inter-class domain discrepancy is maximized to perform class-aware domain alignment.

Note that we independently sample the labeled source data to minimize the cross-entropy loss  $\ell^{ce}$  and those to estimate the CDD  $\hat{\mathcal{D}}_{\mathcal{L}}^{cdd}$ . In this way, we are able to design more efficient sampling strategy (see Section 3.4) to facilitate the mini-batch stochastic optimization with CDD, while not disturbing the conventional optimization with cross-entropy loss on labeled source data.

### 3.4. Optimizing CAN

The framework of CAN is illustrated in Fig. 2. In this section, we mainly focus on discussing how to minimize CDD loss in CAN.

**Alternative optimization (AO).** As shown in Eq. (5), we need to jointly optimize the target label hypothesis  $\hat{y}_{1:n_t}^t$  and the feature representations  $\phi_{1:L}$ . We adopt alternative steps to perform such optimization. In detail, at each loop, given current feature representations, *i.e.* fixing  $\theta$ , we update target labels through clustering. Then, based on the updated target labels  $\hat{y}^t$ , we estimate and minimize CDD to adapt the features, *i.e.* update  $\theta$  through back-propagation.

We employ the input activations  $\phi_1(\cdot)$  of the first task-specific layer to represent a sample. For example, in ResNet, each sample can be represented as the outputs of the global average pooling layer, which are also the inputs of the following task-specific layer. Then the spherical K-means is adopted to perform the clustering of target samples and attach corresponding labels. The number of clusters is the same as the number of underlying classes  $M$ . For each class, the target cluster center  $O^{tc}$  is initialized as the source cluster center  $O^{sc}$ , *i.e.*  $O^{tc} \leftarrow O^{sc}$ , where  $O^{sc} = \sum_{i=1}^{N_s} \mathbf{1}_{y_i^s=c} \frac{\phi_1(\mathbf{x}_i^s)}{\|\phi_1(\mathbf{x}_i^s)\|}$ ,  $\mathbf{1}_{y_i^s=c} \begin{cases} 1 & \text{if } y_i^s = c \\ 0 & \text{otherwise} \end{cases}$  and  $c = \{0, 1, \dots, M-1\}$ . For the metric measuring the distance between points  $\mathbf{a}$  and  $\mathbf{b}$  in the feature space, we apply the cosine dissimilarity, *i.e.*  $dist(\mathbf{a}, \mathbf{b}) = \frac{1}{2}(1 - \frac{\langle \mathbf{a}, \mathbf{b} \rangle}{\|\mathbf{a}\| \|\mathbf{b}\|})$ .

Then the clustering process is iteratively 1) attaching labels for each target samples:  $\hat{y}_i^t \leftarrow \operatorname{argmin}_c dist(\phi_1(\mathbf{x}_i^t), O^{tc})$ , and 2) updating the cluster centers:  $O^{tc} \leftarrow \sum_{i=1}^{N_t} \mathbf{1}_{\hat{y}_i^t=c} \frac{\phi_1(\mathbf{x}_i^t)}{\|\phi_1(\mathbf{x}_i^t)\|}$ , till convergence or reaching the maximum clustering steps.

After clustering, each target sample  $\mathbf{x}_i^t$  is assigned a label  $\hat{y}_i^t$  same as its affiliated clusters. Moreover, ambiguous data, which is far from its affiliated cluster center, is discarded, *i.e.* we select a subset  $\tilde{\mathcal{T}} = \{(\mathbf{x}^t, \hat{y}^t) | dist(\phi_1(\mathbf{x}^t), O^{t(\hat{y}^t)}) < D_0, \mathbf{x}^t \in \mathcal{T}\}$ , where  $D_0 \in [0, 1]$  is a constant.

Moreover, to give a more accurate estimation of the distribution statistics, we assume that the minimum number of samples in  $\tilde{\mathcal{T}}$  assigned to each class, should be guaran-

