

# ECACL: A Holistic Framework for Semi-Supervised Domain Adaptation

Kai Li<sup>1,2\*</sup>, Chang Liu<sup>2</sup>, Handong Zhao<sup>3</sup>, Yulun Zhang<sup>2</sup>, Yun Fu<sup>2</sup>

<sup>1</sup>NEC Laboratories, America, <sup>2</sup>Northeastern University, <sup>3</sup>Adobe Research  
kaili@nec-labs.com, liu.chang6@northeastern.edu, hdzhao@adobe.com  
yulun100@gmail.com, yunfu@ece.neu.edu

## Abstract

*This paper studies Semi-Supervised Domain Adaptation (SSDA), a practical yet under-investigated research topic that aims to learn a model of good performance using unlabeled samples and a few labeled samples in the target domain, with the help of labeled samples from a source domain. Several SSDA methods have been proposed recently, which however fail to fully exploit the value of the few labeled target samples. In this paper, we propose Enhanced Categorical Alignment and Consistency Learning (ECACL), a holistic SSDA framework that incorporates multiple mutually complementary domain alignment techniques. ECACL includes two categorical domain alignment techniques that achieve class-level alignment, a strong data augmentation based technique that enhances the model’s generalizability and a consistency learning based technique that forces the model to be robust with image perturbations. These techniques are applied on one or multiple of the three inputs (labeled source, unlabeled target, and labeled target) and align the domains from different perspectives. ECACL unifies them together and achieves fairly comprehensive domain alignments that are much better than the existing methods: For example, ECACL raises the state-of-the-art accuracy from 68.4 to 81.1 on VisDA2017 and from 45.5 to 53.4 on DomainNet for the 1-shot setting. Our code is available at <https://github.com/kailigo/pacl>.*

## 1. Introduction

Domain adaptation investigates techniques of avoiding severe performance drop when deploying a model on a new domain (target) that has domain gap with the one (source) which the model is trained on. Most existing research focuses on Unsupervised Domain Adaptation (UDA) where a model is trained jointly with unlabeled target data and labeled source data. Many effective UDA approaches have been proposed, from early works that project data from both

domains to a shared feature space [13, 34], to recent ones that are based on adversarial learning [43, 28].

This paper makes a slight diversion from the mainstream UDA research direction and investigates how much it can help if we are further provided with a few (e.g., one sample per class) labeled target samples. We call these scarce labeled target samples as “landmarks”. This is a practical (with minimal labeling effort) yet under-investigated problem and is referred as semi-supervised domain adaptation (SSDA). Preliminary works before the deep learning era use landmarks to more precisely measure the data distribution mismatch between source and target domains, either based on Maximum Mean Discrepancy (MMD) or domain invariant subspace learning [1, 10, 50]. Recent ones revisit this problem and establish new evaluation benchmarks in the deep learning context [38, 18]. However, these prior works have not fully realized the value of the precious landmarks: They are mainly used to optimize the cross-entropy loss along with the labeled source samples that are of a much greater amount; the contribution of the landmarks is significantly diluted and thus the learned model shall still be biased towards the source domain.

In this paper, we propose Enhanced Categorical Alignment and Consistency Learning (ECACL), a SSDA framework which unifies multiple techniques that align domains from different complementary perspectives. As we have access to a few labeled samples from the target domain, i.e., landmarks, we can align the domains in a supervised way by explicitly aligning samples of the same category from the two domains. We propose two techniques to achieve this objective. The first one is based on the prototypical loss [40, 23]. We calculate a target prototype for each class by averaging feature embeddings of the landmarks from that class. Then, source samples are aligned with the target prototype from the same class. The second one is based on the triplet loss. We explicitly push source samples close to the landmarks from the same class and apart from the landmarks from the different classes.

Overfitting would likely occur since the number of landmarks is small. An intuitive approach to this is data aug-

\*Work done when Kai Li was with Northeastern University.

mentation. But rather than employing some commonly used techniques, e.g., image flipping, we harvest a recently proposed one, RandAugment [7] which produces highly perturbed images by applying various image transformations on an image. We apply RandAugment on both labeled source samples and landmarks, which makes the categorical alignment non-trivial to achieve and thus enhances model generalibility.

Consistency learning has recently been proved a very successful solution to various label scarce problems [41, 46, 5]. Inspired by this, we introduce consistency learning to cope the SSDA problem. Specifically, we apply simultaneously a light data augmentation (e.g., image flipping) and a strong augmentation (e.g., RandAugment) on each unlabeled target image, and obtain two versions for each image. We enforce the consistency constraint by producing a pseudo label from the lightly augmented version, and use the pseudo label as the ground truth label for the strongly augmented version for supervised learning. By requiring different perturbed versions of the same image being predicted with the same label, we encourage the model to be robust to changes in the image space and thus be more capable of handling the domain gap. Besides, since unlabeled target samples share the same label space as the labeled source samples, this constraint facilitates label propagation from the labeled source domain to the unlabeled target domain.

Integrating the above techniques that align domains from different perspectives using different combinations of the inputs, we reach the holistic SSDA framework, ECACL. We show in the experiments that ECACL significantly advances the state-of-the-art performance on the common evaluation benchmarks. For example, ECACL lifts the state-of-the-art mean accuracy from 68.4 to 81.1 on *VisDA2017* and from 45.5 to 53.4 on *DomainNet* for the 1-shot setting

In summary, the contributions of this paper are as follows: (1) We propose ECACL, a holistic SSDA framework that incorporates multiple complementary domain alignment techniques. Although each of the incorporated technique is not fundamentally new, we are the first to introduce them to address the SSDA problem and assemble them in a holistic framework. (2) We conduct a comprehensive ablation study and analysis of ECACL, which offer insights on drawing connections among seemingly distinct tasks and identifying contributing techniques. (3) We significantly advance the state-of-the-art performance for SSDA.

## 2. Related Work

**Domain Adaptation (DA).** According to the type of data available in target domain, DA methods can be divided into three categories: Unsupervised Domain Adaptation (UDA), Few-Shot Domain Adaptation (FSDA) and Semi-Supervised Domain Adaptation (SSDA). UDA assumes that target domain data are purely unlabeled. Early methods

in the shallow regime address UDA either by reweighting source instances [16, 12] or projecting samples into a domain invariant feature space [13, 34]. Recent ones are more in the deep regime and approach UDA by moment matching [27, 4, 29] or adversarial learning [43, 11, 30, 28, 24]. FSDA assumes that there is no access to unlabeled samples, but a few labeled ones in target domain. To fully utilize the few labeled target samples, existing methods perform class-wise domain alignment using contrastive loss [33] or triplet loss [49]. SSDA is a hybrid of FSDA and UDA where we have access to both a few labeled samples and many unlabeled samples from target domain. Early works use the extra labeled target samples to help more precisely measure the data distribution mismatch between source and target domains, either based on Maximum Mean Discrepancy (MMD) or domain invariant subspace learning [1, 10, 50, 38, 42]. Saito et al. recently proposed a deep learning based method which alternates between maximizing the classification entropy with respect to the classifier and minimizing it with respect to the feature encoder [38]. Kim et al. extended this work by alleviating the intra-domain discrepancy problem [18]. We approach SSDA in a new way by proposing a general framework into which existing UDA methods can be incorporated as one component for domain alignment along with other novel components.

