

Transferable Semantic Augmentation for Domain Adaptation

Shuang Li¹ Mixue Xie¹ Kaixiong Gong¹ Chi Harold Liu^{1*} Yulin Wang² Wei Li³

¹Beijing Institute of Technology ²Tsinghua University ³Inceptio Tech.

shuangli@bit.edu.cn michellexie102@gmail.com kxgong@bit.edu.cn liuchi02@gmail.com
wang-y119@mails.tsinghua.edu.cn liweimcc@gmail.com

Abstract

Domain adaptation has been widely explored by transferring the knowledge from a label-rich source domain to a related but unlabeled target domain. Most existing domain adaptation algorithms attend to adapting feature representations across two domains with the guidance of a shared source-supervised classifier. However, such classifier limits the generalization ability towards unlabeled target recognition. To remedy this, we propose a Transferable Semantic Augmentation (TSA) approach to enhance the classifier adaptation ability through implicitly generating source features towards target semantics. Specifically, TSA is inspired by the fact that deep feature transformation towards a certain direction can be represented as meaningful semantic altering in the original input space. Thus, source features can be augmented to effectively equip with target semantics to train a more transferable classifier. To achieve this, for each class, we first use the inter-domain feature mean difference and target intra-class feature covariance to construct a multivariate normal distribution. Then we augment source features with random directions sampled from the distribution class-wisely. Interestingly, such source augmentation is implicitly implemented through an expected transferable cross-entropy loss over the augmented source distribution, where an upper bound of the expected loss is derived and minimized, introducing negligible computational overhead. As a light-weight and general technique, TSA can be easily plugged into various domain adaptation methods, bringing remarkable improvements. Comprehensive experiments on cross-domain benchmarks validate the efficacy of TSA.

1. Introduction

Deep learning has achieved remarkable success on various vision tasks, including image recognition [16, 12, 14] and semantic segmentation [55, 49]. However, the recent success of deep learning methods heavily relies on massive

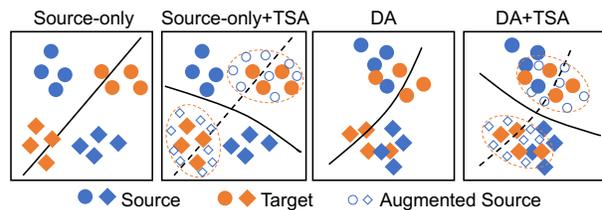


Figure 1. Overview of TSA, which relaxes the assumption that source and target domains share the same classifier. For both traditional deep learning and DA methods, TSA could augment source features towards target semantics to successfully adapt the final classifier from source domain to target domain.

labeled data. In practice, collecting abundant annotated data is expensive [30, 6, 54, 33]. Meanwhile, each domain has its own specific exploratory factors, namely semantics, e.g., the illuminations, colors, visual angles or backgrounds, resulting in the domain shift [28]. Hence, traditional deep models trained on a large dataset usually show poor generalizations on a new domain due to the domain shift issues [24, 10]. To remedy this, one appealing alternative is domain adaptation (DA), which strives to leverage the knowledge of a label-rich source domain to assist the learning in a related but unlabeled target domain [24, 10].

Prior deep DA methods can be roughly categorized as 1) statistical discrepancy minimization based methods [27, 20, 51], which leverage statistical regularizations to explicitly mitigate the cross-domain distribution discrepancy; and 2) adversarial learning based methods [21, 25, 10], which strive to learn domain-invariant representations across two domains via adversarial manners.

Indeed, these DA methods have admittedly yielded promising results, but most of them assume a shared classifier with domain-invariant representations derived. Rare attention has been paid to explicitly enhancing the adaptation ability of the source-supervised classifier, which is also fundamental to DA problems as shown in Fig. 1. To achieve classifier adaptation, Long et al. [26] introduce classifier residual learning to explicitly model the classifier difference across domains. SymNets [52] constructs three classifiers to facilitate joint distribution alignment. However, they

*Corresponding author.

all rely on designing complex network architectures, which may suffer from high computational overhead and hinder the capability and versatility of these methods.

To alleviate aforementioned issues, we propose a Transferable Semantic Augmentation (TSA) approach to implicitly augment source features with target semantic guidance in the deep feature space to facilitate classifier adaptation. Specifically, TSA is motivated by the intriguing property that deep networks excel at disentangling the underlying factors of data variation and linearizing the deep features [41, 45]. There exist many different semantic transformation directions in the deep feature space, and the semantic transformation of one sample can be enforced by translating its deep feature along a certain direction, such as the direction of changing backgrounds. However, it is nontrivial to explicitly discover all kinds of semantic transformation directions. In addition, not all directions are meaningful.

Hence, to effectively explore meaningful transformation directions, we first estimate the inter-domain feature mean difference for each class as the class-wise overall semantic (i.e., an integration of various semantics in one class) bias in the deep feature space. Besides, since the specific semantic information (e.g., different backgrounds, shapes or visual angles) of source and target are different, TSA further estimates target intra-class feature covariance to effectively capture the intra-class semantic variations of target domain. To obtain the accurate estimation, we introduce a memory module to class-wisely calculate feature mean and covariance with pseudo-labeled target samples. At last, we sample semantic transformation directions for source augmentation from a multivariate normal distribution, with the estimated feature mean difference as mean and the target intra-class covariance. In this way, the overall semantic difference between domains and the target intra-class semantic variations can guide source augmented features towards target.

Furthermore, to avoid explicitly generating augmented features and improve the efficiency of TSA, we develop an expected transferable cross-entropy loss over the augmented source distribution with an upper bound derived. By minimizing the upper bound of the expected loss, the source semantic augmentation can be performed and the extra computational overhead is negligible. Then, the trained source classifier can be successfully adapted to target.

Contributions of this work are summarized as follows:

- We propose a novel Transferable Semantic Augmentation (TSA) method for classifier adaptation, which enables source feature augmentation towards target in an implicit manner. Notably, TSA introduces no extra network modules over the backbone network, making it simple to implement and computationally efficient.
- We develop a novel expected transferable cross-entropy loss over the augmented source distribution

for DA, which greatly enhances the classifier adaptation ability. Moreover, as a light-weight and general technique, TSA can be easily plugged into various DA methods to significantly boost their performances.

- Extensive experiments on several cross-domain benchmarks, including Office-Home, Office-31, VisDA-2017 and digits demonstrate that TSA can consistently yield significant performance improvements.

2. Related Work

2.1. Feature Adaptation

Existing domain-invariant feature learning methods mainly fall into two categories. One is the discrepancy-based works [24, 27, 20, 51, 28, 31, 7, 19], which concentrates on mitigating the domain difference by minimizing some statistical discrepancy metrics. To name a few, DAN [24] minimizes maximum mean discrepancy (MMD) on task-specific layers to close the domain gap. Further, JAN in [27] introduces a joint MMD to enforce joint distribution alignment between domains. [51] proposes margin disparity discrepancy (MDD) to facilitate precise alignment with theoretical guarantees. Based on the optimal transport (OT) distance [53], [19] proposes an enhanced transport distance (ETD) to learn the optimal transport plan.

