

# Prototypical Cross-domain Self-supervised Learning for Few-shot Unsupervised Domain Adaptation

Xiangyu Yue<sup>1,\*</sup> Zangwei Zheng<sup>2,\*</sup> Shanghang Zhang<sup>1</sup> Yang Gao<sup>3</sup>  
Trevor Darrell<sup>1</sup> Kurt Keutzer<sup>1</sup> Alberto Sangiovanni Vincentelli<sup>1</sup>  
<sup>1</sup>UC Berkeley <sup>2</sup>Nanjing University <sup>3</sup>Tsinghua University

## Abstract

Unsupervised Domain Adaptation (UDA) transfers predictive models from a fully-labeled source domain to an unlabeled target domain. In some applications, however, it is expensive even to collect labels in the source domain, making most previous works impractical. To cope with this problem, recent work performed instance-wise cross-domain self-supervised learning, followed by an additional fine-tuning stage. However, the instance-wise self-supervised learning only learns and aligns low-level discriminative features. In this paper, we propose an end-to-end Prototypical Cross-domain Self-Supervised Learning (PCS) framework for Few-shot Unsupervised Domain Adaptation (FUDA)<sup>1</sup>. PCS not only performs cross-domain low-level feature alignment, but it also encodes and aligns semantic structures in the shared embedding space across domains. Our framework captures category-wise semantic structures of the data by in-domain prototypical contrastive learning; and performs feature alignment through cross-domain prototypical self-supervision. Compared with state-of-the-art methods, PCS improves the mean classification accuracy over different domain pairs on FUDA by 10.5%, 3.5%, 9.0%, and 13.2% on Office, Office-Home, VisDA-2017, and DomainNet, respectively.

## 1. Introduction

Deep Learning has achieved remarkable performance in various computer vision tasks, such as image classification [30, 32] and semantic segmentation [43, 77, 8]. Despite high accuracy, deep neural networks trained on specific datasets often fail to generalize to new domains owing to the *domain shift* problem [66, 15, 67]. Unsupervised domain adaptation (UDA) transfers predictive models from a fully-labeled source domain to an unlabeled target domain. Although it is challenging with no label information in the

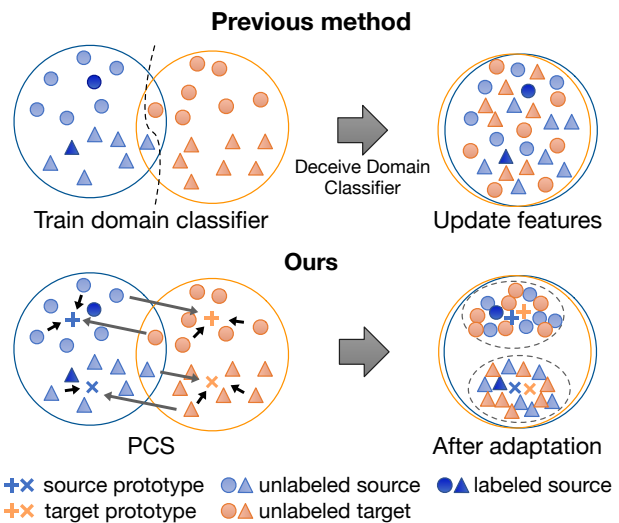


Figure 1: We address the task of few-shot unsupervised domain adaptation. Top: Existing domain-classifier based methods align source and target distributions but fail to extract discriminative features due to lack of labeled data. Bottom: Our method estimates prototypes for in-domain and cross-domain self-supervised learning to extract domain-aligned discriminative features.

target domain, many UDA methods [67, 31, 44, 18] could achieve high accuracy on the target domain using the abundant explicit supervision in source domain, together with the unlabeled target samples for domain alignment.

In some real-world applications, however, providing large-scale annotations even in the source domain is often challenging due to the high cost and difficulty of annotation. Taking medical imaging for instance, each image of the Diabetic Retinopathy dataset [28] is annotated by a panel of 7 or 8 U.S. board-certified ophthalmologists, with a total group of 54 doctors. Thus practically it is too stringent to assume the availability of source data with abundant labels.

In this paper, to cope with the labeling costs of the source domain, we instead consider a few-shot unsupervised domain adaptation (FUDA) setting, where only an extremely small fraction of source samples are labeled, while

\*Equal contribution; correspondence to xyue@berkeley.edu

<sup>1</sup>Project page: <http://xyue.io/pcs-fuda>

all the rest source and target samples remain unlabeled. Most state-of-the-art UDA methods align source and target features by minimizing some form of distribution distances [44, 45, 63, 18], and learn the discriminative representation by minimizing the supervision loss on fully-labeled source domain data. In FUDA, however, since we have a very limited number of labeled source samples, it is much harder to learn discriminative features in the source domain, not to mention in the target domain.

Several recent papers [29, 10, 27, 52, 72] on self-supervised learning (SSL) present promising representation learning results on images from a single domain and [39] further extended to perform SSL across two domains for better domain adaptation performance. Despite the improved performance, the instance-based method in [39] has some fundamental weaknesses. First, the semantic structure of the data is not encoded by the learned structure. This is because the in-domain self-supervision in [39] treats two instances as negative pairs as long as they are from different samples, regardless of the semantic similarity. Consequently, many instances sharing the same semantic are undesirably pushed apart in the feature space. Second, the cross-domain instance-to-instance matching in [39] is very sensitive to abnormal samples. Imagine a case where the embeddings of source and target samples are far apart (*i.e.* the domain gap is large) and one abnormal source sample is mapped closer to all target samples than any other source sample. Then the method in [39] would match all target samples to the same source sample (cf. Figure 3). For a given sample, the matched sample in the other domain may change drastically during the training procedure, making the optimization harder to converge. Third, the two-stage pipeline (*i.e.* SSL followed by domain adaptation) is complicated and experiments show that the optimal DA methods for different datasets are different. As a result, the training is rather cumbersome and it is unclear how to choose the optimal DA method in the second stage for different datasets.

In this paper, we propose *Prototypical Cross-domain Self-supervised learning*, a novel single-stage framework for FUDA that unifies representation learning and domain alignment with few-shot labeled source samples. PCS contains three major components to learn both discriminative and domain-invariant features. First, PCS performs in-domain prototypical self-supervision to implicitly encode the semantic structure of data into the embedding space. This is motivated by [41], but we further leverage the known semantic information of the task and learn better semantic structure in each domain. Second, PCS performs cross-domain instance-to-prototype matching to transfer knowledge from source to the target in a more robust manner. Instead of instance-to-instance matching, matching a sample to a prototype (*i.e.* representative embedding for a group of instances that are semantically similar) is more robust

to abnormal instances in the other domain and makes the optimization converge faster and more smoothly. Third, PCS unifies prototype learning with cosine classifier and update cosine classifier adaptively with source and target prototypes. transfers from source prototypes to target prototypes for better performance on the target domain. In order to further mitigate the effect of cross-domain mismatching, we perform entropy maximization to obtain a more diversified output. We show that together with entropy minimization, this is equivalent to maximizing the mutual information (MI) between input image and the network prediction.

To summarize, our contributions are three-fold:

- We propose a novel Prototypical Cross-domain Self-supervised learning framework (PCS) for few-shot unsupervised Domain Adaptation.
- We propose to leverage prototypes to perform better semantic structure learning, discriminative feature learning, and cross-domain alignment in a unified, unsupervised and adaptive manner.
- While it is hard to choose the optimal domain adaptation method in the complex two-stage framework [39], PCS can be easily trained in an end-to-end manner, and outperforms all state-of-the-art methods by a large margin across multiple benchmark datasets.

