

Re-energizing Domain Discriminator with Sample Relabeling for Adversarial Domain Adaptation

Xin Jin¹ Cuiling Lan^{2*} Wenjun Zeng² Zhibo Chen^{1*}

¹ University of Science and Technology of China ² Microsoft Research Asia, Beijing, China

jinxustc@mail.ustc.edu.cn {culan,wezeng}@microsoft.com chen-zhibo@ustc.edu.cn

Abstract

Many unsupervised domain adaptation (UDA) methods exploit domain adversarial training to align the features to reduce domain gap, where a feature extractor is trained to fool a domain discriminator in order to have aligned feature distributions. The discrimination capability of the domain classifier w.r.t. the increasingly aligned feature distributions deteriorates as training goes on, thus cannot effectively further drive the training of feature extractor. In this work, we propose an efficient optimization strategy named *Reinforceable Adversarial Domain Adaptation (RADA)* which aims to re-energize the domain discriminator during the training by using dynamic domain labels. Particularly, we relabel the well aligned target domain samples as source domain samples on the fly. Such relabeling makes the less separable distributions more separable, and thus leads to a more powerful domain classifier w.r.t. the new data distributions, which in turn further drives feature alignment. Extensive experiments on multiple UDA benchmarks demonstrate the effectiveness and superiority of our RADA.

1. Introduction

The rapid development of deep learning has helped achieve significant improvements for many computer vision tasks. Those learned models usually have high performance on the test data that share similar characteristics as the training data, but suffer from significant performance degradation when deployed in new environments [11, 13, 26, 25]. There are *domain gaps* between the training source data (source domain) and the testing target data (target domain). Besides, annotating the data for the target domain is expensive and time-consuming. To mitigate the notorious domain gap issue without manual annotation efforts, unsupervised domain adaptation (UDA) [11, 13, 19, 22, 64, 27, 9] is recently extensively explored, which aims to learn from labeled source domain and unlabeled target domain.

*Corresponding Author.

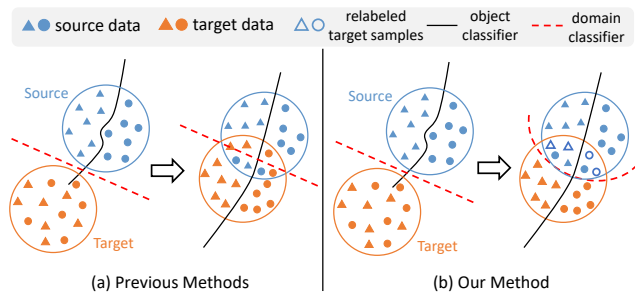


Figure 1: Motivation and our idea. (a) Previous methods align domain distributions by adversarially training the domain classifier with static domain labels of the samples. The discrimination capability of the domain classifier w.r.t. the increasingly aligned feature distributions deteriorates as training goes on (*i.e.*, more samples are inseparable), which in turn provides less driving power to the feature extractor for alignment and prevents effective optimization. (b) We propose a new solution which allows dynamic domain labels. We relabel the “well aligned” target samples as source domain, which makes the two less separable distributions more separable and thus leads to a more powerful domain classifier, which in turn further drives feature alignment.

For domain adaptation, theoretical analysis by Ben-David *et al.* [1] shows that reducing the feature differences between the source and target domains can reduce the upper bound of the target domain error. Many domain adaptation methods tend to learn domain invariant/transferable feature representations [15, 17, 14, 35]. Inspired by the Generative Adversarial Networks (GANs) [16], adversarial learning has been successfully applied for UDA [13, 46, 7, 43, 49, 31, 62, 10]. The core idea of adversarial domain adaptation approaches is to train a domain discriminator/classifier to distinguish between the source and target domains and train the generator/feature extractor to minimize the feature discrepancy between the source and target domain in order to fool the discriminator in a minmax two-player game. High discrimination capability for the domain discriminator is desired in order to be able to drive the feature alignment. Typically, a domain adversarial neural

network (DANN) [13] introduces a gradient reversal layer (GRL) for adversarial training, where the ordinary gradient descent is applied for optimizing the domain classifier and the sign of the gradient is reversed when passing through the GRL to optimize the feature extractor. Conditional domain adversarial network (CDAN) [31] improves DANN by conditioning the domain discriminator on both the object classification predictions/likelihoods and the extracted features.

All these methods train the domain discriminator with the static domain labels of the source and target samples. However, as shown in Figure 1 (a), as the adversarial training goes on, the feature distributions for the source domain and target domain are increasingly aligned. **The discrimination capability of the domain discriminator/classifier w.r.t. the more aligned distributions is weaker than that w.r.t. the earlier less aligned distributions. Such degradation of discrimination capability in turn provides less driving power to the feature extractor for alignment, hindering effective optimization,** even though there are still not aligned samples in the feature space.

To better understand what is happening, on top of a representative baseline scheme CDAN [31], we observe the variation of the discrimination capability of the domain discriminator by calculating the average entropy of domain classification for all the training samples at each training epoch (see Figure 2 (a)), where an epoch is a single pass through the full training set. As we know, a larger entropy of domain classification (*i.e.*, the discriminator has larger ambiguity on the domain classes of the samples) indicates a poorer discrimination capability of the domain classifier. We also observe how well the alignment is by calculating a domain discrepancy measurement Maximum Mean Discrepancy (MMD) [30, 2, 55] (see Figure 2 (b)).

Figure 2 (a) reveals that the average entropy of the baseline scheme goes through a process of fast decreasing and then slowly rises with fluctuation. 1) In the early epochs, the successive training of domain classifier increases its discrimination power and thus the entropy decreases. 2) Meanwhile, as the training goes on, the feature distributions for the source domain and target domain are increasingly aligned, *i.e.*, the domain discrepancy decreases as shown in Figure 2 (b). This could adversarially increase the entropy. 3) The optimization progresses/paces of the feature extractor and the domain discriminator are usually not the same. Around 5 to 15 epochs, the domain discrepancy continues to decrease even when the discrimination capability reduces (*i.e.*, entropy increases). This may be because the optimization of feature extractor lags behind that of the domain discriminator and thus the feature extractor can be further optimized. 4) The discrimination capability of the domain classifier w.r.t. the more aligned distributions is becoming weaker than that w.r.t. the earlier less aligned distributions (the entropy increases). The persistent degradation