Another category is adversarial-based works [39, 44, 10, 25, 21], which learns domain-invariant features in adversarial manners. For example, DANN [10] introduces a domain discriminator to force the features of two domains indistinguishable. Following DANN, CDAN [25] integrates classification information into the domain discriminator to alleviate class-wise mismatching. While MCD [39] employs a new adversarial paradigm where the adversarial manner occurs between the feature extractor and classifiers rather than the feature extractor and the domain discriminator.

In contrast, instead of learning domain-invariant features, TSA harnesses the augmented source features to ameliorate the generalization ability of source classifier, which is also crucial for tackling DA problems.

2.2. Classifier Adaptation

On par with feature adaptation methods, classifier adaptation is also an indispensable part of DA, since the assumption that source and target domains can share one identical classifier is rather restrictive in practical scenarios.

A branch of classifier adaptation methods augments training samples towards target style to adapt the classifier [5, 2, 13, 43, 22, 23, 8]. To be specific, Hoffman et al. [13] exploit pixel cycle-consistency to make transformed images visually close to target domain images and use transformed samples to train the target model. UNIT [23] exploits image-to-image translation framework based

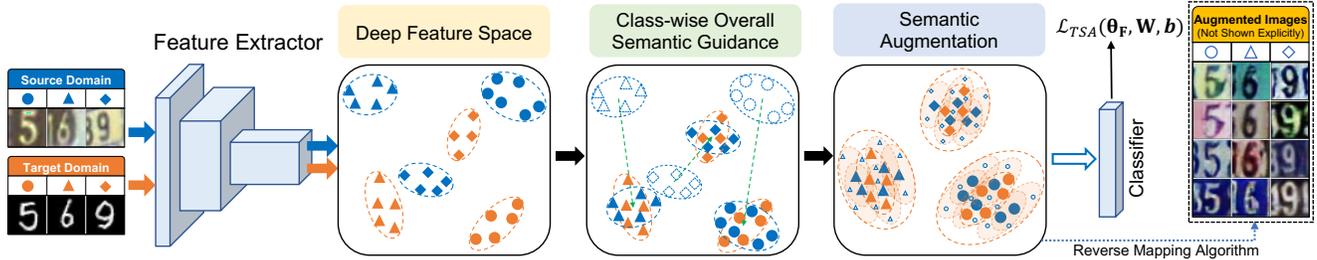


Figure 2. Illustration of TSA, where blue and orange dots represent source and target data respectively. For each class, we exploit the inter-domain feature mean difference which is denoted as the green dashed arrow and target intra-class covariance to augment source features towards target style. Instead of explicitly generating augmented features, we achieve the source augmentation by minimizing the loss \mathcal{L}_{TSA} , which adapts the classifier from source to target. Moreover, we provide the visualization of augmented features using the reverse mapping algorithm in the supplement. The augmented source data with different shapes, colors and backgrounds manifest that TSA can perform meaningful semantic transformations.

on coupled GANs to reconstruct images. TAT [22] generates augmented features with reversed gradients to adversarially train the classifier. [8] leverages the gradually vanishing bridge (GVB) mechanism to generate intermediate features to bridge two domains.

Apart from the generative methods, some methods do not require to explicitly generate samples [26, 18, 52]. For instance, RTN [26] plugs a residual block right after the source classifier to model and mitigate the classifier perturbation. SymNets [52] achieves joint distribution alignment by constructing three different task-specific classifiers.

By contrast, TSA does not need to delicately design auxiliary network modules or explicitly generate training samples. TSA implicitly generates augmented source features with target semantics (e.g., object shapes, colors or backgrounds) to adapt the classifier only by minimizing a robust transferable cross-entropy loss, which is extremely efficient.

3. Transferable Semantic Augmentation

3.1. Motivation and Preliminaries

This work delves into the classifier adaptation for DA problems, which has not yet been fully explored. We propose a Transferable Semantic Augmentation (TSA) approach to effectively augment source features towards target semantics (e.g., backgrounds, view angles or colors) by generating a very large (infinite) number of meaningful semantic transformation directions. The augmented source features will facilitate adapting the classifier from source to target successfully. TSA is inspired by the fact that deep neural networks surprisingly excel at linearizing features of input data, making it linearly separable [41, 45]. Therefore, the relative positions in deep feature space can represent certain semantic transformations in the original input space. Based on this observation, MCF [9] and ISDA [45] have performed semantic augmentation for linear models and deep networks, respectively. However, neither of them considers the crucial domain shift issue in the real world.

To tackle the cross-domain semantic mismatch problems, we need to perform meaningful and target-wise semantic augmentations for source data. In other words, the semantic augmentation directions should effectively guide source features to be target style, such as the shape, color, visual angle or background of target data.

Formally, we denote the source and target domains in DA as $\mathcal{S} = \{(\mathbf{x}_{si}, y_{si})\}_{i=1}^{n_s}$ with n_s labeled samples and $\mathcal{T} = \{\mathbf{x}_{tj}\}_{j=1}^{n_t}$ with n_t unlabeled samples, respectively, where $y_{si} \in \{1, 2, \dots, C\}$ is the label of \mathbf{x}_{si} . Note that, source and target data are sampled from different distributions, such that certain discrepancy may exist between their corresponding ideal classifiers. Our goal is to adapt the source-supervised classifier to target by augmenting the source feature representation $\mathbf{f}_s = F(\mathbf{x}_s, \Theta_F) \in \mathbb{R}^K$, where K is the number of feature embedding dimensions and $F(\cdot)$ is the feature extractor parameterized by Θ_F .

3.2. Cross-Domain Semantic Augmentation

As aforementioned, certain translating directions in deep feature space represent meaningful semantic transformations in the original input space. Thus, we attempt to augment source data towards target in deep feature space. However, manually searching for such cross-domain transformation directions is nontrivial. To address the issue, for each class, we propose to sample random vectors from a multivariate normal distribution with inter-domain feature mean difference as mean and class-conditional target covariance as the covariance matrix. By doing so, rich meaningful cross-domain transformation directions can be discovered. Then, we design an expected transferable cross-entropy loss over the augmented feature distribution to effectively and efficiently adapt the classifier from source to target domain.

3.2.1 Class-wise Overall Semantic Guidance.

Due to the domain shift, for the same class, the means of source and target samples are different in the deep feature space. Such difference reflects the bias of the overall se-

mantics, an integration of various semantics in one class. To conduct useful transformations, we force the mean alignment between augmented source features and target features class-wisely, which effectively bridges the overall semantic bias across two domains. Specifically, for each class c , let μ_s^c and μ_t^c denote the estimated mean feature vector for source and target domains, respectively. Here, pseudo labels for target data are employed as $y'_{tj} = \arg \max_c P_{tj}^c$ to address the label scarce problem in target domain, where P_{tj} is the softmax outputs of target sample \mathbf{x}_{tj} . Then, the inter-domain mean difference $\Delta\mu^c = \mu_t^c - \mu_s^c$ can be exploited to mitigate such overall semantic bias.

3.2.2 Semantic Transformation Direction Learning.

To facilitate meaningful cross-domain semantic augmentations, all possible feature transformation directions towards target style need to be discovered. However, only utilizing the inter-domain mean difference is infeasible to fully explore target semantic variations.

