

Larger Norm More Transferable: An Adaptive Feature Norm Approach for Unsupervised Domain Adaptation

Ruijia Xu¹ Guanbin Li^{1*} Jihan Yang¹ Liang Lin^{1,2}
¹School of Data and Computer Science, Sun Yat-sen University, China
²DarkMatter AI Research

xurj3@mail2.sysu.edu.cn liguanbin@mail.sysu.edu.cn jihanyang13@gmail.com linliang@ieee.org

Abstract

Domain adaptation enables the learner to safely generalize into novel environments by mitigating domain shifts across distributions. Previous works may not effectively uncover the underlying reasons that would lead to the drastic model degradation on the target task. In this paper, we empirically reveal that the erratic discrimination of the target domain mainly stems from its much smaller feature norms with respect to that of the source domain. To this end, we propose a novel parameter-free Adaptive Feature Norm approach. We demonstrate that progressively adapting the feature norms of the two domains to a large range of values can result in significant transfer gains, implying that those task-specific features with larger norms are more transferable. Our method successfully unifies the computation of both standard and partial domain adaptation with more robustness against the negative transfer issue. Without bells and whistles but a few lines of code, our method substantially lifts the performance on the target task and exceeds state-of-the-arts by a large margin (11.5% on Office-Home [45] and 17.1% on VisDA2017 [31]). We hope our simple yet effective approach will shed some light on the future research of transfer learning. Code is available at <https://github.com/jihanyang/AFN>.

1. Introduction

Deep neural networks, driven by numerous labeled samples, have made remarkable progress in a wide range of computer vision tasks. However, those models are very vulnerable to generalize into new application scenarios. Even a subtle deviation from the training regime can lead to a drastic degradation of the model [43, 49]. Therefore, with the strong motivation to safely transfer knowledge from a label-rich source domain to an unlabeled target domain, Unsupervised Domain Adaptation (UDA) attempts to train a classi-

*Corresponding author is Guanbin Li.

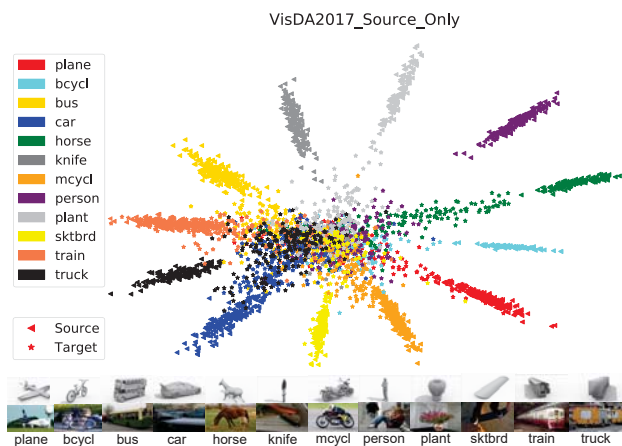


Figure 1: Feature visualization of source and target samples on the Source Only model. This technique is widely used to characterize the feature embeddings under the softmax-related objectives [47, 21, 51]. Specifically, we set the task-specific features to be two-dimensional and retrain the model. Unlike t-SNE [27] whose size of empty space does not account for the similarity between the two data points, this visualization map enables us to interpret the size of feature norms as well as inter-class and intra-class variances. As illustrated, target samples tend to collide in the small-norm (i.e., low-radius) regions which are vulnerable to slight angular variations of the decision boundaries and lead to erratic discrimination.

fier using source samples that can generalize well to the target domain while mitigating the domain shift between the two underlying distributions.

Under the guidance of the theoretical upper bound in [1], the key idea of most existing DA algorithms is to capture not only the task-discriminative but the domain-invariant representations by simultaneously minimizing the source error and some specific statistical discrepancy across the two domains, e.g., \mathcal{H} -divergence [1, 9, 43], $\mathcal{H}\Delta\mathcal{H}$ -divergence [1, 35], Maximum Mean Discrepancy (MMD) [3, 44, 23], correlation distance [39, 40] and etc.

Adversarial domain adaptation [9, 43, 35, 17, 36, 20, 5,

4, 32], which seeks to minimize an approximate domain discrepancy with an adversarial objective, has recently evolved into a dominant method in this field. To the best of our knowledge, RevGrad [9] is the pioneer to empirically measure the \mathcal{H} -divergence by a parametric domain discriminator and adversarially align the features via reverse gradient backpropagation. ADDA [43] instead facilitates the adversarial alignment with GAN-based objective in an asymmetric manner. MCD [35] places a min-max game between the feature generator and the two-branch classifiers to reduce the $\mathcal{H}\Delta\mathcal{H}$ -divergence. On par with the feature-level alignment, generative pixel-level adaptation [17, 36, 20, 4, 32] utilizes Image-to-Image translation techniques to capture the low-level domain shifts.

While the notion of model degradation has been well recognized within the DA community [43, 49], little research work has been published to analyze the underlying cause of this phenomenon. Thus, existing statistic divergences may fail to precisely depict the domain shift and bridging such discrepancies may not guarantee the safe transfer across domains. For example, Shu *et al.* [37] verify that bridging Jensen–Shannon divergence between the two domains does not imply high accuracy on the target task. In this paper, we take a step towards unveiling the nature of model degradation from a solid empirical observation, which is highlighted by Fig. 1. This visualization map suggests that the excessively smaller norms of target-domain features with respect to that of the source domain account for their erratic discrimination. However, there remain two hypothetical interpretations from the current observation:

1) **Misaligned-Feature-Norm Hypothesis:** The domain shift between the source and target domains relies on their misaligned feature-norm expectations. Matching the mean feature norms of the two domains to an arbitrary shared scalar is supposed to yield similar transfer gains.

2) **Smaller-Feature-Norm Hypothesis:** The domain shift substantially relies on the excessively less-informative features with smaller norms for the target task. Despite non-rigorous alignment, adapting the target features far away from the small-norm regions can lead to safe transfer.

With these points in mind, we introduce our parameter-free Adaptive Feature Norm (AFN) approach. First, we propose a simple yet effective statistic distance to characterize the mean-feature-norm discrepancy across domains. Second, we design the Hard AFN to bridge this domain gap by restricting the expected feature norms of the two domains to approximate a shared scalar. It suggests that norm-aligned features can bring effective transfer yet the results can be further improved with a larger scalar. To explore a more sufficient large feature norm in a stable way, we propose the Stepwise AFN to encourage a progressive feature-norm enlargement for each individual sample across domains. As stepwise AFN reveals, the key to achieving successful trans-

fer is to properly lift the target samples towards the large-norm regions while the rigorous alignment is superfluous.

