

Cross-Domain Cross-Set Few-Shot Learning via Learning Compact and Aligned Representations

Wentao Chen^{1,2}, Zhang Zhang^{2,3}, Wei Wang^{2,3}, Liang Wang^{2,3}, Zilei Wang¹,
and Tieniu Tan^{1,2,3}

¹ University of Science and Technology of China, Hefei, China

² Center for Research on Intelligent Perception and Computing,
NLPR, CASIA, Beijing, China

³ University of Chinese Academy of Sciences, Beijing, China
wentao95@mail.ustc.edu.cn, zlwang@ustc.edu.cn
{zzhang, wangwei, wangliang, tnt}@nlpr.ia.ac.cn

Abstract. Few-shot learning (FSL) aims to recognize novel queries with only a few support samples through leveraging prior knowledge from a base dataset. In this paper, we consider the domain shift problem in FSL and aim to address the domain gap between the support set and the query set. Different from previous cross-domain FSL work (CD-FSL) that considers the domain shift between base and novel classes, the new problem, termed cross-domain cross-set FSL (CDSC-FSL), requires few-shot learners not only to adapt to the new domain, but also to be consistent between different domains within each novel class. To this end, we propose a novel approach, namely *stabPA*, to learn prototypical compact and cross-domain aligned representations, so that the domain shift and few-shot learning can be addressed simultaneously. We evaluate our approach on two new CDSC-FSL benchmarks built from the DomainNet and Office-Home datasets respectively. Remarkably, our approach outperforms multiple elaborated baselines by a large margin, e.g., improving 5-shot accuracy by 6.0 points on average on DomainNet. Code is available at <https://github.com/WentaoChen0813/CDSC-FSL>.

Keywords: cross-domain cross-set few-shot learning, prototypical alignment

1 Introduction

Learning a new concept with a very limited number of examples is easy for human beings. However, it is quite difficult for current deep learning models, which usually require plenty of labeled data to learn generalizable and discriminative representations. To bridge the gap between humans and machines, few-shot learning (FSL) has been recently proposed [42,28].

Similar to human beings, most FSL algorithms leverage prior knowledge from known classes to assist recognizing novel concepts. Typically, a FSL algorithm is composed of two phases: (i) pre-train a model on a base set that contains a large number of seen classes (called meta-training phase); (ii) transfer the

pre-trained model to novel classes with a small labeled support set and test it with a query set (meta-testing phase). Despite great progresses on FSL algorithms [9,29,42,37], most previous studies adopt a single domain assumption, where all images in both meta-training and meta-testing phases are from a single domain. Such assumption, however, may be easily broken in real-world applications. Considering a concrete example of online shopping, a clothing retailer commonly shows several high-quality pictures taken by photographers for each fashion product (support set), while customers may use their cellphone photos (query set) to match the displayed pictures of their expected products. In such case, there is a distinct domain gap between the support set and the query set. Similar example can be found in security surveillance: given the low-quality picture of a suspect captured at night (query set), the surveillance system is highly expected to recognize its identity based on a few high-quality registered photos (e.g., ID card). With such domain gap, FSL models will face more challenges besides limited support data.

In this paper, we consider the above problem in FSL and propose a new setting to address the domain gap between the support set and the query set. Following previous FSL work, a large base set from the source domain is available for meta-training. Differently, during meta-testing, only the support set or the query set is from the source domain, while the other is from a different target domain. Some recent studies also consider the cross-domain few-shot learning problem (CD-FSL) [15,39,26]. However, the domain shift in CD-FSL occurs between the meta-training and meta-testing phases. In other words, both the support and query sets in the meta-testing phase are still from the same domain (pictorial illustration is given in Figure 1 (a)). To distinguish the considered setting from CD-FSL, we name this setting as cross-domain cross-set few-shot learning (CDCS-FSL), as the support set and the query set are across different domains. Compared to CD-FSL, the domain gap within each novel class imposes more requirements to learn a well-aligned feature space. Nevertheless, in terms of the above setting, it is nearly intractable to conquer the domain shift due to the very limited samples of the target domain, e.g., the target domain may contain only one support (or query) image. Thus, we follow the CD-FSL literature [26] to use unlabeled auxiliary data from the target domain to assist model training. Note that we do not suppose that the auxiliary data are from novel classes. Therefore, we can collect these data from some common-seen classes (e.g., base classes) without any annotation costs.

One may notice that re-collecting a few support samples from the same domain as the query set can ‘simply’ eliminate the domain gap. However, it may be intractable to re-collect support samples in some real few-shot applications, e.g., re-collecting ID photos for all persons is difficult. Besides, users sometimes not only want to get the class labels, but more importantly they’d like to retrieve the support images themselves (like the high-quality fashion pictures). Therefore, the CDCS-FSL setting can not be simply transferred into previous FSL and CD-FSL settings.

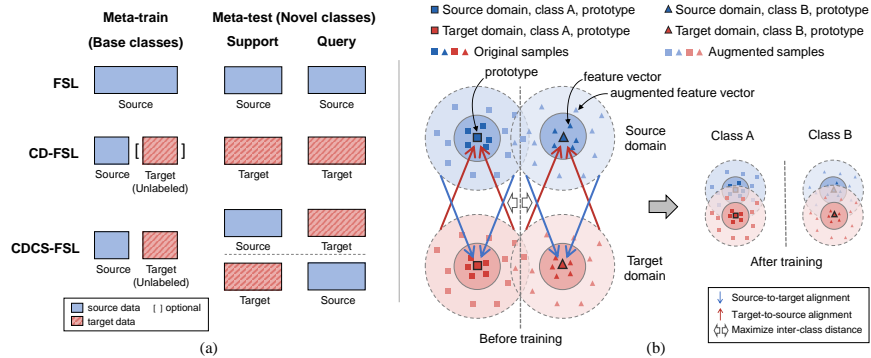


Fig. 1. Problem setup and motivation. (a) CD-FSL considers the domain shift between the meta-training and meta-testing phases, while CDCS-FSL considers the domain shift between the support set and the query set in the meta-testing phase. Following previous CD-FSL work [26], unlabeled target domain data are used in CDCS-FSL to assist model training. (b) We propose a bi-directional prototypical alignment framework to address CDCS-FSL, which pushes feature vectors of one domain to be gathered around the corresponding prototype in the other domain bi-directionally, and separates feature vectors from different classes.

To address the CDCS-FSL problem, we propose a simple but effective bi-directional prototypical alignment framework to learn compact and cross-domain aligned representations, which is illustrated in Figure 1 (b). The main idea of our approach is derived from two intuitive insights: (i) we need aligned representations to alleviate the domain shift between the source and target domains, and (ii) compact representations are desirable to learn a center-clustered class space, so that a small support set can better represent a new class. Specifically, given the labeled base set in the source domain and the unlabeled auxiliary set in the target domain, we first assign pseudo labels to the unlabeled data considering that pseudo labels can preserve the coarse semantic similarity with the visual concepts in source domain. Then, we minimize the point-to-set distance between the prototype (class center) in one domain and the corresponding feature vectors in the other domain bi-directionally. As results, the feature vectors of the source (or target) domain will be gathered around the prototype in the other domain, thus reducing the domain gap and intra-class variance simultaneously. Moreover, the inter-class distances are maximized to attain a more separable feature space. Furthermore, inspired by the fact that data augmentation even with strong transformations generally does not change sample semantics, we augment samples in each domain, and suppose that the augmented samples between different domains should also be aligned. Since these augmented samples enrich the data diversity, they can further encourage to learn the underlying invariance and strengthen the cross-domain alignment.