Figure 2: The variation of (a) the discrimination capability of the domain discriminator (measured by entropy of domain classification) and (b) alignment state (measured by domain discrepancy measure of MMD) in the training. For the baseline scheme CDAN [31] (marked by green), the discrimination capability of the domain discriminator deteriorates w.r.t. the gradually aligned distributions after the initial dip of the entropy, which in turn provides less driving power to the feature extractor for alignment. In contrast, thanks to our strategy of re-labeling the aligned target samples as source samples, our scheme (marked by red) could improve the discrimination capability of the domain discriminator (*i.e.*, preventing the increasing of the entropy) and thus in turn further drives feature alignment. These experiments are conducted on Office-31 of the setting $W \rightarrow A$. Note that more analysis can be found in our **Supplementary**.

in turn provides less driving power to the feature extractor for alignment, hindering effective optimization (where the alignment state cannot be effectively improved after 15 epochs, please refer to the green curve in Figure 2 (b)). **Thus, a mechanism that could enhance/improve on the fly the discrimination capability of the domain classifier to re-energize the adversarial training is highly desired.**

Motivated by this, in this paper, we propose a new optimization strategy named Re-enforceable Adversarial Domain Adaptation (RADA) which aims to re-energize the domain discriminator/classifier during the training. Particularly, instead of training the domain classifier using static domain labels for the source and target samples (see Figure 1 (a)), we propose to exploit dynamic domain labels, where we relabel the “well aligned” target samples, *i.e.*, those close to source domain, as source domain samples, in each mini-batch to be optimized/trained. As illustrated in Figure 1 (b), this makes the two less separable distributions more separable and thus leads to a more powerful domain discriminator. Such re-energized domain discriminator in turn further drives the feature alignment for the feature extractor. From Figure 2, we can observe that after using our proposed data relabeling strategy, the domain classification entropy is much lower than that of the baseline scheme and the domain discrepancy is also smaller.

We summarize our main contributions as follows:

- We pinpoint that the popular adversarial domain adaptation approaches in general face optimization difficulty, which is caused by the deteriorated discrimination capa-

bility of the domain classifier as the feature distributions become increasingly aligned in training.

- To alleviate the deterioration of the discrimination capability of the domain discriminator, we propose an efficient optimization strategy named Re-enforceable Adversarial Domain Adaptation (RADA), which is capable of re-energizing the domain discriminator during training which in turn further drives feature alignment. We achieve this by relabeling the well aligned target samples as source domain samples for online training of the domain discriminator.

We will demonstrate the effectiveness of the proposed RADA on top of multiple widely-used adversarial learning based domain adaptation baselines. RADA significantly outperforms the state-of-the-art UDA approaches. RADA is simple yet effective and can be used as a *plug-and-play* optimization strategy for existing adversarial learning based UDA approaches. Note that we do not make any change to the network architecture of the domain discriminator, which makes the RADA friendly to many adversarial domain adaptation methods.

2. Related Work

Unsupervised Domain Adaptation (UDA). UDA has achieved great progress with the advancement of deep learning [11, 13, 19, 22, 27, 9]. These UDA methods could be mainly categorized into three classes: pixel-level translation based methods [5, 4, 29], explicit domain alignment based methods [17, 30, 32, 44, 55], and adversarial learning based methods [13, 46, 7, 43, 49, 31, 62, 10]. The third category predominates in recent years due to their superior performance and we will focus and elaborate on them.

Inspired by Generative adversarial networks (GANs) [16], adversarial domain adaptation has been widely explored [7, 43, 49]. The core idea of adversarial UDA approaches is to train a domain discriminator to distinguish the source and target domains and train a feature extractor to learn domain-invariant features to fool the discriminator. Domain Adversarial Neural Network (DANN) [13] is a representative work, where a domain classifier is connected to the feature extractor via a gradient reversal layer (GRL). Subsequently, many well-designed adversarial UDA variants are proposed, such as ADDA [46], CyCADA [21], SBADA [40], CDAN [31], MCD [42], MSTN [52] and TADA [50]. Symnets [62] builds a symmetric design of source and target task classifiers, on which the domain discrimination and domain confusion training are based. TADA [50] is also built upon adversarial domain adaptation framework, which adds adversarial alignment constraints on both of transferable local regions and global images through two local/global attention modules. MSTN [52] improves the adversarial UDA framework by additionally leveraging clustering technique to align labeled source cen-

troid and pseudo-labeled target centroid, where features in the same class but different domains are expected to be mapped to nearby. Similarly, PFAN [6] performs feature alignment between category-wise prototypes, with the number of target domain samples progressively increasing. CDAN [31] proposes to condition the domain discriminator on the discriminative information conveyed in the classifier predictions (class likelihood). Besides, for safe transfer, it prioritizes over those easy-to-transfer samples with confident predictions of object classes (measured by the entropy) by reweighting each sample by an entropy-aware weight. Recently, Cui *et al.* [10] equip the adversarial adaptation framework with a Gradually Vanishing Bridge (GVB) mechanism, which reduces the transfer difficulty by gradually reducing the domain-specific characteristics in domain-invariant representations, which idea is similar to [14] where authors exploit the labeled source data landmarks (*i.e.*, that similar to the target domain) to bridge domain gap by constructing provably easier auxiliary domain adaptation tasks. The solutions of those auxiliary tasks form the basis to compose invariant features for the original task. CMSS [56] introduces an agent to learn a dynamic curriculum to softly select/re-weight source samples that are best suited for aligning to the target domain.

As the optimization of the minmax game in domain adversarial adaptation proceeds, the increasingly aligned feature distributions would degrade the discrimination capability of the domain classifier, which reduces the driving power for the feature alignment. In this paper, we propose to relabel the well aligned target samples as source samples to make the less separable distributions more separable, re-energizing the entire adversarial optimization.

Mixup Technique. Mixup, which mixes two samples and their labels to generate a new sample, as a data augmentation/regularization technique, was first introduced in [59] and was widely applied in deep learning pipelines to improve performance [20, 8, 58, 18]. The study of using Mixup for domain adaptation is rare [53, 51] and remains an open problem. Xu *et al.* [53] present adversarial domain adaptation with domain mixup (DM-ADA) to guarantee domain-invariance in a more continuous latent space, where two samples (one from source domain and the other from target domain) and their domain labels are softly mixed to explore inter-domain information. Wu *et al.* [51] propose a dual mixup regularized learning (DMRL) method for UDA, which jointly conducts category and domain mixup regularization on the pixel level to guide the object classifier in enhancing consistent predictions in-between samples and enrich the intrinsic structures of the latent space. These mixup methods aim to explore inter-domain information. In contrast, we aim to enhance the continuity within the source domain after the re-partition of domains, where the samples to be mixed share the same source domain labels.

