

Cluster Alignment with a Teacher for Unsupervised Domain Adaptation

Zhijie Deng, Yucen Luo, Jun Zhu*

Dept. of Comp. Sci. & Tech., Institute for AI, BNRist Lab, THBI Lab, Tsinghua University

{dzj17, luoyc15}@mails.tsinghua.edu.cn, dcszj@tsinghua.edu.cn

Abstract

Deep learning methods have shown promise in unsupervised domain adaptation, which aims to leverage a labeled source domain to learn a classifier for the unlabeled target domain with a different distribution. However, such methods typically learn a domain-invariant representation space to match the marginal distributions of the source and target domains, while ignoring their fine-level structures. In this paper, we propose Cluster Alignment with a Teacher (CAT) for unsupervised domain adaptation, which can effectively incorporate the discriminative clustering structures in both domains for better adaptation. Technically, CAT leverages an implicit ensembling teacher model to reliably discover the class-conditional structure in the feature space for the unlabeled target domain. Then CAT forces the features of both the source and the target domains to form discriminative class-conditional clusters and aligns the corresponding clusters across domains. Empirical results demonstrate that CAT achieves state-of-the-art results in several unsupervised domain adaptation scenarios.

1. Introduction

Deep learning has achieved remarkable performance in a wide variety of computer vision tasks, such as image recognition [15] and object detection [33]. However, classifiers trained on specific datasets cannot always generalize effectively to new datasets owing to the well-known *domain shift* problem [5, 43]. Enabling models to generalize from a source domain to a target domain is usually referred to as domain adaptation (DA) [1]. In many cases, it is expensive or difficult to collect annotations on the target domain. Learning algorithms attempting to tackle the transferring problem from a fully labeled source domain to an unlabeled target domain is called unsupervised domain adaptation (UDA) [10]. UDA is particularly challenging because the target domain cannot provide explicit information to facilitate the adaptation of classifiers.

*Corresponding author.




Figure 1: (Best viewed in color.) Left: The two domains have diverse modes. Right: The two domains have different class imbalance ratios. Existing methods aligning the marginal distributions while ignoring the class-conditional structures cannot perform well in these cases. However, CAT incorporates the discriminative clustering structures in both domains for better adaptation, thus delivers a more reasonable domain-invariant cluster-structure feature space with enhanced discriminative power. See Fig. 3 and Appendix. A for the learned feature space of real data.

Recently, deep models have been developed with promise in unsupervised domain adaptation to learn expressive features [45, 7, 22, 20, 44, 23, 37, 39, 36]. These deep UDA methods mainly focus on matching the source and target domains via adversarial training [7, 44, 2, 20, 23, 49, 37, 13] or kernelized training [22, 23, 24]. The main hypothesis behind them is that the marginal distributions of the two domains can be aligned in some feature space learned by optimizing a deep network, and thus the classifier trained with source data tends to perform well on the target domain. Theoretical analysis [1] also shows that minimizing the divergence between the marginal distributions in the learned feature space is beneficial to reduce the classifier’s error.

However, these methods are not problemless. In classification, as the classes correspond to different semantics and different characteristics, the marginal distribution of the data naturally has a class-conditional multi-modal structure. Moreover, the modes corresponding to the same class in different domains are not always geometrically similar. Thus, it is not sufficient for existing deep UDA methods to only minimize the discrepancy between the marginal distributions while neglecting their structures, and such methods tend to fail in challenging cases, such as those in Fig. 1.

Properly incorporating this fine-grained class-conditional structure has been shown beneficial in various tasks. For example, Shi and Sha [40] make the *discriminative clustering* assumption which helps to adapt the decision boundaries for the source domain to the target domain discriminatively.¹ However, one limitation of [40] is that it adopts a simple linear transformation to learn the feature space, which cannot effectively extract high-order features from raw data (e.g., images) as the deep UDA methods. Another limitation is that [40] builds a nearest neighbor based prediction model, which outputs the prediction of one sample based on all the source data. Then, the training is not compatible with the stochastic training of deep network and has a high complexity.

In this paper, we present *Cluster Alignment with a Teacher* (CAT), a new deep UDA model that incorporates the class-conditional structures for more effective adaptation. CAT conjoins the complimentary advantages of deep learning methods and discriminative clustering methods for UDA. Technically, there are three learning objectives in CAT. At first, CAT minimizes the supervised classification loss on the labeled source data and builds a teacher classifier, *i.e.* an implicit ensemble of the source classifier, to provide pseudo labels for unlabeled target data. The underlying notion is that the golden classifier trained on source domain can perform well on a majority of target samples because of the similarity between the two domains and the teacher-student paradigm is not sensitive to the false pseudo labels [16]. To exploit the fine-grained class-conditional structures in the feature space and address the aforementioned issues suffered by existing deep UDA methods, CAT also includes two objectives which depend on the pseudo labels provided by the teacher classifier. On one hand, for discriminative learning in both domains, CAT deploys a class-conditional clustering loss to force the features from the same class to concentrate together and the features from different classes to be separated. On the other hand, for the class-conditional alignment between the two domains, CAT aligns the clusters which correspond to the same class but come from different domains via a conditional feature matching loss. The prediction models used in CAT are a student deep network and its implicit ensemble (*i.e.*, the teacher classifier), thus CAT can address the training issues of [40] and also enjoy the more flexible ability of feature learning. Furthermore, it is obvious that CAT is compatible to the marginal distribution alignment methods on the tasks where the source data is similarly distributed as the target data. The former can provide a fine-grained class-conditional alignment of domains and the latter can provide a global alignment of them.

¹Besides UDA tasks, previous work [30] has also shown an interesting exploration of the class-conditional structures for learning deep models that are robust against adversarial attacks.

We evaluate the proposed CAT through extensive experiments on both synthetic and real-world datasets. Empirical results show that CAT presents striking performance across various tasks. In addition, we further combine CAT with the existing deep UDA methods, and CAT can bias them successfully to achieve the discriminative alignment between domains, establishing new state-of-the-art baselines on popular benchmarks. In the combined methods, we also propose a confidence-thresholding technique to filter out low-confidence target samples (which are likely to be mapped into incorrect clusters by the marginal distribution alignment methods) to enhance the stability of the training.