**Semi-Supervised Learning (SSL).** Leveraging unlabeled data along with labeled ones in the training process, SSL has boosted performance with a variety of training strategies, including graph-based [19], adversarial [32], generative [8], model-ensemble [20], self-training [26, 41], etc. The difference of SSL and SSDA lies that labeled samples of SSL are from the same domain as the unlabeled ones. In SSDA, labeled samples instead come from two different domains and the majority are out-of-domain (relative to the target). So, compared with SSL, SSDA needs first address the domain shift problem in order to leverage the plenty yet out-of-domain labeled samples. We achieve this by employing off-the-shelf UDA techniques and proposing the categorical domain alignment techniques.

**Few-Shot Learning (FSL).** FSL aims to acquire knowledge of *novel classes* with only a few labeled samples [45, 40, 25, 22]. FSL has very distinct goals from SSDA. FSL emphasizes generalizability of a learned model towards novel classes for which there is no sample (neither labeled nor unlabeled) available during training but a few labeled ones in test. SSDA instead focuses on enhancing generalizability of a model towards unlabeled samples of the classes for which during training there are plenty of labeled samples from source domain, a few labeled ones and many unlabeled ones from target domain. Even with different goals, the way that an FSL method [40] utilizes a few labeled samples to recognize other unlabeled ones inspires us to develop the supervised alignment module which achieves

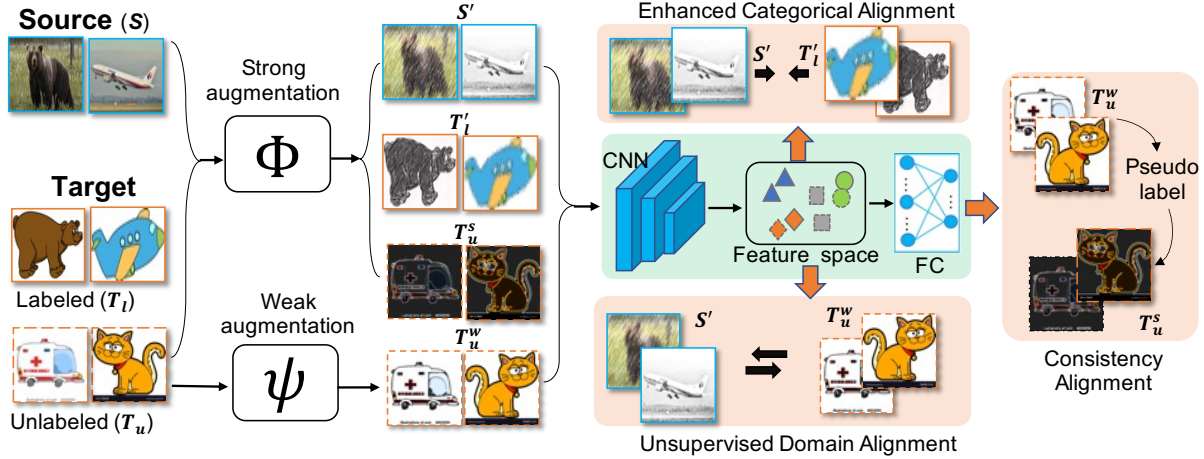


Figure 1. Illustration of the proposed ECACL framework. ECACL includes three modules: (1) The unsupervised domain alignment module that performs domain alignment using labeled source and unlabeled target samples. (2) The enhanced categorical alignment module that conducts class-level alignment by explicitly pushing close cross-domain images that are from the same class, even when the images are strongly perturbed. (3) The consistency alignment module that generates a pseudo label for each unlabeled target sample from its weakly augmented version and applies the pseudo label on its strongly augmented version with supervised learning.

categorical domain alignment.

### 3. Algorithm

Semi-supervised domain adaptation (SSDA) investigates the adaptation from a label-rich source dataset  $\mathcal{S} = \{(s_i, y_i^s)\}_{i=1}^{N_s}$  to a label-scarce target dataset  $\mathcal{T} = \mathcal{T}_l \cup \mathcal{T}_u$ , where  $\mathcal{T}_l = \{(t_i, y_i^t)\}_{i=1}^{N_t}$  is a labeled set and  $\mathcal{T}_u = \{u_i\}_{i=1}^{N_u}$  an unlabeled set.  $\mathcal{S}$  and  $\mathcal{T}$  are drawn from the same label space  $\mathcal{Y} = \{1, 2, \dots, C\}$  but with different data distributions that cause domain shifts. Usually, the number of labeled samples in  $\mathcal{T}_l$  is very small, e.g., one sample per class in the extreme case. We call these labeled target samples as “landmarks”. Our goal is to learn a domain adaptive model using  $\mathcal{S}$ ,  $\mathcal{T}_u$  and landmarks  $\mathcal{T}_l$ . Let the model be  $h = f \circ g$  with parameters  $\theta$ , where  $f$  generates features from images and  $g$  outputs label predictions based on the extracted features.

We can see that the difference of SSDA from UDA is the extra access to landmarks. A naive way to extend an UDA method to SSDA is to optimize

$$\mathcal{L}_{uda} = \mathcal{L}_{ce} + \alpha \mathcal{L}_{ua}, \quad (1)$$

where

$$\mathcal{L}_{ce} = \frac{1}{N_s} \sum_{(s_i, y_i^s) \sim \mathcal{S}} L(h(s_i), y_i^s) + \frac{1}{N_t} \sum_{(t_i, y_i^t) \sim \mathcal{T}_l} L(h(t_i), y_i^t) \quad (2)$$

is the cross-entropy loss over labeled source and target samples.  $\mathcal{L}_{ua}$  is an UDA loss that exploits unlabeled target samples and labeled source samples for domain alignment. It varies in different UDA methods.

Naively merging landmarks into source samples for the cross-entropy loss optimization does not fully release their

potentials, as their contribution would be severely diluted. The learned model would thus still be biased towards the source domain. We solve this by performing *categorical alignment* where we explicitly push source samples towards the landmarks from the same class and apart from the landmarks from the different classes (Section 3.1). This encourages the model to produce features that maintain class discrimination despite of domain shifts. In Section 3.2, we further enhance the categorical alignment technique with a data augmentation strategy where images are heavily perturbed, making the task harder to fulfill and hence improving model generalizability. Besides, we carefully design another data augmentation based technique that is applied on unlabeled target samples. This technique constrains the model to make consistent predictions for different versions of the same image undergone different levels of perturbations. Figure 1 shows our framework.