3. Re-enforceable Adversarial Domain Adaptation (RADA)

We propose an efficient domain adaptation optimization strategy, named Re-enforceable Adversarial Domain Adaptation (RADA), which re-energizes the domain classifier/discriminator during the adversarial training by using sample domain re-labeling strategy. Figure 3 illustrates the main procedure of our RADA optimization. As the adversarial training goes on, the source and target distributions are increasingly aligned. As illustrated in Figure 3 (b), some samples are already aligned such that the domain classifier cannot distinguish them well while some other samples are still not aligned. As revealed in Figure 2, the discrimination capability of the domain classifier (given a fixed model size) w.r.t. the more aligned distributions deteriorates since the aligned samples are inseparable, which in turn results in less driving power for the feature alignment.

We re-energize this domain discriminator by relabeling the well-aligned target domain samples as source domain, leading to a clear boundary for the two new distributions. Particularly, as illustrated in Figure 3 (c), we relabel the well aligned target samples as source domain samples and train the domain classifier under the updated partition. To determine which target samples are to be relabeled in the training, we use a metric to evaluate how well a target sample is aligned to the source domain, by simply taking into account the certainty (entropy) of the prediction of the domain classifier. When the entropy is large, it indicates the sample is close to the domain classification boundary and we will relabel it to the source domain.

For unsupervised domain adaption classification, we denote the source domain as $\mathcal{D}_S = \{(\mathbf{x}_i^s, y_i^s)\}_{i=1}^{N_s}$ with N_s labeled samples covering C classes, and the target domain as $\mathcal{D}_T = \{\mathbf{x}_j^t\}_{j=1}^{N_t}$ with N_t unlabeled samples that belong to the same C classes. UDA aims to train on \mathcal{D}_S and \mathcal{D}_T to achieve high accuracy on the target domain test set.

3.1. Recap of Adversarial Domain Adaptation

Adversarial Domain Adaptation aims to learn domain invariant representations for UDA, by adversarially training the feature extractor and domain discriminator/classifier. The adversarial domain adaptation [13, 31, 62, 10] can be achieved by optimizing object classification loss \mathcal{L}_{cls} and domain adversarial loss \mathcal{L}_{adv} [12, 21, 31]:

$$\mathcal{L}_{cls} = \frac{1}{N_s} \sum_{i=1}^{N_s} \mathcal{L}_{ce}(C(F(\mathbf{x}_i^s))), y_i^s), \quad (1)$$

$$\begin{aligned} \mathcal{L}_{adv} = & -\frac{1}{N_s} \sum_{i=1}^{N_s} \log D(F(\mathbf{x}_i^s)) \\ & -\frac{1}{N_t} \sum_{j=1}^{N_t} \log(1 - D(F(\mathbf{x}_j^t))), \quad (2) \end{aligned}$$

where F, C, D denote the feature extractor, the object classifier, and domain discriminator/classifier, respectively. \mathcal{L}_{ce} is the cross-entropy loss. The total training objective is described as follows:

$$\min_D \mathcal{L}_{adv}, \quad (3)$$

$$\min_{F,C} \mathcal{L}_{cls} - \lambda \mathcal{L}_{adv}, \quad (4)$$

where the domain classifier D is trained to discriminate the domains. Conversely, the feature extractor is trained to extract features that make it difficult for the domain classifier to discriminate the domains, resulting in domain invariant features. λ denotes a hyper-parameter for balancing the losses. Gradient reversal layer (GRL) [12] which connects feature extractor F and the domain discriminator D is usually used to achieve the adversarial function by multiplying the gradient from D by a certain negative constant during the back-propagation to the feature extractor.

3.2. Proposed Training Strategy RADA

During the training, the discrimination capability of the domain classifier w.r.t. to the increasingly aligned distributions *deteriorates*. This in turn negatively influences the effectiveness of the optimization towards feature alignment. Thus, with the adversarial optimization objectives as in (3) and (4), the networks converge to states that are less than optimal. Even though the distributions are not satisfactorily/fully aligned, the deteriorated discrimination capability of the domain classifier cannot further efficiently drive the feature alignment. We propose to re-energize the domain classifier by making the two distributions more separable. As illustrated in Figure 3 (c), we relabel the ‘‘well aligned’’ (ambiguous w.r.t. domain classifier) target domain samples as source domain samples. *Such relabeling makes the less separable distributions more separable and thus leads to a more powerful domain classifier, which in turn further drives feature alignment.* As revealed in Figure 2 (a), the discrimination capability of the domain discriminator after using our relabeling strategy is stronger (*i.e.*, with smaller entropy) than the baseline.

Measurement of Alignment for a Target Sample. To select and relabel the ‘‘well aligned’’ target samples, we need to evaluate whether a target sample is ‘‘well aligned’’ or not. As we know, when a sample is more aligned, it is harder for the domain classifier to identify its domain and has a higher uncertainty w.r.t. to the predicted domain label of this sample. Correspondingly, the entropy of domain classification of this sample is in general higher. Therefore, we measure how well a target sample is aligned with the source domain by simply using the entropy of domain classification,

$$\mathcal{H}(\mathbf{p}) = -p_0 \log p_0 - (1 - p_0) \log(1 - p_0), \quad (5)$$

where $\mathbf{p} = [p_0, 1 - p_0]$, with p_0 and $1 - p_0$ denoting the probability of predicting a sample as target domain and