This innovative discovery inspires us to revisit what features are transferable. We recognize that those task-specific features with larger norms imply more informative transferability. Similar findings are explored in the field of model compression in terms of the smaller-norm-less-informative assumption [48], which suggests that parameters or features with smaller norms play a less informative role during the inference. Like the two sides of a coin, in contrast to the model compression that prunes unnecessary computational elements or paths, we place the larger-norm constraint upon the task-specific features to facilitate the more informative and transferable computation on the target domain.

It is noteworthy that under the partial DA, the negative transfer is caused not only from the unrelated samples within the shared categories but also from the unrelated data from the source outlier categories. To this end, we propose meaningful protocols to evaluate the robustness w.r.t a specific algorithm to defense against these potential risks of negative transfer. With thorough evaluation, it reveals that our fairly novel feature-norm-adaptive manner is more robust to safely transfer knowledge from the source domain.

We summarize our contributions as follows:

- i) We empirically unveil the nature of model degradation from a solid observation that the excessively smaller norms of the target-domain features with respect to that of the source domain account for their erratic discrimination.
- ii) We propose a novel AFN approach for UDA by progressively adapting the feature norms of the two domains to a large range of scalars. Our approach is fairly simple yet effective and is translated into a few lines of code.
- iii) We succeed in unifying the computation for both vanilla and partial DA and the feature-norm-adaptive manner is more robust to defense against the negative transfer.
- iv) Extensive experimental results have demonstrated the promise of our approach by exceeding state-of-the-arts across a wide range of visual DA benchmarks.

2. Related Work

Domain adaptation [30, 33, 41] generalizes the learner across different domains by mitigating the domain shift problem. Supervised DA [42, 29, 28] exploits a few labeled data in the target domain while unsupervised DA has no access to that. We focus on the latter scenario in our paper.

Under the guidance of the theoretical upper bound proposed in [1], existing methods explore domain-invariant structures by minimizing some specific statistic distances between the two domains. For example, Maximum Mean Discrepancy (MMD) [3] based methods [23, 26, 44] learn transferable features by minimizing the MMD of their kernel embeddings. Deep correlation alignment [40] proposes to match the mean and covariance of the two distri-

butions. [1] introduces \mathcal{H} - and $\mathcal{H}\Delta\mathcal{H}$ -divergence to characterize the domain discrepancy, which are further developed into matching the corresponding deep representations by [9, 43] and [35] respectively. Regarding the methodology, kernel-based DA [13, 12, 23, 26] and adversarial DA [9, 43, 35, 17, 36, 20, 4, 32] are widely-used in the field.

Inspired by GANs [14], adversarial DA involves a sub-network as the domain classifier to discriminate features of different domains while the deep learner tries to generate the features that deceive the domain classifier. For example, RevGrad [9] utilize a parametric subnetwork as the domain discriminator and adversarially align the features via reverse gradient backpropagation. ADDA [43] instead facilitates the adversarial alignment with GAN-based objectives in an asymmetric manner. MCD [35] conducts a min-max game between the feature generator and the two-branch classifiers in order to reduce the $\mathcal{H}\Delta\mathcal{H}$ -divergence. On par with the feature-level adversarial alignment, generative pixel-level adaptation [17, 36, 20, 4, 32] utilizes Image-to-Image translation techniques to capture the low-level domain shifts.

In addition, other methods are proposed to learn target-specific structures. DRCN [11] involves a reconstruction penalty on target samples. [34] utilizes tri-training to obtain target pseudo labels. [37] refines the target decision boundary based on the cluster assumption. iCAN [50] iteratively applies sample selection on pseudo-labeled target samples and retrain the network.

Standard domain adaptation assumes that the two domains share the identical label space. [6, 7] further open up the partial setting where source label space subsumes the target one. However, it is not trivial to directly migrate the current models in the standard DA as they are prone to suffer from the negative transfer effect. PADA [7] attempts to alleviate this issue by detecting and down-weighting samples belonging to the source outlier classes.

3. Method

3.1. Preliminaries

Given a source domain $\mathcal{D}_s = \{(x_i^s, y_i^s)\}_{i=1}^{n_s}$ of n_s labeled samples associated with $|\mathcal{C}_s|$ categories and a target domain $\mathcal{D}_t = \{x_i^t\}_{i=1}^{n_t}$ of n_t unlabeled samples associated with $|\mathcal{C}_t|$ categories. DA occurs when the underlying distributions corresponding to the source and target domains in the shared label space are *different but similar* [2] to make sense the transfer. Unsupervised DA considers the scenario that we have no access to any labeled target examples.

Vanilla Setting Under this setting, the source and target domains share the identical label space, i.e., $\mathcal{C}_s = \mathcal{C}_t$.

Partial Setting The source label space subsumes the target one, i.e., $\mathcal{C}_s \supset \mathcal{C}_t$. The source labeled data belonging to the outlier categories $\mathcal{C}_s \setminus \mathcal{C}_t$ are unrelated to the target task.

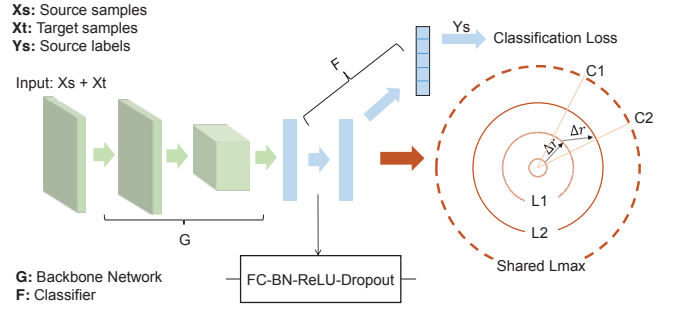


Figure 2: The overall framework of our proposed Adaptive Feature Norm approach. The backbone network G denotes the general feature extraction module. F is employed as the task-specific classifier with l layers, each of which is organized in the FC-BN-ReLU-Dropout order. During each iteration, we apply the feature norm adaptation upon the task-specific features along with the source classification loss as our optimization objective. For the **Hard** variant of AFN, the mean feature norms of source and target samples are constrained to a shared scalar. For the **Stepwise** variant, we encourage a progressive feature-norm enlargement with respect to each individual example at the step size of Δr . To this end, far away from the small-norm regions after the adaptation, the target samples can be correctly classified without any supervision.