## 2. Related Work

**Domain Adaptation.** Unsupervised Domain Adaptation (UDA) [24] addresses the problem of transferring knowledge from a fully-labeled source domain to an unlabeled target domain. Most UDA methods have focused on feature distribution alignment. Discrepancy-based methods explicitly compute the Maximum Mean Discrepancy (MMD) [26] between the source and the target to align the two domains [44, 68, 46]. Long *et al.* [47] proposed to align the joint distributions using the Joint MMD criterion. Sun *et al.* [63] and Zhuo *et al.* [79] further proposed to align second-order statistics of source and target features. With the development of Generative Adversarial Networks [23], additional papers proposed to perform domain alignment in the feature space with adversarial learning [17, 67, 31, 73, 45, 61]. Recently, image translation methods, *e.g.* [78, 42] have been adopted to further improve domain adaptation by performing pixel-level alignment [31, 5, 57, 50, 75, 60, 62]. Instead of explicit feature alignment, Saito *et al.* [59] proposed minimax entropy for adaptation. While these methods have full supervision on the source domain, similar to [39], we focus on a new adaptation setting with few source labels.

**Self-supervised Learning.** Self-supervised learning (SSL) is a subset of unsupervised learning methods where supervision is automatically generated from the data [36, 13, 76, 51, 22, 70]. One of the most common strategies for

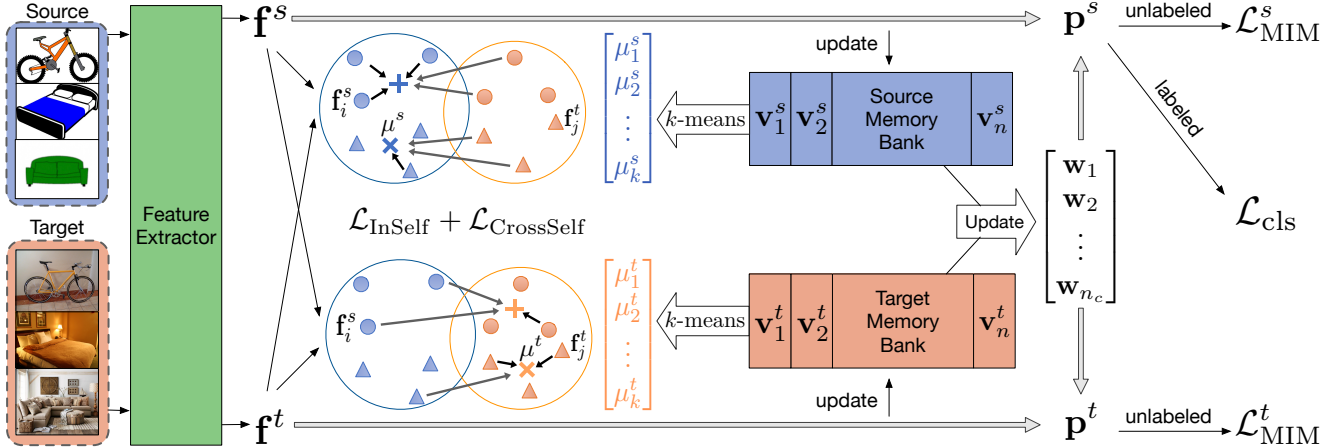


Figure 2: An overview of the PCS framework. In-domain and cross-domain self-supervision are performed between normalized feature vectors  $\mathbf{f}$  and prototypes  $\mu$  computed by clustering vectors  $\mathbf{v}$  in memory banks. Features with confident predictions ( $\mathbf{p}$ ) are used to adaptively update classifier vectors  $\mathbf{w}$ . MI maximization and classification loss are further used to extract discriminative features.

SSL is handcrafting auxiliary pretext tasks predicting future, missing or contextual information. In particular, image colorization [76, 40], patch location prediction [13, 14], image jigsaw puzzle [51], image inpainting [53] and geometric transformations [22, 16] have been shown to be helpful. Currently, contrastive learning [3, 29, 52, 65, 49] has achieved state-of-the-art performance on representation learning [27, 10, 12, 11]. Most contrastive methods are instance-wise, aiming to learn an embedding space where samples from the same instance are pulled closer and samples from different instances are pushed apart [72, 10]. Recently, contrastive learning based on prototypes has shown promising results in representation learning [41, 2, 7, 19].

**Self-supervised Learning for Domain Adaptation.** Self-supervision-based methods incorporate SSL losses into the original task network [20, 21]. Reconstruction was first utilized as self-supervised task in some early works [20, 21], in which source and target share the same encoder to extract domain-invariant features. To capture both domain-specific and shared properties, Bousmalis *et al.* [5] explicitly extracts image representations into two spaces, one private for each domain and one shared across domains. In [6], solving jigsaw puzzle [51] was leveraged as a self-supervision task to solve domain adaptation and generalization. Sun *et al.* [64] further proposed to perform adaptation by jointly learning multiple self-supervision tasks. The feature encoder is shared by both source and target images, and the extracted features are then fed into different self-supervision task heads. Recently, based on instance discrimination [72], Kim *et al.* [39] proposed a cross-domain SSL approach for adaptation with few source labels. SSL has also been incorporated for adaptation in other fields, including point cloud recognition [1], medical imaging [33], action segmentation [9], robotics [34], facial tracking [74], *etc.*

### 3. Approach

In few-shot unsupervised domain adaptation, we are given a very limited number of labeled source images  $\mathcal{D}_s = \{(\mathbf{x}_i^s, y_i^s)\}_{i=1}^{N_s}$ , as well as unlabeled source images  $\mathcal{D}_{su} = \{(\mathbf{x}_i^{su})\}_{i=1}^{N_{su}}$ . In the target domain, we are only given unlabeled target images  $\mathcal{D}_{tu} = \{(\mathbf{x}_i^{tu})\}_{i=1}^{N_{tu}}$ . The goal is to train a model on  $\mathcal{D}_s, \mathcal{D}_{su}$ , and  $\mathcal{D}_{tu}$ ; and evaluate on  $\mathcal{D}_{tu}$ .

The base model consists of a feature encoder  $F$ , a  $\ell_2$  normalization layer, which outputs a normalized feature vector  $\mathbf{f} \in \mathbb{R}^d$  and a cosine similarity-based classifier  $C$ .

#### 3.1. In-domain Prototypical Contrastive Learning

We learn a shared feature encoder  $F$  that extracts discriminative features in both domains. Instance Discrimination [72] is employed in [39] to learn discriminative features. As an instance-wise contrastive learning method, it results in an embedding space where all instances are well separated. Despite promising results, instance discrimination has a fundamental weakness: the semantic structure of the data is not encoded by the learned representations. This is because two instances are treated as negative pairs as long as they are from different samples, regardless of their semantics. For a single domain, ProtoNCE [41] is proposed to learn semantic structure of the data by performing iterative clustering and representation learning. The goal is to drive features within the same cluster to become more aggregated and features in different clusters further apart.

However, naively applying ProtoNCE to  $\mathcal{D}_s \cup \mathcal{D}_{su} \cup \mathcal{D}_{tu}$  in our domain adaptation setting would cause potential problems. Primarily due to the domain shift, images of different classes from different domains can be incorrectly aggregated into the same cluster, and images of the same class from different domains can be mapped into clusters

that are far apart. To mitigate these problems, we propose to perform prototypical contrastive learning separately in  $\mathcal{D}_s \cup \mathcal{D}_{su}$  and in  $\mathcal{D}_{tu}$ . This aims to prevent the incorrect clustering of images across domains and indiscriminative feature learning.