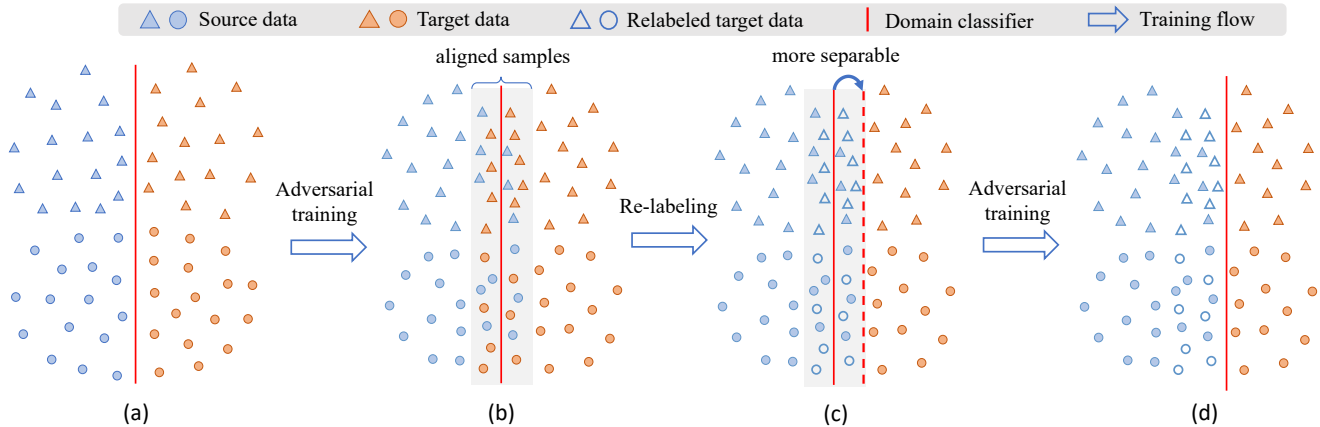


Figure 3: Illustration of our Re-enforceable Adversarial Domain Adaptation (RADA) optimization using toy examples. (a) Initial feature distributions of source and target data sets. (b) As the online adversarial training goes on, the source and target distributions are increasingly aligned. Some samples are already aligned such that the domain classifier cannot distinguish them well while some other samples are still not aligned. The ambiguous samples deteriorate the discrimination capability of the domain classifier. (c) We relabel the well aligned target samples as source domain samples, which makes the less separable distributions more separable. With this strategy, we have a re-enforceable domain classifier, which further re-energizes the domain adversarial optimization in (d).

source domain, respectively. The larger the entropy value, the larger the ambiguity/uncertainty of the discriminator has when identifying which domain it belongs to, and the more aligned the sample is. Note that there may exist other more accurate or advanced measurement metrics, but this is not the focus of this paper and we use this simple one here.

Well aligned Target Sample Selection and Relabeling.

For a given target sample, when it cannot be well distinguished by the domain classifier, it could be considered approximately aligned. In other words, when the entropy of the domain classifier prediction is larger than a threshold τ , we define it as a “well aligned” sample, which is to be relabeled as source domain. τ is a hyper-parameter that controls the strictness of the definition of “well align”. We will study its influence in the experiments in Section 4.

When to start RADA. Intuitively, we could start our relabeling strategy to re-energize the domain discriminator whenever its discrimination capability begins to deteriorate. In general, at the early training stage, the optimization of the domain discriminator persistently improves its discrimination capability as revealed in Figure 2 (a), where the entropy decreases rapidly (in the first 5 epochs). It is thus unnecessary to enable the re-energizing strategy. Inspired by the popular learning rate (lr) adjustment algorithms [39, 3] which adjust the learning rate if no improvement is seen for a ‘patience’ number of epochs (where ‘patience’ is usually set to 2 to 10)¹, we start RADA if no improvement of the discrimination capability is seen for a ‘patience’ number of epochs and we denote this hyper-parameter as K .

Adversarial Training. As is done in previous adversar-

ial domain adaptation methods, we perform mini-batch level optimization where a batch consists of both source and target domain samples. Once RADA is activated, for each mini-batch, we first check whether each target sample should be relabeled as a source sample and relabel them if deemed so. Then the adversarial training is performed under the updated domain labels. The domain classifier is updated by minimizing the domain classification loss (*i.e.*, adversarial loss as in (2)) over the updated source set and target set. Simultaneously, the feature extractor is trained to fool the domain classifier. Note that the domain relabeling has no impact on object classification loss.

Mixup within Updated Source Set. To promote the continuity of feature space in the updated source set, where there may be low density space between the newly absorbed samples and previous source samples, we leverage the mixup technique [59] to softly mix the features between the previous source samples and the newly absorbed samples. In a mini-batch, we randomly select a sample from the original source set and a sample from the target samples that are relabeled as source, where we denote their features as \mathbf{f}_i^s and \mathbf{f}_j^t , respectively. Such mixup generates a new feature $\tilde{\mathbf{f}}^s$ with domain label being source domain:

$$\tilde{\mathbf{f}}^s = \mathcal{M}_\alpha(\mathbf{f}_i^s, \mathbf{f}_j^t) = \alpha \mathbf{f}_i^s + (1 - \alpha) \mathbf{f}_j^t, \quad (6)$$

which is used to train the domain classifier and adversarially train the feature extractor. α is a value randomly selected from a uniform distribution, *i.e.*, $\alpha \sim U(0, 1)$.

¹https://pytorch.org/docs/stable/optim.html?highlight=reduceironplateau#torch.optim.lr_scheduler.ReduceLROnPlateau

4. Experiment

Datasets. 1) **Office-31** [41] is the most widely used dataset for visual domain adaptation (DA), with 4,652 images and 31 categories collected from three distinct domains: Amazon (**A**), Webcam (**W**) and DSLR (**D**). We evaluate all methods on six transfer tasks $A \rightarrow W$, $D \rightarrow W$, $W \rightarrow D$, $A \rightarrow D$, $D \rightarrow A$, and $W \rightarrow A$, respectively. 2) **Office-Home** [48] is a more difficult dataset (with relative large domain discrepancy) than *Office-31*. It consists of 15,500 images of 65 object classes in office and home settings. It has four dissimilar domains: Artistic images (**Ar**), ClipArt (**Cl**), Product images (**Pr**), and Real-World images (**Rw**). Among the four domains, there are a total of 12 DA tasks. 3) **VisDA-2017** [37] is a simulation-to-real dataset for DA with over 280,000 images across 12 categories in the training, validation and testing domains. 4) **DomainNet** [36] is a benchmark for large-scale multi-source domain adaptation, which has six domains (Clipart, Infograph, Painting, Quickdraw, Real and Sketch) and 0.6M images of 345 classes. 5) **Digit-Five** consists of five different digit recognition datasets: MNIST (*mt*) [28], MNIST-M (*mm*) [12], USPS (*up*) [23], SVHN (*sv*) [34], and Synthetic (*syn*) [12].

Following previous works [10, 56, 24], we perform single source to single target adaptation on the first three datasets. We also conduct experiments for multi-source domain adaptation on Digit-Five and DomainNet, following the leave-one-out protocol used in works [33, 36], where one domain is selected as the target domain while the rest of domains are used as source domains.