To summarize, our contributions are three-folds:

- We exploit the discriminative class-conditional structures of distributions in deep UDA and propose CAT to achieve better alignment between the source domain and the target domain.
- CAT is compatible and applicable to the existing UDA methods which rely on marginal distribution alignment.
- Empirically, CAT is not sensitive to hyper-parameters and can boost the marginal distribution alignment approaches significantly, achieving new state-of-the-art across various settings.

2. Related work

Unsupervised domain adaptation has drawn increasing interests, and has been developed mainly in two directions: Maximum Mean Discrepancy (MMD) based approaches and adversarial training based approaches. Tzeng *et al.* [45] and Long *et al.* [22] minimize MMD to match the two domains while [23] proposes to align the joint distributions of them using Joint MMD criterion. Since the development of Generative Adversarial Networks (GANs) [9, 4], adversarial training has been applied into domain adaptation and fruitful works emerge. Ganin *et al.* [7] develop the framework of domain adversarial training and plenty of works are proposed to improve it by aligning source domain and target domain better in the feature space [44, 49, 14] or image space [20]. Zhang *et al.* [50] successfully extend RevGrad [7] to consider each domain's characteristics using collaborative games. Image to image translation approaches [8, 2, 13, 34, 28, 39] also play an important role in the advancement of domain adaptation and demonstrate impressive performance, especially on semantic segmentation tasks. In addition, Saito *et al.* [37] propose to align the two domains using decision boundaries of task-specific classifier. Associative DA [11] proposes an associative loss to reduce discrepancy between domains and SimNet [31] proposes to use a similarity-based classifier in UDA. Though the existing methods match the two domains in different ways, most of them ignore the class-conditional information

in the alignment procedure, thus hard to attain the objective of discriminative learning. Conversely, CAT explicitly discovers classes in the feature space via a teacher model and hence constructs a more reasonable matching procedure. Concurrently, several works [51, 48] analyze the conditional distribution shift and label distribution shift issues in UDA theoretically, but only propose limited solutions. In contrast, we provide a more practical and more powerful way to solve them.

Using a teacher model for labeling data is inspired by the impressive consistency-based methods in semi-supervised learning (SSL) [16, 42]. Recent attempts to apply SSL techniques in UDA include [6, 41, 46]. CAT differs from these previous works in that CAT exploits the discriminative class-conditional structures in both the alignment and classification procedures while they focus on improving the classifier for the target domain by implementing the *cluster* assumption [3]. CAT imposes a much stronger regularization and assists in a better alignment.

3. Methodology

In this section, we first introduce the setting and framework of deep UDA and then present the *Cluster Alignment with a Teacher* (CAT). Finally, we discuss about CAT.

3.1. Deep unsupervised domain adaptation

In an UDA task, we are given a set of source samples $\mathcal{X}_s = \{x_s^i\}_{i=1}^N$ with labels $\mathcal{Y}_s = \{y_s^i\}_{i=1}^N$, $y_s^i \in \{1, 2, \dots, K\}$ and a set of unlabeled target samples $\mathcal{X}_t = \{x_t^i\}_{i=1}^M$. Notably, the two sets of samples are drawn from different distributions which lead to the *domain shift* challenge. Therefore, the UDA algorithms should learn to adapt the classifier trained on the source domain to the unlabeled target domain. Deep learning techniques have been introduced into UDA [7, 44, 2, 20, 23, 37, 22, 24] and they demonstrate remarkable performance across tasks. Generally, in these methods, the classifier h (parameterized by θ) is constructed as $h = g \circ f$ where f maps samples into features in the space \mathcal{F} and g outputs the predictions based on the extracted features. The learning includes simultaneously optimizing the classifier h w.r.t. the labeled source data and minimizing the distance between the marginal distributions of the two domains in the feature space \mathcal{F} , resulting in a domain-invariant feature space. Technically, in the source domain, we minimize the supervised loss as:

$$\min_{\theta} \mathcal{L}_y(\mathcal{X}_s, \mathcal{Y}_s) = \frac{1}{N} \sum_{i=1}^N \ell(h(x_s^i; \theta), y_s^i), \quad (1)$$

where ℓ is a pre-defined loss, *e.g.*, cross-entropy loss. Meanwhile, we minimize the discrepancy loss as:

$$\min_{\theta} \mathcal{L}_d(\mathcal{X}_s, \mathcal{X}_t) = \mathcal{D}(f(\mathcal{X}_s, \theta), f(\mathcal{X}_t, \theta)), \quad (2)$$

where \mathcal{D} is a distance and usually correlated with the $\mathcal{H}\Delta\mathcal{H}$ distance in the error bound theory of DA [1]. The theory reveals that the expected error on target samples of any classifier h drawn from a hypothesis set \mathcal{H} has the following bound [1, 49]:

$$\begin{aligned} \epsilon_t(h) &\leq \epsilon_s(h) + \frac{1}{2} d_{\mathcal{H}\Delta\mathcal{H}}(s, t) \\ &\quad + \min_{\hat{h} \in \mathcal{H}} (\epsilon_s(\hat{h}, l_s) + \epsilon_t(\hat{h}, l_t)) \\ &\leq \epsilon_s(h) + \frac{1}{2} d_{\mathcal{H}\Delta\mathcal{H}}(s, t) + \epsilon_t(l_s, l_t) \\ &\quad + \min_{\hat{h} \in \mathcal{H}} (\epsilon_s(\hat{h}, l_s) + \epsilon_t(\hat{h}, l_s)), \end{aligned} \quad (3)$$

where $\epsilon_s(h)$ denotes the expected error on source samples of h , and l_s and l_t represent the labelling functions [1] for the source and target domains, respectively. $\epsilon_t(l_s, l_t)$ denotes the disagreement between the labelling functions in the target domain. Notably, $\min_{\hat{h} \in \mathcal{H}} (\epsilon_s(\hat{h}, l_s) + \epsilon_t(\hat{h}, l_s))$ can be small enough by optimizing \hat{h} w.r.t. the labeled source data. The supervised loss and discrepancy loss focus on minimizing $\epsilon_s(h)$ and $d_{\mathcal{H}\Delta\mathcal{H}}(s, t)$ respectively to obtain small target domain classification error.