#### 3.1. Categorical Alignment

With labeled source samples  $\mathcal{S} = \{(s_i, y_i^s)\}_{i=1}^{N_s}$  and labeled target samples  $\mathcal{T}_l = \{(t_i, y_i^t)\}_{i=1}^{N_t}$ , we propose two approaches to achieve categorical alignment, one based on the prototypical loss and the other based on the triplet loss.

**Prototypical loss based approach.** This approach learns to minimize the risk of assigning source samples to the landmarks from the same class. Specifically, we calculate a target representation, or a target prototype for each class by averaging the feature embeddings of the landmarks from that class [40]:

$$\mathbf{c}_k = \frac{1}{|\mathcal{T}_k|} \sum_{(t_i, y_i) \in \mathcal{T}_k} f(t_i), \quad (3)$$

where  $\mathcal{T}_k$  is the landmark collection for class  $k$ . With the

target prototypes for all classes  $\{\mathbf{c}_k\}_{k=1}^C$ , we can compute a distribution over classes for a source sample  $(\mathbf{s}_i, y_i^s)$  based on a softmax over distances to the target prototypes in the embedding space:

$$p(y_i^s = y | \mathbf{s}_i) = \frac{\exp(-\|f(\mathbf{s}_i) - \mathbf{c}_y\|_2)}{\sum_{k=1}^C \exp(-\|f(\mathbf{s}_i) - \mathbf{c}_k\|_2)}. \quad (4)$$

Then, we can calculate the prototypical loss over all source samples as

$$L_{pa} = \frac{1}{CN_s} \sum_{(\mathbf{s}_i, y_i^s) \sim \mathcal{S}} \sum_{y \sim \mathcal{Y}} y_i^s \log[-p(y_i^s = y | \mathbf{s}_i)]. \quad (5)$$

Note that although class prototype based solution has been proposed to address UDA [35, 47], we are the first to introduce it for SSDA. Besides, while the class prototypes in the UDA method are calculated based on pseudo labels which are not always reliable [51], we compute class prototypes directly from landmarks.

**Triplet loss based approach.** This approach explicitly optimizes the model to produce features such that cross-domain samples of the same class should be of higher similarity than those from different classes [15]. Specifically, for each landmark  $(\mathbf{t}_i, y_i^t) \in \mathcal{T}_l$ , we find from  $\mathcal{S}$  the least similar source sample  $(\mathbf{s}_p, y_p)$  which also belongs to class  $y_t$  (i.e., the hard positive). Meanwhile, we find from  $\mathcal{S}$  the most similar sample  $(\mathbf{s}_n, y_n)$  which belongs to a class different from  $y_t$  (i.e., the hard negative). With the hardest triplet  $(\mathbf{t}_i, \mathbf{s}_n, \mathbf{s}_p)$ , we optimize the following triplet loss as:

$$L_{ta} = \frac{1}{N_t} \sum_{(\mathbf{t}_i, y_i^t) \sim \mathcal{T}_l} [\|f(\mathbf{t}_i) - f(\mathbf{s}_p)\|_2^2 - \|f(\mathbf{t}_i) - f(\mathbf{s}_n)\|_2^2 + m]_+. \quad (6)$$

The loss makes sure that for any given landmark, its hard positive sample should be closer to it than its hard negative sample by at least a margin  $m$ .

## 3.2. Domain Alignment with Data Augmentation

To further alleviate the scarcity of landmarks, we propose two data augmentation based domain alignment techniques, one applied on labeled samples and the other applied on unlabeled samples.

### 3.2.1 Enhanced Categorical Alignment

It has been shown recently that strong augmentation that creates highly perturbed images brings significant performance gains for supervised learning [6, 7]. Inspired by this, we introduce strong data augmentation to address the DA problem. For each labeled sample from source and target domains,  $(\mathbf{s}_i, y_i^s) \in \mathcal{S}$  or  $(\mathbf{t}_i, y_i^t) \in \mathcal{T}_l$ , we process it with RandAugment [7] by applying random augmentation techniques sampled from a transformation set, including color,

brightness, contrast adjustments, rotation, polarization, etc. This is then followed by the Cutout [9]. Now we obtain  $\mathcal{S}'$  and  $\mathcal{T}'$  that consist of highly perturbed images. With  $\mathcal{S}'$  and  $\mathcal{T}'$ , we can get the enhanced categorical alignment objectives. For the prototypical loss based one, we reach

$$L'_{pa} = \frac{1}{CN_s} \sum_{(\mathbf{s}'_i, y_i^s) \sim \mathcal{S}'} \sum_{y \sim \mathcal{Y}} y_i^s \log[-p(y_i^s = y | \mathbf{s}'_i)], \quad (7)$$

where  $P(y_s = y | \mathbf{s}')$  is calculated with Eq. (4), and the target prototypes are calculated with  $\mathcal{T}'$ . Similarly, for the triplet loss based one, we have

$$L'_{ta} = \frac{1}{N_t} \sum_{(\mathbf{t}'_i, y_i^t) \sim \mathcal{T}'_l} [\|f(\mathbf{t}'_i) - f(\mathbf{s}'_p)\|_2^2 - \|f(\mathbf{t}'_i) - f(\mathbf{s}'_n)\|_2^2 + m]_+, \quad (8)$$

where  $\mathbf{s}'_p$  and  $\mathbf{s}'_n$  are the hard positive sample and hard negative sample mined from  $\mathcal{S}'$ , respectively.

Strong data augmentation produces a wider range of highly perturbed images, which makes the model harder to memorize the few landmarks and therefore enhances the generalizability of the learned model. On the one hand, the model is forced to be insensitive to more diverse changes or perturbations in the image space, which helps domain alignment as these changes model a wide range of factors causing domain shift. On the other hand, the above categorical alignment techniques in essence optimize the model extracting image features that the intra-class ones are of higher similarity than the inter-class ones regardless of domain shifts. It is harder for the model to achieve this optimization objective with highly perturbed images. Thus, the model is encouraged to mine the most discriminative class semantics out of highly perturbed images.