Totally, we summarize all the above steps into one approach termed “**Strongly Augmented Bi-directional Prototypical Alignment**”, or *stabPA*. We evaluate its

effectiveness on two new CDCS-FSL benchmarks built from the DomainNet [25] and Office-Home [40] datasets. Remarkably, the proposed *stabPA* achieves the best performance over both benchmarks and outperforms other baselines with a large margin, e.g., improving 5-shot accuracy by 6.0 points on average on the DomainNet dataset.

In summary, our contributions are three-fold:

- We consider a new FSL setting, CDCS-FSL, where a domain gap exists between the support set and the query set.
- We propose a novel approach, namely *stabPA*, to address the CDCS-FSL problem, of which the key is to learn prototypical compact and domain aligned representations.
- Extensive experiments demonstrate that *stabPA* can learn discriminative and generalizable representations and outperforms all baselines by a large margin on two CDCS-FSL benchmarks.

2 Related Work

FSL aims to learn new classes with very few labeled examples. Most studies follow a meta-learning paradigm [41], where a meta-learner is trained on a series of training tasks (episodes) so as to enable fast adaptation to new tasks. The meta-learner can take various forms, such as an LSTM network [28], initial network parameters [9], or closed-form solvers [29]. Recent advances in pre-training techniques spawn another FSL paradigm. In [4], the authors show that a simple pre-training and fine-tuning baseline can achieve competitive performance with respect to the SOTA FSL models. In [37,5], self-supervised pre-training techniques have proven to be useful for FSL. Our approach also follows the pre-training paradigm, and we further expect the learned representations to be compact and cross-domain aligned to address the CDCS-FSL problem.

CD-FSL [15,39,26,43,20,14,10] considers the domain shift problem between the base classes and the novel classes. Due to such domain gap, [4] show that meta-learning approaches fail to adapt to novel classes. To alleviate this problem, [39] propose a feature-wise transformation layer to learn rich representations that can generalize better to other domains. However, they need to access multiple labeled data sources with extra data collection costs. [26] solve this problem by exploiting additional unlabeled target data with self-supervised pre-training techniques. Alternatively, [14] propose to utilize the semantic information of class labels to minimize the distance between source and target domains. Without the need for extra data or language annotations, [43] augment training tasks in an adversarial way to improve the generalization capability.

Using target domain images to alleviate domain shift is related to the field of domain adaptation (DA). Early efforts align the marginal distribution of each domain by minimizing a pre-defined discrepancy, such as $\mathcal{H}\Delta\mathcal{H}$ -divergence [1] or Maximum Mean Discrepancy (MMD) [13]. Recently, adversarial-based methods adopt a discriminator [12] to approximate the domain discrepancy, and learn domain-invariant distribution at image level [17], feature level [21] or output

level [38]. Another line of studies assign pseudo labels to unlabeled target data [49,47,48], and directly align the feature distribution within each class. Although these DA methods are related to our work, they usually assume that the testing stage shares the same class space as the training stage, which is broken by the setting of FSL. Open-set DA [24,31] and Universal DA [30,45] consider the existence of unseen classes. However, they merely mark them as ‘unknown’. In this work, we are more interested in addressing the domain shift for these unseen novel classes within a FSL assumption.

3 Problem Setup

Formally, a FSL task often adopts a setting of N-way-K-shot classification, which aims to discriminate between N novel classes with K exemplars per class. Given a support set $\mathcal{S} = \{(x_i, y_i)\}_{i=1}^{N \times K}$ where $x_i \in \mathcal{X}_{\mathcal{N}}$ denotes a data sample in novel classes and $y_i \in Y_{\mathcal{N}}$ is the class label, the goal of FSL is to learn a mapping function $\phi : \phi(x_q) \rightarrow y_q$ which classifies a query sample x_q in the query set \mathcal{Q} to the class label $y_q \in Y_{\mathcal{N}}$. Besides \mathcal{S} and \mathcal{Q} , a large labeled dataset $\mathcal{B} \subset \mathcal{X}_{\mathcal{B}} \times \mathcal{Y}_{\mathcal{B}}$ (termed base set) is often provided for meta-training, where $\mathcal{X}_{\mathcal{B}}$ and $\mathcal{Y}_{\mathcal{B}}$ do not overlap with $\mathcal{X}_{\mathcal{N}}$ and $\mathcal{Y}_{\mathcal{N}}$.

Conventional FSL studies assume the three sets \mathcal{S} , \mathcal{Q} and \mathcal{B} are from the same domain. In this paper, we consider the domain gap between the support set and the query set (only one is from the same domain as the base set, namely the source domain \mathcal{D}_s , and the other is from a new target domain \mathcal{D}_t). Specifically, this setting has two situations:

- (i) $\mathcal{D}_s - \mathcal{D}_t$: the support set is from the source domain and the query set is from the target domain, i.e., $\mathcal{S} \subset \mathcal{D}_s$ and $\mathcal{Q} \subset \mathcal{D}_t$.
- (ii) $\mathcal{D}_t - \mathcal{D}_s$: the support set is from the target domain and the query set is from the source domain, i.e., $\mathcal{S} \subset \mathcal{D}_t$ and $\mathcal{Q} \subset \mathcal{D}_s$.

As the support set and the query set are across different domains, we name this setting as cross-domain cross-set few-shot learning (CDCS-FSL). Besides the above three sets, to facilitate crossing the domain gap, an unlabeled auxiliary set \mathcal{U} from the target domain is available in the meta-training phase, where the data from novel classes are manually removed to promise they are not seen in meta-training.

4 Approach

Briefly, our approach contains two stages: 1) In the meta-training stage, we train a feature extractor $f : x_i \rightarrow f(x_i)$ with the base set \mathcal{B} and the unlabeled auxiliary set \mathcal{U} ; 2) In the meta-testing stage, we fix the feature extractor and train a linear classification head $g : f(x_i) \rightarrow y_i$ on the support set \mathcal{S} , and the entire model $\phi = g \circ f$ is used to predict the labels for the query set \mathcal{Q} . The framework of our approach is illustrated in Figure 2.