However, the methods working in the above framework are not problemless. Theoretically, they ignore minimizing $\epsilon_t(l_s, l_t)$ which may lead to a large upper bound of $\epsilon_t(h)$ and result in unsatisfying target domain performance [49]. Empirically, the data in classification naturally has a class-conditional multi-modal structure, thus aligning marginal distributions while ignoring the fine-level discriminative structures of domains may hurt the target classification performance (Fig. 1-left). Moreover, these methods may fail in more practical and challenging problems, *e.g.* the source and target domains have obviously different class imbalance ratios (Fig. 1-right).

3.2. Cluster Alignment with a Teacher

To overcome these issues, we expect to exploit the fine-level structures in the feature space for discriminative learning and match the class-conditional distributions of source and target domains to reduce of the mismatching between l_s and l_t . Therefore, in the deep UDA scenario, we propose *Cluster Alignment with a Teacher* (CAT), a new deep UDA model for more effective adaptation. Specifically, for the objectives of discriminative learning and class-conditional alignment between domains, we propose a discriminative clustering loss \mathcal{L}_c to force the features of both the source and the target domains to form discriminative clusters, and a cluster-based alignment loss \mathcal{L}_a to align the clusters corresponding to the same class in different domains. Given them, we propose to train CAT by solving the following optimization problem:

$$\min_{\theta} \mathcal{L}_y + \alpha(\mathcal{L}_c + \mathcal{L}_a), \quad (4)$$

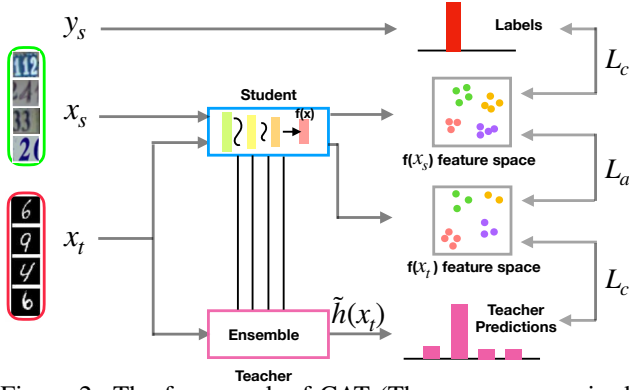


Figure 2: The framework of CAT (The source supervised loss \mathcal{L}_y is omitted for clarity).

where the hyper-parameter α sets a relative trade-off. The whole framework is shown in Fig. 2. We build a teacher classifier, *i.e.* an implicit ensemble of the classifier to be optimized, to provide pseudo labels for the unlabeled target data. These pseudo labels will be used in \mathcal{L}_c and \mathcal{L}_a . We use stochastically sampled mini-batches in the two objectives and the two classifiers make predictions in a forward-propagation way, thus CAT can be trained more efficiently than [40]. Furthermore, \mathcal{L}_c and \mathcal{L}_a optimize the feature space directly and will be more effective than the nearest neighbor based clustering loss in [40]. We elaborate \mathcal{L}_c and \mathcal{L}_a in the following sections.

3.2.1 Discriminative clustering with a teacher

For better classification and alignment, we propose to discover the class-conditional structures in the feature space in both the source and the target domains, and then shape them to be discriminative clusters. In the source domain, the class-conditional structure is obvious because the data is fully labeled. Nevertheless, in the target domain, we cannot obtain the class-conditional structure easily due to the lack of labels. The semantic similarity between the two domains implies that the classifier h trained on the source domain can predict most of target samples correctly. Consequently, using pseudo labels [18] as the annotations for target data and conducting self-training is a direct approach, but it suffers from the error amplification issue which can be detrimental to the learning procedure. To discover the class-conditional structure of the target features in a reliable way, we introduce a teacher classifier \tilde{h} defined as an implicit ensemble of the previous student classifier h [16] to provide pseudo labels for target data.

Based on the pseudo labels given by the teacher, we can explicitly force the target class-conditional structure to be more discriminative using a clustering loss. For the source domain, a similar one can be applied. Formally, resembling the effective SNTG loss in [25], we employ the following discriminative clustering loss (we omit the dependence of f

on θ for clarity, unless stated otherwise):

$$\mathcal{L}_c(\mathcal{X}_s, \mathcal{X}_t) = \mathcal{L}_c(\mathcal{X}_s) + \mathcal{L}_c(\mathcal{X}_t), \quad (5)$$

$$\mathcal{L}_c(\mathcal{X}) = \frac{1}{|\mathcal{X}|^2} \sum_{i=1}^{|\mathcal{X}|} \sum_{j=1}^{|\mathcal{X}|} [\delta_{ij} d(f(x^i), f(x^j)) + (1 - \delta_{ij}) \max(0, m - d(f(x^i), f(x^j)))] , \quad (6)$$

where d is the distance (*e.g.*, squared Euclidean distance) between two features, m is a pre-defined margin, and δ_{ij} is an indicator function which outputs 1 only if x^i and x^j have the same ground truth label (source domain) or teacher-annotated label (target domain). \mathcal{L}_c encourages the features from the same class to concentrate together and pushes the features from different classes far away from each other with a distance m at least. This loss modifies the structures in the representation space gradually, and consequently, it demonstrates a class-conditional cluster structure (as shown in Sec. 4.4). Note that minimizing \mathcal{L}_c is consistent with the *cluster* assumption [25] of classifier and benefits the performance of classification.

It's a common doubt whether the incorrect predictions of the teacher classifier would destroy the training dynamics. However, previous works on semi-supervised learning [16, 42] have validated that this kind of training always leads to good convergence and demonstrates robustness against incorrect labels. Intuitively, the teacher instructs the training of one instance through a bundle of others' predictions which alleviates the negative influence of incorrect predictions notably and aids the student to give better predictions.