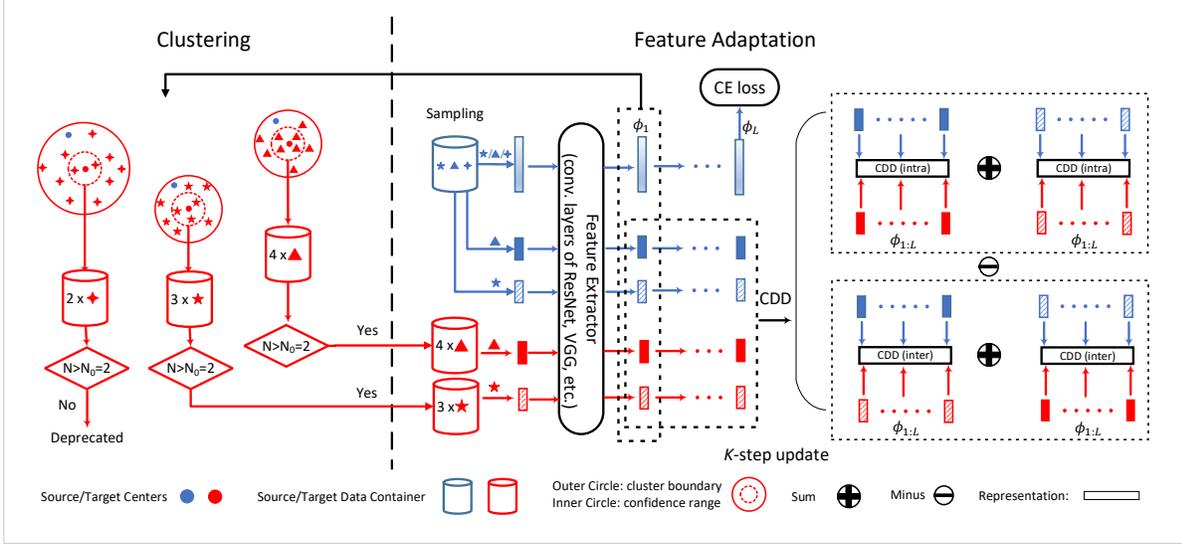


Figure 2. The training process of CAN. To minimize CDD, we perform alternative optimization between updating the target label hypothesis through clustering and adapting feature representations through back-propagation. For the clustering, we apply spherical K-means clustering of target samples based on their current feature representations. The number of clusters equal to that of underlying classes and the initial center of each class cluster is set to the center of source data within the same class. Then ambiguous data (*i.e.* far from the affiliated cluster centers) and ambiguous classes (*i.e.* containing few target samples around affiliated cluster centers) are discarded. For the feature adaptation, the labeled target samples provided by the clustering stage, together with the labeled source samples, pass through the network to achieve their multi-layer feature representations. The features of domain-specific FC layers are adopted to estimate CDD (Eq. (5)). Besides, we apply cross-entropy loss on independently sampled source data. Back-propagating with minimizing CDD and cross-entropy loss (Eq. (8)) adapts the features and provides class-aware alignment. Detailed descriptions can be found in Section 3.4.

teed. The class which doesn't satisfy such condition will not be considered in current loop, *i.e.* at loop  $T_e$ , the selected subset of classes  $\mathcal{C}_{T_e} = \{c | \sum_i^{\tilde{T}} \mathbf{1}_{y_i^t=c} > N_0, c \in \{0, 1, \dots, M-1\}\}$ , where  $N_0$  is a constant.

At the start of training, due to the domain shift, it is more likely to exclude partial classes. However, as training proceeds, more and more classes are included. The reason is two folds: 1) as training proceeds, the model becomes more accurate and 2) benefiting from the CDD penalty, the intra-class domain discrepancy becomes smaller, and the inter-class domain discrepancy becomes larger, so that the hard (*i.e.* ambiguous) classes are able to be taken into account.

**Class-aware Sampling (CAS).** In the conventional training of deep neural networks, a mini-batch of data is usually sampled at each iteration without being differentiated by their classes. However, it will be less efficient for computing the CDD. For example, for class  $C$ , there may only exist samples from one domain (source or target) in the mini-batch, thus the intra-class discrepancy could not be estimated.

We propose to use class-aware sampling strategy to enable the efficient update of network with CDD. It is easy to implement. We randomly select a subset of classes  $\mathcal{C}'_{T_e}$  from  $\mathcal{C}_{T_e}$ , and then sample source data and target data for each class in  $\mathcal{C}'_{T_e}$ . Consequently, in each mini-batch of data during training, we are able to estimate the intra-class dis-

crepancy for each selected class.

**Algorithm.** Algorithm 1 shows one loop of the AO procedure, *i.e.* alternating between a clustering phase (Step 1-4), and a  $K$ -step network update phase (Step 5-11). The loop of AO is repeated multiple times in our experiments. Because the feature adapting process is relatively slower, we asynchronously update the target labels and the network parameters to make the training more stable and efficient.