Specifically, two memory banks  $\mathbf{V}^s$  and  $\mathbf{V}^t$  are maintained for source and target respectively:

$$\mathbf{V}^s = [\mathbf{v}_1^s, \dots, \mathbf{v}_{(N_s+N_{su})}^s], \mathbf{V}^t = [\mathbf{v}_1^t, \dots, \mathbf{v}_{N_{tu}}^t], \quad (1)$$

where  $\mathbf{v}_i$  is the stored feature vector of  $\mathbf{x}_i$ , initialized with  $\mathbf{f}_i$  and updated with a momentum  $m$  after each batch:

$$\mathbf{v}_i \leftarrow m\mathbf{v}_i + (1 - m)\mathbf{f}_i. \quad (2)$$

In order for in-domain prototypical contrastive learning,  $k$ -means clustering is performed on  $\mathbf{V}^s$  and  $\mathbf{V}^t$  to get source clusters  $\mathcal{C}^s = \{C_1^{(s)}, C_2^{(s)}, \dots, C_k^{(s)}\}$  and similarly  $\mathcal{C}^t$  with normalized source prototypes  $\{\mu_j^s\}_{j=1}^k$  and normalized target prototypes  $\{\mu_j^t\}_{j=1}^k$ . Specifically,  $\mu_j^s = \frac{\mathbf{u}_j^s}{\|\mathbf{u}_j^s\|}$ , where  $\mathbf{u}_j^s = \frac{1}{|C_j^{(s)}|} \sum_{\mathbf{v}_i^s \in C_j^{(s)}} \mathbf{v}_i^s$ . We explain only on the source domain for succinct notation, all operations are performed on target similarly.

During training, with the feature encoder  $F$ , we compute a feature vector  $\mathbf{f}_i^s = F(\mathbf{x}_i^s)$ . To perform in-domain prototypical contrastive learning, we compute the similarity distribution vector between  $\mathbf{f}_i^s$  and  $\{\mu_j^s\}_{j=1}^k$  as  $P_i^s = [P_{i,1}^s, P_{i,2}^s, \dots, P_{i,k}^s]$ , with

$$P_{i,j}^s = \frac{\exp(\mu_j^s \cdot \mathbf{f}_i^s / \phi)}{\sum_{r=1}^k \exp(\mu_r^s \cdot \mathbf{f}_i^s / \phi)}, \quad (3)$$

where  $\phi$  is a temperature value determining the level of concentration. Then the in-domain prototypical contrastive loss can be written as:

$$\mathcal{L}_{\text{PC}} = \sum_{i=1}^{N_s+N_{su}} \mathcal{L}_{\text{CE}}(P_i^s, c_s(i)) + \sum_{i=1}^{N_{tu}} \mathcal{L}_{\text{CE}}(P_i^t, c_t(i)) \quad (4)$$

where  $c_s(\cdot)$  and  $c_t(\cdot)$  return the cluster index of the instance. Due to the randomness in clustering, we perform  $k$ -means on the samples  $M$  times with different number of clusters  $\{k_m\}_{m=1}^M$ . Moreover, in the FUDA setting, since the number of classes  $n_c$  is known, we set  $k_m = n_c$  for most  $m$ . The overall loss for in-domain self-supervision is:

$$\mathcal{L}_{\text{InSelf}} = \frac{1}{M} \sum_{m=1}^M \mathcal{L}_{\text{PC}}^{(m)} \quad (5)$$

### 3.2. Cross-domain Instance-Prototype SSL

In order to explicitly enforce learning domain-aligned and more discriminative features in both source and target




Figure 3: Comparison of cross-domain instance-instance (I-I) matching [39] (left) and our cross-domain instance-prototype (I-P) matching (right). Left: I-I incorrectly matches all orange samples to the same blue sample. Right: I-P robustly matches samples to the correct prototypes.

domains, we perform cross-domain instance-prototype self-supervised learning.

Many previous works focus on domain alignment via discrepancy minimization or adversarial learning. However, these methods provide inferior performance or have unstable training. Moreover, most of them focus on distribution matching, without considering semantic similarity matching across domains. Instance-instance matching [39] is proposed to match an instance  $i$  to another instance  $j$  in the other domain with the most similar representation. However, due to the domain gap, instances can be easily mapped to instances of different classes in the other domain. In some cases, if an outlier in one domain is extremely close to the other domain, it will be matched to all the instances in the other domain, as illustrated in Figure 3.

Instead, our method discovers positive matching as well as negative matchings between instance and cluster prototypes in different domains. To find a matching for an instance  $i$ , we perform entropy minimization on the similarity distribution vector between its representation, e.g.  $\mathbf{f}_i^s$  and the centroids of the other domain, e.g.  $\{\mu_j^t\}_{j=1}^k$ .

Specifically, given feature vector  $\mathbf{f}_i^s$  in the source domain, and centroids  $\{\mu_j^t\}_{j=1}^k$  in the target domain, we first compute the similarity distribution vector  $P_i^{s \rightarrow t} = [P_{i,1}^{s \rightarrow t}, \dots, P_{i,k}^{s \rightarrow t}]$ , in which

$$P_{i,j}^{s \rightarrow t} = \frac{\exp(\mu_j^t \cdot \mathbf{f}_i^s / \tau)}{\sum_{r=1}^k \exp(\mu_r^t \cdot \mathbf{f}_i^s / \tau)}. \quad (6)$$

Then we minimize the entropy of  $P_i^{s \rightarrow t}$ , which is:

$$H(P_i^{s \rightarrow t}) = - \sum_{j=1}^k P_{i,j}^{s \rightarrow t} \log P_{i,j}^{s \rightarrow t}. \quad (7)$$

Similarly, we can compute  $H(P_i^{t \rightarrow s})$ , and the final loss for cross-domain instance-prototype SSL is:

$$\mathcal{L}_{\text{CrossSelf}} = \sum_{i=1}^{N_s+N_{su}} H(P_i^{s \rightarrow t}) + \sum_{i=1}^{N_{tu}} H(P_i^{t \rightarrow s}) \quad (8)$$



### 3.3. Adaptive Prototypical Classifier Learning

The goal of this section is to learn a better domain-aligned, discriminative feature encoder  $F$  and more importantly, a cosine classifier  $C$  that could achieve high accuracy on the target domain.

The cosine classifier  $C$  consists of weight vectors  $\mathbf{W} = [\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_{n_c}]$ , where  $n_c$  denotes the total number of classes, and a temperature  $T$ . The output of  $C$ ,  $\frac{1}{T} \mathbf{W}^T \mathbf{f}$  is fed into a softmax layer  $\sigma$  to obtain the final probabilistic output  $\mathbf{p}(\mathbf{x}) = \sigma(\frac{1}{T} \mathbf{W}^T \mathbf{f})$ . With the availability of the labeled set  $\mathcal{D}_s$ , it is straightforward to train  $F$  and  $C$  with a standard cross-entropy loss for classification:

$$\mathcal{L}_{\text{cls}} = \mathbb{E}_{(\mathbf{x}, y) \in \mathcal{D}_s} \mathcal{L}_{CE}(\mathbf{p}(\mathbf{x}), y) \quad (9)$$

However, since  $\mathcal{D}_s$  is quite small under FUDA setting, only training with  $\mathcal{L}_{\text{cls}}$  is hard to get a  $C$  with high performance on the target.

**Adaptive Prototype-Classifier Update (APCU)** Note that for  $C$  to classify samples correctly, the direction of a weight vector  $\mathbf{w}_i$  needs to be representative of features of the corresponding class  $i$ . This indicates that the meaning of  $\mathbf{w}_i$  coincide with the ideal cluster prototype of class  $i$ . We propose to use an estimate of the ideal cluster prototypes to update  $\mathbf{W}$ . Yet the computed  $\{\mu_j^s\}$  and  $\{\mu_j^t\}$  cannot be naively used for this purpose, not only because the correspondence between  $\{\mathbf{w}_i\}$  and  $\{\mu_j\}$  is unknown, but also the  $k$ -means result may contain very impure clusters, leading to non-representative prototypes.