3.2.2 Cluster alignment via conditional feature matching

Once the feature space presents discriminative cluster structure, the classifier is expected to make more accurate predictions. However, the label predictor g trained on source domain features may fail due to the geometrical mismatching between the clusters which correspond to the same class in different domains. This kind of mismatching is brought by the individual characteristics of each domain. As a result, it is necessary to impose a class-conditional alignment of two domains to learn better domain-invariant features and adjust the target feature space to be more suitable for classification. Naturally, we expect to minimize the divergence between the corresponding clusters in source domain and the teacher-annotated target domain:

$$\min_{\theta} \mathcal{D}(\mathcal{F}_{s,k} || \mathcal{F}_{t,k}), \quad (7)$$

where $\mathcal{F}_{s,k}$ ($\mathcal{F}_{t,k}$) denotes the set consisting of all the features belonging to class k of domain s (domain t). Extensive previous works [9, 19] have been proposed to minimize

the distance between two sets of samples but we expect to achieve this in a more simple and efficient way by exploiting the separable and tight clusters in the feature space. Drawing inspiration from feature matching GANs [38] which optimizes the distance between the first-order statistics of distributions and demonstrates striking results on SSL tasks, we choose to extend it to work in a conditional way. Formally, we introduce the following cluster alignment loss:

$$\mathcal{L}_a(\mathcal{X}_s, \mathcal{X}_t) = \frac{1}{K} \sum_{k=1}^K \|\lambda_{s,k} - \lambda_{t,k}\|_2^2, \quad (8)$$

where $\lambda_{s,k}$ and $\lambda_{t,k}$ are calculated by

$$\lambda_{s,k} = \frac{1}{|\mathcal{X}_{s,k}|} \sum_{x_s^i \in \mathcal{X}_{s,k}} f(x_s^i), \quad \lambda_{t,k} = \frac{1}{|\mathcal{X}_{t,k}|} \sum_{x_t^i \in \mathcal{X}_{t,k}} f(x_t^i) \quad (9)$$

where $\mathcal{X}_{s,k}$ is the subset of \mathcal{X}_s containing all the source samples whose ground-truth labels are k and $\mathcal{X}_{t,k}$ is the subset of \mathcal{X}_t including all the target samples annotated as class k by the teacher classifier h . This loss is slightly different from the original feature matching loss: it matches the statistics of the representation space \mathcal{F} which totally determine the predictions instead of those produced by an extra critic network. Arguably, the objective has a local optima where class-conditional distributions are matched thoroughly. The cluster alignment loss and the discriminative clustering loss work together to align the class-conditional structures of the two domains in a discriminative way.

3.2.3 Improved marginal distribution alignment

In fact, the source domain and target domain in the existing popular UDA tasks (e.g., digits adaptation and *Office-31*) have analogous marginal distributions. Therefore, in these experiments, we combine CAT with the marginal distribution alignment methods, and CAT contributes to bias them to match the cluster-based marginal distributions. The negative effects of these methods of ignoring the discriminability may hurt the stability of training and the capability of converged models. For example, several target circle samples in Fig. 1-left will be misclassified by them.

We are dedicated to delivering a technique to improve these models given the observation that in the early stages of training, a portion of target samples lie around the decision boundaries of the adapted classifier, i.e., they have low classification confidence (the largest output probability) and are likely to be misclassified. Therefore, these samples are possible to be mapped into the incorrect clusters in the marginal alignment process and the training falls into local optima. To solve this, we propose to use confidence-thresholding method to hold out uncertain data points with confidence less than p , while aligning the confident instances which

are more geometrically close to the source domain with the source data. Formally, we instantiate this technique in the typical and brief RevGrad [7] and propose robust RevGrad (rRevGrad) which optimizes the following loss:

$$\min_{\theta} \max_{\phi} \mathcal{L}_d(\mathcal{X}_s, \mathcal{X}_t) = \frac{1}{N} \sum_{i=1}^N [\log c(f(x_s^i, \theta); \phi)] + \frac{1}{\tilde{M}} \sum_{i=1}^{\tilde{M}} [\log (1 - c(f(x_t^i, \theta); \phi)) \gamma_i], \quad (10)$$

where c is the critic model parameterized by ϕ and γ_i is an indicator function which outputs 1 only if teacher's confidence of x_t^i is greater than p . With the divergence between the two domains decreasing, more and more target samples are selected into the domain adversarial training. Gradually, almost all the target samples are included in the training which avoids the lost of target information. We empirically observed that rRevGrad improves the classification performance on the target domain and enhances the stability of training (see Sec.4).

3.3. Discussion

Comparison with SSL based deep UDA methods [41, 36, 6]. CAT not only implements of *cluster* assumption for better classification but also imposes a class-conditional alignment between domains which is more principal in UDA. However, [41, 36, 6] focus on improving classifier to make it more consistent and robust for the target domain based on *cluster* assumption. Thus CAT is compatible to these methods (see Sec. 4).

Comparison with MSTN [49]. The cluster alignment loss using conditional feature matching technique is similar to the semantic loss in MSTN [49]. However, in MSTN, minimizing distance between centers is necessary but not sufficient to achieve semantic alignment. In CAT, we regularize the features to form separable and tight clusters, so the feature matching based loss can match the clusters naturally. They are also different in implementation.

Mini-batch stochastic training of CAT. We implement the two objectives in CAT using stochastically sampled mini-batches as \mathcal{X}_s and \mathcal{X}_t . Specifically, \mathcal{L}_c is an instance-wise loss and can work well. The class-conditional expectation $\lambda_{s,k}$ or $\lambda_{t,k}$ in \mathcal{L}_c could be none when there is no points belonging to class k . At this time, we remove the term corresponding to class k in Eq. 8 and calculate the mean of the other terms. We empirically find CAT needs only $0.05 \times$ more training time on a GTX 1080Ti when combining with existing methods.

Teacher-student paradigm. First, using teacher as labeling function on target domain avoids the error amplification issue. Furthermore, once the classifier becomes more accurate on the target domain, the teacher classifier per-

forms better as well. Then the pseudo labels used in \mathcal{L}_c and \mathcal{L}_a are more likely to be correct which in turn enhances the classifier. Consequently, a boosting cycle between them is formed.