Therefore, to address this problem, we further exploit the intra-class covariance of target features to capture target semantic variations. Specifically, we randomly sample semantic transformation directions from a multivariate distribution $N(\Delta\mu^c, \Sigma_t^c)$ for each class c , and Σ_t^c is the corresponding target intra-class covariance, which contains rich target semantic variations. Here, we develop the normal distribution class-wisely, since the semantic knowledge in each class may differ from each other vastly. Notably, $\Delta\mu^c$ attends to mitigating the overall semantic bias for class c , and Σ_t^c focuses on providing abundant target intra-class semantic variation knowledge. By adding the sampled transformation vectors, the augmented source features will be close to target domain and vary along target semantic variations.

3.2.3 Memory Module.

To efficiently implement TSA in an end-to-end training manner, for each class c , we propose to estimate $\Delta\mu^c$ and Σ_t^c according to a memory module that stores all the latest features. By contrast, ISDA [45] proposes an iterative manner by accumulating statistics of batches from first to current batch. However, the weight distribution of network in early stage will vastly differ from that in later stage. Thus, due to its accumulation property, such iterative manner may encounter the issue that out-of-date features will bias the estimation of mean and covariance. To avoid this, we use a memory module to cache all features of two domains and update in each batch. By doing so, we can discard those out-of-date features and replace them with the latest ones at the negligible cost of memory. Formally, in each iteration i , we will update a batch of features and corresponding labels in memory module \mathbb{M} :

$$\mathbf{f}_j^{\mathbb{M}} \leftarrow \mathbf{f}_j^{(i)}, y_j^{\mathbb{M}} \leftarrow y_j^{(i)}, j \in \mathbf{B}^{(i)}, \quad (1)$$

where j is the sample index within a batch $\mathbf{B}^{(i)}$ and $\mathbf{f}_j^{\mathbb{M}}/y_j^{\mathbb{M}}$ is the feature/label stored in memory module \mathbb{M} . Comparisons of the two manners will be shown in the experiment.

3.2.4 Sampling Strategy.

In the early training stage, the predictions of target samples are inevitably not that accurate. Thus, the estimation of target mean and covariance may not be quite informative as expected. To alleviate this issue, for each class c , we sample semantic transformation directions from $N(\lambda\Delta\mu^c, \lambda\Sigma_t^c)$, where λ is a positive parameter to control the augmentation strength. Since the target predictions will become more and more accurate as the training progresses, here we set $\lambda = (t/T) \times \lambda_0$, where t and T are the current and maximum iterations respectively, and λ_0 is a hyper-parameter. Consequently, as the training goes on, λ will gradually grow from 0 to λ_0 . And this sampling strategy will reduce the impact of less accurate estimations of mean and covariance at the beginning of the training stage.

3.3. Transferable Cross-Entropy Loss Learning

Once C sampling distributions are constructed, each source deep feature \mathbf{f}_{si} can conduct various semantic transformations along the random directions sampled from $N(\lambda\Delta\mu^{y_{si}}, \lambda\Sigma_t^{y_{si}})$ to generate the augmented feature $\hat{\mathbf{f}}_{si}$, i.e., $\hat{\mathbf{f}}_{si} \sim N(\mathbf{f}_{si} + \lambda\Delta\mu^{y_{si}}, \lambda\Sigma_t^{y_{si}})$. Considering a naive method, we can augment each \mathbf{f}_{si} for M times with its label preserved, which will result in an augmented feature set $\{(\mathbf{f}_{si}^1, y_{si}), (\mathbf{f}_{si}^2, y_{si}), \dots, (\mathbf{f}_{si}^M, y_{si})\}_{i=1}^{n_s}$. Based on this, the source network can be trained with a traditional cross-entropy loss on the augmented feature set:

$$\mathcal{L}_M(\Theta_F, \mathbf{W}, \mathbf{b}) = \frac{1}{n_s} \sum_{i=1}^{n_s} \frac{1}{M} \sum_{m=1}^M -\log \left(\frac{e^{\mathbf{w}_{y_{si}}^\top \hat{\mathbf{f}}_{si}^m + b_{y_{si}}}}{\sum_{c=1}^C e^{\mathbf{w}_c^\top \hat{\mathbf{f}}_{si}^m + b_c}} \right), \quad (2)$$

where $\mathbf{W} = [\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_C]^\top \in \mathbb{R}^{C \times K}$ and $\mathbf{b} = [b_1, b_2, \dots, b_C]^\top \in \mathbb{R}^C$ are the weight matrix and bias vector of the last fully connected layer, respectively.

To achieve desired performance, M is usually large, resulting in unexpected computational cost. In TSA, instead of explicitly generating the augmented source features for M times, we intend to implicitly generate infinite augmented source features. When M approaches infinity, we can derive an upper-bound loss according to the Law of Large Numbers over the augmented source distribution.

Specifically, in the case of $M \rightarrow \infty$, our proposed expected transferable cross-entropy loss over the augmented source feature distribution is defined as:

$$\begin{aligned} \lim_{M \rightarrow \infty} \mathcal{L}_M &= \frac{1}{n_s} \sum_{i=1}^{n_s} \mathbb{E}_{\hat{\mathbf{f}}_{si}} \left[-\log \left(\frac{e^{\mathbf{w}_{y_{si}}^\top \hat{\mathbf{f}}_{si} + b_{y_{si}}}}{\sum_{c=1}^C e^{\mathbf{w}_c^\top \hat{\mathbf{f}}_{si} + b_c}} \right) \right] \\ &= \frac{1}{n_s} \sum_{i=1}^{n_s} \mathbb{E}_{\hat{\mathbf{f}}_{si}} \left[\log \left(\sum_{c=1}^C e^{(\mathbf{w}_c - \mathbf{w}_{y_{si}})^\top \hat{\mathbf{f}}_{si} + (b_c - b_{y_{si}})} \right) \right]. \end{aligned} \quad (3)$$

However, it is infeasible to tackle Eq (3) directly. According to the Jensen’s inequality [3] and the concave property of logarithmic function $\log(\cdot)$, we have $\mathbb{E}[\log(X)] \leq \log(\mathbb{E}[X])$. Thus, we can derive the upper bound of the expected loss as follows:

$$\begin{aligned} \lim_{M \rightarrow \infty} \mathcal{L}_M &\leq \frac{1}{n_s} \sum_{i=1}^{n_s} \log \left(\mathbb{E}_{\tilde{\mathbf{f}}_{si}} \left[\sum_{c=1}^C e^{(\mathbf{w}_c^\top - \mathbf{w}_{y_{si}}^\top) \tilde{\mathbf{f}}_{si} + (b_c - b_{y_{si}})} \right] \right) \\ &= \frac{1}{n_s} \sum_{i=1}^{n_s} \log \left(\sum_{c=1}^C \mathbb{E}_{\tilde{\mathbf{f}}_{si}} \left[e^{(\mathbf{w}_c^\top - \mathbf{w}_{y_{si}}^\top) \tilde{\mathbf{f}}_{si} + (b_c - b_{y_{si}})} \right] \right). \end{aligned} \quad (4)$$