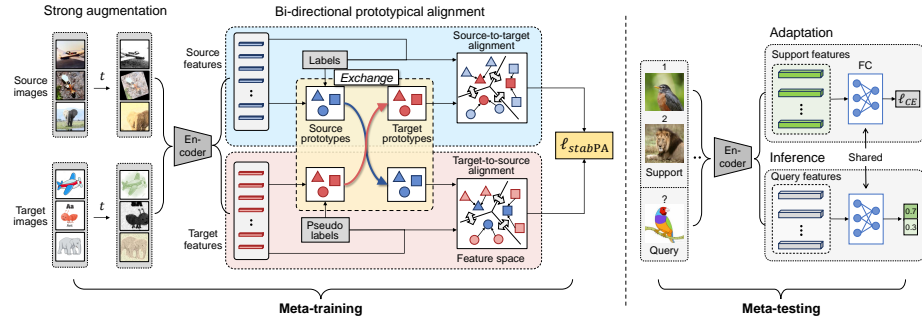


Fig. 2. Framework. In the meta-training stage, we train a feature extractor within the bi-directional prototypical alignment framework to learn compact and aligned representations. In the meta-testing stage, we fix the feature extractor and train a new classification head with the support set, and then evaluate the model on the query set.

4.1 Bi-directional Prototypical Alignment

A straightforward way to align feature distributions is through estimating class centers (prototypes) in both source and target domains. With the labeled base data, it is easy to estimate prototypes for the source domain. However, it is difficult to estimate prototypes in the target domain with only unlabeled data available. To address this issue, we propose to assign pseudo labels to the unlabeled data and then use the pseudo labels to approximate prototypes. The insight is that the pseudo labels can preserve the coarse semantic similarity even under domain or category shift (e.g., a painting tiger could be more likely to be pseudo-labeled as a cat rather than a tree). Aggregating samples with the same pseudo label can extract the shared semantics across different domains.

Specifically, given the source domain base set \mathcal{B} and the target domain unlabeled set \mathcal{U} , we first assign pseudo labels to the unlabeled samples with an initial classifier ϕ_0 trained on the base set and obtain $\hat{\mathcal{U}} = \{(x_i, \hat{y}_i) | x_i \in \mathcal{U}\}$, where $\hat{y}_i = \phi_0(x_i)$ is the pseudo label. Then, we obtain the source prototypes $\{p_k^s\}_{k=1}^{|\mathcal{Y}^s|}$ and the target prototypes $\{p_k^t\}_{k=1}^{|\mathcal{Y}^s|}$ by averaging the feature vectors with the same label (or pseudo label). It should be noted that the prototypes are estimated on the entire datasets \mathcal{B} and $\hat{\mathcal{U}}$, and adjusted together with the update of the feature extractor and pseudo labels (details can be found below).

With the obtained prototypes, directly minimizing the *point-to-point* distance between two prototypes p_k^s and p_k^t can easily reduce the domain gap for the class k . However, this may make the feature distribution of different classes mix together and the discrimination capability of the learned representations is still insufficient. To overcome these drawbacks, we propose to minimize the *point-to-set* distance across domains in a *bi-directional* way. That is, we minimize the distance between the *prototype in one domain* and the corresponding *feature vectors in the other domain*, and meanwhile maximize the feature distance between different classes. In this way, we can not only align features across

domains, but also simultaneously obtain compact feature distributions for both domains to suit the requirement of few-shot learning.

Concretely, for a source sample $(x_i^s, y_i^s) \in \mathcal{B}$ of the q -th class (i.e., $y_i^s = q$), we minimize its feature distance to the prototype p_q^t in the target domain, and meanwhile maximize its distances to the prototypes of other classes. Here, a softmax loss function for the source-to-target alignment is formulated as:

$$\ell_{s-t}(x_i^s, y_i^s) = -\log \frac{\exp(-\|f(x_i^s) - p_q^t\|^2/\tau)}{\sum_{k=1}^{|\mathcal{Y}_{\mathcal{B}}|} \exp(-\|f(x_i^s) - p_k^t\|^2/\tau)}, \quad (1)$$

where τ is a temperature factor. To get a better feature space for the target domain, a similar target-to-source alignment loss is applied for each target sample $(x_i^t, \hat{y}_i^t) \in \hat{\mathcal{U}}$ with $\hat{y}_i^t = q$:

$$\ell_{t-s}(x_i^t, \hat{y}_i^t) = -\log \frac{\exp(-\|f(x_i^t) - p_q^s\|^2/\tau)}{\sum_{k=1}^{|\mathcal{Y}_{\mathcal{B}}|} \exp(-\|f(x_i^t) - p_k^s\|^2/\tau)}. \quad (2)$$

Since the initial pseudo labels are more likely to be incorrect, we gradually increase the weights of these two losses following the principle of curriculum learning [2]. For the source-to-target alignment, the loss weight starts from zero and converges to one, formulated as:

$$w(t) = \frac{2}{1 + \exp(-t/T_{max})} - 1, \quad (3)$$

where t is the current training step and T_{max} is the maximum training step. For the target-to-source alignment, since the pseudo labels become more confident along with the training process, a natural curriculum is achieved by setting a confidence threshold to filter out the target samples with low confidence pseudo labels [33].

Together, the total loss for the bi-directional prototypical alignment is

$$\ell_{bPA} = \frac{1}{|\mathcal{B}|} \sum_{i=1}^{|\mathcal{B}|} w(t) \ell_{s-t}(x_i^s, y_i^s) + \frac{1}{|\hat{\mathcal{U}}|} \sum_{i=1}^{|\hat{\mathcal{U}}|} \mathbb{1}(p(\hat{y}_i^t) > \beta) \ell_{t-s}(x_i^t, \hat{y}_i^t), \quad (4)$$

where $p(\cdot)$ is the confidence of a pseudo label, and β is the confidence threshold below which the data samples will be dropped.

Updating Pseudo Label. The pseudo labels are initially predicted by a classifier ϕ_0 pre-trained on the base set \mathcal{B} . As the representations are updated, we update the pseudo labels by re-training a classifier $\phi_t = h \circ f$ based on the current feature extractor f , where h is a linear classification head for the base classes. The final pseudo labels are updated by linear interpolation between the predictions of the initial classifier ϕ_0 and the online updated classifier ϕ_t :

$$\hat{y}_i = \arg \max_{k \in \mathcal{Y}_{\mathcal{B}}} \lambda \phi_0(k|x_i) + (1 - \lambda) \phi_t(k|x_i), \quad (5)$$

where λ is the interpolation coefficient. The combination of these two classifiers makes it possible to rectify the label noise of the initial classifier, and meanwhile

inhibit the rapid change of pseudo labels of online classifier especially in the early training stage.

Generating Prototypes. Note that we are intended to estimate the prototypes on the entire dataset and update them with representation learning. For the source domain, instead of calculating the mean value of intra-class samples in the feature space, a cheaper way is to approximate prototypes with the normalized weights of the classification head h , as the classifier weights tend to align with class centers in order to reduce classification errors [27]. Specifically, we set the source prototypes as $p_k^s = W_k$, where W_k is the normalized classification weight for the k -th class. For the target domain, we adopt the momentum technique to update prototypes. The prototypes are initialized as zeros. At each training step, we first estimate the prototypes using target samples in current batch with their pseudo labels. Then, we update the target prototype p_k^t as:

$$p_k^t \leftarrow mp_k^t + (1 - m) \frac{1}{n_k} \sum_{i=1}^{|\hat{\mathcal{U}}_b|} \mathbb{1}(\hat{y}_i^t = k) f(x_i^t), \quad (6)$$

where n_k is the number of the target samples classified into the k -th class in a target batch $\hat{\mathcal{U}}_b$, and m is the momentum term controlling the update speed.