4. Experiments

To demonstrate the effectiveness of CAT, we evaluate it through various experiments on synthetic imbalanced dataset and three challenging UDA tasks: *SVHN-MNIST-USPS*, *Office-31* [35] and *ImageCLEF-DA*². We show that CAT considers and exploits the fine-level class-conditional structures of the source and target domains, and makes the learned feature space discriminative and aligned, thus yielding improved performance on the target domain.

Datasets and configurations. *SVHN-MNIST-USPS* is a challenging adaptation task of digits between three datasets *SVHN* [29], *MNIST* [17] and *USPS*. We conduct experiments in three directions: *SVHN*→*MNIST*, *MNIST*→*USPS* and *USPS*→*MNIST*. We follow the protocol in [44]: we use the whole training sets for the adaptation from *SVHN* to *MNIST* and randomly sample 2000 images from *MNIST* and 1800 images from *USPS* for the adaptation between the two datasets. Following MSTN [49], the images are cast to $28 \times 28 \times 1$ when using LeNet [17] as classifier. When combining with MCD [37] and VADA [41], We take the identical settings as the original methods.

Imbalanced SVHN-MNIST-USPS. We randomly sample 1000 instances from class 0 and 100 instances from class 1 from the original source domain and construct a new one. Then we sample 100 instances from class 0 and 1000 instances from class 1 from the target domain to form a new target. Therefore, the synthetic adaptation dataset contains several imbalanced two-class adaptation tasks. The experiment settings are the same as those of *SVHN-MNIST-USPS*.

Office-31 and *ImageCLEF-DA* are two real-world datasets which are widely used in domain adaptation research. *Office-31* is composed of three domains: Amazon (A), DSLR (D) and Webcam (W), containing 2817, 498 and 795 images from 31 categories, respectively. *ImageCLEF-DA* includes three domains: Caltech-256 (C), ImageNet ILSVRC 2012 (I) and Pascal VOC 2012 (P), containing 600 images from 12 classes, respectively. We use data augmentation such as random flipping and cropping in training for fair comparison with the baselines.

Implementation. In synthetic and digits experiments using LeNet, we set $m = 30$ according to the performance of CAT on the synthetic dataset and we forward-propagate a target sample twice under different perturbations (*i.e.*, dropout) and use the latter as the prediction of the teacher for its simplicity (similar with the II model of [16]). In all the other experiments, we fix $m = 3$ and deploy a tem-

Method	<i>SVHN to MNIST</i>	<i>MNIST to USPS</i>	<i>USPS to MNIST</i>
RevGrad [7]	27.4 ± 6.3	26.7 ± 2.0	17.9 ± 1.4
MSTN [49]	25.8 ± 3.6	30.3 ± 1.0	29.4 ± 0.5
CAT	100.0 ± 0.05	100.0 ± 0.0	99.9 ± 0.2

Table 1: Summary of domain adaptation results on the imbalanced digits datasets in terms of test accuracy (%).

poral ensemble [16] of previous predictions of h as teacher (the accumulation decay constant is set to 0.6). We design a ramp-up function similar with that of [16] to update α in experiments using LeNet and set $\alpha = \frac{2}{1+\exp(-10t)} - 1$ suggested by RevGrad [7] in which t increases linearly from 0 to 1 in the others. We set $p = 0.9$ in all the experiments without tuning. Refer to Appendix. E for more details of the used architectures and optimization settings.

4.1. Experiments on imbalanced SVHN-MNIST-USPS

We first test CAT on the *imbalanced SVHN-MNIST-USPS* dataset, a challenging task where the source domains have 10 : 1 ratio of class imbalance while the target domains have 1 : 10. We implement CAT and RevGrad [7] based on the official codes of MSTN [49] using LeNet [17]. The results are shown in Table 1. We repeat each task 3 times and report the averaged test accuracy and standard deviation.

It is notable that RevGrad [7] and MSTN [49] fail thoroughly owing to their obsession of matching the marginal distributions. In contrast, CAT gives almost completely correct predictions for the target domains. This experiment verifies that existing methods through aligning marginal distributions are restrictive and require the modes corresponding to the same class but different domains to be geometrically similar. They are sensitive and fragile in practical tasks.

4.2. SVHN-MNIST-USPS digits datasets

We apply CAT to the popular digits adaptation task *SVHN-MNIST-USPS* and compare to the state-of-the-art approaches in Table 3 (all baseline results are taken from related literature). CAT, RevGrad+CAT and rRevGrad+CAT follow the settings of MSTN [49] using the LeNet [17]. We implement MCD+CAT and VADA+CAT based on the official codes of MCD [37] and VADA [41] using their architectures instead of LeNet for fair comparison. We only integrate CAT with the first stage algorithm VADA in DIRT-T [41] while discarding the fine-tuning stage for simpleness.

There are several conclusions we can make. First, CAT reveals strikingly improved test accuracy on *SVHN* to *MNIST* task without tuning the hyper-parameters, and CAT even outperforms MCD [37] and VADA [41] which use much wider and deeper neural networks thanks to the class-conditional discriminative alignment between the source and target domains. This task is the most challenging one among the three because of the complex samples and the

²Source coda is at <https://github.com/thudzj/CAT>.