3.2. L_2 -preserved Dropout

In this part, we first prove that the standard Dropout operator is L_1 -preserved. As our algorithm is computed based on the L_2 norms of the hidden features, we introduce the following L_2 -preserved Dropout operation to meet our goal.

Dropout is a widely-used regularization technique in deep neural networks [38, 18]. Given a d -dimensional input vector \mathbf{x} , in the training phase, we randomly zero the element $x_k, k = 1, 2, \dots, d$ with probability p by samples $a_k \sim P$ that are generated from the Bernoulli distribution:

$$P(a_k) = \begin{cases} p, & a_k = 0 \\ 1 - p, & a_k = 1 \end{cases} \quad (1)$$

To compute an identity function in the evaluation stage, the outputs are further scaled by a factor of $\frac{1}{1-p}$ and thus

$$\hat{x}_k = a_k \frac{1}{1-p} x_k, \quad (2)$$

which implicitly preserves the L_1 -norm in both of the training and evaluation phases since x_k and a_k are independent:

$$\mathbb{E}[|\hat{x}_k|] = \mathbb{E}[|a_k \frac{1}{1-p} x_k|] = \frac{1}{1-p} \mathbb{E}[a_k] \mathbb{E}[|x_k|] = \mathbb{E}[|x_k|]. \quad (3)$$

However, as we are in pursuit of adaptive L_2 feature norm, we instead scale the output by a factor of $\frac{1}{\sqrt{1-p}}$ and obtain

$$\hat{x}_k = a_k \frac{1}{\sqrt{1-p}} x_k, \quad (4)$$

which satisfies

$$\mathbb{E}[|\hat{x}_k|^2] = \mathbb{E}[|a_k \frac{1}{\sqrt{1-p}} x_k|^2] = \frac{1}{1-p} \mathbb{E}[a_k^2] \mathbb{E}[|x_k|^2] = \mathbb{E}[|x_k|^2]. \quad (5)$$

3.3. Framework

As indicated in Fig. 2, our framework consists of a backbone network G and a classifier F . Existing findings reveal that deep features eventually transit from general to specific along the network and feature transferability significantly drops in higher layers [49]. In our case, G is regarded as the general feature extraction module that inherits from the prevailing network architecture such as ResNet [16]. F represents the task-specific classifier that has l fully-connected layers. We denote the first $l-1$ layers of the classifier as F_f , which results in the so-called bottleneck feature embeddings \mathbf{f} [7, 26]. Those features computed by F_f depend greatly on the specific domain and are not safely transferable to a novel domain. Eventually, we calculate the class probabilities along the last layer F_y , which is followed by a softmax operator. We denote the parameters of G , F_f , F_y with θ_g , θ_f and θ_y respectively. Our intention is to explore an adaptive algorithm to compute the domain-transferable features $\mathbf{f} = F_f(G(\cdot))$ using only source domain supervision. On the other side, as we are unifying the computation with respect to both vanilla and partial DA, it raises an interleaving challenge to defense against the negative transfer effect caused by the outlier categories in the source domain.

3.4. Hard Adaptive Feature Norm

Based on our Misaligned-Feature-Norm Hypothesis, we propose the Maximum Mean Feature Norm Discrepancy (MMFND) to characterize the mean-feature-norm distance between the two distributions and verify whether bridging this statistical domain gap can result in appreciable transfer gains. MMFND is defined by Eq. (6), where the function class \mathcal{H} is the combination of all the possible functions composited by the L_2 -norm operator with the deep representation module, i.e., $h(x) = (\|\cdot\|_2 \circ F_f \circ G)(x)$.

$$\text{MMFND}[\mathcal{H}, \mathcal{D}_s, \mathcal{D}_t] := \sup_{h \in \mathcal{H}} \left(\frac{1}{n_s} \sum_{x_i \in \mathcal{D}_s} h(x_i) - \frac{1}{n_t} \sum_{x_i \in \mathcal{D}_t} h(x_i) \right). \quad (6)$$

Intuitively, the functions class \mathcal{H} are rich enough to contain substantial positive real valued functions on the input \mathbf{x} and the upper bound will greatly deviate from zero if there is no restriction on the function type. In order to avoid this happening, we place a restrictive scalar R to match the corresponding mean feature norms. By restricting both the mean feature norms of the two domains respectively converging towards the shared equilibrium R , the domain gap in terms of MMFND will vanish to zero. We implement this via the Hard Adaptive Feature Norm (HAFN) algorithm,

which is illustrated by Eq. (7).

$$C_1(\theta_g, \theta_f, \theta_y) = \frac{1}{n_s} \sum_{(x_i, y_i) \in \mathcal{D}_s} L_y(x_i, y_i) + \lambda \left(L_d\left(\frac{1}{n_s} \sum_{x_i \in \mathcal{D}_s} h(x_i), R\right) + L_d\left(\frac{1}{n_t} \sum_{x_i \in \mathcal{D}_t} h(x_i), R\right) \right). \quad (7)$$

The optimization objective consists of two terms: the source classification loss L_y in order to obtain the task-discriminative features by minimizing the softmax cross entropy on the source labeled samples, which is indicated by Eq. (8), where $p = [p_1, \dots, p_{|C_s|}]$ is the softmax of the activations predicted by the classifier, i.e., $p = \text{softmax}(F(G(x)))$; the feature-norm penalty in order to obtain the domain-transferable features by minimizing the feature-norm discrepancy between the two domains, where $L_d(\cdot, \cdot)$ is taken as the L_2 -distance and λ is a hyperparameter to trade off the two objectives.

$$L_y(x_i^s, y_i^s; \theta_g, \theta_f, \theta_y) = - \sum_{k=1}^{|C_s|} \mathbb{1}_{[k=y_i^s]} \log p_k. \quad (8)$$

Simple yet effective, MMFND appears to be a novel and superior statistical distance to characterize the cross-domain shift. And by bridging this feature-norm discrepancy with only source-domain supervision through executing HAFN, we can finally achieve the task-discriminative as well as domain-transferable features.