Implementation Details. *Baseline Setup:* We build our baseline networks based on two representative domain adversarial domain adaptation frameworks: DANN [13] and CDAN [31]. Following [31], we could also use entropy conditioning (E) regularization as proposed in CDAN [31], which reweighs the samples based on their entropy of the object classification predictions to prioritize the easy-to-transfer samples for easy optimization. We denote the schemes with entropy conditioning (E) regularization as CDAN+E, DANN+E. We employ stochastic gradient descent (SGD) as optimizer with an initial learning rate of $1e-3$ and momentum of 0.9 to train all the models.

Our RADA Setup: As a plug-and-play optimization strategy, we apply our RADA on top of the above two representative UDA baselines (DANN and CDAN) for validation. Batch size is set as 36 for Office-31 and VisDA-2017. More details are presented in **Supplementary**. As discussed in Section 3.2, we perform the “well aligned” target sample selection, relabeling, and domain discriminator and feature extractor updating with the updated domain partitions at the mini-batch level. We train our RADA method for 100 epochs on Office-31 and Office-Home, and 150 epochs for the larger VisDA-2017 dataset. We set threshold τ to 0.35 and K to 5 by default, which is investigated in our ablation study.

Table 1: Performance (classification accuracy %) comparisons of our schemes and the corresponding baselines. “w/o MU” denotes without using mixup within the updated source set. All schemes use ResNet-50 as backbone.

Methods	Office-31	Office-Home	VisDA-2017
Baseline (DANN [13])	83.42	60.05	61.23
DANN+RADA w/o MU	85.24	63.14	65.91
DANN+RADA	86.79	64.81	67.29
Baseline (CDAN [31])	87.90	68.11	70.82
CDAN+RADA w/o MU	89.58	70.25	75.62
CDAN+RADA	91.08	71.37	76.28

4.1. Ablation Study

We perform comprehensive ablation studies to demonstrate the effectiveness of our proposed Re-enforceable Adversarial Domain Adaptation strategy. We conduct experiments on the two representative adversarial domain adaptation frameworks *DANN* [13] and *CDAN* [31].

Effectiveness of RADA. Table 1 shows the averaged domain adaptive classification results on three benchmarks. We have the following observations:

- 1) On top of the baseline *DANN*, **our DANN+RADA w/o MU** even *without using mixup* (for the updated source set) brings significant improvement of **1.82%/3.09%/4.68%** in accuracy on Office-31/Office-Home/VisDA-2017, respectively. On top of the stronger baseline *CDAN*, **our CDAN+RADA w/o MU** still brings significant improvements, *i.e.*, **1.68%/2.14%/4.80%** in accuracy on Office-31/Office-Home/VisDA-2017, respectively.
- 2) Mixup within the updated source set promotes the continuity in the updated source set. It brings additional gain of 1.55%/1.67%/1.38% in accuracy on Office-31/Office-Home/VisDA-2017 on top of *DANN+RADA w/o MU*. Note that on VisDA-2017, the majority of gain comes from our relabeling strategy, *i.e.*, 4.68%/4.80% on DANN/CDAN, whereas the mixup specific to our updated source set brings additional gain of 1.38%/0.66%.
- 3) The improvement of our *RADA* over *Baseline* on VisDA-2017 is larger than that on the other two datasets, which reveals that our method is effective on the challenging dataset with large domain gap (synthetic→real).

Table 2: Comparison with re-weighting related methods for multi-source domain adaptation on digits classification.

Methods	Digit-Five					Avg.
	<i>mt</i>	<i>mm</i>	<i>sv</i>	<i>sy</i>	<i>up</i>	
DANN [11]	97.9	70.8	68.5	87.3	93.4	83.6
DANN + E [31]	98.4	73.7	71.5	87.8	94.7	85.2
DANN + inverse-E [31]	98.7	74.4	72.4	88.6	96.4	86.1
DANN+IWAN [60]	98.2	74.2	72.9	88.9	95.8	86.0
DANN+RADA w/o MU	99.0	76.6	88.9	94.7	98.4	91.5
DANN+RADA w/ MU	99.3	78.2	90.0	95.2	98.4	92.2

Comparison with Sample Re-weighting based Methods. We compare RADA with some sample (re)weighting

based schemes, including entropy-based re-weighting (E) [31], IWAN [60]. Entropy-based re-weighting (E) aims to prioritize the easy-to-transfer samples to ease the optimization, sharing a similar idea in curriculum learning. IWAN [60] re-weights the samples based on the outputs of domain discriminator, which assigns small weights to those source samples that are easy to be distinguished by the domain discriminator to exclude the outlier classes in the source domain. *DANN+inverse-E* denotes the scheme that after half of the optimization iterations, we turn the focus (large weight) from easy-to-transfer samples to those hard-to-aligned samples to enable hard-mining, by adjusting entropy-based reweighting strategy from $\omega(\text{ent}(\cdot)) = 1 + e^{-\text{ent}(\cdot)}$ to $\omega(\text{ent}(\cdot)) = 1/(1 + e^{-\text{ent}(\cdot)})$, where $\text{ent}(\cdot)$ denotes the entropy of the predicted object classes.

We conduct experiments under the multi-source setup on Digit-Five. To guarantee fairness of comparison, all competitors use the same network architecture (the feature extractor is composed of three convolutional layers and two FC layers (Cov_3FC_2)), and are built upon the representative domain adversarial training framework DANN [13].

Table 2 shows the results. We can see that all sample re-weighting strategies bring performance gain. Our RADA strategy outperforms the performance of all these strategies. In essence, the re-weighting based methods adjust the levels of importance of samples during optimization (from the perspective of curriculum learning or hard mining). They all use static domain labels in optimization. Giving larger weights to the “well aligned” samples actually provides more aligned distributions to the domain discriminator, which cannot effectively enhance the discrimination capability of the domain classifier. Giving smaller weights to the “well aligned” samples may result in less effective use of them for training the domain discriminator. *In contrast, we allow dynamic domain labels by relabeling those “well aligned” target domain samples as source domain samples. On one hand, this provides more separable distributions (vs. more aligned distributions as in assigning larger weights) to re-energize the domain discriminator. On the other hand, this enables the full use of “well aligned” samples (vs. less effective use as in assigning smaller weights).*

4.2. Design Choices

For clear analysis, we do experiments on our scheme *RADA w/o MU* (i.e., without using mixup for the updated source set) for design choices study on VisDA-2017.

Influence of Threshold τ . As described in Section 3.2, we employ a hyper-parameter τ as the threshold to determine whether a target sample is “well aligned”. We study its influence in Figure 4 (a). We can see that a superior performance is achieved when τ ranges from 0.3 to 0.35 on two schemes. Similar trends are observed on other datasets. We set $\tau = 0.35$ on all datasets. Note that we set $K = 5$ for all




Figure 4: Influence of (a) threshold τ , and (b) K .