Method	A to W	D to W	W to D	A to D	D to A	W to A	Avg
ResNet-50 [12]	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
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
RevGrad [7]	82.0 \pm 0.4	96.9 \pm 0.2	99.1 \pm 0.1	79.4 \pm 0.4	68.2 \pm 0.4	67.4 \pm 0.5	82.2
JAN [23]	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
SimNet [32]	88.6 \pm 0.5	98.2 \pm 0.2	99.7 \pm 0.2	85.3 \pm 0.3	73.4 \pm 0.8	71.8 \pm 0.6	86.2
GenToAdapt [39]	89.5 \pm 0.5	97.9 \pm 0.3	99.8 \pm 0.4	87.7 \pm 0.5	72.8 \pm 0.3	71.4 \pm 0.4	86.5
CAT	91.1 \pm 0.2	98.6 \pm 0.6	99.6 \pm 0.1	90.6 \pm 1.0	70.4 \pm 0.7	66.5 \pm 0.4	86.1
JAN+CAT	94.0 \pm 0.4	96.6 \pm 0.6	100.0 \pm 0.0	88.1 \pm 1.0	68.9 \pm 0.7	69.4 \pm 0.5	86.2
rRevGrad+CAT	94.4 \pm 0.1	98.0 \pm 0.2	100.0 \pm 0.0	90.8 \pm 1.8	72.2 \pm 0.6	70.2 \pm 0.1	87.6
AlexNet [15]	61.6 \pm 0.5	95.4 \pm 0.3	99.0 \pm 0.2	63.8 \pm 0.5	51.1 \pm 0.6	49.8 \pm 0.4	70.1
DDC [45]	61.8 \pm 0.4	95.0 \pm 0.5	98.5 \pm 0.4	64.4 \pm 0.3	52.1 \pm 0.6	52.2 \pm 0.4	70.6
DRCN [8]	68.7 \pm 0.3	96.4 \pm 0.3	99.0 \pm 0.2	66.8 \pm 0.5	56.0 \pm 0.5	54.9 \pm 0.5	73.6
RevGrad [7]	73.0 \pm 0.5	96.4 \pm 0.3	99.2 \pm 0.3	72.3 \pm 0.3	53.4 \pm 0.4	51.2 \pm 0.5	74.3
JAN [23]	74.9 \pm 0.3	96.6 \pm 0.2	99.5 \pm 0.2	71.8 \pm 0.2	58.3 \pm 0.3	55.0 \pm 0.4	76.0
MSTN [49]	80.5 \pm 0.4	96.9 \pm 0.1	99.9 \pm 0.1	74.5 \pm 0.4	62.5 \pm 0.4	60.0 \pm 0.6	79.1
CAT	77.4 \pm 0.1	97.4 \pm 0.1	99.9 \pm 0.1	74.7 \pm 0.1	63.4 \pm 0.2	60.8 \pm 0.6	78.9
JAN+CAT	78.4 \pm 0.5	97.2 \pm 0.2	100.0 \pm 0.0	74.5 \pm 0.5	63.6 \pm 0.4	61.2 \pm 0.6	79.2
rRevGrad+CAT	80.7 \pm 1.6	97.6 \pm 0.1	100.0 \pm 0.0	76.4 \pm 0.6	63.7 \pm 0.5	62.2 \pm 0.4	80.1

Table 2: Accuracy on the *Office-31* datasets in terms of test accuracy (%) (ResNet-50 and AlexNet).

Method	SVHN to MNIST	MNIST to USPS	USPS to MNIST
Source Only	60.1 \pm 1.1	75.2 \pm 1.6	57.1 \pm 1.7
DDC [45]	68.1 \pm 0.3	79.1 \pm 0.5	66.5 \pm 3.3
CoGAN [20]	-	91.2 \pm 0.8	89.1 \pm 0.8
DRCN [8]	82.0 \pm 0.1	91.8 \pm 0.09	73.7 \pm 0.04
ADDA [44]	76.0 \pm 1.8	89.4 \pm 0.2	90.1 \pm 0.8
LEL [26]	81.0 \pm 0.3	-	-
AssocDA [11]	97.6	-	-
MSTN [49]	91.7 \pm 1.5	92.9 \pm 1.1	-
CAT	98.1 \pm 1.3	90.6 \pm 2.3	80.9 \pm 3.1
RevGrad [7]	73.9	77.1 \pm 1.8	73.0 \pm 2.0
RevGrad+CAT	98.0 \pm 0.8	93.7 \pm 1.1	95.7 \pm 1.3
rRevGrad+CAT	98.8 \pm 0.02	94.0 \pm 0.7	96.0 \pm 0.9
MCD [37]	96.2 \pm 0.4	94.2 \pm 0.7	94.1 \pm 0.3
MCD+CAT	97.1 \pm 0.2	96.3 \pm 0.5	95.2 \pm 0.4
VADA [41]	94.5	-	-
VADA+CAT	95.2	-	-

Table 3: Summary of domain adaptation results on the digits datasets in terms of test accuracy (%).

internal class imbalance of *SVHN*. Second, CAT does not perform well enough on the other two tasks but when combining with rRevGrad and MCD [37], CAT outperforms the strong baselines MSTN [49] and MCD [37] with obvious margins. Third, applying CAT into RevGrad [7], MCD [37] and VADA [41] can enhance the base methods significantly, especially on the typical and simple RevGrad [7]. Finally, rRevGrad+CAT displays higher test accuracy and lower variance than those of RevGrad+CAT and the advantage is

particularly obvious when the two domains have different class-conditional structures (e.g., *SVHN* to *MNIST*), so we utilize rRevGrad+CAT on more challenging tasks.

4.3. Experiments on *Office-31* and *ImageCLEF-DA*

We evaluate CAT using two sets of extensive experiments on the widely used *Office-31* and *ImageCLEF-DA*. They contain more realistic and high-dimensional images, providing a good complement to the digits adaptation task. The results are provided in Table 2 and Table 4, respectively. We integrate CAT with rRevGrad and JAN [23] (using ResNet-50 [12] and AlexNet [15] as the classifiers), which is sufficient to testify the effectiveness of discriminative cluster-based alignment and teacher-student paradigm.

We observe that CAT can boost rRevGrad and JAN [23] significantly, especially on the difficult A to W, A to D and D to A tasks in *Office-31*, and the combined models surpass the strong baselines RevGrad [7] and JAN [23] by obvious margins. CAT based methods also outperform MSTN [49] on various tasks substantially which proves the class-conditional discriminative alignment is superior to the semantic alignment used by MSTN [49]. The improvement of test accuracy on most tasks of *ImageCLEF-DA* shows that CAT can still work well when the domains are small containing only 600 images. We further confirm that CAT can deliver a discriminative and aligned feature space by visualizing the learned features in the Appendix. A.