Due to $\tilde{\mathbf{f}}_{si} \sim N(\mathbf{f}_{si} + \lambda \Delta \boldsymbol{\mu}^{y_{si}}, \lambda \Sigma_t^{y_{si}})$, we can derive that $(\mathbf{w}_c^\top - \mathbf{w}_{y_{si}}^\top) \tilde{\mathbf{f}}_{si} + (b_c - b_{y_{si}}) \sim N((\mathbf{w}_c^\top - \mathbf{w}_{y_{si}}^\top)(\mathbf{f}_{si} + \lambda \Delta \boldsymbol{\mu}^{y_{si}}) + (b_c - b_{y_{si}}), \sigma_{si}^c)$, where $\sigma_{si}^c = \lambda(\mathbf{w}_c^\top - \mathbf{w}_{y_{si}}^\top) \Sigma_t^{y_{si}} (\mathbf{w}_c - \mathbf{w}_{y_{si}})$. Leveraging the moment-generating function $\mathbb{E}[e^{aX}] = e^{a\mu + \frac{1}{2}a^2\sigma}$, $X \sim N(\mu, \sigma)$, we can obtain

$$\lim_{M \rightarrow \infty} \mathcal{L}_M \leq \mathcal{L}_\infty = -\frac{1}{n_s} \sum_{i=1}^{n_s} \log \frac{e^{Z_{si}^{y_{si}}}}{\sum_{c=1}^C e^{Z_{si}^c}}, \quad (5)$$

where $Z_{si}^c = \hat{y}_{si}^c + \lambda(\mathbf{w}_c^\top - \mathbf{w}_{y_{si}}^\top) \Delta \boldsymbol{\mu}^{y_{si}} + \frac{\sigma_{si}^c}{2}$ and \hat{y}_{si}^c is the c -th element of logits output of \mathbf{x}_{si} .

Essentially, Eq (5) gives us a surrogate loss \mathcal{L}_∞ for optimizing the expected transferable cross-entropy loss, which could effectively adapt the source classifier to target with the augmented source features. Moreover, the proposed novel loss allows us to utilize TSA as a plug-in module for other DA methods to further improve their transferability.

Discussion. Although the derivation of the upper bound is kind of similar with ISDA [45], the motivation of our method is essentially different. We aim to enhance the classifier adaptation for DA where target domain is unlabeled and large domain shift exists between domains. Thus, the proposed $\Delta \boldsymbol{\mu}^c$ and Σ_t^c are crucial to successfully generate transformation directions that can really bridge the domain gap. While ISDA is a completely supervised algorithm and ignores the crucial domain shift issues in reality. In the ablation study, we show that the performance of directly applying ISDA in DA context is inferior. Besides, different from the iterative estimation manner in ISDA, our memory module manner can achieve more accurate estimations of the mean and covariance, which is shown in the experiment.

3.4. Overall Formulation

In information theory, mutual information $I(X; Y)$ measures how related two random variables X and Y are. Actually, strong correlations between target features and predictions will benefit our semantic augmentations, because the extracted features will be more informative and contain more important semantics for predictions, ignoring trivial semantics. Thus, we employ the mutual information maximization on target data, i.e., minimizing the loss in Eq (6).

$$\mathcal{L}_{MI} = \sum_{c=1}^C \hat{P}^c \log \hat{P}^c - \frac{1}{n_t} \sum_{j=1}^{n_t} \sum_{c=1}^C P_{tj}^c \log P_{tj}^c, \quad (6)$$

where $\hat{P} = \frac{1}{n_t} \sum_{j=1}^{n_t} P_{tj}$. Since target domain is unlabeled, we use the average of target predictions to approximate the ground-truth distribution on target domain.

To sum up, the overall objective function of TSA is:

$$\mathcal{L}_{TSA} = \mathcal{L}_\infty + \beta \mathcal{L}_{MI}, \quad (7)$$

where β is a trade-off parameter. The effects of different parts in TSA will be analyzed in details in ablation study.

3.5. Theoretical Insight

To theoretically understand TSA, we introduce the domain adaptation theory proposed by [1], which reveals the ingredients of target generalization error ϵ_t . Formally, let \mathcal{H} denote the hypothesis space and $h \in \mathcal{H}$ denote the classifier, we can formulate the upper bound of ϵ_t as:

$$\epsilon_t(h) \leq \epsilon_s(h) + \frac{1}{2} d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{S}, \mathcal{T}) + \lambda^*, \forall h \in \mathcal{H}, \quad (8)$$

where $\epsilon_s(h)$ is the source generalization error of h , $d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{S}, \mathcal{T})$ is the $\mathcal{H}\Delta\mathcal{H}$ -distance between source and target domains, and $\lambda^* = \epsilon_s(h^*) + \epsilon_t(h^*)$ denotes the error of an ideal joint hypothesis h^* on source and target domains.

With the supervision of labeled source data, $\epsilon_s(h)$ is well bounded. Besides, TSA generates transformation directions from the constructed multivariate distribution to augment source features towards target domain. Since these augmented source features can fill the domain gap, TSA further bounds the $d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{S}, \mathcal{T})$. Moreover, TSA implicitly generates infinite augmented source features that are close to target domain class-wisely, enabling the classifier to jointly minimize $\epsilon_s(h^*)$ and $\epsilon_t(h^*)$ of the shared error λ^* on the augmented training set. To sum up, TSA complies well with the theory, thus further enhancing the transferability.

4. Experiment

4.1. Datasets

Office-Home [42] is a challenging benchmark for domain adaptation, which contains 15,500 images in 65 classes drawn from 4 distinct domains: Artistic (Ar), Clip Art (Cl), Product (Pr), and Real-World (Rw).

Office-31 [38] is a classical cross-domain benchmark, including images in 31 classes drawn from 3 distinct domains: Amazon (A), Webcam (W) and DSLR (D).

VisDA-2017 [35] is a large-scale dataset for visual domain adaptation, containing over 280K images across 12 categories. Following [39], we choose the synthetic domain with 152,397 images as source and the realistic domain with 55,388 images as target.

Table 1. Accuracy (%) on Office-Home for UDA (ResNet-50).