4.2 stabPA

Strong data augmentation has proved to be effective for learning generalizable representations, especially in self-supervised representation learning studies [16,3]. Given a sample x , strong data augmentation generates additional data points $\{\tilde{x}_i\}_{i=1}^n$ by applying various intensive image transformations. The assumption behind strong data augmentation is that the image transformations do not change the semantics of original samples.

In this work, we further hypothesize that strongly augmented intra-class samples in different domains can also be aligned. It is expected that strong data augmentation can further strengthen the learning of cross-domain representations, since stronger augmentation provides more diverse data samples and makes the learned aligned representations more robust for various transformations in both the source and target domains.

Following this idea, we extend the bi-directional prototypical alignment with strong data augmentation and the entire framework is termed *stabPA*. Specifically, for a source sample (x_i^s, y_i^s) and a target sample (x_i^t, \hat{y}_i^t) , we generate their augmented versions (\tilde{x}_i^s, y_i^s) and $(\tilde{x}_i^t, \hat{y}_i^t)$. Within the bi-directional prototypical alignment framework, we minimize the feature distance of a strongly augmented image to its corresponding prototype in the other domain, and maximize its distances to the prototypes of other classes. Totally, the *stabPA* loss is

$$\ell_{stabPA} = \frac{1}{|\tilde{\mathcal{B}}|} \sum_{i=1}^{|\tilde{\mathcal{B}}|} w(t) \ell_{s-t}(\tilde{x}_i^s, y_i^s) + \frac{1}{|\tilde{\mathcal{U}}|} \sum_{i=1}^{|\tilde{\mathcal{U}}|} \mathbb{1}(p(\hat{y}_i^t) > \beta) \ell_{t-s}(\tilde{x}_i^t, \hat{y}_i^t), \quad (7)$$

where $\tilde{\mathcal{B}}$ and $\tilde{\mathcal{U}}$ are the augmented base set and unlabeled auxiliary set, respectively.

To perform strong data augmentation, we apply random crop, Cutout [8], and RandAugment [7]. RandAugment comprises 14 different transformations and randomly selects a fraction of transformations for each sample. In our implementation, the magnitude for each transformation is also randomly selected, which is similar to [33].

5 Experiments

5.1 Datasets

DomainNet. DomainNet [25] is a large-scale multi-domain image dataset. It contains 345 classes in 6 different domains. In experiments, we choose the *real* domain as the source domain and choose one domain from *painting*, *clipart* and *sketch* as the target domain. We randomly split the classes into 3 parts: base set (228 classes), validation set (33 classes) and novel set (65 classes), and discard 19 classes with too few samples.

Office-Home. Office-Home [40] contains 65 object classes usually found in office and home settings. We randomly select 40 classes as the base set, 10 classes as the validation set, and 15 classes as the novel set. There are 4 domains for each class: *real*, *product*, *clipart* and *art*. We set the source domain as *real* and choose the target domain from the other three domains.

In both datasets, we construct the unlabeled auxiliary set by collecting data from the base and validation sets of the target domain and removing their labels. These unlabeled data combined with the labeled base set are used for meta-training. The validation sets in both domains are used to tune hyper-parameters. Reported results are averaged across 600 test episodes from the novel set.

5.2 Comparison Results

We compare our approach to a broad range of related methods. Methods in the first group [32,34,19,37,46,39,43,23,6] do not use the unlabeled auxiliary data during meta-training, while methods in the second group [11,35,36,33,26,18] utilize the unlabeled target images to facilitate crossing the domain gap. Note that methods in the second group are only different in representation learning, and adopt the same evaluation paradigm as ours, i.e., training a linear classifier on the support set. We also implement a baseline method, termed *stabPA*⁻, where we do not apply domain alignment and only train the feature extractor on augmented source images, which is also equivalent to applying strong augmentation to Tian et al. [37]. We set $\beta = 0.5$, $\lambda = 0.2$ and $m = 0.1$ as default for our approach. All compared methods are implemented with the same backbone and optimizer. Implementation details including augmentation techniques can be found in supplementary materials.

The comparison results are shown in Tables 1 and 2. The ‘r-r’ setting denotes all images are from the source domain, and thus is not available for methods in

Table 1. Comparison to baselines on the DomainNet dataset. We denote ‘r’ as *real*, ‘p’ as *painting*, ‘c’ as *clipart* and ‘s’ as *sketch*. We report 5-way 1-shot and 5-way 5-shot accuracies with 95% confidence interval.