We use the few-shot labeled data as well as samples with high-confident predictions to estimate the prototype for each class. Formally, we define  $\mathcal{D}_s^{(i)} = \{\mathbf{x} | (\mathbf{x}, y) \in \mathcal{D}_s, y = i\}$  and denote by  $\mathcal{D}_{su}^{(i)}$  and  $\mathcal{D}_{tu}^{(i)}$  the set of samples with high-confident label  $i$  in source and target, respectively. With  $\mathbf{p}(\mathbf{x}) = [\mathbf{p}(\mathbf{x})_1, \dots, \mathbf{p}(\mathbf{x})_{n_c}]$ ,  $\mathcal{D}_{su}^{(i)} = \{\mathbf{x} | \mathbf{x} \in \mathcal{D}_{su}, \mathbf{p}(\mathbf{x})_i > t\}$ , where  $t$  is a confidence threshold; and similarly for  $\mathcal{D}_{tu}^{(i)}$ . Then the estimate of  $\mathbf{w}_i$  from source and target domain can be computed as:

$$\hat{\mathbf{w}}_i^s = \frac{1}{|\mathcal{D}_{s+}^{(i)}|} \sum_{\mathbf{x} \in \mathcal{D}_{s+}^{(i)}} \mathbf{V}^s(\mathbf{x}); \hat{\mathbf{w}}_i^t = \frac{1}{|\mathcal{D}_{tu}^{(i)}|} \sum_{\mathbf{x} \in \mathcal{D}_{tu}^{(i)}} \mathbf{V}^t(\mathbf{x}) \quad (10)$$

where  $\mathcal{D}_{s+}^{(i)} = \mathcal{D}_s^{(i)} \cup \mathcal{D}_{su}^{(i)}$  and  $\mathbf{V}(\mathbf{x})$  returns the representation in memory bank corresponding to  $\mathbf{x}$ .

With only few labeled samples in source, it is hard to learn a representative prototype shared across domains. Instead of directly employing a global prototype for a class  $i$ , we further propose to update  $\mathbf{w}_i$  in an domain adaptive manner, with  $\hat{\mathbf{w}}_i^s$  during early training stage and with  $\hat{\mathbf{w}}_i^t$  at later stage. This is because that  $\hat{\mathbf{w}}_i^s$  is more robust in early training stage due to the few labeled source samples,

while  $\hat{\mathbf{w}}_i^t$  would be more representative later for target domain to get better adaptation performance. Specifically, we use  $|\mathcal{D}_{tu}^{(i)}|$  to determine whether  $\hat{\mathbf{w}}_i^t$  is robust to use:

$$\mathbf{w}_i = \begin{cases} \text{unit}(\hat{\mathbf{w}}_i^s) & \text{if } |\mathcal{D}_{tu}^{(i)}| < t_w \\ \text{unit}(\hat{\mathbf{w}}_i^t) & \text{otherwise} \end{cases} \quad (11)$$

where  $\text{unit}(\cdot)$  normalizes the input vector and  $t_w$  is a threshold hyper-parameter.

**Mutual Information Maximization** In order for the above unified prototype-classifier learning paradigm to work well, the network is desired to have enough confident predictions, e.g.  $|\mathcal{D}^{(i)}| > t_w$ , for all classes to get robust  $\hat{\mathbf{w}}_i^s$  and  $\hat{\mathbf{w}}_i^t$  for  $i = 1, \dots, n_c$ . First, to promote the network to have diversified outputs over the dataset, we maximize the entropy of expected network prediction  $\mathcal{H}(\mathbb{E}_{\mathbf{x} \in \mathcal{D}} [p(y|\mathbf{x}; \theta)])$ , where  $\theta$  denotes learnable parameters in both  $F$  and  $C$ , and  $\mathcal{D} = \mathcal{D}_s \cup \mathcal{D}_{su} \cup \mathcal{D}_{tu}$ . Second, in order to get high-confident prediction for each sample, we leverage entropy minimization on the network output which has shown efficacy in label-scarce scenarios [25, 4]. These two desired behaviors turn out to be equivalent to maximizing the mutual information between input and output:

$$\mathcal{I}(y; \mathbf{x}) = \mathcal{H}(\mathbf{p}_0) - \mathbb{E}_{\mathbf{x}}[\mathcal{H}(p(y|\mathbf{x}; \theta))], \quad (12)$$

where the prior distribution  $\mathbf{p}_0$  is given by  $\mathbb{E}_{\mathbf{x}}[p(y|\mathbf{x}; \theta)]$ , and the detailed derivation is presented in the supplementary material. We can get the objective as:

$$\mathcal{L}_{\text{MIM}} = -\mathcal{I}(y; \mathbf{x}) \quad (13)$$

### 3.4. PCS Learning for FUDA

The PCS learning framework performs in-domain prototypical contrastive learning, cross-domain instance-prototype self-supervised learning, and unified adaptive prototype-classifier learning. Together with APCU in Eq. 11, the overall learning objective is:

$$\mathcal{L}_{PCS} = \mathcal{L}_{\text{cls}} + \lambda_{\text{in}} \cdot \mathcal{L}_{\text{InSelf}} + \lambda_{\text{cross}} \cdot \mathcal{L}_{\text{CrossSelf}} + \lambda_{\text{mim}} \cdot \mathcal{L}_{\text{MIM}} \quad (14)$$

## 4. Experiments

### 4.1. Experimental Setting

**Datasets.** We evaluate our approach on four public datasets and choose labeled images in source domain based on previous work [39]. **Office** [58] is a real-world dataset for domain adaptation tasks. It contains 3 domains (Amazon, DSLR, Webcam) with 31 classes. Experiments are conducted with 1-shot and 3-shots source labels per class in this dataset. **Office-Home** [69] is a more difficult dataset

Table 1: Adaptation accuracy (%) comparison on 1-shot and 3-shots per class on the Office dataset.