4.4. Analysis

Visualization of feature space. We visualize the features of the two domains learned by the powerful

Method	I to P	P to I	I to C	C to I	C to P	P to C	Avg
ResNet-50 [12]	74.8 \pm 0.3	83.9 \pm 0.1	91.5 \pm 0.3	78.0 \pm 0.2	65.5 \pm 0.3	91.2 \pm 0.3	80.7
DAN [22]	74.5 \pm 0.4	82.2 \pm 0.2	92.8 \pm 0.2	86.3 \pm 0.4	69.2 \pm 0.4	89.8 \pm 0.4	82.5
RevGrad [7]	75.0 \pm 0.6	86.0 \pm 0.3	96.2 \pm 0.4	87.0 \pm 0.5	74.3 \pm 0.5	91.5 \pm 0.6	85.0
JAN [23]	76.8 \pm 0.4	88.0 \pm 0.2	94.7 \pm 0.2	89.5 \pm 0.3	74.2 \pm 0.3	91.7 \pm 0.3	85.8
CAT	76.7 \pm 0.2	89.0 \pm 0.7	94.5 \pm 0.4	89.8 \pm 0.3	74.0 \pm 0.2	93.7 \pm 1.0	86.3
JAN+CAT	76.3 \pm 0.8	89.2 \pm 0.8	95.3 \pm 0.7	89.3 \pm 0.3	75.9 \pm 1.1	92.2 \pm 1.3	86.4
rRevGrad+CAT	77.2 \pm 0.2	91.0 \pm 0.3	95.5 \pm 0.3	91.3 \pm 0.3	75.3 \pm 0.6	93.6 \pm 0.5	87.3
AlexNet [15]	66.2 \pm 0.2	70.0 \pm 0.2	84.3 \pm 0.2	71.3 \pm 0.4	59.3 \pm 0.5	84.5 \pm 0.3	73.9
RTN [24]	67.4 \pm 0.3	81.3 \pm 0.3	89.5 \pm 0.4	78.0 \pm 0.2	62.0 \pm 0.2	89.1 \pm 0.1	77.9
RevGrad [7]	66.5 \pm 0.5	81.8 \pm 0.4	89.0 \pm 0.5	79.8 \pm 0.5	63.5 \pm 0.4	88.7 \pm 0.4	78.2
JAN [23]	67.2 \pm 0.5	82.8 \pm 0.4	91.3 \pm 0.5	80.0 \pm 0.5	63.5 \pm 0.4	91.0 \pm 0.4	79.3
MSTN [49]	67.3 \pm 0.3	82.8 \pm 0.2	91.5 \pm 0.1	81.7 \pm 0.3	65.3 \pm 0.2	91.2 \pm 0.2	80.0
CAT	68.3 \pm 0.5	83.6 \pm 0.7	91.3 \pm 0.3	79.1 \pm 0.5	64.0 \pm 0.7	90.9 \pm 0.3	79.5
JAN+CAT	67.8 \pm 0.2	83.7 \pm 0.4	92.3 \pm 0.6	80.8 \pm 0.3	65.8 \pm 0.8	91.1 \pm 0.2	80.3
rRevGrad+CAT	68.6 \pm 0.1	84.6 \pm 0.5	91.9 \pm 0.4	80.8 \pm 0.3	65.6 \pm 0.6	92.5 \pm 0.2	80.7

Table 4: Accuracy on the *ImageCLEF-DA* datasets in terms of test accuracy (%) (ResNet-50 and AlexNet).




Figure 3: (Best viewed in color.) (a) Feature space learned by RevGrad. (b) Feature space learned by rRevGrad+CAT. The features are projected to 2-D using t-SNE. Blue violet denotes the source domain and the other colors denotes classes of target domain. See Appendix. A for more results.




Figure 4: Summary of clustering accuracy(%).

rRevGrad+CAT and RevGrad [7] on the *SVHN* to *MNIST* task using t-SNE [27]. The results are shown in Fig. 3. As expected, using CAT (Fig. 3b), the features are concentrated and form tight clusters and those from different classes are separated. In contrast, the features learned by RevGrad [7] (Fig. 3a) are more overlapping and less discriminative.

Clustering in the feature space. We further examine

the feature space shaped by CAT and other baselines by conducting K-means [21] clustering using the aligned features. We utilize the trained models to infer the hidden features of all the images from two domains. Then the features are clustered into k components by K-means in scikit-learn [47]. We set k as the number of categories. We greedily set the label of a cluster as the most frequent label in it to calculate clustering accuracy as shown in Fig. 4. As expected, the feature spaces learned by rRevGrad+CAT demonstrate more discriminative cluster structure and this is consistent with the classification results. Appendix. B, C and D provide more analyses.

5. Conclusion

In this paper we address the challenges of making better alignment between domains and advocate to exploit the discriminative class-conditional structures for effective adaptation in deep UDA. We propose *Cluster Alignment with a Teacher* (CAT) to achieve the objectives of discriminative learning and class-conditional alignment via a discriminative clustering loss and a cluster-based alignment loss. CAT produces a domain-invariant feature space with improved discriminative power and enhances the performance significantly. CAT establishes new state-of-the-art baselines on benchmarks and additional analyses testify its effectiveness.

Acknowledgements

This work was supported by the NSFC Projects (Nos. 61621136008, 61620106010), Beijing NSF Project (No. L172037), Tiangong Institute for Intelligent Computing, Beijing Academy of Artificial Intelligence (BAAI), the NVIDIA NVAIL Program with GPU/DGX Acceleration and the JP Morgan Faculty Research Program.