---

**Algorithm 1:** Optimization of CAN at loop  $T_e$ .

---

- Input:**  
source data:  $\mathcal{S} = \{(\mathbf{x}_1^s, y_1^s), \dots, (\mathbf{x}_{N_s}^s, y_{N_s}^s)\}$ ,  
target data:  $\mathcal{T} = \{\mathbf{x}_1^t, \dots, \mathbf{x}_{N_t}^t\}$
- Procedure:**
- 1 Forward  $\mathcal{S}$  and compute the  $M$  cluster centers  $O^{sc}$ ;
  - 2 Initialize  $O^{tc}$ :  $O^{tc} \leftarrow O^{sc}$ ;
  - 3 Cluster target samples  $\mathcal{T}$  using spherical K-means;
  - 4 Filter the ambiguous target samples and classes;
  - 5 **for** ( $k \leftarrow 1; k \leq K; k \leftarrow k + 1$ ) **do**
  - 6     Class-aware sampling based on  $\mathcal{C}'_{T_e}$ ,  $\tilde{\mathcal{T}}$ , and  $\mathcal{S}$ ;
  - 7     Compute  $\hat{D}_{\mathcal{L}}^{cdd}$  using Eq. (6);
  - 8     Sample from  $\mathcal{S}$  and compute  $\ell^{ce}$  using Eq. (7);
  - 9     Back-propagate with the objective  $\ell$  (Eq.(8));
  - 10    Update network parameters  $\theta$ .
  - 11 **end**
-



Figure 3. The gallery of Office-31 and VisDA-2017 datasets.

## 4. Experiments

### 4.1. Setups

**Datasets:** We validate our method on two public benchmarks. *Office-31* [30] is a common dataset for real-world domain adaptation tasks. It consists of 4,110 images belonging to 31 classes. This dataset contains three distinct domains, *i.e.*, images which are collected from the 1) Amazon website (Amazon domain), 2) web camera (Webcam domain), and 3) digital SLR camera (DSLR domain) under different settings, respectively. The dataset is imbalanced across domains, with 2,817 images in **A** domain, 795 images in **W** domain, and 498 images in **D** domain.

*VisDA-2017* [29] is a challenging testbed for UDA with the domain shift from synthetic data to real imagery. In this paper, we validate our method on its classification task. In total there are  $\sim 280k$  images from 12 categories. The images are split into three sets, *i.e.* a training set with 152,397 synthetic images, a validation set with 55,388 real-world images, and a test set with 72,372 real-world images. The gallery of two datasets is shown in Fig. 3

**Baselines:** We compare our method with class-agnostic discrepancy minimization methods: RevGrad [10, 11], DAN [22], and JAN [25]. Moreover, we compare our method with the ones which explicitly or implicitly take the class information or decision boundary into consideration to learn more discriminative features: MADA [28], MCD [32], and ADR [31]. The descriptions of these methods can be found in Section 2. We implement DAN and JAN using the released code<sup>1</sup>. For a comparison under optimal parameter setting, we cite the performance of MADA, RevGrad, MCD and ADR reported in their corresponding papers [28, 31, 32, 10].

**Implementation details:** We use ResNet-50 and ResNet-101 [14, 15] pretrained on ImageNet [7] as our backbone networks. We replace the last FC layer with the task-specific FC layer, and finetune the model with labeled source domain data and unlabeled target domain data. All

<sup>1</sup><https://github.com/thuml/Xlearn>

the network parameters are shared between the source domain and target domain data other than those of the batch normalization layers which are domain-specific. The hyperparameters are selected following the same protocol as described in [22], *i.e.* we train a domain classifier and perform selection on a validation set (of labeled source samples and unlabeled target samples) by jointly evaluating the test errors of the source classifier and the domain classifier.

We use mini-batch stochastic gradient descent (SGD) with momentum of 0.9 to train the network. We follow the same learning rate schedule as described in [10, 22, 25], *i.e.* the learning rate  $\eta_p$  is adjusted following  $\eta_p = \frac{\eta_0}{(1+ap)^b}$ , where  $p$  linearly increases from 0 to 1. The  $\eta_0$  is the initial learning rate, *i.e.* 0.001 for the convolutional layers and 0.01 for the task-specific FC layer. For Office-31,  $a = 10$  and  $b = 0.75$ , while for VisDA-2017,  $a = 10$  and  $b = 2.25$ . The  $\beta$  selected is 0.3. The thresholds ( $D_0, N_0$ ) are set to (0.05, 3) for Office-31 tasks **A**  $\rightarrow$  **W** and **A**  $\rightarrow$  **D**. And we don't filter target samples and classes for other tasks during training.