Method	Office: Target Acc. on 1-shot / 3-shots						
	A→D	A→W	D→A	D→W	W→A	W→D	Avg
SO	27.5 / 49.2	28.7 / 46.3	40.9 / 55.3	65.2 / 85.5	41.1 / 53.8	62.0 / 86.1	44.2 / 62.7
MME [59]	21.5 / 51.0	12.2 / 54.6	23.1 / 60.2	60.9 / 89.7	14.0 / 52.3	62.4 / 91.4	32.3 / 66.5
CDAN [45]	11.2 / 43.7	6.2 / 50.1	9.1 / 65.1	54.8 / 91.6	10.4 / 57.0	41.6 / 89.8	22.2 / 66.2
SPL [71]	12.0 / 77.1	7.7 / 80.3	7.3 / 74.2	7.2 / 93.5	7.2 / 64.4	10.2 / 91.6	8.6 / 80.1
CAN [38]	25.3 / 48.6	26.4 / 45.3	23.9 / 41.2	69.4 / 78.2	21.2 / 39.3	67.3 / 82.3	38.9 / 55.8
MDDIA [35]	45.0 / 62.9	54.5 / 65.4	55.6 / 67.9	84.4 / 93.3	53.4 / 70.3	79.5 / 93.2	62.1 / 75.5
CDS [39]	33.3 / 57.0	35.2 / 58.6	52.0 / 67.6	59.0 / 86.0	46.5 / 65.7	57.4 / 81.3	47.2 / 69.3
DANN + ENT [18]	32.5 / 57.6	37.2 / 54.1	36.9 / 54.1	70.1 / 87.4	43.0 / 51.4	58.8 / 89.4	46.4 / 65.7
MME + ENT	37.6 / 69.5	42.5 / 68.3	48.6 / 66.7	73.5 / 89.8	47.2 / 63.2	62.4 / 95.4	52.0 / 74.1
CDAN + ENT	31.5 / 68.3	26.4 / 71.8	39.1 / 57.3	70.4 / 88.2	37.5 / 61.5	61.9 / 93.8	44.5 / 73.5
CDS + ENT	40.4 / 61.2	44.7 / 66.7	66.4 / 73.1	71.6 / 90.6	58.6 / 71.8	69.3 / 86.1	58.5 / 74.9
CDS + MME + ENT	39.4 / 61.6	43.6 / 66.3	66.0 / 74.5	75.7 / 92.1	53.1 / 73.0	70.9 / 90.6	58.5 / 76.3
CDS + CDAN + ENT	52.6 / 65.1	55.2 / 68.8	65.7 / 71.2	76.6 / 88.1	59.7 / 71.0	73.3 / 87.3	63.9 / 75.3
CDS / MME + ENT <sup>†</sup>	55.4 / 75.7	57.2 / 77.2	62.8 / 69.7	84.9 / 92.1	62.6 / 69.9	77.7 / 95.4	65.3 / 80.0
CDS / CDAN + ENT <sup>†</sup>	53.8 / 78.1	65.6 / 79.8	59.5 / 70.7	83.0 / 93.2	57.4 / 64.5	77.1 / 97.4	66.1 / 80.6
PCS (Ours)	<b>60.2 / 78.2</b>	<b>69.8 / 82.9</b>	<b>76.1 / 76.4</b>	<b>90.6 / 94.1</b>	<b>71.2 / 76.3</b>	<b>91.8 / 96.0</b>	<b>76.6 / 84.0</b>
Improvement	<b>+4.8 / +0.1</b>	<b>+4.2 / +3.1</b>	<b>+9.7 / +1.9</b>	<b>+5.7 / +0.9</b>	<b>+8.6 / +3.3</b>	<b>+14.1 / -1.4</b>	<b>+10.5 / +3.4</b>

<sup>†</sup> Two-stage training results reported in [39].

Table 2: Performance contribution of each part in PCS framework on Office.

Method	Office: Target Acc. on 1-shot / 3-shots						
	A→D	A→W	D→A	D→W	W→A	W→D	Avg
$\mathcal{L}_{cls}$	27.5 / 49.2	28.7 / 46.3	40.9 / 55.3	65.2 / 85.5	41.1 / 53.8	62.0 / 86.1	44.2 / 62.7
+ $\mathcal{L}_{InSelf}$	39.0 / 55.6	38.6 / 55.1	47.2 / 68.5	71.7 / 89.4	50.9 / 68.4	65.1 / 90.6	52.1 / 71.3
+ $\mathcal{L}_{CrossSelf}$	47.2 / 71.1	52.7 / 70.6	59.0 / 75.5	76.4 / 90.3	58.5 / 74.1	66.9 / 91.8	60.1 / 78.9
+ $\mathcal{L}_{MIM}$	52.8 / 73.5	57.5 / 71.2	67.2 / 76.3	78.9 / 91.4	64.2 / 74.3	68.7 / 92.2	64.9 / 79.8
+APCU (PCS)	<b>60.2 / 78.2</b>	<b>69.8 / 82.9</b>	<b>76.1 / 76.4</b>	<b>90.6 / 94.1</b>	<b>71.2 / 76.3</b>	<b>91.8 / 96.0</b>	<b>76.6 / 84.0</b>
PCS w/o MIM	59.0 / 75.9	58.6 / 76.5	76.2 / 76.4	87.8 / 93.2	68.7 / 74.7	89.8 / 95.0	73.5 / 82.0

than Office, which consists of 4 domains (Art, Clipart, Product, Real) in 65 classes. Following [39], we look into the settings with 3% and 6% labeled source images per class, which means each class has 2 to 4 labeled images on average. **VisDA-2017** [55] is a challenging simulation-to-real dataset containing over 280K images across 12 classes. We validate our model on settings with 0.1% and 1% labeled source images per class as suggested in [39]. **DomainNet** [54] is a large-scale domain adaptation benchmark. Since some domains and classes are noisy, we follow [59] and use a subset containing four domains (Clipart, Real, Painting, Sketch) with 126 classes. We show results on settings with 1-shot and 3-shots source labels on this dataset.

**Implementation Details.** We use ResNet-101 [30] (for DomainNet) and ResNet-50 (for other datasets) pre-trained on ImageNet [56] as our backbones. To enable a fair comparison with [39], we replaced the last FC layer with a 512-D randomly initialized linear layer. L2-normalizing are performed on the output features. We use  $k$ -means GPU implementation in faiss [37] for efficient clustering. We use SGD

with momentum of 0.9, a learning rate of 0.01, a batch size of 64. More implementation details can be found in the supplementary material.

## 4.2. Results on FUDA

**Baselines.** **SO** is a model only trained using the labeled source images. **CDAN** [45] and **MDDIA** [35] are both state-of-the-art methods in UDA with a domain classifier to perform domain alignment. **MME** [59] minimizes the conditional entropy of unlabeled target data with respect to the feature extractor and maximizes it with respect to the classifier. **CAN** [38] uses clustering information to contrast discrepancy of source and target domain. **CDS** [39] is a instance-based cross-domain self-supervised pre-training, which can be used for other domain adaptation methods and form *two-stage* methods, such as **CDS / CDAN** and **CDS / MME**. We re-implement CDS into an end-to-end version by adding losses in two stage together and tuning the weight for different losses. We also investigate the *one-stage* version of the methods above (**CDS + CDAN**, **CDS + MME**). Following [39], entropy minimization (**ENT**) on source is

Table 3: Adaptation accuracy (%) comparison on 3% and 6% labeled samples per class on the Office-Home dataset.