Method	Ar→Cl	Ar→Pr	Ar→Rw	Cl→Ar	Cl→Pr	Cl→Rw	Pr→Ar	Pr→Cl	Pr→Rw	Rw→Ar	Rw→Cl	Rw→Pr	Avg
JAN [27]	45.9	61.2	68.9	50.4	59.7	61.0	45.8	43.4	70.3	63.9	52.4	76.8	58.3
TAT [22]	51.6	69.5	75.4	59.4	69.5	68.6	59.5	50.5	76.8	70.9	56.6	81.6	65.8
TPN [31]	51.2	71.2	76.0	65.1	72.9	72.8	55.4	48.9	76.5	70.9	53.4	80.4	66.2
ETD [19]	51.3	71.9	85.7	57.6	69.2	73.7	57.8	51.2	79.3	70.2	57.5	82.1	67.3
SymNets [52]	47.7	72.9	78.5	64.2	71.3	74.2	64.2	48.8	79.5	74.5	52.6	82.7	67.6
BNM [7]	52.3	73.9	80.0	63.3	72.9	74.9	61.7	49.5	79.7	70.5	53.6	82.2	67.9
MDD [51]	54.9	73.7	77.8	60.0	71.4	71.8	61.2	53.6	78.1	72.5	60.2	82.3	68.1
GSP [47]	56.8	75.5	78.9	61.3	69.4	74.9	61.3	52.6	79.9	73.3	54.2	83.2	68.4
GVB-GD [8]	57.0	74.7	79.8	64.6	74.1	74.6	65.2	55.1	81.0	74.6	59.7	84.3	70.4
ResNet-50 [12]	34.9	50.0	58.0	37.4	41.9	46.2	38.5	31.2	60.4	53.9	41.2	59.9	46.1
+TSA	53.6	75.1	78.3	64.4	73.7	72.5	62.3	49.4	77.5	72.2	58.8	82.1	68.3
DANN [10]	45.6	59.3	70.1	47.0	58.5	60.9	46.1	43.7	68.5	63.2	51.8	76.8	57.6
+TSA	56.5	71.5	78.7	62.2	73.5	72.9	61.8	55.6	80.6	72.6	61.3	82.0	69.1
CDAN [25]	50.7	70.6	76.0	57.6	70.0	70.0	57.4	50.9	77.3	70.9	56.7	81.6	65.8
+TSA	56.7	75.3	80.5	63.9	75.6	75.5	63.8	55.5	81.6	74.8	60.2	84.4	70.7
BSP [4]	52.0	68.6	76.1	58.0	70.3	70.2	58.6	50.2	77.6	72.2	59.3	81.9	66.3
+TSA	57.6	75.8	80.7	64.3	76.3	75.1	66.7	55.7	81.2	75.7	61.9	83.8	71.2

Digital Datasets contain 4 standard digital datasets: MNIST [17], USPS [15], Street View House Numbers (SVHN) [29] and synthetic digits dataset (SYN) [11]. All of these datasets provide number images from 0 to 9. We construct four transfer tasks: MNIST to USPS (M → U), USPS to MNIST (U → M), SVHN to MNIST (SV → M) and SYN to MNIST (SY → M).

4.2. Implementation Details

For a fair comparison, we use ResNet [12] pre-trained on ImageNet [37] as the backbone network for datasets: Office-Home, Office-31 and VisDA-2017. As for digital datasets, we employ the same network structures in [21] and train the networks from scratch. In this paper, all experiments are implemented via PyTorch [32]. We adopt mini-batch SGD optimizer with momentum 0.9 for network optimization, and deep embedded validation [50] to select hyper-parameters λ_0 from $\{0.1, 0.25, 0.5, 0.75, 1.0\}$ and β from $\{0.01, 0.05, 0.1, 0.15, 0.2\}$, and we found $\lambda_0 = 0.25, \beta = 0.1$ works well on all datasets.

We evaluate our approach by applying TSA to the source-only model and several mainstream DA methods, i.e., DANN [10], CDAN [25], BSP [4] and JADA [21], based on their open-source codes. For each task, we report the average accuracy of 3 random trials. Code of TSA is available at <https://github.com/BIT-DA/TSA>.

4.3. Results

Results on Office-Home are reported in Table 1. Office-Home is a challenging dataset with large domain discrepancy for DA. TSA consistently improves the generalization ability of the four based methods. Specifically, TSA empowers BSP with 4.9% improvement, achieving the highest accuracy 71.2% on average. Based on these promising re-

Table 2. Accuracy (%) on Office-31 for UDA (ResNet-50).

Method	A→W	D→W	W→D	A→D	D→A	W→A	Avg
ADDA [40]	86.2	96.2	98.4	77.8	69.5	68.9	82.9
JAN [27]	85.4	97.4	99.8	84.7	68.6	70.0	84.3
ETD [19]	92.1	100.0	100.0	88.0	71.0	67.8	86.2
MCD [39]	88.6	98.5	100.0	92.2	69.5	69.7	86.5
BNM [7]	91.5	98.5	100.0	90.3	70.9	71.6	87.1
DMRL [46]	90.8	99.0	100.0	93.4	73.0	71.2	87.9
SymNets [52]	90.8	98.8	100.0	93.9	74.6	72.5	88.4
TAT [22]	92.5	99.3	100.0	93.2	73.1	72.1	88.4
MDD [51]	94.5	98.4	100.0	93.5	74.6	72.2	88.9
GVB-GD [8]	94.8	98.7	100.0	95.0	73.4	73.7	89.3
GSP [47]	92.9	98.7	99.8	94.5	75.9	74.9	89.5
ResNet-50 [12]	68.4	96.7	99.3	68.9	62.5	60.7	76.1
+TSA	94.8	99.1	100.0	92.6	74.9	74.4	89.3
DANN [10]	82.0	96.9	99.1	79.7	68.2	67.4	82.2
+TSA	94.9	98.5	100.0	92.0	76.3	74.6	89.4
CDAN [25]	94.1	98.6	100.0	92.9	71.0	69.3	87.7
+TSA	95.1	98.7	100.0	95.0	76.3	75.9	90.2
BSP [4]	93.3	98.2	100.0	93.0	73.6	72.6	88.5
+TSA	96.0	98.7	100.0	95.4	76.7	76.8	90.6

sults, we can infer that TSA can stably enhance the transferability of classifiers on this difficult cross-domain dataset.

Results on Office-31 are presented in Table 2. Compared with feature adaptation methods (e.g., JAN and DANN), ResNet-50+TSA surpasses them by a large margin, indicating that classifier adaptation is also indispensable for DA. In particular, TSA brings a large improvement of 7.2% to DANN, revealing that TSA is complementary to previous DA methods. Besides, ResNet-50+TSA also surpasses recent classifier adaptation methods, e.g., SymNets and TAT, revealing that TSA can explore truly useful semantic information to better adapt the classifier.

Results on VisDA-2017 are summarized in Table 3. We can observe that ResNet-101+TSA dramatically outper-

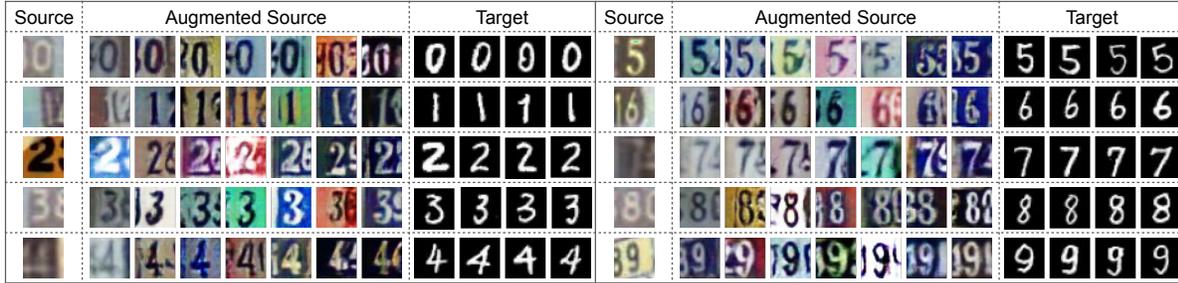


Figure 3. Visualization of semantically augmented images for task SVHN to MNIST. “Source” and “Augmented Source” columns present the images from SVHN and their corresponding augmented images visualized by using our designed reverse mapping algorithm (presented in the supplement), respectively. “Target” column provides several images representing corresponding classes from MNIST.

forms other augmentation methods, like DMRL [46] and DM-ADA [48]. This is mainly due to the fact that TSA exploits the mean difference and target covariance to fully capture the meaningful semantic information class-wisely, achieving better augmentation results. Besides, TSA also brings remarkable improvements to other baseline methods, which validates the effectiveness and versatility of TSA.