Figure 5: Visualization of t-SNE distributions on Digit-Five. Compared with CDAN, the features of different domains are better aligned and the different classes/digits are clearly separated for our CDAN+RADA.

these comparison experiments.

Influence of K . As described in Section 3.2, we start our RADA if no improvement of the discrimination capability of the domain classifier is seen for K number of epochs. Figure 4 (b) shows the influence of K . We observe that superior performance is achieved when K ranges from 4 to 6. Similar trends are observed on other datasets. We set $K = 5$ for all datasets. When K is too small, the judgement of no improvement is not reliable which is sensitive to noise. When K is too large, this optimization strategy cannot be fully utilized.

4.3. Visualization

Visualization of Feature Distributions. In Figure 5, we visualize the distributions of the features using t-SNE [47] in $mm, mt, sv, syn \rightarrow up$ setting on Digit-Five. We compare the feature distribution with the baseline CDAN [31], and observe that the features of different domains are better aligned and the different classes/digits are clearly separated for our scheme CDAN+RADA.

4.4. Comparison with State-of-the-Arts

We compare our scheme CDAN [31]+RADA with state-of-the-art unsupervised single/multi-source domain adaptation methods. We report the results from their original papers if available. For fair comparison, we also report the result of our baseline scheme CDAN run by us. SRDC [45], CAN [27], and Symnets [62] tend to encourage the category/group-level consistency of distributions to alleviate the damage of intrinsic discrimination of data when minimizing the domain discrepancy. GVB [10] and CMSS [56] explore to reduce domain discrepancy in a curriculum manner. Ours is conceptually complementary to

Table 3: Performance (%) comparisons with the state-of-the-art UDA approaches on Office-31. All experiments are conducted based on ResNet-50 pre-trained on ImageNet.

Method	Venue	A→D	A→W	D→W	W→D	D→A	W→A	Avg.
DANN [13]	JMLR'16	79.7±0.4	82.0±0.4	96.9±0.2	99.1±0.1	68.2±0.4	67.4±0.5	82.2
Simnet [38]	CVPR'18	85.3±0.3	88.6±0.5	98.2±0.2	99.7±0.2	73.4±0.8	71.8±0.6	86.2
MCD [42]	CVPR'18	92.2±0.2	88.6±0.2	98.5±0.1	100.0±.0	69.5±0.1	69.7±0.3	86.5
CDAN [31]	NIPS'18	92.9±0.2	93.1±0.1	98.6±0.1	100.0±.0	71.0±0.3	69.3±0.3	87.5
TADA [50]	AAAI'19	91.6±0.3	94.3±0.3	98.7±0.1	99.8±0.2	72.9±0.2	73.0±0.3	88.4
Symnets [62]	CVPR'19	93.9±0.5	90.8±0.1	98.8±0.3	100.0±.0	74.6±0.6	72.5±0.5	88.4
CAN [27]	CVPR'19	95.0±0.3	94.5±0.3	99.1±0.2	99.8±0.2	78.0±0.3	77.0±0.3	90.6
GVB [10]	CVPR'20	95.0±0.4	94.8±0.5	98.7±0.3	100.0±.0	73.4±0.3	73.7±0.4	89.3
CDAN-GVB [10]	CVPR'20	93.7±0.2	94.0±0.2	98.6±0.1	100.0±.0	73.4±0.3	73.0±0.2	88.8
SRDC [45]	CVPR'20	95.8±0.2	95.7±0.2	99.2±0.1	100.0±.0	76.7±0.3	77.1±0.1	90.8
CDAN (Baseline) [31]	NIPS'18	90.8±0.3	94.0±0.5	98.1±0.3	100.0±.0	72.4±0.4	72.1±0.3	87.9
CDAN+RADA	This work	96.1±0.4	96.2±0.4	99.3±0.1	100.0±.0	77.5±0.1	77.4±0.3	91.1

Table 4: Performance (%) comparisons with the state-of-the-art UDA approaches on Office-Home, VisDA-2017, Digit-Five, and DomainNet. For all these approaches, ResNet-50 is taken as backbone for Office-Home, VisDA-2017, Cov_3FC_2 is taken as backbone for Digit-Five, ResNet-101 is taken as backbone for DomainNet. See **Supplementary** for more results.

(a) Comparison on Office-Home.			(b) Comparison on VisDA-2017.			(c) Comparison on Digit-Five.			(d) Comparison on DomainNet.		
Method	Venue	Avg.	Method	Venue	Avg.	Methods	Venue	Avg.	Methods	Venue	Avg.
DANN [13]	JMLR'16	57.6	DAN [30]	ICML'15	61.6	DANN [13]	JMLR'16	83.6	DANN [13]	JMLR'16	32.7
MCD [42]	CVPR'18	64.1	DANN [13]	JMLR'16	57.4	IWAN [57]	ICML'19	86.0	DCTN [54]	CVPR'18	38.3
CDAN [31]	NIPS'18	65.8	GTA [43]	CVPR'18	69.5	MDAN [63]	NIPS'18	86.7	MCD [42]	CVPR'18	38.5
Symnets [62]	CVPR'19	67.2	CDAN [31]	NIPS'18	70.0	MCD [42]	CVPR'18	86.1	MDAN [63]	NIPS'18	42.8
TADA [50]	AAAI'19	67.6	MDD [61]	ICML'19	74.6	DCTN [54]	CVPR'18	88.6	M3SDA [36]	ICCV'19	42.7
BNM [9]	CVPR'20	67.9	CDAN-GVB [10]	CVPR'20	74.9	M3SDA [36]	ICCV'19	87.6	FAR [24]	Arxiv'19	45.5
CDAN-GVB [10]	CVPR'20	69.0	GVB [10]	CVPR'20	75.3	CMSS [56]	ECCV'20	90.8	CMSS [56]	ECCV'20	46.5
GVB [10]	CVPR'20	70.4									
SRDC [45]	CVPR'20	71.3									
CDAN (Baseline)	NIPS'18	68.1	CDAN (Baseline)	NIPS'18	70.8	CDAN (Baseline)	NIPS'18	88.7	CDAN (Baseline)	NIPS'18	45.2
CDAN+RADA	This work	71.4	CDAN+RADA	This work	76.3	CDAN+RADA	This work	93.2	CDAN+RADA	This work	47.5

these methods. They in general ignore the degradation of the discrimination capability of the domain discriminator and its side effect. To address this, we propose RADA which could re-energize the domain discriminator by exploring dynamic domain labels, without making any change to the network architecture of the domain discriminator.