Method	Office-Home: Target Acc. (%)												Avg
	Ar → Cl	Ar → Pr	Ar → Rw	Cl → Ar	Cl → Pr	Cl → Rw	Pr → Ar	Pr → Cl	Pr → Rw	Rw → Ar	Rw → Cl	Rw → Pr	
<b>3% labeled source</b>													
SO	24.4	38.3	43.1	26.4	34.7	33.7	27.5	26.5	42.6	41.2	29.0	52.3	35.0
MME [59]	4.5	15.4	25.0	28.7	34.1	37.0	25.6	25.4	44.9	39.3	29.0	52.0	30.1
CDAN [45]	5.0	8.4	11.8	20.6	26.1	27.5	26.6	27.0	40.3	38.7	25.5	44.9	25.2
MDDIA [35]	21.7	37.3	42.8	29.4	43.9	44.2	37.7	29.5	51.0	47.1	29.2	56.4	39.1
CAN [38]	17.1	30.5	33.2	22.5	34.5	36.0	18.5	19.4	41.3	28.7	18.6	43.2	28.6
CDS [39]	33.5	41.1	41.9	45.9	46.0	49.3	44.7	37.8	51.0	51.6	35.7	53.8	44.4
DANN + ENT [18]	19.5	30.2	38.1	18.1	21.8	24.2	31.6	23.5	48.1	40.7	28.1	50.2	31.2
MME + ENT	31.2	35.2	40.2	37.3	39.5	37.4	48.7	42.9	60.9	59.3	46.4	58.6	44.8
CDAN + ENT	20.6	31.4	41.2	20.6	24.9	30.6	33.5	26.5	56.7	46.9	29.5	48.4	34.2
CDS + ENT	39.2	46.1	47.8	49.9	50.7	54.1	48.0	43.5	59.3	58.6	44.3	59.3	50.1
CDS + MME + ENT	39.4	48.0	52.1	50.0	52.3	56.4	47.8	44.2	60.6	61.1	45.3	62.1	51.6
CDS + CDAN + ENT	43.8	55.5	60.2	51.5	56.4	59.6	51.3	46.4	64.5	62.2	52.4	70.2	56.2
CDS / MME + ENT <sup>†</sup>	41.7	49.4	57.8	51.8	52.3	55.9	54.3	46.2	69.0	65.6	52.2	68.2	55.4
CDS / CDAN + ENT <sup>†</sup>	37.7	49.2	56.5	49.8	51.9	55.9	50.0	42.3	68.1	63.1	48.7	67.5	53.4
PCS (Ours)	<b>42.1</b>	<b>61.5</b>	<b>63.9</b>	<b>52.3</b>	<b>61.5</b>	<b>61.4</b>	<b>58.0</b>	<b>47.6</b>	<b>73.9</b>	<b>66.0</b>	<b>52.5</b>	<b>75.6</b>	<b>59.7</b>
Improvement	-1.7	<b>+6.0</b>	<b>+6.1</b>	<b>+3.7</b>	<b>+5.1</b>	<b>+1.8</b>	<b>+3.7</b>	<b>+1.2</b>	<b>+4.9</b>	<b>+0.4</b>	<b>+0.1</b>	<b>+5.4</b>	<b>+3.5</b>
<b>6% labeled source</b>													
SO	28.7	45.7	51.2	31.9	39.8	44.1	37.6	30.8	54.6	49.9	36.0	61.8	42.7
MME [59]	27.6	43.2	49.5	41.1	46.6	49.5	43.7	30.5	61.3	54.9	37.3	66.8	46.0
CDAN [45]	26.2	33.7	44.5	34.8	42.9	44.7	42.9	36.0	59.3	54.9	40.1	63.6	43.6
MDDIA [35]	25.1	44.5	51.9	35.6	46.7	50.3	48.3	37.1	64.5	58.2	36.9	68.4	50.3
CAN [38]	20.4	34.7	44.7	29.0	40.4	38.6	33.3	21.1	53.4	36.8	19.1	58.0	35.8
CDS [39]	38.8	51.7	54.8	53.2	53.3	57.0	53.4	44.2	65.2	63.7	45.3	68.6	54.1
DANN + ENT [18]	22.4	32.9	43.5	23.2	30.9	33.3	33.2	26.9	54.6	46.8	32.7	55.1	36.3
MME + ENT	37.2	42.4	50.9	46.1	46.6	49.1	53.5	45.6	67.2	63.4	48.1	71.2	51.8
CDAN + ENT	23.1	35.5	49.2	26.1	39.2	43.8	44.7	33.8	61.7	55.1	34.7	67.9	42.9
CDS + ENT	42.9	55.5	59.5	55.2	55.1	59.1	54.3	46.9	68.1	65.7	50.6	71.5	57.0
CDS + MME + ENT	41.7	58.1	61.7	55.7	56.2	61.3	54.6	47.3	68.6	66.4	50.3	72.1	57.8
CDS + CDAN + ENT	45.4	60.4	65.5	54.9	59.2	63.8	55.4	49.0	71.6	66.6	54.1	75.4	60.1
CDS / MME + ENT <sup>†</sup>	44.1	51.6	63.3	53.9	55.2	62.0	56.5	46.6	70.9	67.7	54.7	74.7	58.4
CDS / CDAN + ENT <sup>†</sup>	39.0	51.3	63.1	51.0	55.0	63.6	57.8	45.9	72.8	65.8	50.4	73.5	57.4
PCS (Ours)	<b>46.1</b>	<b>65.7</b>	<b>69.2</b>	<b>57.1</b>	<b>64.7</b>	<b>66.2</b>	<b>61.4</b>	<b>47.9</b>	<b>75.2</b>	<b>67.0</b>	<b>53.9</b>	<b>76.6</b>	<b>62.6</b>
Improvement	<b>+0.7</b>	<b>+5.3</b>	<b>+3.7</b>	<b>+2.2</b>	<b>+5.5</b>	<b>+2.4</b>	<b>+3.6</b>	-1.1	<b>+2.4</b>	-0.7	-0.8	<b>+1.2</b>	<b>+2.5</b>

<sup>†</sup> Two-stage training results reported in [39].

Table 4: Adaptation accuracy (%) comparison on 0.1% and 1% labeled samples per class on the VisDA-2017 dataset.

Method	VisDA: Target Acc. (%)	
	0.1% Labeled	1% Labeled
SO	47.9	51.4
MME [59]	55.6	69.4
CDAN [45]	58.0	61.5
MDDIA [35]	68.9	71.3
CAN [38]	51.3	57.2
CDS [39]	34.2	67.5
DANN + ENT [18]	44.5	50.2
MME + ENT	54.0	66.1
CDAN + ENT	57.7	58.1
CDS + ENT	49.8	75.3
CDS + ENT + MME	60.0	78.3
CDS / MME + ENT <sup>†</sup>	62.5	69.4
CDS / CDAN + ENT <sup>†</sup>	69.0	69.1
PCS (Ours)	<b>78.0</b>	<b>79.0</b>
Improvement	<b>+9.0</b>	<b>+0.7</b>

<sup>†</sup> Two-stage training results reported in [39].

added to previous DA methods to obtain better baseline performance.

We compare the proposed PCS with state-of-the-art methods on FUDA (adaptation with few source labels). Ex-

Table 5: Adaptation accuracy (%) comparison on 1-shot and 3-shots per class on the DomainNet dataset.

Method	DomainNet: Target Acc. (%)							
	R-C	R-P	R-S	P-C	P-R	C-S	S-P	Avg
<b>1-shot labeled source</b>								
SO	18.4	30.6	16.7	16.2	28.9	12.7	10.5	19.1
MME [59]	13.8	29.2	9.7	16.0	26.0	13.4	14.4	17.5
CDAN [45]	16.0	25.7	12.9	12.6	19.5	7.2	8.0	14.6
MDDIA [35]	18.0	30.6	15.9	15.4	27.4	9.3	10.2	18.1
CAN [38]	18.3	22.1	16.7	13.2	23.9	11.1	12.1	16.8
CDS [39]	16.7	24.4	11.1	14.1	15.9	13.4	19.0	16.4
CDS + ENT	21.7	30.1	18.2	17.4	20.5	18.6	22.7	21.5
CDS + MME + ENT	21.2	28.8	15.5	15.8	17.6	19.0	20.7	19.8
PCS (Ours)	<b>39.0</b>	<b>51.7</b>	<b>39.8</b>	<b>26.4</b>	<b>38.8</b>	<b>23.7</b>	<b>23.6</b>	<b>34.7</b>
Improvement	<b>+17.3</b>	<b>+21.1</b>	<b>+21.6</b>	<b>+9.0</b>	<b>+9.9</b>	<b>+4.7</b>	<b>+0.9</b>	<b>+13.2</b>
<b>3-shots labeled source</b>								
SO	30.2	44.2	25.7	24.6	49.8	24.2	23.2	31.7
MME [59]	22.8	46.5	14.5	25.1	50.0	20.1	24.9	29.1
CDAN [45]	30.0	40.1	21.7	21.4	40.8	17.1	19.7	27.3
MDDIA [35]	41.4	50.7	37.4	31.4	52.9	23.1	24.1	37.3
CAN [38]	28.1	33.5	25	24.7	46.9	23.3	20.1	28.8
CDS [39]	35.0	43.8	36.7	34.1	36.8	31.1	34.5	36.0
CDS + ENT	44.5	52.2	40.9	40.0	47.2	33.0	40.1	42.5
CDS + MME + ENT	43.8	54.9	41.1	38.9	45.9	32.8	38.7	42.3
PCS (Ours)	<b>45.2</b>	<b>59.1</b>	<b>41.9</b>	<b>41.0</b>	<b>66.6</b>	31.9	37.4	<b>46.1</b>
Improvement	<b>+0.7</b>	<b>+6.9</b>	<b>+0.8</b>	<b>+1.0</b>	<b>+13.7</b>	-0.9	-2.7	<b>+3.6</b>

tensive experiments are conducted on Office, Office-Home, VisDA-2017 and DomainNet, with the results presented in