### 3.2.2 Consistency Alignment

Inspired by the recent success of consistency learning in semi-supervised learning [41, 46, 2, 3], we introduce consistency learning to address SSDA and propose the Consistency Alignment (CONA) module. For each unlabeled target sample  $\mathbf{u}_i \in \mathcal{T}_u$ , we apply weak augmentation  $\psi$  and strong augmentation  $\Phi$ :

$$\mathbf{u}_i^w = \psi(\mathbf{u}_i), \quad (9)$$

$$\mathbf{u}_i^s = \Phi(\mathbf{u}_i). \quad (10)$$

The weak augmentation  $\psi$  includes the widely-used image flipping and image translation. Same as the practice for labeled samples, we use RandAugment [7] and Cutout [9] as our strong data augmentation  $\Phi$ . We feed  $\mathbf{u}_i^s$  and  $\mathbf{u}_i^w$  to the model  $h$ , and optimize the following objective function:

$$L_{cona} = \sum_{\mathbf{u}_i \sim \mathcal{U}} [\mathbb{1}(\max(\mathbf{p}_w) \geq \sigma) H(\tilde{\mathbf{p}}_w, \mathbf{p}_s)], \quad (11)$$

where  $\mathbf{p}_s$  and  $\mathbf{p}_w$  are the classification probabilities of augmented samples  $\mathbf{u}_i^s$  and  $\mathbf{u}_i^w$ , respectively.  $\tilde{\mathbf{p}}_w = \arg \max(\mathbf{p}_w)$  returns a one-hot vector for the prediction;  $H(\cdot, \cdot)$  is the cross-entropy of two possibility distributions;  $\mathbb{1}(\cdot)$  is an indicator function and  $\max(\mathbf{p}_w)$  returns the highest possibility score.

In essence, the above CONA module computes a pseudo label for an unlabeled sample from its weakly-augmented version and applies the pseudo label on its strongly-augmented version for the cross-entropy loss optimization. To mitigate the impact of incorrect pseudo labels, only the samples with confident predictions (the highest probability scores are above a threshold) are used for loss computation. This introduces a form of consistency regularization, encouraging the model to be insensitive to image perturbations and hence stronger in classifying unlabeled images.

Pseudo labeling (or self-training) has been investigated before for domain adaptation [35, 51], but our method is clearly distinct from the previous ones. Existing methods usually perform stage-wise pseudo labeling: Each stage consists of a number of training epochs and the latest model is applied on unlabeled samples in the end of each stage. The confidently predicted samples are selected for model training in the next stage usually in the same way as labeled samples from source domain. Within all training epochs in each stage, the pseudo-labeled samples remain unchanged. Our method instead performs mini-batch-wise pseudo labeling: In each mini-batch, we compute a pseudo label for every sample from its weakly-augmented version and apply the pseudo label on the strongly-augmented one. We discard all of the pseudo labels after each mini-batch, which alleviates the harmful impact of incorrect pseudo labels.

The overall learning objective of our method is a weighted combination of the UDA loss, the enhanced categorical alignment loss, and the consistency alignment loss:

$$L = L_{uda} + \lambda_1 L_{cata} + \lambda_2 L_{cona}, \quad L_{cata} = \{L'_{pa}, L'_{ta}\} \quad (12)$$

where  $\lambda_1$  and  $\lambda_2$  are the hyper-parameters.  $L_{cata} = L'_{pa}$  and  $L_{cata} = L'_{ta}$  correspond to the two variants of our ECACL framework which performs categorical alignment based on the prototypical loss and the triplet loss, respectively. We refer these two variants as ECACL-P and ECACL-T, respectively.

**Algorithm 1** outlines the main steps of the proposed framework.

## 4. Experiments

**Datasets and evaluation protocols.** We conduct experiments on three commonly used datasets, namely, *VisDA2017* [37], *DomainNet* [36], and *Office-Home* [44].

*DomainNet* consists of 6 domains of 345 categories. Following [38], we select the *Real* (R), *Clipart* (C), *Painting*

---

### Algorithm 1. Proposed ECACL framework

---

**Input:** Source set  $\mathcal{S} = \{\mathcal{X}_s, \mathcal{Y}_s\}$ , labeled target set  $\mathcal{T}_l = \{\mathcal{X}_t, \mathcal{Y}_t\}$  and unlabeled target set  $\mathcal{T}_u$ .

**Output:** Domain adaptive model  $h$ .

---

**while not done do**

1. Sample from  $\mathcal{S} \cup \mathcal{T}_l$  labeled images  $\mathcal{B}_l = \{\{\mathbf{s}_{i,j}\}_{i=1}^{N_s}, \{\mathbf{t}_{i,j}\}_{i=1}^{N_t}, y_j\}_{j=1}^M$ . Sample unlabeled images  $\mathcal{B}_u = \{\mathbf{u}_i\}_{i=1}^{N_u}$  from  $\mathcal{U}$ . Denote  $\mathcal{B}_l = \mathcal{B}_s \cup \mathcal{B}_t$  where  $\mathcal{B}_s = \{\{\mathbf{s}_{i,j}\}_{i=1}^{N_s}, y_j\}_{j=1}^M$  and  $\mathcal{B}_t = \{\{\mathbf{t}_{i,j}\}_{i=1}^{N_t}, y_j\}_{j=1}^M$ .
2. Apply strong augmentation on  $\mathcal{B}_s$ ,  $\mathcal{B}_t$  and  $\mathcal{B}_u$ , and get  $\mathcal{B}'_s = \Phi(\mathcal{B}_s)$ ,  $\mathcal{B}'_t = \Phi(\mathcal{B}_t)$  and  $\mathcal{B}'_u = \Phi(\mathcal{B}_u)$ .
3. Apply weak augmentation on  $\mathcal{B}_u$  and get  $\mathcal{B}^w_u = \psi(\mathcal{B}_u)$ .
4. Calculate the cross-entropy loss with  $\mathcal{B}'_s$  and  $\mathcal{B}'_t$  using Eq. (2).
5. Calculate the unsupervised alignment loss with  $\mathcal{B}'_s$  and  $\mathcal{B}^w_u$  using an existing UDA method (e.g., MME [38]).
6. Calculate the enhanced categorical alignment loss with  $\mathcal{B}'_s$  and  $\mathcal{B}'_t$  by using Eq. (7) (for the ECACL-P variant), or Eq. (8) (for the ECACL-T variant).
7. Calculate the consistency alignment loss with  $\mathcal{B}'_u$  and  $\mathcal{B}^w_u$  using Eq. (11).
8. Optimize the model  $h$  using Eq. (12).

**end while**

---

(P), and *Sketch* (S) as the 4 evaluation domains and perform the following cross-domain evaluations: R→C (adaptation from source *Real* to target *Clipart*), R→P, P→C, C→S, S→P, R→S, and P→R. For each set of cross-domain experiment, we evaluate both 1-shot and 3-shot settings, where there are 1 and 3 labeled target samples, respectively. The labeled samples are randomly selected and we use the provided splits for experiments. We evaluate the classification accuracy for all the 7 sets of experiments and also report the mean of the accuracies.

*VisDA2017* includes 152,397 synthetic and 55,388 real images from 12 categories. For SSDA evaluation, we randomly select 1 and 3 real images from each of the 12 classes as the landmarks, which correspond to the 1-shot and 3-shot evaluation settings, respectively. Following previous method [48], we report the per-class classification accuracy and also the Mean Class Accuracy (MCA) over all classes.