However, the preference setting of R still remains unsettled. As the Misaligned-Feature-Norm Hypothesis suggests, matching feature-norm expectations of the two domains to an arbitrary shared positive real value is supposed to yield similar transfer gains. But this assertion is found not to be true by our empirical results. Specifically, although restricting the mean feature norms of the two domains to even a fairly small value (e.g., $R = 1$, that is, feature normalization) has shown effective results, however, with R gradually increases, the obtained models are still prone to achieve higher accuracies on the target task. To this end, it is natural to explore a sufficiently large R and verify whether the rigorous alignment between the feature-norm expectations is necessary, which is revealed by our Smaller-Feature-Norm Hypothesis. In fact, it is unfortunate that HAFN fails to set an extremely large R as the gradients generated by the feature-norm penalty may eventually lead to an explosion.

3.5. Stepwise Adaptive Feature Norm

To break the aforementioned bottleneck, we introduce an improved variant called Stepwise Adaptive Feature Norm (SAFN) in order to encourage the model to learn task-specific features with larger norms in a progressive manner,

Table 1: Accuracy (%) on *Office-Home* under vanilla setting (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
ResNet [16]	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
DAN [23]	43.6	57.0	67.9	45.8	56.5	60.4	44.0	43.6	67.7	63.1	51.5	74.3	56.3
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
JAN [26]	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
CDAN* [24]	49.0	69.3	74.5	54.4	66.0	68.4	55.6	48.3	75.9	68.4	55.4	80.5	63.8
HAFN	50.2 \pm 0.2	70.1 \pm 0.2	76.6 \pm 0.3	61.1 \pm 0.4	68.0 \pm 0.1	70.7 \pm 0.2	59.5 \pm 0.2	48.4 \pm 0.3	77.3 \pm 0.2	69.4 \pm 0.0	53.0 \pm 0.6	80.2 \pm 0.3	65.4
SAFN	52.0 \pm 0.1	71.7 \pm 0.6	76.3 \pm 0.3	64.2 \pm 0.3	69.9 \pm 0.6	71.9 \pm 0.6	63.7 \pm 0.4	51.4 \pm 0.2	77.1 \pm 0.0	70.9 \pm 0.4	57.1 \pm 0.1	81.5 \pm 0.0	67.3
SAFN*	54.4	73.3	77.9	65.2	71.5	73.2	63.6	52.6	78.2	72.3	58.0	82.1	68.5

Table 2: Accuracy (%) on *VisDA2017* under vanilla setting (ResNet-101)

Method	plane	bicycl	bus	car	horse	knife	mcycl	person	plant	sktbrd	train	truck	Per-class
ResNet [16]	55.1	53.3	61.9	59.1	80.6	17.9	79.7	31.2	81.0	26.5	73.5	8.5	52.4
DAN [23]	87.1	63.0	76.5	42.0	90.3	42.9	85.9	53.1	49.7	36.3	85.8	20.7	61.1
DANN [10]	81.9	77.7	82.8	44.3	81.2	29.5	65.1	28.6	51.9	54.6	82.8	7.8	57.4
MCD [35]	87.0	60.9	83.7	64.0	88.9	79.6	84.7	76.9	88.6	40.3	83.0	25.8	71.9
HAFN	92.7 \pm 0.7	55.4 \pm 4.1	82.4 \pm 2.6	70.9 \pm 1.2	93.2 \pm 0.9	71.2 \pm 3.9	90.8 \pm 0.5	78.2 \pm 1.3	89.1 \pm 0.7	50.2 \pm 2.4	88.9 \pm 0.8	24.5 \pm 0.5	73.9
SAFN	93.6 \pm 0.2	61.3 \pm 4.0	84.1 \pm 0.5	70.6 \pm 2.2	94.1 \pm 0.5	79.0 \pm 4.1	91.8 \pm 0.5	79.6 \pm 1.3	89.9 \pm 0.7	55.6 \pm 3.4	89.0 \pm 0.3	24.4 \pm 2.9	76.1

which is indicated by Eq. (9) as follows:

$$C_2(\theta_g, \theta_f, \theta_y) = \frac{1}{n_s} \sum_{(x_i, y_i) \in D_s} L_y(x_i, y_i) + \frac{\lambda}{n_s + n_t} \sum_{x_i \in D_s \cup D_t} L_d(h(x_i; \theta_0) + \Delta r, h(x_i; \theta)), \quad (9)$$

where $\theta = \theta_g \cup \theta_f$. θ_0 and θ represent the updated and updating model parameters in the last and current iterations respectively. Δr denotes the positive residual scalar to control the feature-norm enlargement. During each iteration, the second penalty in SAFN encourages a feature-norm enlargement at the step size of Δr with respect to individual examples, based on their feature norms calculated by the past model parameters in the last iteration. Instead of assigning a hard value, SAFN enables the optimization process more stable and fairly easy to trade off between the two objectives. To this end, executing SAFN can lead to higher accuracies on the target task by generating more informative features with larger norms. It is noteworthy that SAFN does not rigorously bridge the mean-feature-norm discrepancy, yet one can alternatively place a terminal R to restrict the endless enlargement, which is indicated by Eq. (10). However, our empirical results revealed that Eq. (10) has slightly different result as to replace the second term in Eq. (9). As the Smaller-Feature-Norm hypothesis suggests, once we properly adapt the target samples towards the large-norm regions, the rigorous alignment becomes superfluous.

$$L_d(\max(h(x_i; \theta_0) + \Delta r, R), h(x_i; \theta)). \quad (10)$$

3.6. Model Robustness Evaluation

Though the notion of *negative transfer* has been well recognized within the DA community [30], its rigorous definition is still unclear [46]. A widely accepted description of

negative transfer [30] is stated as “*transferring knowledge from the source can have a negative impact on the target learner*”. While intuitive, how to evaluate it still remains open. Inspired by [46], we propose meaningful protocols to evaluate the robustness of a given algorithm especially under the more general partial setting. It is noteworthy that in this setting, the negative transfer is caused not only from the unrelated samples within the shared categories but also from the unrelated data from the source outlier classes. Let $A_{\mathcal{T}_{|c_t|}}^{l\%}$, $A_{S_{|c_t|} \rightarrow \mathcal{T}_{|c_t|}}$ and $A_{S_{|c_s|} \rightarrow \mathcal{T}_{|c_t|}}$ denote the accuracies by using just $l\%$ target labeled data, transferring without and with source outlier classes w.r.t an identical algorithm. We define i) $A_{\mathcal{T}_{|c_t|}}^{l\%} - A_{S_{|c_t|} \rightarrow \mathcal{T}_{|c_t|}}$ (Closed Negative Gap, **CNG**): the negative impact occurs if the algorithm cannot obtain more transfer gains over the negative influences from another domain than even just labeling a few (e.g., 1%) target data, which is valueless when deployed in the wild. ii) $A_{S_{|c_t|} \rightarrow \mathcal{T}_{|c_t|}} - A_{S_{|c_s|} \rightarrow \mathcal{T}_{|c_t|}}$ (Outlier Negative Gap, **ONG**): especially measures the negative influences that are caused by the source unrelated categories. iii) $A_{\mathcal{T}_{|c_t|}}^{l\%} - A_{S_{|c_s|} \rightarrow \mathcal{T}_{|c_t|}}$ (Partial Negative Gap, **PNG**): reveals whether it is valuable for an algorithm to access and transfer from those available large domains with the potential risks of CNG and ONG. We say that the negative effect exceeds the positive gains once the gap value is positive and vice versa. The larger absolute value suggests more desperate negative influences or more encouraging positive gains.