### 4.2. Comparison with the state-of-the-art

Table 1 shows the classification accuracy on six tasks of Office-31. All domain adaptation methods yield notable improvement over the ResNet model (first row) which is fine-tuned on labeled source data only. CAN outperforms other baseline methods across all tasks, achieving the state-of-the-art performance. On average, it boosts the accuracy of JAN by a absolute 6.3% and that of MADA by 5.4%.

We visualize the distribution of learned features by t-SNE [27]. Fig. 4 illustrates a representative task **W**  $\rightarrow$  **A**. Compared to JAN, as expected, the target data representations learned by CAN demonstrate higher intra-class compactness and much larger inter-class margin. This suggests that our CDD produces more discriminative features for the target domain and substantiates our improvement in Table 1.

Table 2 lists the accuracy over 12 classes on VisDA-2017 with the validation set as the target domain. Our method outperforms the other baseline methods. The mean accuracy of our model (87.2%) outperforms the self-ensembling (SE) method [9] (84.3%) which wins the first place in the VisDA-2017 competition, by 2.9%. It is worth noting that SE mainly deals with UDA by ensemble and data augmentation, which is orthogonal to the topic of this paper and thus can be easily combined to boost the performance further.

Moreover, we also perform adaptation on the VisDA-2017 test set (as the target domain), and submit our predictions to official evaluation server. Our goal is to evaluate the effectiveness of our proposed technique based on a vanilla backbone (ResNet-101). We choose not to use ensemble or additional data augmentation which is commonly used to boost the performance in the competition. Anyhow, our single model achieves a very competitive accuracy of 87.4%, which is comparable to the method which ranks at

Method	A $\rightarrow$ W	D $\rightarrow$ W	W $\rightarrow$ D	A $\rightarrow$ D	D $\rightarrow$ A	W $\rightarrow$ A	Average
Source-finetune	68.4 $\pm$ 0.2	96.7 $\pm$ 0.1	99.3 $\pm$ 0.1	68.9 $\pm$ 0.2	62.5 $\pm$ 0.3	60.7 $\pm$ 0.3	76.1
RevGrad [10, 11]	82.0 $\pm$ 0.4	96.9 $\pm$ 0.2	99.1 $\pm$ 0.1	79.7 $\pm$ 0.4	68.2 $\pm$ 0.4	67.4 $\pm$ 0.5	82.2
DAN [22]	80.5 $\pm$ 0.4	97.1 $\pm$ 0.2	99.6 $\pm$ 0.1	78.6 $\pm$ 0.2	63.6 $\pm$ 0.3	62.8 $\pm$ 0.2	80.4
JAN [25]	85.4 $\pm$ 0.3	97.4 $\pm$ 0.2	99.8 $\pm$ 0.2	84.7 $\pm$ 0.3	68.6 $\pm$ 0.3	70.0 $\pm$ 0.4	84.3
MADA [28]	90.0 $\pm$ 0.2	97.4 $\pm$ 0.1	99.6 $\pm$ 0.1	87.8 $\pm$ 0.2	70.3 $\pm$ 0.3	66.4 $\pm$ 0.3	85.2
Ours (intra only)	93.2 $\pm$ 0.2	98.4 $\pm$ 0.2	99.8 $\pm$ 0.2	92.9 $\pm$ 0.2	76.5 $\pm$ 0.3	76.0 $\pm$ 0.3	89.5
Ours (CAN)	<b>94.5 <math>\pm</math> 0.3</b>	<b>99.1 <math>\pm</math> 0.2</b>	<b>99.8 <math>\pm</math> 0.2</b>	<b>95.0 <math>\pm</math> 0.3</b>	<b>78.0 <math>\pm</math> 0.3</b>	<b>77.0 <math>\pm</math> 0.3</b>	<b>90.6</b>

Table 1. Classification accuracy (%) for all the six tasks of Office-31 dataset based on ResNet-50 [14, 15]. Our methods named “intra only” and “CAN” are trained with intra-class domain discrepancy and contrastive domain discrepancy, respectively.