*Office-Home* contains images of 65 categories that are from 4 different domains, namely, *Real* (R), *Clipart* (C), *Art* (A), and *Product* (P). We use the same 1-shot and 3-shot splits as [38] and evaluate the adaptation performance for all 12 pairs of domains. We report the classification accuracies for all the experimental sets and the accuracy mean.

**Implementation details.** The proposed ECACL-P and ECACL-T are general SSDA frameworks that can incorporate most existing UDA methods and leverage landmark samples to improve adaptation performance. Depending on the UDA method built upon, we can get different variants. But for ease and fairness of evaluation, we conduct most

	Net	R→C		R→P		P→C		C→S		S→P		R→S		P→R		Mean	
		1-shot	3-shot	1-shot	3-shot	1-shot	3-shot	1-shot	3-shot	1-shot	3-shot	1-shot	3-shot	1-shot	3-shot	1-shot	3-shot
ST	AlexNet	43.3	47.1	42.4	45.0	40.1	44.9	33.6	36.4	35.7	38.4	29.1	33.3	55.8	58.7	40.0	43.4
DANN	AlexNet	43.3	46.1	41.6	43.8	39.1	41.0	35.9	36.5	36.9	38.9	32.5	33.4	53.6	57.3	40.4	42.4
ADR	AlexNet	43.1	46.2	41.4	44.4	39.3	43.6	32.8	36.4	33.1	38.9	29.1	32.4	55.9	57.3	39.2	42.7
CDAN	AlexNet	46.3	46.8	45.7	45.0	38.3	42.3	27.5	29.5	30.2	33.7	28.8	31.3	56.7	58.7	39.1	41.0
ENT	AlexNet	37.0	45.5	35.6	42.6	26.8	40.4	18.9	31.1	15.1	29.6	18.0	29.6	52.2	60.0	29.1	39.8
MME	AlexNet	48.9	55.6	48.0	49.0	46.7	51.7	36.3	39.4	39.4	43.0	33.3	37.9	56.8	60.7	44.2	48.2
Meta-MME	AlexNet	-	56.4	-	50.2	-	51.9	-	39.6	-	43.7	-	38.7	-	60.7	-	48.7
BiAT	AlexNet	54.2	58.6	49.2	50.6	44.0	52.0	37.7	41.9	39.6	42.1	37.2	42.0	56.9	58.8	45.5	49.4
FAN	AlexNet	47.7	54.6	49.0	50.5	46.9	52.1	38.5	42.6	38.5	42.2	33.8	38.7	57.5	61.4	44.6	48.9
ECACL-T	AlexNet	<b>56.8</b>	<b>62.9</b>	<b>54.8</b>	58.9	<b>56.3</b>	<b>60.5</b>	<b>46.6</b>	<b>51.0</b>	<b>54.6</b>	<b>51.2</b>	<b>45.4</b>	<b>48.9</b>	<b>62.8</b>	<b>67.4</b>	<b>53.4</b>	<b>57.7</b>
ECACL-P	AlexNet	55.8	62.6	54.0	<b>59.0</b>	56.1	<b>60.5</b>	46.1	50.6	<b>54.6</b>	50.3	45.0	48.4	62.3	<b>67.4</b>	52.8	57.6
ST	ResNet-34	55.6	60.0	60.6	62.2	56.8	59.4	50.8	55.0	56.0	59.5	46.3	50.1	71.8	73.9	56.9	60.0
DANN	ResNet-34	58.2	59.8	61.4	62.8	56.3	59.6	52.8	55.4	57.4	59.9	52.2	54.9	70.3	72.2	58.4	60.7
ADR	ResNet-34	57.1	60.7	61.3	61.9	57.0	60.7	51.0	54.4	56.0	59.9	49.0	51.1	72.0	74.2	57.6	60.4
CDAN	ResNet-34	65.0	69.0	64.9	67.3	63.7	68.4	53.1	57.8	63.4	65.3	54.5	59.0	73.2	78.5	62.5	66.5
ENT	ResNet-34	65.2	71.0	65.9	69.2	65.4	71.1	54.6	60.0	59.7	62.1	52.1	61.1	75.0	78.6	62.6	67.6
MME	ResNet-34	70.0	72.2	67.7	69.7	69.0	71.7	56.3	61.8	64.8	66.8	61.0	61.9	76.1	78.5	66.4	68.9
MME	ResNet-34	-	73.5	-	70.3	-	72.8	-	62.8	-	68.0	-	63.8	-	79.2	-	70.1
BiAT	ResNet-34	73.0	74.9	68.0	68.8	71.6	74.6	57.9	61.5	63.9	67.5	58.5	62.1	77.0	78.6	67.1	69.7
FAN	ResNet-34	70.4	76.6	70.8	72.1	72.9	76.7	56.7	63.1	64.5	66.1	63.0	67.8	76.6	79.4	67.6	71.7
ECACL-T	ResNet-34	73.5	76.4	72.8	74.3	72.8	75.9	<b>65.1</b>	65.3	70.3	72.2	64.8	68.6	78.3	79.7	71.1	73.2
ECACL-P	ResNet-34	<b>75.3</b>	<b>79.0</b>	<b>74.1</b>	<b>77.3</b>	<b>75.3</b>	<b>79.4</b>	65.0	<b>70.6</b>	<b>72.1</b>	<b>74.6</b>	<b>68.1</b>	<b>71.6</b>	<b>79.7</b>	<b>82.4</b>	<b>72.8</b>	<b>76.4</b>

Table 1. Results on the *DomainNet* dataset. Best results are in **bold**.

of our experiments with the variants based on MME [38]<sup>1</sup>. Note that although MME is proposed for SSDA, it can be viewed as an UDA method that naively merges labeled target samples into labeled source samples for the cross-entropy loss optimization. Since we do this in the same way, MME is still compatible with our framework. We adopt exactly the same training procedures and hyper-parameters as MME, except the way to sample labeled data in each mini-batch. Rather than naively sampling data across the whole labeled set, we perform class-balanced sampling: In each mini-batch, we randomly sample  $M$  classes with  $N_s$  and  $N_t$  images for each class from source and target domains, respectively. We set  $M = 10$ ,  $N_s = 10$ ,  $N_t = 1$  for 1-shot setting, and  $N_t = 3$  for 3-shot setting. We set the balancing hyper-parameters of different losses as  $\lambda_1 = 0.1$  and  $\lambda_2 = 1$  in Eq. (12), and the confident threshold  $\sigma = 0.8$  for pseudo labeling (Eq. (11)) for all experiments. For ECACL-T, we set the margin parameter  $m = 1.0$  for all experiments.