4. Experiment

4.1. Setup

VisDA2017 [31] is the challenging large-scale benchmark that attempts to bridge the significant synthetic-to-

Table 3: Accuracy (%) on *ImageCLEF-DA* in vanilla setting

Method	I→P	P→I	I→C	C→I	C→P	P→C	Avg
ResNet-50 [16]	74.8	83.9	91.5	78.0	65.5	91.2	80.7
DAN [23]	74.5	82.2	92.8	86.3	69.2	89.8	82.5
DANN [10]	75.0	86.0	96.2	87.0	74.3	91.5	85.0
JAN [26]	76.8	88.0	94.7	89.5	74.2	91.7	85.8
CDAN* [24]	76.7	90.6	97.0	90.5	74.5	93.5	87.1
HAFN	76.9	89.0	94.4	89.6	74.9	92.9	86.3
	±0.4	±0.4	±0.1	±0.6	±0.2	±0.1	
SAFN	78.0	91.7	96.2	91.1	77.0	94.7	88.1
	±0.4	±0.5	±0.1	±0.3	±0.5	±0.3	
SAFN+ENT	79.3	93.3	96.3	91.7	77.6	95.3	88.9
	±0.1	±0.4	±0.4	±0.0	±0.1	±0.1	
SAFN+ENT*	80.2	93.8	96.7	92.8	78.4	95.7	89.6

Table 4: Accuracy (%) on *Office-31* under vanilla setting

Method	A→W	D→W	W→D	A→D	D→A	W→A	Avg
ResNet-50 [16]	68.4	96.7	99.3	68.9	62.5	60.7	76.1
DAN [23]	80.5	97.1	99.6	78.6	63.6	62.8	80.4
DANN [10]	82.0	96.9	99.1	79.7	68.2	67.4	82.2
ADDA [43]	86.2	96.2	98.4	77.8	69.5	68.9	82.9
JAN [26]	85.4	97.4	99.8	84.7	68.6	70.0	84.3
GTA [36]	89.5	97.9	99.8	87.7	72.8	71.4	86.5
CDAN* [24]	93.1	98.2	100.0	89.8	70.1	68.0	86.6
HAFN	83.4	98.3	99.7	84.4	69.4	68.5	83.9
	±0.7	±0.1	±0.1	±0.7	±0.5	±0.3	
SAFN	88.8	98.4	99.8	87.7	69.8	69.7	85.7
	±0.4	±0.0	±0.0	±1.3	±0.4	±0.2	
SAFN+ENT	90.1	98.6	99.8	90.7	73.0	70.2	87.1
	±0.8	±0.2	±0.0	±0.5	±0.2	±0.3	
SAFN+ENT*	90.3	98.7	100.0	92.1	73.4	71.2	87.6

real domain gap with over 280K images across 12 object categories. The source domain has 152,397 synthetic images generated by rendering from 3D models. The target domain has 55,388 real object images collected from *Microsoft COCO* [19]. Under the partial setting, we follow [7] to choose (in alphabetic order) the first 6 categories as target categories and conduct the Synthetic-12 → Real-6 task.

Office-Home [45] is another challenging dataset that collects images of everyday objects to form four domains: Artistic images (**Ar**), Clipart images (**Cl**), Product images (**Pr**) and Real-World images (**Rw**). Each domain contains 65 object categories and they amount to around 15,500 images. Under the partial setting, we follow [7] to choose (in alphabetic order) the first 25 categories as target categories.

Office-31 [33] is a widely-used benchmark for visual DA. It contains 4,652 images of 31 office environment categories from three domains: *Amazon* (**A**), *DSLr* (**D**) and *Webcam* (**W**), which correspond to online website, digital SLR camera and web camera images respectively.

ImageCLEF-DA is built for ImageCLEF 2014 domain adaptation challenge¹ and consists of 12 common categories shared by three public datasets: *Caltech-256* (**C**), *ImageNet*

¹<http://imageclef.org/2014/adaptation>

ILSVRC2012 (**I**) and *Pascal VOC 2012* (**P**). There are 50 images in each category and 600 images in each domain.

Implementation Details We follow the standard protocol [23, 35, 9, 10, 7, 26, 43] to utilize all labeled source data and unlabeled target data that belongs to their own label spaces. We implement our experiments on the widely-used **PyTorch**² platform. For fair comparison, our backbone network is identical to the competitive approaches and is also fine-tuned from the ImageNet [8] pre-trained model. We adopted a unified set of hyper-parameters throughout the *Office-Home*, *Office-31* and *ImageCLEF-DA* datasets under both settings, where $\lambda = 0.05$, $R = 25$ in HAFN and $\Delta r = 1.0$ in SAFN. Since the synthetic domain on *VisDA2017* is easy to converge, we applied a slightly smaller λ and Δr that equal to 0.01 and 0.3 respectively. We used mini-batch SGD optimizer with learning rate 1.0×10^{-3} on all benchmarks. We used center-crop images for the reported results. For each transfer task, we reported the average accuracy over three random repeats. For fairer comparison with those methods [7, 24] which used ten-crop images at the evaluation phase with the best-performing models, we also included our corresponding results with the notion of $\{\text{method}\}^*$ to benefit the future comparisons.

4.2. Result Analysis

Results on *Office-Home*, *VisDA2017*, *ImageCLEF-DA* and *Office-31* under the vanilla setting are reported in Table 1, 2, 3, 4 respectively. Results on *Office-Home* and *VisDA2017* under the partial setting are reported in Table 5, 7. Robustness evaluations in terms of CNG, ONG and PNG are shown in Table 6. As illustrated, our methods significantly outperform the state-of-the-arts throughout all experiments, where SAFN is the top-performing variant.