Results on Digital Datasets are reported in Table 4. It is noteworthy that DIFA can be seen as DANN combined with extra GAN-based modules which are designed for generating training samples. Thus, it is fair to compare DIFA with DANN+TSA. DANN+TSA exceeds DIFA on average and especially surpasses it by a large margin on $U \rightarrow M$. Besides, TSA does not require to introduce any auxiliary networks, which is extremely easy to implement. For CDAN and JADA, TSA also boosts their performance on average. These encouraging results validate that our semantic augmentation approach can derive a robust classifier efficiently.

Table 3. Accuracy (%) on VisDA-2017 for UDA (ResNet-101).

Method	Synthetic \rightarrow Real
DAN [24]	61.1
MCD [39]	71.9
SimNet [36]	72.9
DMRL [46]	75.5
DM-ADA [48]	75.6
TPN [31]	80.4
ResNet-101 [12]	52.4
+TSA	78.6 (26.2 \uparrow)
DANN [10]	57.4
+TSA	79.6 (22.2 \uparrow)
CDAN [25]	73.7
+TSA	81.6 (7.9 \uparrow)
BSP [4]	76.9
+TSA	82.0 (5.1 \uparrow)

4.4. Analysis

Ablation Study. For each class c , the multivariate normal distribution for sampling transformation directions consists of two important parts: 1) inter-domain feature mean difference $\Delta\mu^c$ and 2) target intra-class covariance matrix Σ_t^c . Only using Σ_t^c to build sampling distribution

Table 4. Accuracy (%) on Digital Datasets for UDA.

Method	M \rightarrow U	U \rightarrow M	SV \rightarrow M	SY \rightarrow M	Avg
ADDA [40]	89.4 \pm 0.2	90.1 \pm 0.8	76.0 \pm 1.8	96.3 \pm 0.4	88.0
PixelDA [2]	95.9 \pm 0.7	-	-	-	-
DIFA [43]	92.3 \pm 0.1	89.7 \pm 0.5	89.7 \pm 2.0	-	-
UNIT [23]	95.9 \pm 0.3	93.6 \pm 0.2	90.5 \pm 0.3	-	-
CyCADA [13]	95.6 \pm 0.2	96.5 \pm 0.1	90.4 \pm 0.4	-	-
TPN [31]	92.1 \pm 0.2	94.1 \pm 0.1	93.0 \pm 0.3	-	-
DM-ADA [48]	96.7 \pm 0.5	94.2 \pm 0.9	95.5 \pm 1.1	-	-
MCD [39]	96.5 \pm 0.3	94.1 \pm 0.3	96.2 \pm 0.4	-	-
ETD [19]	96.4 \pm 0.3	96.3 \pm 0.1	97.9 \pm 0.4	-	-
DMRL [46]	96.1 \pm 0.2	99.0\pm0.1	96.2 \pm 0.4	-	-
Source-only	79.4 \pm 0.4	63.4 \pm 0.3	67.1 \pm 0.5	89.7 \pm 0.2	74.9
+TSA	90.9 \pm 0.3	96.9 \pm 0.2	99.2\pm0.1	97.8 \pm 0.2	96.2
DANN [10]	85.1 \pm 0.5	73.0 \pm 0.2	71.1 \pm 0.4	90.2 \pm 0.2	79.9
+TSA	93.2 \pm 0.3	97.7 \pm 0.2	97.1 \pm 0.3	98.0 \pm 0.2	96.5
CDAN [25]	95.6 \pm 0.1	98.0 \pm 0.1	89.2 \pm 0.3	-	-
+TSA	95.1 \pm 0.2	98.7 \pm 0.1	98.0 \pm 0.1	98.2 \pm 0.2	97.5
JADA [21]	97.6 \pm 0.2	97.1 \pm 0.3	96.4 \pm 0.2	98.6 \pm 0.2	97.4
+TSA	98.0\pm0.1	98.3 \pm 0.3	98.7 \pm 0.2	99.2\pm0.1	98.5

Table 5. Ablation Study of TSA on Office-31.

method	A \rightarrow W	D \rightarrow W	W \rightarrow D	A \rightarrow D	D \rightarrow A	W \rightarrow A	Avg
DANN [10]	82.0	96.9	99.1	79.7	68.2	67.4	82.2
+TSA (w/o $\Delta\mu^c$)	92.5	98.5	100.0	90.8	70.4	69.4	86.9
+TSA (w/o Σ_t^c)	92.3	98.8	100.0	91.1	71.8	70.9	87.5
+ \mathcal{L}_{MI}	91.6	98.2	100.0	89.8	69.7	68.9	86.4
+TSA (w/o \mathcal{L}_{MI})	92.7	98.5	100.0	90.6	70.7	70.6	87.2
+ISDA	90.3	97.9	100.0	87.1	68.5	68.8	85.4
+TSA	94.9	98.5	100.0	92.0	76.3	74.6	89.4

$N(0, \lambda\Sigma_t^c)$ will suffer from overall semantic bias due to the natural dissimilarity between two domains. While only using $\Delta\mu^c$ to construct the distribution $N(\lambda\Delta\mu^c, 0)$ cannot fully capture semantic variations underlying target domain.

Therefore, to investigate the effects of different components in the sampling distribution, we conduct extensive ablation studies of TSA based on DANN [10] in Table 5. The results manifest that our strategy achieves the best performance, indicating that our sampling distribution can yield meaningful transformation directions to guide the augmented source features towards target.

Fixing the sampling distribution $N(\lambda\Delta\mu^c, \lambda\Sigma_t^c)$, we further explore the impacts of mutual information maxi-

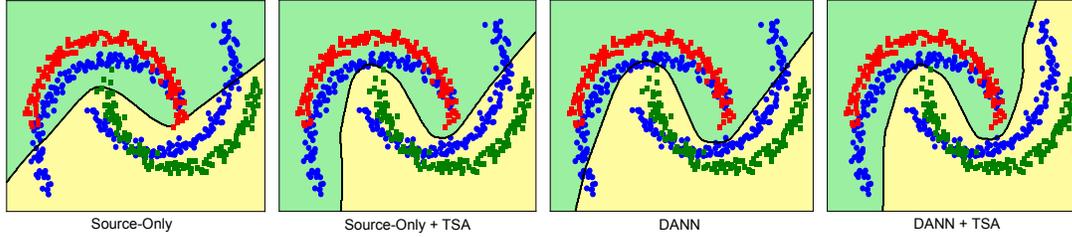


Figure 4. Red and green dots denote source data of class 0 and 1. Blue dots are target data, generated via rotating source data by 30° . Dots in green and yellow areas are respectively classified as 0 and 1. Black solid lines are classification boundaries.

mization loss \mathcal{L}_{MI} in TSA. In Table 5, TSA seems to be more effective when incorporating \mathcal{L}_{MI} , since mutual information maximization helps achieve more accurate target predictions and more informative target features, which is beneficial for capturing the underlying target semantics. Besides, DANN+TSA (w/o \mathcal{L}_{MI}) achieves much better results than DANN+ISDA. The reason is that ISDA [45] does not consider the domain shift problem and fails to capture target semantics, resulting in poor classifier adaptability.