Method	5-way 1-shot						
	r-r	r-p	p-r	r-c	c-r	r-s	s-r
ProtoNet [32]	63.43±0.90	45.36±0.81	45.25±0.97	44.65±0.81	47.50±0.95	39.28±0.77	42.85±0.89
RelationNet [34]	59.49±0.91	42.69±0.77	43.04±0.97	44.12±0.81	45.86±0.95	36.52±0.73	41.29±0.96
MetaOptNet [19]	61.12±0.89	44.02±0.77	44.31±0.94	42.46±0.80	46.15±0.98	36.37±0.72	40.27±0.95
Tian et al. [37]	67.18±0.87	46.69±0.86	46.57±0.99	48.30±0.85	49.66±0.98	40.23±0.73	41.90±0.86
DeepEMD [46]	67.15±0.87	47.60±0.87	47.86±1.04	49.02±0.83	50.89±1.00	42.75±0.79	46.02±0.93
ProtoNet+FWT [39]	62.38±0.89	44.40±0.80	45.32±0.97	43.95±0.80	46.32±0.92	39.28±0.74	42.18±0.95
ProtoNet+ATA [43]	61.97±0.87	45.59±0.84	45.90±0.94	44.28±0.83	47.69±0.90	39.87±0.81	43.64±0.95
S2M2 [23]	67.07±0.84	46.84±0.82	47.03±0.95	47.75±0.83	48.27±0.91	39.78±0.76	40.11±0.91
Meta-Baseline [6]	69.46±0.91	48.76±0.85	48.90±1.12	49.96±0.85	52.67±1.08	43.08±0.80	46.22±1.04
<i>stabPA⁻ (Ours)</i>	68.48±0.87	48.65±0.89	49.14±0.88	45.86±0.85	48.31±0.92	41.74±0.78	42.17±0.95
DANN [11]	-	45.94±0.84	46.85±0.97	47.31±0.86	50.02±0.94	42.44±0.79	43.66±0.92
PCT [35]	-	47.14±0.89	47.31±1.04	50.04±0.85	49.83±0.98	39.10±0.76	39.92±0.95
Mean Teacher [36]	-	46.92±0.83	46.84±0.96	48.48±0.81	49.60±0.97	43.39±0.81	44.52±0.89
FixMatch [33]	-	48.86±0.87	49.15±0.93	48.70±0.82	49.18±0.93	44.48±0.80	45.97±0.95
STARTUP [26]	-	47.53±0.88	47.58±0.98	49.24±0.87	51.32±0.98	43.78±0.82	45.23±0.96
DDN [18]	-	48.83±0.84	48.11±0.91	48.25±0.83	48.46±0.93	43.60±0.79	43.99±0.91
<i>stabPA (Ours)</i>	-	53.86±0.89	54.44±1.00	56.12±0.83	56.57±1.02	50.85±0.86	51.71±1.01
Method	5-way 5-shot						
	r-r	r-p	p-r	r-c	c-r	r-s	s-r
ProtoNet [32]	82.79±0.58	57.23±0.79	65.60±0.95	58.04±0.81	65.91±0.78	51.68±0.81	59.46±0.85
RelationNet [34]	77.68±0.62	52.63±0.74	61.18±0.90	57.24±0.80	62.65±0.81	47.32±0.75	56.39±0.88
MetaOptNet [19]	80.93±0.60	56.34±0.76	63.20±0.89	57.92±0.79	63.51±0.82	48.20±0.79	55.65±0.85
Tian et al. [37]	84.50±0.55	56.87±0.84	63.90±0.95	59.67±0.84	65.33±0.80	50.41±0.80	56.95±0.84
DeepEMD [46]	82.79±0.56	56.62±0.78	63.86±0.93	60.43±0.82	67.46±0.78	51.66±0.80	60.39±0.87
ProtoNet+FWT [39]	82.42±0.55	57.18±0.77	65.64±0.93	57.42±0.77	65.11±0.83	50.69±0.77	59.58±0.84
ProtoNet+ATA [43]	81.96±0.57	57.69±0.83	64.96±0.93	56.90±0.84	64.08±0.86	51.67±0.80	60.78±0.86
S2M2 [23]	85.79±0.52	58.79±0.81	65.67±0.90	60.63±0.83	63.57±0.88	49.43±0.79	54.45±0.89
Meta-Baseline [6]	83.74±0.58	56.07±0.79	65.70±0.99	58.84±0.80	67.89±0.91	50.27±0.76	61.88±0.94
<i>stabPA⁻ (Ours)</i>	85.98±0.51	59.92±0.85	67.10±0.93	57.10±0.88	62.90±0.83	51.03±0.85	57.11±0.93
DANN [11]	-	56.83±0.86	64.29±0.94	59.42±0.84	66.87±0.78	53.47±0.75	60.14±0.81
PCT [35]	-	56.38±0.87	64.03±0.99	61.15±0.80	66.19±0.82	46.77±0.74	53.91±0.90
Mean Teacher [36]	-	57.74±0.84	64.97±0.94	61.54±0.84	67.39±0.89	54.57±0.79	60.04±0.86
FixMatch [33]	-	61.62±0.79	67.46±0.89	61.94±0.82	66.72±0.81	55.26±0.83	62.46±0.87
STARTUP [26]	-	58.13±0.82	65.27±0.92	61.51±0.86	67.95±0.78	54.89±0.81	61.97±0.88
DDN [18]	-	61.98±0.82	67.69±0.88	61.07±0.84	65.58±0.79	54.35±0.83	60.37±0.88
<i>stabPA (Ours)</i>	-	65.65±0.74	73.63±0.82	67.32±0.80	74.41±0.76	61.37±0.82	68.93±0.87

the second group. In Table 1, we can see that the performance of conventional FSL methods drops quickly when there is a domain shift between support and query sets. The proposed *stabPA* leveraging unlabeled target images for domain alignment can alleviate this problem, improving the previous best FSL baseline [6] by 7.05% across 6 CSCS-FSL situations. Similar results can be found on the Office-Home dataset in Table 2, where the *stabPA* outperforms the previous best FSL method, S2M2 [23], by 3.90% on average. When comparing our approach with methods in the second group, we find that the *stabPA* outperforms them in all situations, improving 5-shot accuracy by 5.98% over the previous best method FixMatch [33] on DomainNet. These improvements indicate that

Table 2. Comparison results on Office-Home. We denote ‘r’ as *real*, ‘p’ as *product*, ‘c’ as *clipart* and ‘a’ as *art*. Accuracies are reported with 95% confidence intervals.