Method	airplane	bicycle	bus	car	horse	knife	motorcycle	person	plant	skateboard	train	truck	Average
Source-finetune	72.3	6.1	63.4	<b>91.7</b>	52.7	7.9	80.1	5.6	90.1	18.5	78.1	25.9	49.4
RevGrad [10, 11]	81.9	77.7	82.8	44.3	81.2	29.5	65.1	28.6	51.9	54.6	82.8	7.8	57.4
DAN [22]	68.1	15.4	76.5	87.0	71.1	48.9	82.3	51.5	88.7	33.2	88.9	42.2	62.8
JAN [25]	75.7	18.7	82.3	86.3	70.2	56.9	80.5	53.8	92.5	32.2	84.5	54.5	65.7
MCD [32]	87.0	60.9	83.7	64.0	88.9	79.6	84.7	76.9	88.6	40.3	83.0	25.8	71.9
ADR [31]	87.8	79.5	83.7	65.3	92.3	61.8	88.9	73.2	87.8	60.0	85.5	32.3	74.8
SE [9]	95.9	<b>87.4</b>	<b>85.2</b>	58.6	96.2	95.7	90.6	80.0	94.8	90.8	88.4	47.9	84.3
Ours (intra only)	96.5	72.1	80.9	70.8	94.6	<b>98.0</b>	<b>91.7</b>	<b>84.2</b>	90.3	89.8	<b>89.4</b>	47.9	83.9
Ours (CAN)	<b>97.0</b>	87.2	82.5	74.3	<b>97.8</b>	96.2	90.8	80.7	<b>96.6</b>	<b>96.3</b>	87.5	<b>59.9</b>	<b>87.2</b>

Table 2. Classification accuracy (%) on the VisDA-2017 validation set based on ResNet-101 [14, 15]. Our methods named “intra only” and “CAN” are trained with intra-class domain discrepancy and contrastive domain discrepancy, respectively.

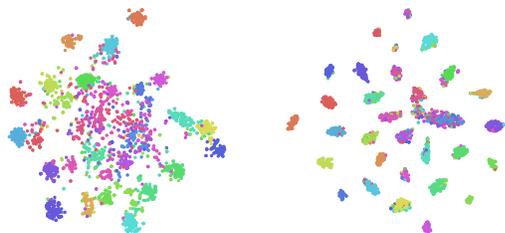


Figure 4. Visualization with t-SNE for different adaptation methods (bested viewed in color). **Left:** t-SNE of JAN. **Right:** CAN. The input activations of the last FC layer are used for the computation of t-SNE. The results are on Office-31 task W  $\rightarrow$  A.

the second place on the leaderboard (87.7%).

From Table 1 and 2, we have two observations: 1) Taking class information/decision boundary into account is beneficial for the adaptation. It can be seen that MADA, MCD, ADR and our method achieve better performance than class-agnostic methods, *e.g.* RevGrad, DAN, JAN, *etc.* 2) Our way of exploiting class information is more effective. We achieve better accuracy than MADA (+5.4%), ADR (+12.4%), and MCD (+15.3%).

### 4.3. Ablation studies

**Effect of inter-class domain discrepancy.** We compare our method (“CAN”) with that trained using intra-class discrepancy only (“intra only”), to verify the merits of intro-

ducing inter-class domain discrepancy measure. The results are shown in the last two rows in Table 1 and 2. It can be seen that introducing the inter-class domain discrepancy improves the adaptation performance. We believe the reason is that it is impossible to completely eliminate the intra-class domain discrepancy, maximizing the inter-class domain discrepancy may alleviate the possibility of the model overfitting to the source data and benefits the adaptation.

**Effect of alternative optimization and class-aware sampling.** Table 3 examines two key components of CAN, *i.e.* alternative optimization (or “AO”), and class-aware sampling (or “CAS”). We perform ablation study by leaving-one-component-out of our framework at a time. In Table 3, the method “w/o. AO” directly employs the outputs of the network at each iteration as pseudo target labels to estimate CDD and back-propagates to update the network. It can be regarded as updating the feature representations and pseudo target labels simultaneously. The method “w/o. CAS” uses conventional class-agnostic sampling instead of CAS. The comparisons to these two special cases verify the contributions of AO and CAS in our method.

Interestingly, even without alternative optimization, the method “w/o. AO” improves over class-agnostic methods, *e.g.* DAN, JAN, *etc.* This suggests our proposed CDD in itself is robust to the label noise to some extent, and MMD is a suitable metric to establish CDD (see Section 3.2).