Results on Office-31, Office-Home and VisDA-2017. For the single source to single target adaptation, Table 3, Table 4a and Table 4b show the comparisons on Office-31, Office-Home, and VisDA-2017. We can see that our CDAN+RADA significantly outperforms the baseline CDAN [31] by **3.2%**, **3.3%**, and **5.5%** on three datasets. CDAN+RADA also achieves the best performance.

Results on Digit-Five and DomainNet. We also conduct experiments on the multi-source to single target adaptation settings on Digit-Five and DomainNet. On Table 4c and Table 4d show that our CDAN+RADA outperforms our baseline CDAN by 4.5% and 2.3% on Digit-Five and DomainNet, respectively. Our CDAN+RADA also achieves the state-of-the-art performance. More experimental results, including various sub-settings on different single/multi-source UDA datasets, can be found in **Supplementary**.

5. Conclusion

In this paper, we propose an effective optimization strategy for adversarial domain adaptation, termed as Re-enforceable Adversarial Domain Adaptation (RADA), which aims to re-energize the domain classifier during the training and in turn further drives feature alignment. We achieve this by relabeling the “well aligned” target samples as source domain, which enables the exploration of dynamic domain labels, for re-energizing the domain classifier on the fly. Extensive experiments on multiple benchmarks, including single- and multi-source domain adaptation, and different adversarial domain adaptation networks, have demonstrated the effectiveness and generalization capability of our RADA strategy for adversarial domain adaptation. We hope to inspire more works to incorporate more dynamics to the adversarial domain adaptation for effective training.

6. Acknowledgments

This work was supported in part by NSFC under Grant U1908209, 61632001, 62021001 and the National Key Research and Development Program of China 2018AAA0101400.

References

- [1] Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for domain adaptation. In *NeurIPS*, pages 137–144, 2007. 1
- [2] Karsten M Borgwardt, Arthur Gretton, Malte J Rasch, Hans-Peter Kriegel, Bernhard Schölkopf, and Alex J Smola. Integrating structured biological data by kernel maximum mean discrepancy. *Bioinformatics*, 22(14):e49–e57, 2006. 2
- [3] Léon Bottou. Stochastic gradient descent tricks. In *Neural networks: Tricks of the trade*, pages 421–436. Springer, 2012. 5
- [4] Konstantinos Bousmalis, Nathan Silberman, David Dohan, Dumitru Erhan, and Dilip Krishnan. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *CVPR*, pages 3722–3731, 2017. 3
- [5] Konstantinos Bousmalis, George Trigeorgis, Nathan Silberman, Dilip Krishnan, and Dumitru Erhan. Domain separation networks. In *NeurIPS*, pages 343–351, 2016. 3
- [6] Chaoqi Chen, Weiping Xie, Wenbing Huang, Yu Rong, Xinghao Ding, Yue Huang, Tingyang Xu, and Junzhou Huang. Progressive feature alignment for unsupervised domain adaptation. In *CVPR*, pages 627–636, 2019. 3
- [7] Qingchao Chen, Yang Liu, Zhaowen Wang, Ian Wassell, and Kevin Chetty. Re-weighted adversarial adaptation network for unsupervised domain adaptation. In *CVPR*, pages 7976–7985, 2018. 1, 3
- [8] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation strategies from data. In *CVPR*, pages 113–123, 2019. 3
- [9] Shuhao Cui, Shuhui Wang, Junbao Zhuo, Liang Li, Qingming Huang, and Qi Tian. Towards discriminability and diversity: Batch nuclear-norm maximization under label insufficient situations. In *CVPR*, pages 3941–3950, 2020. 1, 3, 8
- [10] Shuhao Cui, Shuhui Wang, Junbao Zhuo, Chi Su, Qingming Huang, and Qi Tian. Gradually vanishing bridge for adversarial domain adaptation. In *CVPR*, pages 12455–12464, 2020. 1, 3, 4, 6, 7, 8
- [11] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. *ICML*, 2015. 1, 3, 6
- [12] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, pages 1180–1189. PMLR, 2015. 4, 6
- [13] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1):2096–2030, 2016. 1, 2, 3, 4, 6, 7, 8
- [14] Boqing Gong, Kristen Grauman, and Fei Sha. Connecting the dots with landmarks: Discriminatively learning domain-invariant features for unsupervised domain adaptation. In *ICML*, pages 222–230, 2013. 1, 3
- [15] Boqing Gong, Yuan Shi, Fei Sha, and Kristen Grauman. Geodesic flow kernel for unsupervised domain adaptation. In *CVPR*, pages 2066–2073. IEEE, 2012. 1
- [16] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, pages 2672–2680, 2014. 1, 3
- [17] Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(Mar):723–773, 2012. 1, 3
- [18] Hongyu Guo, Yongyi Mao, and Richong Zhang. Mixup as locally linear out-of-manifold regularization. In *AAAI*, volume 33, pages 3714–3722, 2019. 3
- [19] Philip Haeusser, Thomas Frerix, Alexander Mordvintsev, and Daniel Cremers. Associative domain adaptation. In *ICCV*, pages 2765–2773, 2017. 1, 3
- [20] Tong He, Zhi Zhang, Hang Zhang, Zhongyue Zhang, Junyuan Xie, and Mu Li. Bag of tricks for image classification with convolutional neural networks. In *CVPR*, pages 558–567, 2019. 3
- [21] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *ICML*, pages 1989–1998. PMLR, 2018. 3, 4
- [22] Lanqing Hu, Meina Kan, Shiguang Shan, and Xilin Chen. Duplex generative adversarial network for unsupervised domain adaptation. In *CVPR*, pages 1498–1507, 2018. 1, 3
- [23] Jonathan J. Hull. A database for handwritten text recognition research. *TPAMI*, 16(5):550–554, 1994. 6
- [24] Xin Jin, Cuiling Lan, Wenjun Zeng, and Zhibo Chen. Feature alignment and restoration for domain generalization and adaptation. *arXiv preprint arXiv:2006.12009*, 2020. 6, 8
- [25] Xin Jin, Cuiling Lan, Wenjun Zeng, and Zhibo Chen. Global distance-distributions separation for unsupervised person re-identification. In *ECCV*, pages 735–751. Springer, 2020. 1
- [26] Xin Jin, Cuiling Lan, Wenjun Zeng, Zhibo Chen, and Li Zhang. Style normalization and restitution for generalizable person re-identification. In *CVPR*, pages 3143–3152, 2020. 1
- [27] Guoliang Kang, Lu Jiang, Yi Yang, and Alexander G Hauptmann. Contrastive adaptation network for unsupervised domain adaptation. In *CVPR*, pages 4893–4902, 2019. 1, 3, 7, 8
- [28] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 6
- [29] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *NeurIPS*, pages 700–708, 2017. 3
- [30] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *ICML*, pages 97–105, 2015. 2, 3, 8
- [31] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. In *NeurIPS*, pages 1640–1650, 2018. 1, 2, 3, 4, 6, 7, 8
- [32] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Deep transfer learning with joint adaptation networks. In *ICML*, pages 2208–2217, 2017. 3