References

- [1] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79(1-2):151–175, 2010.
- [2] Konstantinos Bousmalis, Nathan Silberman, David Dohan, Dumitru Erhan, and Dilip 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.
- [3] Olivier Chapelle, Bernhard Scholkopf, and Alexander Zien. Semi-supervised learning (chapelle, o. et al., eds.; 2006)[book reviews]. *IEEE Transactions on Neural Networks*, 20(3):542–542, 2009.
- [4] Zhijie Deng, Hao Zhang, Xiaodan Liang, Luona Yang, Shizhen Xu, Jun Zhu, and Eric P Xing. Structured generative adversarial networks. In *Advances in Neural Information Processing Systems*, pages 3899–3909, 2017.
- [5] Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. In *International conference on machine learning*, pages 647–655, 2014.
- [6] Geoffrey French, Michal Mackiewicz, and Mark Fisher. Self-ensembling for visual domain adaptation. *ICLR*, 2017.
- [7] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning-Volume 37*, pages 1180–1189. JMLR.org, 2015.
- [8] Muhammad Ghifary, W Bastiaan Kleijn, Mengjie Zhang, David Balduzzi, and Wen Li. Deep reconstruction-classification networks for unsupervised domain adaptation. In *European Conference on Computer Vision*, pages 597–613. Springer, 2016.
- [9] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [10] Raghuraman Gopalan, Ruonan Li, and Rama Chellappa. Domain adaptation for object recognition: An unsupervised approach. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 999–1006. IEEE, 2011.
- [11] Philip Haeusser, Thomas Frerix, Alexander Mordvintsev, and Daniel Cremers. Associative domain adaptation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2765–2773, 2017.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [13] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei A Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. *arXiv preprint arXiv:1711.03213*, 2017.
- [14] Weixiang Hong, Zhenzhen Wang, Ming Yang, and Junsong Yuan. Conditional generative adversarial network for structured domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1335–1344, 2018.
- [15] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [16] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242*, 2016.
- [17] 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.
- [18] Dong-Hyun Lee. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on Challenges in Representation Learning, ICML*, volume 3, page 2, 2013.
- [19] Yujia Li, Kevin Swersky, and Rich Zemel. Generative moment matching networks. In *International Conference on Machine Learning*, pages 1718–1727, 2015.
- [20] Ming-Yu Liu and Oncl Tuzel. Coupled generative adversarial networks. In *Advances in neural information processing systems*, pages 469–477, 2016.
- [21] Stuart Lloyd. Least squares quantization in pcm. *IEEE transactions on information theory*, 28(2):129–137, 1982.
- [22] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael I Jordan. Learning transferable features with deep adaptation networks. *arXiv preprint arXiv:1502.02791*, 2015.
- [23] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Deep transfer learning with joint adaptation networks. *arXiv preprint arXiv:1605.06636*, 2016.
- [24] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Unsupervised domain adaptation with residual transfer networks. In *Advances in Neural Information Processing Systems*, pages 136–144, 2016.
- [25] Yucen Luo, Jun Zhu, Mengxi Li, Yong Ren, and Bo Zhang. Smooth neighbors on teacher graphs for semi-supervised learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [26] Zelun Luo, Yuliang Zou, Judy Hoffman, and Li F Fei-Fei. Label efficient learning of transferable representations across domains and tasks. In *Advances in Neural Information Processing Systems*, pages 165–177, 2017.
- [27] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.
- [28] Zak Murez, Soheil Kolouri, David Kriegman, Ravi Ramamoorthi, and Kyungnam Kim. Image to image translation for domain adaptation. *arXiv preprint arXiv:1712.00479*, 13, 2017.
- [29] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bisacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. In *NIPS workshop on deep learning and unsupervised feature learning*, volume 2011, page 5, 2011.

- [30] Tianyu Pang, Chao Du, and Jun Zhu. Max-mahalanobis linear discriminant analysis networks. In *International Conference on Machine Learning*, pages 4013–4022, 2018.
- [31] Pedro O Pinheiro. Unsupervised domain adaptation with similarity learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8004–8013, 2018.
- [32] Pedro O. Pinheiro. Unsupervised domain adaptation with similarity learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [33] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [34] Paolo Russo, Fabio Maria Carlucci, Tatiana Tommasi, and Barbara Caputo. From source to target and back: symmetric bi-directional adaptive gan. *arXiv preprint arXiv:1705.08824*, 3, 2017.
- [35] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *European conference on computer vision*, pages 213–226. Springer, 2010.
- [36] Kuniaki Saito, Yoshitaka Ushiku, and Tatsuya Harada. Asymmetric tri-training for unsupervised domain adaptation. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2988–2997. JMLR. org, 2017.
- [37] Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [38] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in neural information processing systems*, pages 2234–2242, 2016.
- [39] Swami Sankaranarayanan, Yogesh Balaji, Carlos D. Castillo, and Rama Chellappa. Generate to adapt: Aligning domains using generative adversarial networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [40] Yuan Shi and Fei Sha. Information-theoretical learning of discriminative clusters for unsupervised domain adaptation. *arXiv preprint arXiv:1206.6438*, 2012.
- [41] Rui Shu, Hung H Bui, Hirokazu Narui, and Stefano Ermon. A dirt-t approach to unsupervised domain adaptation. *ICLR*, 2018.
- [42] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Advances in neural information processing systems*, pages 1195–1204, 2017.
- [43] Antonio Torralba and Alexei A Efros. Unbiased look at dataset bias. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 1521–1528. IEEE, 2011.
- [44] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *Computer Vision and Pattern Recognition (CVPR)*, volume 1, page 4, 2017.
- [45] Eric Tzeng, Judy Hoffman, Ning Zhang, Kate Saenko, and Trevor Darrell. Deep domain confusion: Maximizing for domain invariance. *arXiv preprint arXiv:1412.3474*, 2014.
- [46] Riccardo Volpi, Pietro Morerio, Silvio Savarese, and Vittorio Murino. Adversarial feature augmentation for unsupervised domain adaptation. *arXiv preprint arXiv:1711.08561*, 2017.
- [47] Ian H Witten, Eibe Frank, Mark A Hall, and Christopher J Pal. *Data Mining: Practical machine learning tools and techniques*. Morgan Kaufmann, 2016.
- [48] Yifan Wu, Ezra Winston, Divyansh Kaushik, and Zachary Lipton. Domain adaptation with asymmetrically-relaxed distribution alignment. In *International Conference on Machine Learning*, pages 6872–6881, 2019.
- [49] Shaoan Xie, Zibin Zheng, Liang Chen, and Chuan Chen. Learning semantic representations for unsupervised domain adaptation. In *International Conference on Machine Learning*, pages 5419–5428, 2018.
- [50] Weichen Zhang, Wanli Ouyang, Wen Li, and Dong Xu. Collaborative and adversarial network for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3801–3809, 2018.
- [51] Han Zhao, Remi Tachet Des Combes, Kun Zhang, and Geoffrey Gordon. On learning invariant representations for domain adaptation. In *International Conference on Machine Learning*, pages 7523–7532, 2019.