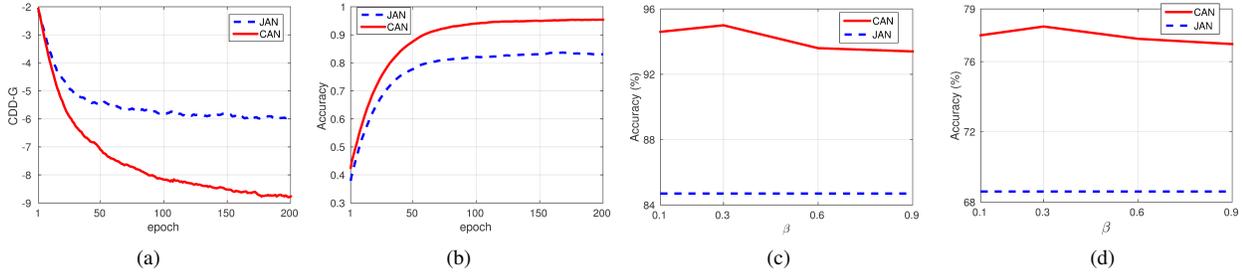


Figure 5. (a-b) The curve of CDD and accuracy during training on task  $A \rightarrow D$  of the Office-31 dataset. The “CDD-G” denotes the contrastive domain discrepancy computed with ground-truth target labels. (c-d) The sensitivity of accuracy of CAN to  $\beta$ . The results for  $A \rightarrow D$  (Left) and  $D \rightarrow A$  (Right) are illustrated as examples. The trends for other tasks are similar.

Dataset	w/o. AO	w/o. CAS	CAN
Office-31	88.1	89.1	90.6
VisDA-2017	77.5	81.6	87.2

Table 3. The effect of alternative optimization (AO) and CAS. The mean accuracy over six tasks on Office-31 and the mean accuracy over 12 classes on VisDA-2017 validation set are reported.

Method	$A \rightarrow W$	$A \rightarrow D$	$D \rightarrow A$	$W \rightarrow A$	Average
pseudo <sub>0</sub>	85.8	86.3	74.9	72.3	79.8
pseudo <sub>1</sub>	90.2 ± 1.6	92.5 ± 0.4	75.7 ± 0.2	75.3 ± 0.6	83.4
CAN	94.5 ± 0.3	95.0 ± 0.3	78.0 ± 0.3	77.0 ± 0.3	86.1

Table 4. Comparison with different ways of utilizing pseudo target labels. The “pseudo<sub>0</sub>” means training with pseudo target labels (achieved by our initial clustering) directly. The “pseudo<sub>1</sub>” is to alternatively update target labels through clustering and minimize the cross-entropy loss on pseudo labeled target data. In “pseudo<sub>1</sub>”, the cross-entropy loss on source data is also minimized.

**Ways of using pseudo target labels.** The estimates for the target labels can be achieved through clustering, which enables various ways to train a model. In Table 4, we compare our method with two different ways of training with pseudo target labels achieved by the clustering. One way (“pseudo<sub>0</sub>”) is to fix these pseudo labels to train a model directly. The other (“pseudo<sub>1</sub>”) is to update the pseudo target labels during training, which is the same as CAN, but to train the model based on the cross-entropy loss over pseudo labeled target data rather than estimating the CDD.

As shown in Table 4, “pseudo<sub>0</sub>” leads to a model whose accuracy exactly matches with that of the initial clustering, due to the large capacity of deep neural networks. The “pseudo<sub>1</sub>” achieves significantly better results than “pseudo<sub>0</sub>”, but is still worse than our CAN, which verifies that our way of explicitly modeling the class-aware domain discrepancy makes the model better adapted and less likely to be affected by the label noise.

**CDD value during training.** In our training, we generate target label hypothesis to estimate CDD. We expect that the underlying metric computed with the ground-truth target labels would decrease steadily during training until

convergence. To do so, during training, we evaluate the *ground-truth* CDD (denoted by CDD-G) for JAN and CAN with the ground-truth target labels. The trend of CDD and the test accuracy during training are plotted in Fig. 5.

As we see, for JAN (the blue curve), the ground-truth CDD rapidly becomes stable at a high level after a short decrease. This indicates that JAN cannot minimize the contrastive domain discrepancy effectively. For CAN (the red curve), although we can only estimate the CDD using inaccurate target label hypothesis, its CDD value steadily decreases along training. The result illustrates our estimation works as a good proxy of ground-truth contrastive domain discrepancy. And from the accuracy curve illustrated in Fig. 5, we see that minimizing CDD leads to notable accuracy improvement of CAN, compared to JAN.