- [33] Massimiliano Mancini, Lorenzo Porzi, Samuel Rota Bulò, Barbara Caputo, and Elisa Ricci. Boosting domain adaptation by discovering latent domains. In *CVPR*, June 2018. 6
- [34] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bis-sacco, Bo Wu, and Andrew Y. Ng. Reading digits in natural images with unsupervised feature learning. In *NeurIPS-W*, 2011. 6
- [35] Jie Ni, Qiang Qiu, and Rama Chellappa. Subspace interpolation via dictionary learning for unsupervised domain adaptation. In *CVPR*, pages 692–699, 2013. 1
- [36] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *ICCV*, pages 1406–1415, 2019. 6, 8
- [37] Xingchao Peng, Ben Usman, Neela Kaushik, Judy Hoffman, Dequan Wang, and Kate Saenko. Visda: The visual domain adaptation challenge. *arXiv preprint arXiv:1710.06924*, 2017. 6
- [38] Pedro O Pinheiro. Unsupervised domain adaptation with similarity learning. In *CVPR*, pages 8004–8013, 2018. 8
- [39] Aaditya Prakash, James Storer, Dinei Florencio, and Cha Zhang. Repr: Improved training of convolutional filters. In *CVPR*, pages 10666–10675, 2019. 5
- [40] Paolo Russo, Fabio M Carlucci, Tatiana Tommasi, and Barbara Caputo. From source to target and back: symmetric bi-directional adaptive gan. In *CVPR*, pages 8099–8108, 2018. 3
- [41] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *ECCV*, pages 213–226. Springer, 2010. 6
- [42] Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In *CVPR*, pages 3723–3732, 2018. 3, 8
- [43] Swami Sankaranarayanan, Yogesh Balaji, Carlos D Castillo, and Rama Chellappa. Generate to adapt: Aligning domains using generative adversarial networks. In *CVPR*, pages 8503–8512, 2018. 1, 3, 8
- [44] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *ECCV*, pages 443–450, 2016. 3
- [45] Hui Tang, Ke Chen, and Kui Jia. Unsupervised domain adaptation via structurally regularized deep clustering. In *CVPR*, pages 8725–8735, 2020. 7, 8
- [46] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *CVPR*, pages 7167–7176, 2017. 1, 3
- [47] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008. 7
- [48] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *CVPR*, pages 5018–5027, 2017. 6
- [49] Riccardo Volpi, Pietro Morerio, Silvio Savarese, and Vittorio Murino. Adversarial feature augmentation for unsupervised domain adaptation. In *CVPR*, pages 5495–5504, 2018. 1, 3
- [50] Ximei Wang, Liang Li, Weirui Ye, Mingsheng Long, and Jianmin Wang. Transferable attention for domain adaptation. In *AAAI*, volume 33, pages 5345–5352, 2019. 3, 8
- [51] Yuan Wu, Diana Inkpen, and Ahmed El-Roby. Dual mixup regularized learning for adversarial domain adaptation. In *ECCV*, pages 540–555. Springer, 2020. 3
- [52] Shaoan Xie, Zibin Zheng, Liang Chen, and Chuan Chen. Learning semantic representations for unsupervised domain adaptation. In *ICML*, pages 5423–5432. PMLR, 2018. 3
- [53] Minghao Xu, Jian Zhang, Bingbing Ni, Teng Li, Chengjie Wang, Qi Tian, and Wenjun Zhang. Adversarial domain adaptation with domain mixup. *AAAI*, 2020. 3
- [54] Ruijia Xu, Ziliang Chen, Wangmeng Zuo, Junjie Yan, and Liang Lin. Deep cocktail network: Multi-source unsupervised domain adaptation with category shift. In *CVPR*, pages 3964–3973, 2018. 8
- [55] Hongliang Yan, Yukang Ding, Peihua Li, Qilong Wang, Yong Xu, and Wangmeng Zuo. Mind the class weight bias: Weighted maximum mean discrepancy for unsupervised domain adaptation. In *CVPR*, pages 2272–2281, 2017. 2, 3
- [56] Luyu Yang, Yogesh Balaji, Ser-Nam Lim, and Abhinav Shrivastava. Curriculum manager for source selection in multi-source domain adaptation. *ECCV*, 2020. 3, 6, 7, 8
- [57] Kaichao You, Ximei Wang, Mingsheng Long, and Michael Jordan. Towards accurate model selection in deep unsupervised domain adaptation. In *ICML*, pages 7124–7133, 2019. 8
- [58] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *ICCV*, pages 6023–6032, 2019. 3
- [59] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *ICLR*, 2018. 3, 5
- [60] Jing Zhang, Zewei Ding, Wanqing Li, and Philip Ogunbona. Importance weighted adversarial nets for partial domain adaptation. In *CVPR*, pages 8156–8164, 2018. 6, 7
- [61] Yuchen Zhang, Tianle Liu, Mingsheng Long, and Michael I Jordan. Bridging theory and algorithm for domain adaptation. *ICML*, 2019. 8
- [62] Yabin Zhang, Hui Tang, Kui Jia, and Mingkui Tan. Domain-symmetric networks for adversarial domain adaptation. In *CVPR*, pages 5031–5040, 2019. 1, 3, 4, 7, 8
- [63] Han Zhao, Shanghang Zhang, Guanhang Wu, José MF Moura, Joao P Costeira, and Geoffrey J Gordon. Adversarial multiple source domain adaptation. In *NeurIPS*, pages 8559–8570, 2018. 8
- [64] Junbao Zhuo, Shuhui Wang, Shuhao Cui, and Qingming Huang. Unsupervised open domain recognition by semantic discrepancy minimization. In *CVPR*, pages 750–759, 2019. 1