Method	5-way 1-shot						
	r-r	r-p	p-r	r-c	c-r	r-a	a-r
ProtoNet [32]	35.24 \pm 0.63	30.72 \pm 0.62	30.27 \pm 0.62	28.52 \pm 0.58	28.44 \pm 0.63	26.80 \pm 0.47	27.31 \pm 0.58
RelationNet [34]	34.86 \pm 0.63	28.28 \pm 0.62	27.59 \pm 0.56	27.66 \pm 0.58	25.86 \pm 0.60	25.98 \pm 0.54	27.83 \pm 0.63
MetaOptNet [19]	36.77 \pm 0.65	33.34 \pm 0.69	33.28 \pm 0.65	28.78 \pm 0.53	28.70 \pm 0.64	29.45 \pm 0.69	28.36 \pm 0.64
Tian et al. [37]	39.53 \pm 0.67	33.88 \pm 0.69	33.98 \pm 0.67	30.44 \pm 0.60	30.86 \pm 0.66	30.26 \pm 0.57	30.30 \pm 0.62
DeepEMD [46]	41.19 \pm 0.71	34.27 \pm 0.72	35.19 \pm 0.71	30.92 \pm 0.62	31.82 \pm 0.70	31.05 \pm 0.59	31.07 \pm 0.63
ProtoNet+FWT [39]	35.43 \pm 0.64	32.18 \pm 0.67	30.92 \pm 0.61	28.75 \pm 0.62	27.93 \pm 0.63	27.58 \pm 0.52	28.37 \pm 0.65
ProtoNet+ATA [43]	35.67 \pm 0.66	31.56 \pm 0.68	30.40 \pm 0.62	27.20 \pm 0.56	26.61 \pm 0.62	27.88 \pm 0.55	28.48 \pm 0.65
S2M2 [23]	41.92 \pm 0.68	35.46\pm0.74	35.21 \pm 0.70	31.84\pm0.66	31.96\pm0.66	30.36 \pm 0.59	30.88 \pm 0.65
Meta-Baseline [6]	38.88 \pm 0.67	33.44 \pm 0.72	33.73 \pm 0.68	30.41 \pm 0.61	30.43 \pm 0.67	30.00 \pm 0.58	30.31 \pm 0.64
<i>stabPA</i> ⁻ (<i>Ours</i>)	43.43\pm0.69	35.16 \pm 0.72	35.74\pm0.68	31.16 \pm 0.66	30.44 \pm 0.64	32.09\pm0.62	31.71\pm0.67
DANN [11]	-	33.41 \pm 0.71	33.60 \pm 0.66	30.98 \pm 0.64	30.81 \pm 0.70	31.67 \pm 0.60	32.07 \pm 0.64
PCT [35]	-	35.53 \pm 0.73	35.58 \pm 0.71	28.83 \pm 0.58	28.44 \pm 0.67	31.56 \pm 0.58	31.59 \pm 0.65
Mean Teacher [36]	-	33.24 \pm 0.70	33.13 \pm 0.67	31.34 \pm 0.62	30.91 \pm 0.67	30.98 \pm 0.60	31.57 \pm 0.61
FixMatch [33]	-	36.05\pm0.73	35.83\pm0.76	33.79\pm0.64	33.20\pm0.74	31.81 \pm 0.60	32.32 \pm 0.66
STARTUP [26]	-	34.62 \pm 0.74	34.80 \pm 0.68	30.70 \pm 0.63	30.17 \pm 0.68	32.06\pm0.59	32.40\pm0.66
<i>stabPA</i> (<i>Ours</i>)	-	38.02\pm0.76	38.09\pm0.82	35.44\pm0.76	34.74\pm0.76	34.81\pm0.69	35.18\pm0.72
Method	5-way 5-shot						
	r-r	r-p	p-r	r-c	c-r	r-a	a-r
ProtoNet [32]	49.21 \pm 0.59	39.74 \pm 0.64	38.98 \pm 0.64	34.81 \pm 0.59	35.85 \pm 0.59	34.56 \pm 0.58	36.27 \pm 0.66
RelationNet [34]	47.02 \pm 0.57	33.95 \pm 0.60	32.78 \pm 0.59	33.58 \pm 0.60	30.15 \pm 0.55	30.44 \pm 0.55	35.42 \pm 0.70
MetaOptNet [19]	52.00 \pm 0.59	43.21 \pm 0.69	42.97 \pm 0.63	36.48 \pm 0.57	36.56 \pm 0.65	36.75 \pm 0.63	38.48 \pm 0.68
Tian et al. [37]	56.89 \pm 0.61	45.79 \pm 0.69	44.27 \pm 0.63	38.27 \pm 0.64	38.99 \pm 0.63	38.80 \pm 0.61	41.56 \pm 0.72
DeepEMD [46]	58.76 \pm 0.61	47.47 \pm 0.71	45.39 \pm 0.65	38.87 \pm 0.63	40.06 \pm 0.66	39.20 \pm 0.58	41.62 \pm 0.72
ProtoNet+FWT [39]	51.40 \pm 0.61	41.50 \pm 0.68	40.32 \pm 0.60	36.07 \pm 0.62	35.80 \pm 0.60	34.60 \pm 0.56	37.36 \pm 0.67
ProtoNet+ATA [43]	51.19 \pm 0.63	41.19 \pm 0.68	38.06 \pm 0.61	32.74 \pm 0.56	33.98 \pm 0.67	35.36 \pm 0.56	36.87 \pm 0.68
S2M2 [23]	60.82 \pm 0.58	47.84 \pm 0.70	46.32\pm0.67	40.09\pm0.66	41.63\pm0.64	40.01 \pm 0.60	42.68 \pm 0.67
Meta-Baseline [6]	55.75 \pm 0.60	45.33 \pm 0.73	42.62 \pm 0.63	37.29 \pm 0.60	38.21 \pm 0.66	38.35 \pm 0.62	41.54 \pm 0.71
<i>stabPA</i> ⁻ (<i>Ours</i>)	61.87\pm0.57	48.02\pm0.73	46.27 \pm 0.67	38.22 \pm 0.66	39.88 \pm 0.63	41.75\pm0.59	44.09\pm0.69
DANN [11]	-	45.09 \pm 0.48	42.71 \pm 0.65	39.11 \pm 0.61	39.49 \pm 0.69	41.40 \pm 0.59	43.68 \pm 0.73
PCT [35]	-	48.06 \pm 0.68	46.25 \pm 0.64	34.10 \pm 0.58	35.59 \pm 0.66	40.85 \pm 0.58	43.30 \pm 0.74
Mean Teacher [36]	-	44.80 \pm 0.69	43.16 \pm 0.61	39.30 \pm 0.61	39.37 \pm 0.66	39.98 \pm 0.60	42.50 \pm 0.68
FixMatch [33]	-	48.45\pm0.70	47.17\pm0.68	43.13\pm0.67	43.20\pm0.69	41.48 \pm 0.60	44.68 \pm 0.72
STARTUP [26]	-	47.18 \pm 0.71	45.00 \pm 0.64	38.10 \pm 0.62	38.84 \pm 0.70	41.94\pm0.63	44.71\pm0.73
<i>stabPA</i> (<i>Ours</i>)	-	49.83\pm0.67	50.78\pm0.74	44.02\pm0.71	45.55\pm0.70	45.64\pm0.63	48.97\pm0.69

the proposed bi-directional prototypical alignment is an effective approach to leveraging unlabeled images to reduce domain gap for CDCS-FSL.

5.3 Analysis

Has *stabPA* learned compact and aligned representations? To verify whether *stabPA* indeed learns compact and aligned representations, we visualize the feature distributions through the meta-training process using t-SNE [22]. From Figure 3 (a)-(d), it can be seen that in the beginning, samples from different classes are heavily mixed. There are no distinct classification boundaries between classes. Besides, samples from two domains are far away from each other, indicating the existence of a considerable domain shift (such as the classes in green and orange). However, as training continues, samples from the same class begin

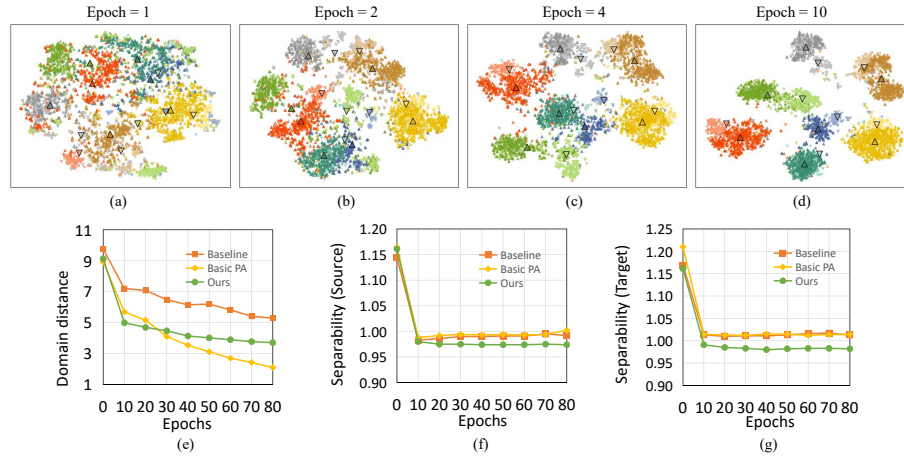


Fig. 3. (a)-(d) t-SNE visualization of feature distribution at different training epochs. Samples of the same class are painted in similar colors, where darker triangles represent source samples and lighter reverted triangles represent target samples (best viewed in color). Class centers are marked in black border. (e) Domain distance on novel classes. (f)-(g) Separability among novel classes in the source and target domains. Separability is represented by the average distance ratio, the lower the better.

to aggregate together, and the margins between different classes are increasing. In other words, compact feature representation can be obtained by the *stabPA*. Moreover, we can see that samples from different domains are grouped into their ground-truth classes, even though no label information is given for the target domain data. These observations demonstrate that *stabPA* is indeed capable to learn compact and aligned representations.