Results on *VisDA2017* reveal some interesting observations: i) Adversarial based models such as DANN may not effectively learn a diverse transformation across domains on this extremely large-scale transfer task and is prone to suffer from the risk of mode mismatch. However, our encouraging results prove the efficacy of AFN to work reasonably on this large-scale dataset and bridge the significant synthetic-to-real domain gap. ii) Note that existing methods usually mix and optimize multiple learning objectives, and it is not always easy to get an optimal solution. For example, MCD [35] incorporates another class balance objective to align target samples in a balanced way. Nevertheless, our method yields superior performance on most categories, revealing that it is robust to the unbalanced issue without any other auxiliary constraint. iii) Our model is parameter-free thus is more lightweight than the compared methods.

As indicated in Table 1, 3 and 4, our methods achieve new state-of-the-arts on these three benchmarks, and with larger rooms of improvement for those hard transfer tasks,

²<https://pytorch.org/>

Table 5: Accuracy (%) on *Office-Home* under partial setting (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
ResNet [16]	38.57	60.78	75.21	39.94	48.12	52.90	49.68	30.91	70.79	65.38	41.79	70.42	53.71
DAN [23]	44.36	61.79	74.49	41.78	45.21	54.11	46.92	38.14	68.42	64.37	45.37	68.85	54.48
DANN [10]	44.89	54.06	68.97	36.27	34.34	45.22	44.08	38.03	68.69	52.98	34.68	46.50	47.39
PADA* [7]	51.95	67.00	78.74	52.16	53.78	59.03	52.61	43.22	78.79	73.73	56.60	77.09	62.06
HAFN	53.35	72.66	80.84	64.16	65.34	71.07	66.08	51.64	78.26	72.45	55.28	79.02	67.51
	±0.44	±0.53	±0.50	±0.48	±0.30	±1.04	±0.68	±0.42	±0.51	±0.13	±0.37	±0.19	
SAFN	58.93	76.25	81.42	70.43	72.97	77.78	72.36	55.34	80.40	75.81	60.42	79.92	71.83
	±0.50	±0.33	±0.27	±0.46	±1.39	±0.52	±0.31	±0.46	±0.78	±0.37	±0.83	±0.20	
SAFN*	60.00	77.92	83.32	72.91	74.57	79.13	73.09	57.43	82.11	77.87	61.79	82.58	73.56

Table 6: Evaluation on the Robustness

Method	Ar → Rw (5%)			Cl → Rw (5%)			Pr → Rw (5%)			VisDA2017 (1%)		
	CNG	ONG	PNG	CNG	ONG	PNG	CNG	ONG	PNG	CNG	ONG	PNG
DAN [23]	-14.6	13.8	-0.8	-4.9	24.5	19.6	-12.0	17.3	5.3	18.7	25.0	43.7
DANN [10]	-17.2	21.9	4.7	-10.9	39.4	28.5	-14.5	19.5	5.0	13.6	26.7	40.3
JAN [26]	-15.3	13.1	-2.2	-7.1	19.1	12.0	-13.4	12.2	-1.2	20.4	24.5	44.9
PADA [7]	-17.2	12.2	-5.0	-10.9	25.6	14.7	-14.5	9.4	-5.1	13.6	24.2	37.8
SAFN	-18.1	8.5	-9.6	-13.6	8.2	-5.4	-16.1	7.7	-8.4	5.0	15.6	20.6

Table 7: Accuracy (%) on *VisDA2017* under partial setting

Method	Synthetic-12→Real-6
ResNet-50 [16]	45.26
DAN [23]	47.60
DANN [10]	51.01
PADA* [7]	53.53
HAFN	65.06±0.90
SAFN	67.65±0.51
SAFN*	70.67

e.g., $D \rightarrow A$, $A \rightarrow D$, $Cl \rightarrow Pr$, $Cl \rightarrow Rw$ and etc, where the source and target domains are substantially different.

As illustrated in Table 5 and 7, our models obtain substantial improvements for partial DA, with 11.5% gain on *Office-Home* and 17.1% gain on *VisDA2017*. Plain domain-adversarial networks, e.g., DANN, seriously suffer from the mismatch from the source outlier classes and perform even worse than the Source Only variant. An intuitive solution, e.g., PADA [7], is to detect and down-weight the outlier categories during the domain alignment. However, without any heuristic reweighting mechanism, our feature-norm-adaptive manner exhibits stronger robustness against the unrelated data from the source domain. We testify this point via more thorough evaluation in Table 6. Besides, our method works stably and does not require to adjust different hyper-parameters for different subtasks within the same dataset as was done in PADA.

We carefully conduct robustness evaluation for the most challenging transfer tasks, e.g., $Ar65 \rightarrow Rw25$, Synthetic-12 \rightarrow Real-6 and etc. As described in Section 3.6, the positive gap implies more negative impacts over the positive gains and vice versa. The target labeled ratio is 5% and 1% for the two benchmarks. Results in Table 6 reveal some interesting observations: i) Throughout all evaluation met-

rics on all transfer tasks, we can either achieve the largest transfer gains or smallest negative influences. ii) All the methods, including ours, are inevitable to the positive ONG under the more challenging partial setting, while SAFN alleviates the outlier negative impact to the utmost extent. iii) For the $Cl \rightarrow Rw$ transfer task, the comparison methods all have positive PNG, suggesting that they are unable to obtain more transfer gains from the Cl65 domain than using only 5% $Rw25$ labeled samples. However, we still derive encouraging result in this task. iv) It is noteworthy that on the most challenging *VisDA2017* dataset with the significant synthetic-to-real gap, current approaches, including ours, all fail to distill more positive knowledge from the synthetic domain than just labeling 1% real samples. It remains a big challenge for the future development of DA community.

5. Ablation Study

Feature Visualization: Although testifying the efficacy of a DA algorithm via t-SNE [27] embeddings is considered over-interpreted,³ we still follow the *de facto* practice to provide the intuitive understanding. We randomly select 2000 samples across 12 categories from the source and target domains on *VisDA2017* and visualize their task-specific features by t-SNE. As shown in Fig. 4(a), the ResNet features of target samples collide into a mess because of extremely large synthetic-to-real domain gap. After adaptation, as illustrated in Fig. 4(b), our method succeeded in separating target domain samples and better aligning them to the corresponding source domain clusters.