#### 4.1. Comparative Results

We compare with the following methods, DANN [11], ADR [39], CDAN [28], ENT [14], MME [38], FAN [18], BiAT [17], and Meta-MME [21]. All these methods are either specifically designed or tailored to address the SSDA problem. We also report results of the baseline method “ST” which trains models with labeled samples from source and

target domains, without domain alignment.

**DomainNet.** We first compare the two variants of ECACL. We can see from Table 1 that while ECACL-T performs slightly better than ECACL-P with the AlexNet backbone, ECACL-P gets much better performance than ECACL-T with the ResNet-34 backbone. This shows that the prototype based categorical domain alignment technique is more effective than the triplet loss based one. One possible reason is that the former considers discrimination over all classes, which is more beneficial than the latter where only the triplet-wise relationship is modeled. For ease of evaluation, we only evaluate ECACL-P for the rest experiments.

Comparing with other methods, we can see that ECACL-P shows significant advantages for all the experimental settings. Specifically, with AlexNet as the feature extraction model, ECACL-P attains 8.6 and 9.4 point gains for the 1-shot and 3-shot settings, respectively, over MME which ECACL-P is based on. With ResNet-34, the improvements are 6.4 and 7.5 for the 1-shot and 3-shot settings, respectively. ECACL-P also performs significantly better than the most recent methods in both settings with both backbone networks. These results substantiate the effectiveness of ECACL-P on mitigating domain shifts by comprehensively exploring landmarks.

**VisDA2017.** We can see from Table 2 that ECACL-P, built upon on MME, reaches significant gains over MME. The average gains are 12.7 for the 1-shot setting and 12.1 for the 3-shot setting. These tremendous improvements convinc-

<sup>1</sup>Unless otherwise specified, we use “ECACL-P” and “ECACL-T” to represent the variants in short.

Method	plane	bcycl	bus	car	horse	knife	mcycl	person	plant	sktbrd	train	truck	MCA
<b>1-shot</b>													
ST	82.6	52.8	75.0	57.6	72.7	39.7	80.5	53.3	59.0	64.1	77.5	12.6	60.6
MME	86.6	60.1	80.8	61.9	84.0	69.6	87.0	72.4	73.0	50.9	79.4	<b>14.7</b>	68.4
ECACL-P	<b>94.9</b>	<b>81.5</b>	<b>88.9</b>	<b>81.3</b>	<b>95.9</b>	<b>92.4</b>	<b>92.2</b>	<b>83.3</b>	<b>95.2</b>	<b>77.4</b>	<b>88.4</b>	2.3	<b>81.1</b>
<b>3-shot</b>													
ST	74.0	71.7	71.2	64.7	78.5	71.8	69.6	51.4	73.7	49.4	80.8	19.8	64.7
MME	87.2	67.3	74.9	64.5	86.9	85.5	78.8	75.8	84.4	48.0	80.8	<b>19.9</b>	71.2
ECACL-P	<b>95.9</b>	<b>82.9</b>	<b>88.6</b>	<b>84.9</b>	<b>95.9</b>	<b>92.1</b>	<b>93.3</b>	<b>83.7</b>	<b>95.4</b>	<b>79.3</b>	<b>88.0</b>	19.5	<b>83.3</b>

Table 2. Results on the *VisDA2017* dataset.

	R to C	R to P	R to A	P to R	P to C	P to A	A to P	A to C	A to R	C to R	C to A	C to P	Mean
<b>One-shot</b>													
S+T	37.5	63.1	44.8	54.3	31.7	31.5	48.8	31.1	53.3	48.5	33.9	50.8	44.1
DANN	42.5	64.2	45.1	56.4	36.6	32.7	43.5	34.4	51.9	51.0	33.8	49.4	45.1
ADR	37.8	63.5	45.4	53.5	32.5	32.2	49.5	31.8	53.4	49.7	34.2	50.4	44.5
CDAN	36.1	62.3	42.2	52.7	28.0	27.8	48.7	28.0	51.3	41.0	26.8	49.9	41.2
ENT	26.8	65.8	45.8	56.3	23.5	21.9	47.4	22.1	53.4	30.8	18.1	53.6	38.8
MME	42.0	69.6	48.3	58.7	37.8	34.9	52.5	36.4	57.0	54.1	39.5	59.1	49.2
ECACL-P	<b>50.3</b>	<b>70.71</b>	<b>52.2</b>	<b>61.4</b>	<b>41.2</b>	<b>39.3</b>	<b>57.8</b>	<b>39.1</b>	<b>59.1</b>	<b>55.8</b>	<b>41.7</b>	<b>59.9</b>	<b>52.4</b>
<b>Three-shot</b>													
S+T	44.6	66.7	47.7	57.8	44.4	36.1	57.6	38.8	57.0	54.3	37.5	57.9	50.0
DANN	47.2	66.7	46.6	58.1	44.4	36.1	57.2	39.8	56.6	54.3	38.6	57.9	50.3
ADR	45.0	66.2	46.9	57.3	38.9	36.3	57.5	40.0	57.8	53.4	37.3	57.7	49.5
CDAN	41.8	69.9	43.2	53.6	35.8	32.0	56.3	34.5	53.5	49.3	27.9	56.2	46.2
ENT	44.9	70.4	47.1	60.3	41.2	34.6	60.7	37.8	60.5	58.0	31.8	63.4	50.9
MME	51.2	73.0	50.3	61.6	47.2	40.7	63.9	43.8	61.4	59.9	44.7	64.7	55.2
FAN	51.9	74.6	51.2	61.6	47.9	42.1	65.5	44.5	60.9	58.1	44.3	64.8	55.6
ECACL-P	<b>55.4</b>	<b>75.7</b>	<b>56.0</b>	<b>67.0</b>	<b>52.5</b>	<b>46.4</b>	<b>67.4</b>	<b>48.5</b>	<b>66.3</b>	<b>60.8</b>	<b>45.9</b>	<b>67.3</b>	<b>59.1</b>

Table 3. Results on the *Office-Home* dataset.

CA		✓		✓	✓	✓	✓	✓
SA			✓	✓	✓	✓	✓	✓
CONA			✓	✓	✓	✓	✓	✓
	37.9	39.2	39.3	43.5	42.2	45.2	44.1	48.4

Table 4. Ablation study for the adaptation from *Real* to *Sketch* on the *DomainNet* dataset for the 3-shot setting. The second column shows the baseline result obtained by MME.

ingly evidence the effectiveness of ECACL-P.

**Office-Home.** We can see from Table 3 that ECACL-P attains remarkable performance boosts over existing methods as well, although the gains are not as significant as those on the other two datasets. A possible reason is that this dataset is harder than the other two such that improvements are more difficult to attain.