**Hyper-parameter sensitivity.** We study the sensitivity of CAN to the important balance weight  $\beta$  on two example tasks  $A \rightarrow D$  and  $D \rightarrow A$  in Fig. 5. Generally, our model is less sensitive to the change of  $\beta$ . In a vast range, the performance of CAN outperforms the baseline method with a large margin (the blue dashed curve). As the  $\beta$  gets larger, the accuracy steadily increases before decreasing. The bell-shaped curve illustrates the regularization effect of CDD.

## 5. Conclusion

In this paper, we proposed Contrastive Adaptation Network to perform class-aware alignment for UDA. The intra-class and inter-class domain discrepancy are explicitly modeled and optimized through end-to-end mini-batch training. Experiments on real-world benchmarks demonstrate the superiority of our model compared with the strong baselines.

**Acknowledgement.** This work was supported in part by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/Interior Business Center (DOI/IBC) contract number D17PC00340. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes not withstanding any copyright annotation/herein. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DOI/IBC, or the U.S. Government.

## References

- [1] S. Ben-David, J. Blitzer, K. Crammer, A. Kulesza, F. Pereira, and J. W. Vaughan. A theory of learning from different domains. *Machine learning*, 79(1-2):151–175, 2010. [1](#)
- [2] S. Ben-David, J. Blitzer, K. Crammer, and F. Pereira. Analysis of representations for domain adaptation. In *Advances in neural information processing systems*, pages 137–144, 2007. [1](#)
- [3] K. Bousmalis, N. Silberman, D. Dohan, D. Erhan, and D. Krishnan. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, page 7, 2017. [1](#)
- [4] K. Bousmalis, G. Trigeorgis, N. Silberman, D. Krishnan, and D. Erhan. Domain separation networks. In *Advances in Neural Information Processing Systems*, pages 343–351, 2016. [2](#)
- [5] L. Bruzzone and M. Marconcini. Domain adaptation problems: A dasvm classification technique and a circular validation strategy. *IEEE transactions on pattern analysis and machine intelligence*, 32(5):770–787, 2010. [1](#)
- [6] D. Cheng, Y. Gong, S. Zhou, J. Wang, and N. Zheng. Person re-identification by multi-channel parts-based cnn with improved triplet loss function. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1335–1344, 2016. [2](#)
- [7] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. [1](#), [4](#), [6](#)
- [8] X. Dong and Y. Yang. Searching for a robust neural architecture in four gpu hours. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [1](#)
- [9] G. French, M. Mackiewicz, and M. Fisher. Self-ensembling for domain adaptation. *ICLR*, 2018. [6](#), [7](#)
- [10] Y. Ganin and V. Lempitsky. Unsupervised domain adaptation by backpropagation. *ICML*, 2015. [2](#), [6](#), [7](#)
- [11] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1):2096–2030, 2016. [6](#), [7](#)
- [12] R. Hadsell, S. Chopra, and Y. LeCun. Dimensionality reduction by learning an invariant mapping. In *Computer vision and pattern recognition, 2006 IEEE computer society conference on*, volume 2, pages 1735–1742. IEEE, 2006. [2](#)
- [13] P. Haeusser, T. Frerix, A. Mordvintsev, and D. Cremers. Associative domain adaptation. In *International Conference on Computer Vision (ICCV)*, volume 2, page 6, 2017. [2](#)
- [14] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [4](#), [6](#), [7](#)
- [15] K. He, X. Zhang, S. Ren, and J. Sun. Identity mappings in deep residual networks. In *European Conference on Computer Vision*, pages 630–645. Springer, 2016. [4](#), [6](#), [7](#)
- [16] A. Hermans, L. Beyer, and B. Leibe. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*, 2017. [2](#)
- [17] J. Hoffman, E. Tzeng, T. Park, J.-Y. Zhu, P. Isola, K. Saenko, A. A. Efros, and T. Darrell. Cycada: Cycle-consistent adversarial domain adaptation. *arXiv preprint arXiv:1711.03213*, 2017. [1](#)
- [18] J. Hoffman, D. Wang, F. Yu, and T. Darrell. Fcns in the wild: Pixel-level adversarial and constraint-based adaptation. *arXiv preprint arXiv:1612.02649*, 2016. [1](#)
- [19] L. Jiang, D. Meng, T. Mitamura, and A. G. Hauptmann. Easy samples first: Self-paced reranking for zero-example multimedia search. In *Proceedings of the 22nd ACM international conference on Multimedia*, pages 547–556. ACM, 2014. [1](#)
- [20] G. Kang, J. Li, and D. Tao. Shakeout: A new approach to regularized deep neural network training. *IEEE transactions on pattern analysis and machine intelligence*, 40(5):1245–1258, 2018. [1](#)
- [21] G. Kang, L. Zheng, Y. Yan, and Y. Yang. Deep adversarial attention alignment for unsupervised domain adaptation: the benefit of target expectation maximization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 401–416, 2018. [2](#)
- [22] M. Long, Y. Cao, J. Wang, and M. I. Jordan. Learning transferable features with deep adaptation networks. *arXiv preprint arXiv:1502.02791*, 2015. [1](#), [2](#), [4](#), [6](#), [7](#)
- [23] M. Long, J. Wang, G. Ding, J. Sun, and S. Y. Philip. Transfer feature learning with joint distribution adaptation. In *Computer Vision (ICCV), 2013 IEEE International Conference on*, pages 2200–2207. IEEE, 2013. [2](#)
- [24] M. Long, H. Zhu, J. Wang, and M. I. Jordan. Unsupervised domain adaptation with residual transfer networks. In *Advances in Neural Information Processing Systems*, pages 136–144, 2016. [2](#)
- [25] M. Long, H. Zhu, J. Wang, and M. I. Jordan. Deep transfer learning with joint adaptation networks. *ICML*, 2017. [1](#), [2](#), [4](#), [6](#), [7](#)
- [26] Y. Luo, L. Zheng, T. Guan, J. Yu, and Y. Yang. Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [1](#)
- [27] L. v. d. Maaten and G. Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008. [6](#)
- [28] Z. Pei, Z. Cao, M. Long, and J. Wang. Multi-adversarial domain adaptation. In *AAAI Conference on Artificial Intelligence*, 2018. [2](#), [6](#), [7](#)
- [29] X. Peng, B. Usman, N. Kaushik, J. Hoffman, D. Wang, and K. Saenko. Visda: The visual domain adaptation challenge. *arXiv preprint arXiv:1710.06924*, 2017. [2](#), [6](#)
- [30] K. Saenko, B. Kulis, M. Fritz, and T. Darrell. Adapting visual category models to new domains. In *European conference on computer vision*, pages 213–226. Springer, 2010. [1](#), [2](#), [6](#)
- [31] K. Saito, Y. Ushiku, T. Harada, and K. Saenko. Adversarial dropout regularization. *ICLR*, 2018. [2](#), [6](#), [7](#)
- [32] K. Saito, K. Watanabe, Y. Ushiku, and T. Harada. Maximum classifier discrepancy for unsupervised domain adaptation. *CVPR*, 2018. [2](#), [6](#), [7](#)