Sample Size of Target Domain: In this part, we empirically demonstrate that our approach is scalable and data-driven with respect to the increase of **unlabeled** target sam-

³<https://interpretablevision.github.io/>

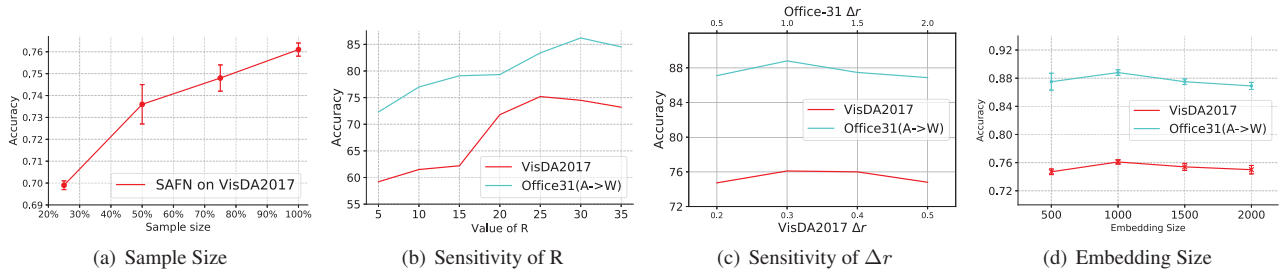


Figure 3: Analysis of (a) varying unlabeled target sample size; (b)(c) parameter sensitivity of R and Δr ; (d) varying embedding size.

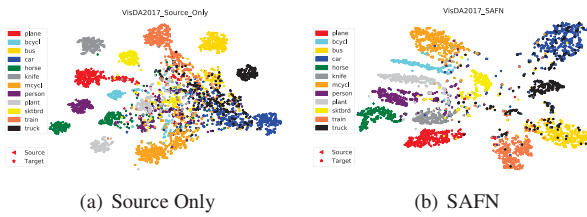


Figure 4: (a) and (b) correspond to the t-SNE embedding visualization of the Source Only and SAFN models on *VisDA2017*. The triangle and star markers denote the source and target samples respectively. Different colors indicate different categories.

ples, which exposes the appealing capacity in practice. It is not necessarily intuitive for adversarial learning based methods to optimize and obtain a diverse transformation upon large volumes of unlabeled target samples. Specifically, we shuffle the target domain on *VisDA2017* and sequentially access the top 25%, 50%, 75% and 100% of the dataset. We train and evaluate our approach on these four subsets. As illustrated in Fig. 3(a), with the sample size gradually increases, the classification accuracy of the corresponding target domain grows accordingly. It shows that the more unlabeled target data are involved in the feature norm adaptation, the more transferable classifier can be obtained.

Complementary with Other Methods: In this part, we demonstrate that our approach can be used in combination with other DA techniques. Because of limited space, we particularly exploit ENTropy minimization (ENT) [15], a low-density separation technique, for demonstration. ENT is widely applied in DA community [25, 37, 22] to encourage the decision boundary to pass through the target low-density regions by minimizing the conditional entropy of target samples. We conduct this case study on *ImageCLEF-DA* and *Office-31* datasets and report the accuracies in Table 3 and 4 respectively. As indicated, with ENT to fit the target-specific structure, we further boost the recognition performance by 0.8% and 1.4% on these two datasets.

Sensitivity of R and Δr : We conduct case studies to investigate the sensitivity of parameter R in HAFN and parameter Δr in SAFN. We select *VisDA2017* and task A→W for demonstration. The results are shown in Fig. 3(b) and Fig. 3(c), by varying $R \in \{5, 10, 15, 20, 25, 30, 35\}$ on

both datasets, $\Delta r \in \{0.2, 0.3, 0.4, 0.5\}$ on *VisDA2017* and $\Delta r \in \{0.5, 1.0, 1.5, 2.0\}$ on task A→W. For parameter R , the accuracy first gradually increases with larger values of R and then begins to decrease as the feature-norm penalty in HAFN may explode. As shown in Fig. 3(c), the accuracies stay almost the same as parameter Δr varies, revealing that SAFN works reasonably stable on these two tasks.

Sensitivity of Embedding Size: We investigate the sensitivity of embedding size of the task-specific features as it plays a significant role in norm computation. We conduct this case study on both *VisDA2017* and A→W transfer tasks. We report the average accuracy over three random repeats for those embedding sizes varying in $\{500, 1000, 1500, 2000\}$. As illustrated in Fig. 3(d), the accuracy stays almost the same and achieves slightly higher when the embedding size is set to 1000, indicating that our approach is robust to a wide range of feature space dimensions.

6. Conclusion

We have presented an innovative discovery for UDA, revealing that the model degradation on the target domain mainly stems from its much smaller feature norms with respect to that of the source domain. To this end, we demonstrated that progressively adapting the feature norms of the two domains to a large range of values can result in significant transfer gains, implying that those task-specific features with larger norms are more transferable. Our method is parameter-free, easy to implement and performs stably. In addition, we successfully unify the computation of both standard and partial DA, and thorough evaluations revealed that our feature-norm-adaptive manner is more robust against the negative transfer. Extensive experimental results have validated the virtue of our proposed approach.

Acknowledgements

This work was supported in part by the State Key Development Program under Grant 2016YFB1001004, in part by the National Natural Science Foundation of China under Grant No.U1811463 and No.61702565, in part by Guangdong “Climbing Program” Special Funds (pdjhb0009). This work was also sponsored by SenseTime Research Fund.