## 4.2. Additional Empirical Analysis

**Ablation study.** Built upon existing UDA methods, ECACL-P includes the following new modules/techniques to address the SSDA problem, namely, the Prototypical Alignment (PA) module, the Strong Augmentation (SA) technique that aims to enhance PA and the Consistency Alignment (CONA) module. Table 4 shows the ablation study on the adaptation from *Real* to *Sketch* on the *DomainNet* dataset for the 3-shot setting. We can see that all the

Method	Net	Setting	MCA	
Source-only	ResNet-101	UDA	52.4	
HAFN	ResNet-101	UDA	73.9	
SAFN	ResNet-101	UDA	76.1	
			1-shot	3-shot
HAFN + ST	ResNet-101	SSDA	77.0	79.3
SAFN + ST	ResNet-101	SSDA	77.5	79.2
ECACL-P (HAFN)	ResNet-101	SSDA	<b>83.9</b>	<b>85.3</b>
ECACL-P (SAFN)	ResNet-101	SSDA	83.3	84.5
ST	ResNet-34	SSDA	60.6	64.7
MME	ResNet-34	SSDA	68.4	71.2
ECACL-P (MME)	ResNet-34	SSDA	<b>81.1</b>	<b>83.3</b>

Table 5. Flexibility analysis of ECACL-P on the *VisDA2017* dataset. “HAFN + ST” and “SAFN + ST” denote the naive extensions of the methods HAFN and SAFN from UDA to SSDA, by further including labeled target data for the cross-entropy loss optimization. “ECACL-P (MME)”, “ECACL-P (HAFN)”, and “ECACL-P (SAFN)” are the methods corresponding to different UDA methods incorporated into our framework.

new modules/techniques contribute to the ultimate performance promotions, thus verifying the efficacy.

**Plug-and-play evaluation.** As mentioned above, ECACL-P is agnostic to the UDA methods built upon. To evaluate this, we apply ECACL-P as a plug-and-play component on




Figure 2. (a): Accuracy with different numbers of labeled samples per class in target domain. (b): Sensitivity analysis with respect to the confident threshold  $\sigma$  (Eq. (11)). (c): t-SNE visualization of the learned features of 3 randomly selected classes. Different colors represent different classes. Shapes “Y”, “★” and “o” represent source samples, labeled target samples, and unlabeled target samples, respectively.

different existing UDA methods and see how much adaptation performance can be improved. Table 5 shows the results of three UDA methods before and after empowered by ECACL-P, namely HAFN [48], SAFN [48], and MME [38]. We can see from Table 5 that ECACL-P indeed significantly promotes the performance of existing UDA methods and their naive SSDA extensions. For example, “ECACL-P (HAFN)” raises the result of the UDA method HAFN from 73.9 to 83.9 with one sample per class labeled and further to 85.3 with 3 samples per class labeled. These significant and consistent improvements with different UDA methods, different backbone networks and different numbers of landmarks convincingly substantiate the effectiveness of ECACL-P on boosting adaptation performance with minimal labeling effort.

**Impact of the number of landmarks.** We have shown the superior performance of the standard 1-shot and 3-shot settings above. We wonder how performance changes with the number of landmarks increased. We study this on the adaptation from *Real* to *Sketch* on *DomainNet*. We can see from Fig. 2 (a) that all the three methods enjoy performance boosts with more target samples labeled. Comparatively speaking, ECACL-P consistently reaches the best performance for all the cases. This substantiates the benefit of our method for flexibly exploiting different amount of labeled target samples to help address the domain shift problem.

**Parameter analysis.** Our consistency alignment module involves an important hyper-parameter  $\sigma$  which is the confident threshold for pseudo labeling. We analyze the sensitivity regarding  $\sigma$  on the adaptation from *Real* to *Sketch* on the *DomainNet* dataset under the 3-shot setting. The result is shown in Figure 2 (b). We see that ECACL-P is not very sensitive to  $\sigma$  and maintains good performance when  $\sigma$  is fairly high (over 0.65).

**Results with different splits.** We follow the prior method [38] and use the provided labeled/unlabeled splits of target domains for experiments. To study the impact of randomness on the performance, we regenerate the splits used for the adaptation from the *real* domain to *sketch* domain in the

		Split 1	Split 2	Split 3
AlexNet	MME	37.9	41.2	43.0
	ECACL-P	48.4	48.9	51.6
ResNet-34	MME	61.9	65.3	65.2
	ECACL-P	71.6	72.5	71.8

Table 6. Variance analysis with different dataset splits.

*DomainNet* dataset for the 3-shot settings. Table 6 shows the results with three different splits. We can see that the proposed ECACL-P consistently improves MME with different splits.

**Feature visualization.** To qualitatively evaluate the alignment results, we plot the t-SNE [31] visualization of the features produced by ECACL-P for the adaptation from *Real* to *Sketch* on *DomainNet* in the 3-shot setting. We can see from Figure 2(c) that the learned features exhibit favorable clustering structure. Features from different domains are close if they belong to the same classes and apart otherwise. This plot further supports that both feature discriminability and domain alignment are achieved by the learned model.

## 5. Conclusion

We propose in this paper a novel semi-supervised domain adaptation (SSDA) framework within which existing unsupervised domain adaptation (UDA) methods can effectively utilize a few labeled samples from the target domain to further mitigate domain shifts. The proposed framework includes two categorical alignment techniques, both of which are further enhanced by a data augmentation based technique that produces highly perturbed images to mitigate overfitting. A consistency alignment module is incorporated into the framework which enforces consistency regularization on the learned model. Ablation study verifies the contributing roles of all the above alignment components. Experiments show that the proposed framework reaches state-of-the-art SSDA performance and consistently promotes the adaptation performance of various UDA methods for different numbers of labeled target samples.