- [33] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015. 2
- [34] D. Sejdinovic, B. Sriperumbudur, A. Gretton, and K. Fukumizu. Equivalence of distance-based and rkhs-based statistics in hypothesis testing. *The Annals of Statistics*, pages 2263–2291, 2013. 3
- [35] O. Sener, H. O. Song, A. Saxena, and S. Savarese. Learning transferrable representations for unsupervised domain adaptation. In *Advances in Neural Information Processing Systems*, pages 2110–2118, 2016. 2
- [36] B. Sun and K. Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *European Conference on Computer Vision*, pages 443–450. Springer, 2016. 2
- [37] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell. Adversarial discriminative domain adaptation. In *Computer Vision and Pattern Recognition (CVPR)*, volume 1, page 4, 2017. 1
- [38] E. Tzeng, J. Hoffman, N. Zhang, K. Saenko, and T. Darrell. Deep domain confusion: Maximizing for domain invariance. *arXiv preprint arXiv:1412.3474*, 2014. 2
- [39] J. Wang, W. Feng, Y. Chen, H. Yu, M. Huang, and P. S. Yu. Visual domain adaptation with manifold embedded distribution alignment. In *2018 ACM Multimedia Conference on Multimedia Conference*, pages 402–410. ACM, 2018. 2
- [40] L. Zhu, Z. Xu, and Y. Yang. Bidirectional multirate reconstruction for temporal modeling in videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2653–2662, 2017. 1