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. 1, 2, 3
- [2] Shai Ben-David, Tyler Lu, Teresa Luu, and Dávid Pál. Impossibility theorems for domain adaptation. In *International Conference on Artificial Intelligence and Statistics*, pages 129–136, 2010. 3
- [3] Karsten M Borgwardt, Arthur Gretton, Malte J Rasch, Hans-Peter Kriegel, Bernhard Schölkopf, and Alex J Smola. Integrating structured biological data by kernel maximum mean discrepancy. *Bioinformatics*, 22(14):e49–e57, 2006. 1, 2
- [4] Konstantinos Bousmalis, Nathan Silberman, David Dohan, Dumitru Erhan, and Dilip Krishnan. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3722–3731, 2017. 2, 3
- [5] Konstantinos Bousmalis, George Trigeorgis, Nathan Silberman, Dilip Krishnan, and Dumitru Erhan. Domain separation networks. In *Advances in Neural Information Processing Systems*, pages 343–351, 2016. 2
- [6] Zhangjie Cao, Mingsheng Long, Jianmin Wang, and Michael I. Jordan. Partial transfer learning with selective adversarial networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 3
- [7] Zhangjie Cao, Lijia Ma, Mingsheng Long, and Jianmin Wang. Partial adversarial domain adaptation. In *European Conference on Computer Vision*, pages 139–155. Springer, 2018. 3, 4, 6, 7
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. 6
- [9] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International Conference on Machine Learning*, pages 1180–1189, 2015. 1, 2, 3, 6
- [10] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1):2096–2030, 2016. 5, 6, 7
- [11] 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. 3
- [12] Boqing Gong, Kristen Grauman, and Fei Sha. Learning kernels for unsupervised domain adaptation with applications to visual object recognition. *International Journal of Computer Vision*, 109(1-2):3–27, 2014. 3
- [13] Boqing Gong, Yuan Shi, Fei Sha, and Kristen Grauman. Geodesic flow kernel for unsupervised domain adaptation. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 2066–2073. IEEE, 2012. 3
- [14] 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. 3
- [15] Yves Grandvalet and Yoshua Bengio. Semi-supervised learning by entropy minimization. In *Advances in neural information processing systems*, pages 529–536, 2005. 8
- [16] 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. 4, 5, 6, 7
- [17] 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 *Proceedings of the 35th International Conference on Machine Learning*, 2018. 2, 3
- [18] Xiang Li, Shuo Chen, Xiaolin Hu, and Jian Yang. Understanding the disharmony between dropout and batch normalization by variance shift. *arXiv preprint arXiv:1801.05134*, 2018. 3
- [19] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 6
- [20] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *Advances in Neural Information Processing Systems*, pages 700–708, 2017. 2, 3
- [21] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Spheraface: Deep hypersphere embedding for face recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, page 1, 2017. 1
- [22] Mingsheng Long, Yue Cao, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Transferable representation learning with deep adaptation networks. *IEEE transactions on pattern analysis and machine intelligence*, 2018. 8
- [23] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *International Conference on Machine Learning*, pages 97–105, 2015. 1, 2, 3, 5, 6, 7
- [24] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. In *Advances in Neural Information Processing Systems*, pages 1647–1657, 2018. 5, 6
- [25] 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. 8
- [26] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Deep transfer learning with joint adaptation networks. In *International Conference on Machine Learning*, pages 2208–2217, 2017. 2, 3, 4, 5, 6, 7
- [27] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008. 1, 7
- [28] Saeid Motiian, Quinn Jones, Seyed Iranmanesh, and Gianfranco Doretto. Few-shot adversarial domain adaptation. In

Advances in Neural Information Processing Systems, pages 6670–6680, 2017. [2](#)

- [29] Saeid Motiian, Marco Piccirilli, Donald A. Adjeroh, and Gianfranco Doretto. Unified deep supervised domain adaptation and generalization. In *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. [2](#)
- [30] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2010. [2](#), [5](#)
- [31] Xingchao Peng, Ben Usman, Neela Kaushik, Judy Hoffman, Dequan Wang, and Kate Saenko. Visda: The visual domain adaptation challenge. *arXiv preprint arXiv:1710.06924*, 2017. [1](#), [5](#)
- [32] Paolo Russo, Fabio M Carlucci, Tatiana Tommasi, and Barbara Caputo. From source to target and back: symmetric bi-directional adaptive gan. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8099–8108, 2018. [2](#), [3](#)
- [33] 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. [2](#), [6](#)
- [34] Kuniaki Saito, Yoshitaka Ushiku, and Tatsuya Harada. Asymmetric tri-training for unsupervised domain adaptation. In *International Conference on Machine Learning*, pages 2988–2997, 2017. [3](#)
- [35] 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. [1](#), [2](#), [3](#), [5](#), [6](#)
- [36] Swami Sankaranarayanan, Yogesh Balaji, Carlos D Castillo, and Rama Chellappa. Generate to adapt: Aligning domains using generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8503–8512, 2018. [2](#), [3](#), [6](#)
- [37] Rui Shu, Hung H Bui, Hirokazu Narui, and Stefano Ermon. A dirt-t approach to unsupervised domain adaptation. In *Proc. 6th International Conference on Learning Representations*, 2018. [2](#), [3](#), [8](#)
- [38] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014. [3](#)
- [39] Baochen Sun, Jiashi Feng, and Kate Saenko. Return of frustratingly easy domain adaptation. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016. [1](#)
- [40] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *European Conference on Computer Vision*, pages 443–450. Springer, 2016. [1](#), [2](#)
- [41] Antonio Torralba and Alexei A Efros. Unbiased look at dataset bias. In *Proceedings of the 2011 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1521–1528. IEEE, 2011. [2](#)
- [42] Eric Tzeng, Judy Hoffman, Trevor Darrell, and Kate Saenko. Simultaneous deep transfer across domains and tasks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4068–4076, 2015. [2](#)
- [43] 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. [1](#), [2](#), [3](#), [6](#)
- [44] Eric Tzeng, Judy Hoffman, Ning Zhang, Kate Saenko, and Trevor Darrell. Deep domain confusion: Maximizing for domain invariance. *arXiv preprint arXiv:1412.3474*, 2014. [1](#), [2](#)
- [45] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, pages 5385–5394. IEEE, 2017. [1](#), [6](#)
- [46] Zirui Wang, Zihang Dai, Barnabás Póczos, and Jaime Carbonell. Characterizing and avoiding negative transfer. *arXiv preprint arXiv:1811.09751*, 2018. [5](#)
- [47] Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. A discriminative feature learning approach for deep face recognition. In *European Conference on Computer Vision*, pages 499–515. Springer, 2016. [1](#)
- [48] Jianbo Ye, Xin Lu, Zhe Lin, and James Z Wang. Rethinking the smaller-norm-less-informative assumption in channel pruning of convolution layers. *arXiv preprint arXiv:1802.00124*, 2018. [2](#)
- [49] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? In *Advances in neural information processing systems*, pages 3320–3328, 2014. [1](#), [2](#), [4](#)
- [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. [3](#)
- [51] Yutong Zheng, Dipan K Pal, and Marios Savvides. Ring loss: Convex feature normalization for face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5089–5097, 2018. [1](#)