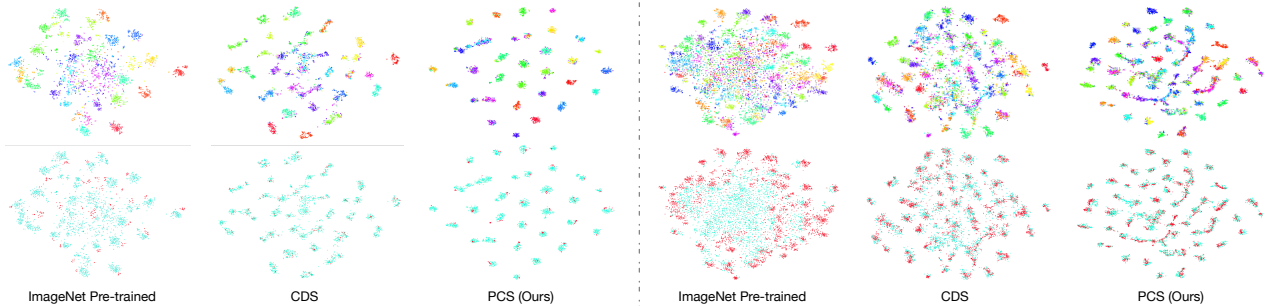


Figure 4: t-SNE visualization of ours and baselines on Office (left) and Office-Home (right). Top row: Coloring represents the class of each sample. Features with PCS are more discriminative than the ones with other methods. Bottom row: Cyan represents source features and Red represents target features. Feature from PCS are better-aligned between domains compared to other methods.

Table 1, 3, 4, and 5, respectively. We can see that PCS outperforms the best state-of-the-arts in all the benchmarks, with large improvements: 10.5% and 3.4% on Office, 4.3% and 4.2% on Office-Home, 9.0% and 0.7% on VisDA, 13.2% and 3.6% on DomainNet. If we look at the result of each domain pair in each dataset (e.g. D  $\rightarrow$  A in Office), PCS outperforms previous best in 47 out of 52 settings. Finally, as the number of labeled samples decreases, PCS shows larger performance improvements against the previous best methods, which demonstrates PCS is extremely beneficial in label-scarce adaptation scenarios.

### 4.3. Ablation Study and Analysis

Next, we investigate the effectiveness of each component in PCS on Office. Table 2 shows that adding each component contributes to the final results without any performance degradation. Comparing the last row in Table 2 and Table 1, we can see even without MIM, PCS still outperforms all previous methods.

We plot the learned features with t-SNE [48] on the DSLR-to-Amazon setting in Office, and Real-to-Clipart in Office-Home respectively in left and right of Figure 4. In the top row, the color represents the class of each sample; while in the bottom row, cyan represents source samples and red represents target samples. Compared to ImageNet pre-training and CDS, it qualitatively shows that PCS well clusters samples with the same class in the feature space; thus, PCS favors more discriminative features. Also, the features from PCS are more closely aggregated than ImageNet pre-training and CDS, which demonstrates that PCS learns a better semantic structure of the datasets.

### 4.4. Sample Efficiency

We compare our method with other state-of-the-art methods on Office dataset (DSLR as source and Amazon as target) with a varying number of source labels. From Figure 5, we can see that PCS outperforms all SOTA methods in all settings with different number of labeled samples. Moreover, our method is very label-efficient: a) For 1-shot

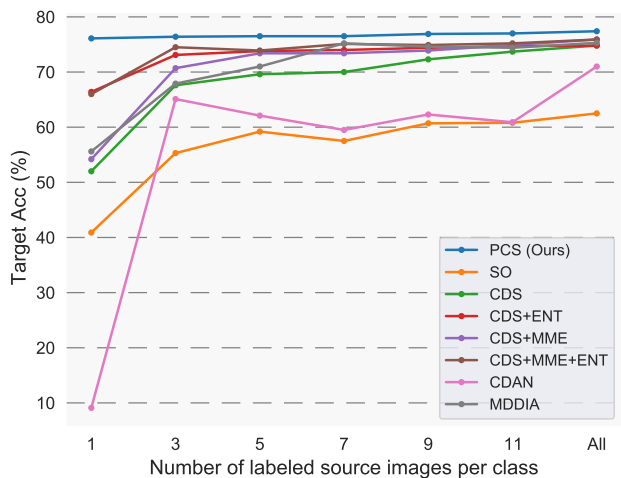


Figure 5: Sample efficiency comparison from DSLR to Amazon in Office dataset.

image per class (31 labeled source images in total), PCS can achieve **76.1%** accuracy. b) For the fully-labeled setting (498 labeled source images in total), PCS can achieve **77.4%** accuracy. c) With **94%** less labeled source images, the performance degradation of our method is only **1.3%**. In short, *with less labeled source data, PCS outperforms other works by a larger margin.*

## 5. Conclusion

In this paper, we investigated Few-shot Unsupervised Domain Adaptation where only few labeled samples are available in the source domain, and no labeled samples in the target domain. We proposed a novel Prototypical Cross-domain Self-supervised learning (PCS) framework that performs both in-domain and cross-domain prototypical self-supervised learning, as well as adaptive prototype-classifier learning. We perform extensive experiments on multiple benchmark datasets, which demonstrates the superiority of PCS over previous best methods. PCS sets a new state of the art for Few-shot Unsupervised Domain Adaptation.



## References

- [1] Idan Achituve, Haggai Maron, and Gal Chechik. Self-supervised learning for domain adaptation on point-clouds. *arXiv preprint arXiv:2003.12641*, 2020.
- [2] Yuki M. Asano, Christian Rupprecht, and Andrea Vedaldi. Self-labelling via simultaneous clustering and representation learning. In *International Conference on Learning Representations (ICLR)*, 2020.
- [3] Philip Bachman, R Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. In *Advances in Neural Information Processing Systems*, pages 15535–15545, 2019.
- [4] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. In *Advances in Neural Information Processing Systems*, pages 5049–5059, 2019.
- [5] 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.
- [6] Fabio M Carlucci, Antonio D’Innocente, Silvia Bucci, Barbara Caputo, and Tatiana Tommasi. Domain generalization by solving jigsaw puzzles. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2229–2238, 2019.
- [7] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in Neural Information Processing Systems*, 33, 2020.
- [8] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017.
- [9] Min-Hung Chen, Baopu Li, Yingze Bao, Ghassan Al-Regib, and Zsolt Kira. Action segmentation with joint self-supervised temporal domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9454–9463, 2020.
- [10] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020.
- [11] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. Big self-supervised models are strong semi-supervised learners. *arXiv preprint arXiv:2006.10029*, 2020.
- [12] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020.
- [13] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE international conference on computer vision*, pages 1422–1430, 2015.
- [14] Carl Doersch and Andrew Zisserman. Multi-task self-supervised visual learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2051–2060, 2017.
- [15] Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. In *International conference on machine learning*, pages 647–655, 2014.
- [16] Alexey Dosovitskiy, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with convolutional neural networks. In *Advances in neural information processing systems*, pages 766–774, 2014.
- [17] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189. PMLR, 2015.
- [18] 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.
- [19] Vivien Sainte Fare Garnot and Loic Landrieu. Metric-guided prototype learning. *arXiv preprint arXiv:2007.03047*, 2020.
- [20] Muhammad Ghifary, W Bastiaan Kleijn, Mengjie Zhang, and David Balduzzi. Domain generalization for object recognition with multi-task autoencoders. In *Proceedings of the IEEE international conference on computer vision*, pages 2551–2559, 2015.
- [21] 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.
- [22] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*, 2018.
- [23] 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.
- [24] Raghuraman Gopalan, Ruonan Li, and Rama Chellappa. Domain adaptation for object recognition: An unsupervised approach. In *2011 international conference on computer vision*, pages 999–1006. IEEE, 2011.
- [25] Yves Grandvalet and Yoshua Bengio. Semi-supervised learning by entropy minimization. In *Advances in neural information processing systems*, pages 529–536, 2005.
- [26] Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012.
- [27] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent: A