Can *stabPA* learn generalizable representations for novel classes? To validate the generalization capability of the representations learned by *stabPA*, we propose two quantitative metrics, Prototype Distance (PD) and Average Distance Ratio (ADR), which indicate the domain distance and class separability among novel classes, respectively. A small PD value means the two domains are well aligned to each other, and an ADR less than 1 indicates most samples are classified into their ground-truth classes. Detailed definitions about these two metrics can be found in the supplementary materials.

We compare *stabPA* with a FSL baseline [37] that does not leverage target images, and the BasicPA which aligns two domains by simply minimizing the point-to-point distance between prototypes in two domains [44]. The results are presented in Figure 3 (e)-(g). It is noticed that all these methods can achieve lower domain distance as training progresses, and BasicPA gets the lowest domain distance at the end. However, BasicPA does not improve the class separability as much as our approach, as shown in Figure 3 (f)-(g). The inferior class separability can be understood that BasicPA merely aims to reduce the feature distance

Table 3. The influence of the number of unlabeled samples and the number of base classes in the auxiliary set. We report average accuracy on DomainNet over 6 situations.

	FixMatch [33]	number of samples				number of base classes				
		10%	40%	70%	100%	0%	10%	40%	70%	100%
1-shot	47.72	51.76	52.97	53.42	53.92	50.74	51.59	52.48	53.24	53.92
5-shot	62.58	65.96	67.56	67.96	68.55	65.04	65.68	67.07	67.87	68.55

Table 4. Pseudo label accuracy on DomainNet real-painting.

	fixed	epoch=0	10	20	30	40	50
Top-1	23.5	4.9	24.2	30.8	34.4	35.9	37.2
Top-5	40.0	14.4	41.9	48.2	51.4	52.8	53.9

between two domains, without taking account of the intra-class variance and inter-class distances in each domain. Instead of aligning centers adopted by BasicPA, the proposed *stabPA* considers the feature-to-prototype distances across different domains and classes, so the domain alignment and class separability can be improved at the same time.

Number of unlabeled target data. To test the robustness to the number of unlabeled samples, we gradually drop data from the auxiliary set in two ways: (i) randomly drop samples from the auxiliary set, (ii) select a subset of base classes and then manually remove samples that are from the selected classes. Table 3 shows the average accuracy of *stabPA* on DomainNet over 6 situations. Unsurprisingly, decreasing the number of samples will lead to performance drop (about 2.4 points from 100% to 10%). However, with only 10% samples remained, our approach still outperforms FixMatch which uses 100% auxiliary data. We can also see that removing whole classes leads to more performance drop than randomly removing samples, probably due to the class imbalance problem caused by the former. Nevertheless, the difference is very small (about 0.3 points), indicating that our approach is robust to the number of base classes.

Pseudo label accuracy. In Table 4, we show the pseudo label accuracies of the target domain images obtained by the fixed classifier and the online classifier during the training process. We can see that the fixed classifier is better than the online classifier at the early training epochs. However, as the training goes on, the online classifier gets more accurate and outperforms the fixed classifier. This is because the online classifier is updated along the representation alignment process and gradually fits the data distribution of the target domain. After training with 50 epochs, the online classifier achieves 53.9% top-5 accuracy. To further improve the reliability of pseudo labels, we set a threshold to filter out pseudo labels with low confidence. Therefore, the actual pseudo label accuracy is higher than 53.9%.

Table 5. Ablation studies on DomainNet with 95% confidence interval.

ℓ_{s-t}	ℓ_{t-s}	aug	real-sketch		sketch-real	
			1-shot	5-shot	1-shot	5-shot
×	×	×	40.23 \pm 0.73	50.41 \pm 0.80	41.90 \pm 0.86	56.95 \pm 0.84
×	×	✓	41.74 \pm 0.78	51.03 \pm 0.85	42.17 \pm 0.95	57.11 \pm 0.93
✓	×	×	42.86 \pm 0.78	52.16 \pm 0.78	44.83 \pm 0.95	60.87 \pm 0.91
×	✓	×	44.20 \pm 0.77	54.83 \pm 0.79	44.45 \pm 0.92	61.97 \pm 0.90
✓	✓	×	47.01 \pm 0.84	56.68 \pm 0.81	47.59 \pm 1.00	64.32 \pm 0.86
✓	✓	✓	50.85\pm0.86	61.37\pm0.82	51.71\pm1.01	68.93\pm0.87

Ablation studies. We conduct ablation studies on key components of the proposed *stabPA*. The results on DomainNet are shown in Table 5. As all key components are removed (the first row), our approach is similar to Tian et al. [37] that trains feature extractor with only the source data. When the unlabeled target data are available, applying either source-to-target alignment or target-to-source alignment can improve the performance evidently. Interestingly, we can see that the target-to-source alignment is more effective than the source-to-target alignment (about 1.2 points on average). This is probably because the source prototypes estimated by the ground truth labels are more accurate than the target prototypes estimated by the pseudo labels. Improving the quality of target prototypes may reduce this gap. Combing these two alignments together, we can get better results, indicating that the two alignments are complementary to each other. Finally, the best results are obtained by combining the strong data augmentation techniques, verifying that strong data augmentation can further strengthen the cross-domain alignment.

6 Conclusions

In this work, we have investigated a new problem in FSL, namely CDCS-FSL, where a domain gap exists between the support set and query set. To tackle this problem, we have proposed *stabPA*, a prototype-based domain alignment framework to learn compact and cross-domain aligned representations. On two widely-used multi-domain datasets, we have compared our approach to multiple elaborated baselines. Extensive experimental results have demonstrated the advantages of our approach. Through more in-depth analysis, we have validated the generalization capability of the representations learned by *stabPA* and the effectiveness of each component of the proposed model.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under Grants 61721004, 61976214, 62076078, 62176246 and in part by the CAS-AIR.