Visualization for Source Augmentation. To intuitively validate that TSA can generate meaningful augmented features, we design a reverse mapping algorithm (presented in the supplement) to search images corresponding to the augmented features. From Fig. 3, we can observe that TSA can alter the semantics of source images, such as backgrounds, colors and shapes of numbers. Moreover, some augmented images closely resemble target images. This verifies that TSA can indeed generate diverse, meaningful and even target-specific augmented features. More augmented images are shown in the supplement.

Inter Twinning Moons 2D Problems. The goal of TSA is to learn transferable classifiers. We carry out experiments on inter twinning moons 2D problem [34] to illustrate the adaptation of classification boundaries brought by TSA. In Fig. 4, source-only model and DANN correctly classify almost all source samples, but misclassify many target samples. By contrast, the classification boundaries of source-only+TSA and DANN+TSA are adapted to target data clearly. The apparent classifier movement reveals that TSA can enhance the classifier transferability notably.

Estimation Bias Comparisons. In this section, we compare the estimation bias of the iterative manner in ISDA [45] and our memory module manner. Specifically, in each epoch, we calculate the average Euclidean distances between ideal and practical estimations of the mean and covariance for all C classes, respectively. The ideal estimation denotes using all features that are freshly produced by feeding the whole dataset into current network to estimate $\Delta\mu^c$ and Σ_t^c for each class c ; the practical estimation denotes leveraging the iterative manner in ISDA or our memory module manner to class-wisely calculate $\Delta\mu^c$ and Σ_t^c . The results are shown in Fig. 5. We observe that the mean and covariance estimations by our memory module man-

ner better approximate the ideal estimation. Based on the results, we can infer that our manner better contributes to deriving more precise mean and covariance estimations.

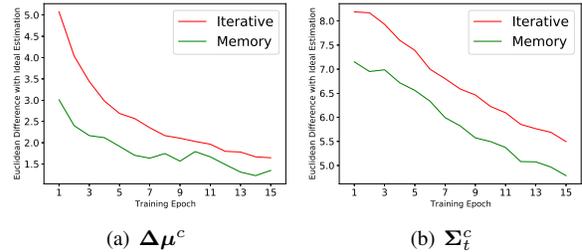


Figure 5. Estimation bias of the iterative manner in [45] and our memory module manner on D→A in different epochs. (a) and (b) present the average estimations bias of inter-domain mean difference $\Delta\mu^c$ and target intra-class covariance Σ_t^c for all C classes.

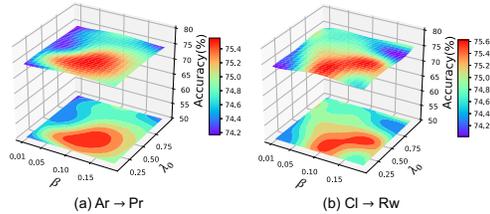


Figure 6. Hyper-parameter sensitivity analysis of TSA on tasks Ar→Pr and Cl→Rw (Office-Home) based on CDAN.

Hyper-parameter Sensitivity. Hyper-parameter λ_0 controls the transformation strength and β balances the two losses in TSA. To check the parameter sensitivity of TSA, we conduct experiments on two random tasks Ar→Pr and Cl→Rw by varying $\lambda_0 \in \{0.1, 0.25, 0.5, 0.75, 1.0\}$ and $\beta \in \{0.01, 0.05, 0.1, 0.15, 0.2\}$. Fig. 6 shows that TSA is not that sensitive to λ_0 and β , and can achieve competitive results under a wide range of hyper-parameter values.

5. Conclusion

This paper presents a transferable semantic augmentation (TSA) approach to ameliorate the adaptation ability of the classifier by optimizing a novel transferable loss over the implicitly augmented source distribution, introducing negligible computational costs. TSA is applicable to various DA methods and can yield significant improvements. Comprehensive experiments on several cross-domain datasets have demonstrated the efficacy and versatility of TSA.

References

- [1] Shai Bendavid, 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. 5
- [2] Konstantinos Bousmalis, Nathan Silberman, David Dohan, Dumitru Erhan, and Dilip Krishnan. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *CVPR*, pages 95–104, 2017. 2, 7
- [3] David Chandler. Introduction to modern statistical. *Mechanics. Oxford University Press, Oxford, UK*, 1987. 5
- [4] Xinyang Chen, Sinan Wang, Mingsheng Long, and Jianmin Wang. Transferability vs. discriminability: Batch spectral penalization for adversarial domain adaptation. In *ICML*, pages 1081–1090, 2019. 6, 7
- [5] Jaehoon Choi, Tae-Kyung Kim, and Changick Kim. Self-ensembling with gan-based data augmentation for domain adaptation in semantic segmentation. In *ICCV*, pages 6829–6839, 2019. 2
- [6] Gabriela Csurka. Domain adaptation for visual applications: A comprehensive survey. *arXiv:1702.05374*, 2017. 1
- [7] Shuhao Cui, Shuhui Wang and] Junbao Zhuo, Liang Li, Qingming Huang, and Qi Tian. Towards discriminability and diversity: Batch nuclear-norm maximization under label insufficient situations. In *CVPR*, pages 3940–3949, 2020. 2, 6
- [8] Shuhao Cui, Shuhui Wang, Junbao Zhuo, Chi Su, Qingming Huang, and Qi Tian. Gradually vanishing bridge for adversarial domain adaptation. In *CVPR*, June 2020. 2, 3, 6
- [9] Laurens Van Der Maaten, Minmin Chen, Stephen Tyree, and Kilian Q Weinberger. Learning with marginalized corrupted features. In *ICML*, pages 410–418, 2013. 3
- [10] Yaroslav Ganin and Victor S. Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, pages 1180–1189, 2015. 1, 2, 6, 7
- [11] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor S. Lempitsky. Domain-adversarial training of neural networks. *JMLR*, 17:59:1–59:35, 2016. 6
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 1, 6, 7
- [13] 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 1994–2003, 2018. 2, 7
- [14] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *CVPR*, pages 4700–4708, 2017. 1
- [15] Jonathan. J. Hull. A database for handwritten text recognition research. *TPAMI*, 16(5):550–554, 1994. 6
- [16] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, pages 1097–1105, 2012. 1
- [17] Yann Lecun, Leon Bottou, Y. Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proc. IEEE*, 86:2278 – 2324, 1998. 6
- [18] Chen-Yu Lee, Tanmay Batra, Mohammad Haris Baig, and Daniel Ulbricht. Sliced wasserstein discrepancy for unsupervised domain adaptation. In *CVPR*, pages 10285–10295, 2019. 3
- [19] Mengxue Li, Yi-Ming Zhai, You-Wei Luo, Peng-Fei Ge, and Chuan-Xian Ren. Enhanced transport distance for unsupervised domain adaptation. In *CVPR*, June 2020. 2, 6, 7
- [20] Shuang Li, Chi Harold Liu, Qiuxia Lin, Qi Wen, Limin Su, Gao Huang, and Zhengming Ding. Deep residual correction network for partial domain adaptation. *TPAMI*, pages 1–1, 2020. 1, 2
- [21] Shuang Li, Chi Harold Liu, Binhui Xie, Limin Su, Zhengming Ding, and Gao Huang. Joint adversarial domain adaptation. In *ACM MM*, pages 729–737. ACM, 2019. 1, 2, 6, 7
- [22] Hong Liu, Mingsheng Long, Jianmin Wang, and Michael Jordan. Transferable adversarial training: A general approach to adapting deep classifiers. In *ICML*, pages 4013–4022, 2019. 2, 3, 6
- [23] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *NeurIPS*, page 700–708, 2017. 2, 7
- [24] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael I Jordan. Learning transferable features with deep adaptation networks. In *ICML*, pages 97–105. ACM, 2015. 1, 2, 7
- [25] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. In *NeurIPS*, pages 1647–1657, 2018. 1, 2, 6, 7
- [26] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Unsupervised domain adaptation with residual transfer networks. In *NeurIPS*, pages 136–144, 2016. 1, 3
- [27] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Deep transfer learning with joint adaptation networks. In *ICML*, pages 2208–2217. ACM, 2017. 1, 2, 6
- [28] You-Wei Luo, Chuan-Xian Ren, Pengfei Ge, Ke-Kun Huang, and Yu-Feng Yu. Unsupervised domain adaptation via discriminative manifold embedding and alignment. In *AAAI*, pages 5029–5036, 2020. 1, 2
- [29] Yuval Netzer, Tiejie Wang, Adam Coates, Alessandro Bisacco, Baolin Wu, and Andrew Y. Ng. Reading digits in natural images with unsupervised feature learning. In *NeurIPS*, 2011. 6
- [30] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *TKDE*, 22(10):1345–1359, 2010. 1
- [31] Yingwei Pan, Ting Yao, Yehao Li, Yu Wang, Chong-Wah Ngo, and Tao Mei. Transferrable prototypical networks for unsupervised domain adaptation. In *CVPR*, June 2019. 2, 6, 7
- [32] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, pages 8024–8035, 2019. 6