**Acknowledgement:** The work was partially supported by the National Science Foundation Award ECCS-1916839.



## References

- [1] Shuang Ao, Xiang Li, and Charles X Ling. Fast generalized distillation for semi-supervised domain adaptation. In *AAAI*, 2017. 1, 2
- [2] David Berthelot, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. Remixmatch: Semi-supervised learning with distribution matching and augmentation anchoring. In *ICLR*, 2019. 4
- [3] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin Raffel. Mixmatch: A holistic approach to semi-supervised learning. *arXiv preprint arXiv:1905.02249*, 2019. 4
- [4] Fabio Maria Carlucci, Lorenzo Porzi, Barbara Caputo, Elisa Ricci, and Samuel Rota Bulò. Autodial: Automatic domain alignment layers. In *ICCV*, 2017. 2
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020. 2
- [6] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation strategies from data. In *CVPR*, 2019. 4
- [7] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical data augmentation with no separate search. *arXiv preprint arXiv:1909.13719*, 2019. 2, 4
- [8] Zihang Dai, Zhilin Yang, Fan Yang, William W. Cohen, and Ruslan Salakhutdinov. Good semi-supervised learning that requires a bad GAN. In *NIPS*, 2017. 2
- [9] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017. 4
- [10] Jeff Donahue, Judy Hoffman, Erik Rodner, Kate Saenko, and Trevor Darrell. Semi-supervised domain adaptation with instance constraints. In *CVPR*, 2013. 1, 2
- [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. *J. Mach. Learn. Res.*, 17:59:1–59:35, 2016. 2, 6
- [12] Boqing Gong, Kristen Grauman, and Fei Sha. Connecting the dots with landmarks: Discriminatively learning domain-invariant features for unsupervised domain adaptation. In *ICML*, 2013. 2
- [13] Boqing Gong, Yuan Shi, Fei Sha, and Kristen Grauman. Geodesic flow kernel for unsupervised domain adaptation. In *CVPR*, 2012. 1, 2
- [14] Yves Grandvalet and Yoshua Bengio. Semi-supervised learning by entropy minimization. In *NeurIPS*, 2005. 6
- [15] Alexander Hermans, Lucas Beyer, and Bastian Leibe. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*, 2017. 4
- [16] Jiayuan Huang, Alexander J. Smola, Arthur Gretton, Karsten M. Borgwardt, and Bernhard Schölkopf. Correcting sample selection bias by unlabeled data. In *NeurIPS*, 2006. 2
- [17] Pin Jiang, Aming Wu, Yahong Han, Yunfeng Shao, Meiyu Qi, and Bingshuai Li. Bidirectional adversarial training for semi-supervised domain adaptation. In *IJCAI*, 2020. 6
- [18] Taekyung Kim and Changick Kim. Attract, perturb, and explore: Learning a feature alignment network for semi-supervised domain adaptation. In *ECCV*, 2020. 1, 2, 6
- [19] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *CoRR*, 2016. 2
- [20] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. *CoRR*, abs/1610.02242, 2016. 2
- [21] Da Li and Timothy Hospedales. Online meta-learning for multi-source and semi-supervised domain adaptation. In *ECCV*, 2020. 6
- [22] Kai Li, Martin Renqiang Min, Bing Bai, Yun Fu, and Hans Peter Graf. On novel object recognition: A unified framework for discriminability and adaptability. In *CIKM*, 2019. 2
- [23] Kai Li, Martin Renqiang Min, and Yun Fu. Rethinking zero-shot learning: A conditional visual classification perspective. In *ICCV*, 2019. 1
- [24] Kai Li, Curtis Wigington, Chris Tensmeyer, Handong Zhao, Nikolaos Barmpalios, Vlad I Morariu, Varun Manjunatha, Tong Sun, and Yun Fu. Cross-domain document object detection: Benchmark suite and method. In *CVPR*, 2020. 2
- [25] Kai Li, Yulun Zhang, Kunpeng Li, and Yun Fu. Adversarial feature hallucination networks for few-shot learning. In *CVPR*, 2020. 2
- [26] Xinzhe Li, Qianru Sun, Yaoyao Liu, Qin Zhou, Shibao Zheng, Tat-Seng Chua, and Bernt Schiele. Learning to self-train for semi-supervised few-shot classification. In *NeurIPS*, 2019. 2
- [27] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *ICML*, 2015. 2
- [28] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. In *NeurIPS*, 2018. 1, 2, 6
- [29] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I. Jordan. Deep transfer learning with joint adaptation networks. In *ICML*, 2017. 2
- [30] Zelun Luo, Yuliang Zou, Judy Hoffman, and Fei-Fei Li. Label efficient learning of transferable representations across domains and tasks. In *NeurIPS*, 2017. 2
- [31] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008. 8
- [32] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, Ken Nakae, and Shin Ishii. Distributional smoothing by virtual adversarial examples. In *ICLR*, 2016. 2
- [33] Saeid Motiian, Marco Piccirilli, Donald A Adjeroh, and Gianfranco Doretto. Unified deep supervised domain adaptation and generalization. In *ICCV*, 2017. 2
- [34] Sinno Jialin Pan, Ivor W. Tsang, James T. Kwok, and Qiang Yang. Domain adaptation via transfer component analysis. *IEEE Trans. Neural Networks*, 22(2):199–210, 2011. 1, 2
- [35] Yingwei Pan, Ting Yao, Yehao Li, Yu Wang, Chong-Wah Ngo, and Tao Mei. Transferrable prototypical networks for unsupervised domain adaptation. In *CVPR*, 2019. 4, 5

- [36] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *ICCV*, 2019. 5
- [37] Xingchao Peng, Ben Usman, Neela Kaushik, Judy Hoffman, Dequan Wang, and Kate Saenko. Visda: The visual domain adaptation challenge. *arXiv preprint arXiv:1710.06924*, 2017. 5
- [38] Kuniaki Saito, Donghyun Kim, Stan Sclaroff, Trevor Darrell, and Kate Saenko. Semi-supervised domain adaptation via minimax entropy. In *ICCV*, 2019. 1, 2, 5, 6, 8
- [39] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. Adversarial dropout regularization. In *ICLR*, 2017. 6
- [40] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *NeurIPS*, 2017. 1, 2, 3
- [41] Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Han Zhang, and Colin Raffel. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In *NeurIPS*, 2020. 2, 4
- [42] Ajinkya Tejankar and Hamed Pirsiavash. A simple baseline for domain adaptation using rotation prediction. *CoRR*, abs/1912.11903, 2019. 2
- [43] Eric Tzeng, Judy Hoffman, Trevor Darrell, and Kate Saenko. Simultaneous deep transfer across domains and tasks. In *ICCV*, 2015. 1, 2
- [44] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *CVPR*, 2017. 5
- [45] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *NeurIPS*, 2016. 2
- [46] Qizhe Xie, Zihang Dai, Eduard Hovy, Minh-Thang Luong, and Quoc V Le. Unsupervised data augmentation for consistency training. In *NeurIPS*, 2020. 2, 4
- [47] Minghao Xu, Hang Wang, Bingbing Ni, Qi Tian, and Wenjun Zhang. Cross-domain detection via graph-induced prototype alignment. In *CVPR*, 2020. 4
- [48] Ruijia Xu, Guanbin Li, Jihan Yang, and Liang Lin. Larger norm more transferable: An adaptive feature norm approach for unsupervised domain adaptation. In *ICCV*, 2019. 5, 8
- [49] Xiang Xu, Xiong Zhou, Ragav Venkatesan, Gurumurthy Swaminathan, and Orchid Majumder. d-sne: Domain adaptation using stochastic neighborhood embedding. In *CVPR*, 2019. 2
- [50] Ting Yao, Yingwei Pan, Chong-Wah Ngo, Houqiang Li, and Tao Mei. Semi-supervised domain adaptation with subspace learning for visual recognition. In *CVPR*, 2015. 1, 2
- [51] Yang Zou, Zhiding Yu, Xiaofeng Liu, BVK Kumar, and Jinsong Wang. Confidence regularized self-training. In *CVPR*, 2019. 4, 5