References

1. Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., Vaughan, J.W.: A theory of learning from different domains. *Machine Learning* **79**(1), 151–175 (2010)
2. Bengio, Y., Louradour, J., Collobert, R., Weston, J.: Curriculum learning. In: *International Conference on Machine Learning* (2009)
3. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: *International Conference on Machine Learning*. pp. 1597–1607. PMLR (2020)
4. Chen, W.Y., Liu, Y.C., Kira, Z., Wang, Y.C.F., Huang, J.B.: A closer look at few-shot classification. In: *International Conference on Learning Representations* (2019)
5. Chen, W., Si, C., Wang, W., Wang, L., Wang, Z., Tan, T.: Few-shot learning with part discovery and augmentation from unlabeled images. *arXiv preprint arXiv:2105.11874* (2021)
6. Chen, Y., Liu, Z., Xu, H., Darrell, T., Wang, X.: Meta-baseline: Exploring simple meta-learning for few-shot learning. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 9062–9071 (2021)
7. Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V.: Randaugment: Practical automated data augmentation with a reduced search space. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. pp. 702–703 (2020)
8. DeVries, T., Taylor, G.W.: Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552* (2017)
9. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: *International Conference on Machine Learning*. pp. 1126–1135. PMLR (2017)
10. Fu, Y., Fu, Y., Jiang, Y.G.: Meta-fdmixup: Cross-domain few-shot learning guided by labeled target data. In: *ACMMM* (2021)
11. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., Lempitsky, V.: Domain-adversarial training of neural networks. *The Journal of Machine Learning Research* **17**(1), 2096–2030 (2016)
12. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. *Advances in Neural Information Processing Systems* **27** (2014)
13. Gretton, A., Borgwardt, K., Rasch, M., Schölkopf, B., Smola, A.: A kernel method for the two-sample-problem. *NeurIPS* (2006)
14. Guan, J., Zhang, M., Lu, Z.: Large-scale cross-domain few-shot learning. In: *ACCV* (2020)
15. Guo, Y., Codella, N.C., Karlinsky, L., Codella, J.V., Smith, J.R., Saenko, K., Rosing, T., Feris, R.: A broader study of cross-domain few-shot learning. In: *Proceedings of the European conference on computer vision (ECCV)*. pp. 124–141. Springer (2020)
16. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9729–9738 (2020)
17. Hoffman, J., Tzeng, E., Park, T., Zhu, J.Y., Isola, P., Saenko, K., Efros, A., Darrell, T.: Cycada: Cycle-consistent adversarial domain adaptation. In: *International Conference on Machine Learning*. pp. 1989–1998. PMLR (2018)

18. Islam, A., Chen, C.F.R., Panda, R., Karlinsky, L., Feris, R., Radke, R.J.: Dynamic distillation network for cross-domain few-shot recognition with unlabeled data. *Advances in Neural Information Processing Systems* **34**, 3584–3595 (2021)
19. Lee, K., Maji, S., Ravichandran, A., Soatto, S.: Meta-learning with differentiable convex optimization. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2019)
20. Liang, H., Zhang, Q., Dai, P., Lu, J.: Boosting the generalization capability in cross-domain few-shot learning via noise-enhanced supervised autoencoder. In: *ICCV* (2021)
21. Long, M., CAO, Z., Wang, J., Jordan, M.I.: Conditional adversarial domain adaptation. In: *Advances in Neural Information Processing Systems* (2018)
22. Van der Maaten, L., Hinton, G.: Visualizing data using t-sne. *Journal of Machine Learning Research* **9**(11) (2008)
23. Mangla, P., Kumari, N., Sinha, A., Singh, M., Krishnamurthy, B., Balasubramanian, V.N.: Charting the right manifold: Manifold mixup for few-shot learning. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 2218–2227 (2020)
24. Panareda Busto, P., Gall, J.: Open set domain adaptation. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 754–763 (2017)
25. Peng, X., Bai, Q., Xia, X., Huang, Z., Saenko, K., Wang, B.: Moment matching for multi-source domain adaptation. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 1406–1415 (2019)
26. Phoo, C.P., Hariharan, B.: Self-training for few-shot transfer across extreme task differences. In: *International Conference on Learning Representations* (2021)
27. Qiao, S., Liu, C., Shen, W., Yuille, A.L.: Few-shot image recognition by predicting parameters from activations. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 7229–7238 (2018)
28. Ravi, S., Larochelle, H.: Optimization as a model for few-shot learning. In: *International Conference on Learning Representations* (2017)
29. Rusu, A.A., Rao, D., Sygnowski, J., Vinyals, O., Pascanu, R., Osindero, S., Hadsell, R.: Meta-learning with latent embedding optimization. In: *International Conference on Learning Representations* (2019)
30. Saito, K., Kim, D., Sclaroff, S., Saenko, K.: Universal domain adaptation through self supervision. *arXiv preprint arXiv:2002.07953* (2020)
31. Saito, K., Yamamoto, S., Ushiku, Y., Harada, T.: Open set domain adaptation by backpropagation. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 153–168 (2018)
32. Snell, J., Swersky, K., Zemel, R.: Prototypical networks for few-shot learning. In: *Advances in Neural Information Processing Systems*. vol. 30 (2017)
33. Sohn, K., Berthelot, D., Li, C.L., Zhang, Z., Carlini, N., Cubuk, E.D., Kurakin, A., Zhang, H., Raffel, C.: Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *arXiv preprint arXiv:2001.07685* (2020)
34. Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P.H., Hospedales, T.M.: Learning to compare: Relation network for few-shot learning. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1199–1208 (2018)
35. Tanwisuth, K., FAN, X., Zheng, H., Zhang, S., Zhang, H., Chen, B., Zhou, M.: A prototype-oriented framework for unsupervised domain adaptation. In: *NeurIPS* (2021)
36. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In: *Advances in Neural Information Processing Systems*. vol. 30 (2017)

37. Tian, Y., Wang, Y., Krishnan, D., Tenenbaum, J.B., Isola, P.: Rethinking few-shot image classification: a good embedding is all you need? In: Proceedings of the European Conference on Computer Vision (ECCV) (2020)
38. Tsai, Y.H., Hung, W.C., Schuler, S., Sohn, K., Yang, M.H., Chandraker, M.: Learning to adapt structured output space for semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7472–7481 (2018)
39. Tseng, H.Y., Lee, H.Y., Huang, J.B., Yang, M.H.: Cross-domain few-shot classification via learned feature-wise transformation. In: ICLR (2020)
40. Venkateswara, H., Eusebio, J., Chakraborty, S., Panchanathan, S.: Deep hashing network for unsupervised domain adaptation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5018–5027 (2017)
41. Vilalta, R., Drissi, Y.: A perspective view and survey of meta-learning. *Artificial Intelligence Review* **18**(2), 77–95 (2002)
42. Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al.: Matching networks for one shot learning. *Advances in neural information processing systems* **29**, 3630–3638 (2016)
43. Wang, H., Deng, Z.H.: Cross-domain few-shot classification via adversarial task augmentation. In: IJCAI (2021)
44. Xie, S., Zheng, Z., Chen, L., Chen, C.: Learning semantic representations for unsupervised domain adaptation. In: International Conference on Machine Learning. pp. 5423–5432. PMLR (2018)
45. You, K., Long, M., Cao, Z., Wang, J., Jordan, M.I.: Universal domain adaptation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2720–2729 (2019)
46. Zhang, C., Cai, Y., Lin, G., Shen, C.: Deepemd: Few-shot image classification with differentiable earth mover’s distance and structured classifiers. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (June 2020)
47. Zhang, P., Zhang, B., Zhang, T., Chen, D., Wang, Y., Wen, F.: Prototypical pseudo label denoising and target structure learning for domain adaptive semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12414–12424 (2021)
48. Zhang, Q., Zhang, J., Liu, W., Tao, D.: Category anchor-guided unsupervised domain adaptation for semantic segmentation. In: *Advances in Neural Information Processing Systems*. vol. 32 (2019)
49. Zheng, Z., Yang, Y.: Rectifying pseudo label learning via uncertainty estimation for domain adaptive semantic segmentation. *International Journal of Computer Vision* **129**(4), 1106–1120 (2021)