- [33] Vishal M. Patel, Raghuraman Gopalan, Ruonan Li, and Rama Chellappa. Visual domain adaptation: A survey of recent advances. *IEEE Signal Process. Mag.*, 32(3):53–69, 2015. [1](#)
- [34] Fabian Pedregosa, Gael Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, Edouard Duchesnay, and Gilles Louppe. Scikit-learn: Machine learning in python. *JMLR*, 12, 01 2012. [8](#)
- [35] Xingchao Peng, Ben Usman, Neela Kaushik, Judy Hoffman, Dequan Wang, and Kate Saenko. Visda: The visual domain adaptation challenge. *arXiv:1710.06924*, 2017. [5](#)
- [36] Pedro O Pinheiro. Unsupervised domain adaptation with similarity learning. In *CVPR*, pages 8004–8013, 2018. [7](#)
- [37] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, and Michael Bernstein. Imagenet large scale visual recognition challenge. *IJCV*, 115(3):211–252, 2015. [6](#)
- [38] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *ECCV*, pages 213–226, 2010. [5](#)
- [39] Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In *CVPR*, pages 3723–3732, 2018. [2](#), [5](#), [6](#), [7](#)
- [40] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *CVPR*, volume 1, page 4, 2017. [6](#), [7](#)
- [41] Paul Upchurch, Jacob Gardner, Geoff Pleiss, Robert Pless, Noah Snavely, Kavita Bala, and Kilian Weinberger. Deep feature interpolation for image content changes. In *CVPR*, page 6090–6099, July 2017. [2](#), [3](#)
- [42] Hemant Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *CVPR*, pages 5018–5027, 2017. [5](#)
- [43] Riccardo Volpi, Pietro Morerio, Silvio Savarese, and Vittorio Murino. Adversarial feature augmentation for unsupervised domain adaptation. In *CVPR*, pages 5495–5504, 2017. [2](#), [7](#)
- [44] Ximei Wang, Liang Li, Weirui Ye, Mingsheng Long, and Jianmin Wang. Transferable attention for domain adaptation. In *AAAI*, pages 5345–5352, 2019. [2](#)
- [45] Yulin Wang, Xuran Pan, Shiji Song, Honghua Zhang, Cheng Wu, and Gao Huang. Implicit semantic data augmentation for deep networks. In *NeurIPS*, pages 12614–12623, 2019. [2](#), [3](#), [4](#), [5](#), [8](#)
- [46] Yuan Wu, Diana Inkpen, and Ahmed El-Roby. Dual mixup regularized learning for adversarial domain adaptation. In *ECCV*, volume 12374, pages 540–555, 2020. [6](#), [7](#)
- [47] Haifeng Xia and Zhengming Ding. Structure preserving generative cross-domain learning. In *CVPR*, June 2020. [6](#)
- [48] Minghao Xu, Jian Zhang, Bingbing Ni, Teng Li, Chengjie Wang, Qi Tian, and Wenjun Zhang. Adversarial domain adaptation with domain mixup. In *AAAI*, pages 6502–6509, 2020. [7](#)
- [49] Yonghao Xu, Du Bo, Lefei Zhang, Qian Zhang, Guoli Wang, and Liangpei Zhang. Self-ensembling attention networks: Addressing domain shift for semantic segmentation. In *AAAI*, volume 33, pages 5581–5588, 2019. [1](#)
- [50] Kaichao You, Ximei Wang, Mingsheng Long, and Michael Jordan. Towards accurate model selection in deep unsupervised domain adaptation. In *ICML*, pages 7124–7133, 2019. [6](#)
- [51] Yuchen Zhang, Tianle Liu, Mingsheng Long, and Michael Jordan. Bridging theory and algorithm for domain adaptation. In *ICML*, pages 7404–7413, 2019. [1](#), [2](#), [6](#)
- [52] Yabin Zhang, Hui Tang, Kui Jia, and Minghui Tan. Domain-symmetric networks for adversarial domain adaptation. In *CVPR*, pages 5031–5040, 2019. [1](#), [3](#), [6](#)
- [53] Zhen Zhang, Mianzhi Wang, and Arye Nehorai. Optimal transport in reproducing kernel hilbert spaces: Theory and applications. *TPAMI*, 42(7):1741–1754, 2020. [2](#)
- [54] Sicheng Zhao, Xiangyu Yue, Shanghang Zhang, Bo Li, Han Zhao, Bichen Wu, Ravi Krishna, Joseph E. Gonzalez, Alberto L. Sangiovanni-Vincentelli, Sanjit A. Seshia, and Kurt Keutzer. A review of single-source deep unsupervised visual domain adaptation. *TNNLS*, 2020. [1](#)
- [55] Mingmin Zhen, Jinglu Wang, Lei Zhou, Tian Fang, and Long Quan. Learning fully dense neural networks for image semantic segmentation. In *AAAI*, volume 33, pages 9283–9290, 2019. [1](#)