- new approach to self-supervised learning. *arXiv preprint arXiv:2006.07733*, 2020.
- [28] Varun Gulshan, Lily Peng, Marc Coram, Martin C Stumpe, Derek Wu, Arunachalam Narayanaswamy, Subhashini Venugopalan, Kasumi Widner, Tom Madams, Jorge Cuadros, et al. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *Jama*, 316(22):2402–2410, 2016.
- [29] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020.
- [30] 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.
- [31] 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 *International conference on machine learning*, pages 1989–1998. PMLR, 2018.
- [32] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- [33] Sontje Ihler, Felix Kuhnke, Max-Heinrich Laves, and Tobias Ortmaier. Self-supervised domain adaptation for patient-specific, real-time tissue tracking. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 54–64. Springer, 2020.
- [34] Rae Jeong, Yusuf Aytar, David Khosid, Yuxiang Zhou, Jackie Kay, Thomas Lampe, Konstantinos Bousmalis, and Francesco Nori. Self-supervised sim-to-real adaptation for visual robotic manipulation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2718–2724. IEEE, 2020.
- [35] Xiang Jiang, Qicheng Lao, Stan Matwin, and Mohammad Havaei. Implicit class-conditioned domain alignment for unsupervised domain adaptation. In *International Conference on Machine Learning*, 2020.
- [36] Longlong Jing and Yingli Tian. Self-supervised visual feature learning with deep neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [37] Jeff Johnson, Matthijs Douze, and Hervé Jégou. Billion-scale similarity search with gpus. *arXiv preprint arXiv:1702.08734*, 2017.
- [38] Guoliang Kang, Lu Jiang, Yi Yang, and Alexander G Hauptmann. Contrastive adaptation network for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4893–4902, 2019.
- [39] Donghyun Kim, Kuniaki Saito, Tae-Hyun Oh, Bryan A Plummer, Stan Sclaroff, and Kate Saenko. Cross-domain self-supervised learning for domain adaptation with few source labels. *arXiv preprint arXiv:2003.08264*, 2020.
- [40] Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Learning representations for automatic colorization. In *European conference on computer vision*, pages 577–593. Springer, 2016.
- [41] Junnan Li, Pan Zhou, Caiming Xiong, Richard Socher, and Steven CH Hoi. Prototypical contrastive learning of unsupervised representations. *arXiv preprint arXiv:2005.04966*, 2020.
- [42] Ming-Yu Liu and Oncl Tuzel. Coupled generative adversarial networks. In *Advances in neural information processing systems*, pages 469–477, 2016.
- [43] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
- [44] 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. PMLR, 2015.
- [45] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. In *Advances in Neural Information Processing Systems*, pages 1640–1650, 2018.
- [46] 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.
- [47] 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. PMLR, 2017.
- [48] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.
- [49] Ishan Misra and Laurens van der Maaten. Self-supervised learning of pretext-invariant representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6707–6717, 2020.
- [50] Zak Murez, Soheil Kolouri, David Kriegman, Ravi Ramamoorthi, and Kyungnam Kim. Image to image translation for domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4500–4509, 2018.
- [51] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European Conference on Computer Vision*, pages 69–84. Springer, 2016.
- [52] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [53] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2536–2544, 2016.
- [54] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source

- domain adaptation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1406–1415, 2019.
- [55] 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.
- [56] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- [57] 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.
- [58] 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.
- [59] Kuniaki Saito, Donghyun Kim, Stan Sclaroff, Trevor Darrell, and Kate Saenko. Semi-supervised domain adaptation via minimax entropy. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8050–8058, 2019.
- [60] 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.
- [61] Jian Shen, Yanru Qu, Weinan Zhang, and Yong Yu. Wasserstein distance guided representation learning for domain adaptation. *arXiv preprint arXiv:1707.01217*, 2017.
- [62] Ashish Shrivastava, Tomas Pfister, Oncel Tuzel, Joshua Susskind, Wenda Wang, and Russell Webb. Learning from simulated and unsupervised images through adversarial training. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2107–2116, 2017.
- [63] Baochen Sun, Jiashi Feng, and Kate Saenko. Correlation alignment for unsupervised domain adaptation. In *Domain Adaptation in Computer Vision Applications*, pages 153–171. Springer, 2017.
- [64] Yu Sun, Eric Tzeng, Trevor Darrell, and Alexei A Efros. Unsupervised domain adaptation through self-supervision. *arXiv preprint arXiv:1909.11825*, 2019.
- [65] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. *arXiv preprint arXiv:1906.05849*, 2019.
- [66] Antonio Torralba and Alexei A Efros. Unbiased look at dataset bias. In *CVPR 2011*, pages 1521–1528. IEEE, 2011.
- [67] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7167–7176, 2017.
- [68] Eric Tzeng, Judy Hoffman, Ning Zhang, Kate Saenko, and Trevor Darrell. Deep domain confusion: Maximizing for domain invariance. *arXiv preprint arXiv:1412.3474*, 2014.
- [69] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5018–5027, 2017.
- [70] Hanchen Wang, Qi Liu, Xiangyu Yue, Joan Lasenby, and Matthew J Kusner. Pre-training by completing point clouds. *arXiv preprint arXiv:2010.01089*, 2020.
- [71] Qian Wang and Toby Breckon. Unsupervised domain adaptation via structured prediction based selective pseudo-labeling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 6243–6250, 2020.
- [72] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3733–3742, 2018.
- [73] Shaoan Xie, Zibin Zheng, Liang Chen, and Chuan Chen. Learning semantic representations for unsupervised domain adaptation. In *International Conference on Machine Learning*, pages 5423–5432, 2018.
- [74] Jae Shin Yoon, Takaaki Shiratori, Shoou-I Yu, and Hyun Soo Park. Self-supervised adaptation of high-fidelity face models for monocular performance tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4601–4609, 2019.
- [75] Xiangyu Yue, Yang Zhang, Sicheng Zhao, Alberto Sangiovanni-Vincentelli, Kurt Keutzer, and Boqing Gong. Domain randomization and pyramid consistency: Simulation-to-real generalization without accessing target domain data. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2100–2110, 2019.
- [76] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *European conference on computer vision*, pages 649–666. Springer, 2016.
- [77] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017.
- [78] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017.
- [79] Junbao Zhuo, Shuhui Wang, Weigang Zhang, and Qingming Huang. Deep unsupervised convolutional domain adaptation. In *Proceedings of the 25th ACM international conference on Multimedia*, pages 261–269, 2017.