

CNG-SFDA: Clean-and-Noisy Region Guided Online-Offline Source-Free Domain Adaptation

Hyeonwoo Cho¹, Chanmin Park², Dong-Hee Kim¹, Jinyoung Kim¹, and Won Hwa Kim³

¹ VUNO Inc, ² POSCO DX, ³ POSTECH
hyeonwoo.cho@gmail.com

Abstract. Domain shift occurs when training (source) and test (target) data diverge in their distribution. Source-Free Domain Adaptation (SFDA) addresses this domain shift problem, aiming to adopt a trained model on the source domain to the target domain in a scenario where only a well-trained source model and unlabeled target data are available. In this scenario, handling false labels in the target domain is crucial because they negatively impact the model performance. To deal with this problem, we propose to update cluster prototypes (i.e., centroid of each sample cluster) and their structure in the target domain formulated by the source model in online manners. In the feature space, samples in different regions have different pseudo-label distribution characteristics affected by the cluster prototypes, and we adopt distinct training strategies for these samples by defining clean and noisy regions: we selectively train the target with clean pseudo-labels in the clean region, whereas we introduce mix-up inputs representing intermediate features between clean and noisy regions to increase the compactness of the cluster. We conducted extensive experiments on multiple datasets in online/offline SFDA settings, whose results demonstrate that our method, CNG-SFDA, achieves state-of-the-art for most cases. Code is available at <https://github.com/hyeonwoocho7/CNG-SFDA>.

Keywords: SFDA · Clean-and-Noisy Region · Online Prototype

1 Introduction

Domain shift occurs when training (source) and testing (target) data diverge in their distributions, and a trained model fails to generalize on the target data. This problem is typically dealt with Unsupervised Domain Adaptation (UDA) methods [6, 16, 35, 46, 54], assuming that one has access to both the source and target data to align their distributions. However, such a scenario is not feasible in many real-world cases where privacy issues or regulations prevent access to the source data (e.g., healthcare data or autonomous vehicles). Consequently, Source-Free Domain Adaptation (SFDA) [4, 25, 31, 64] has been proposed and studied, which utilizes the knowledge from a pre-trained source model and unlabeled target data.

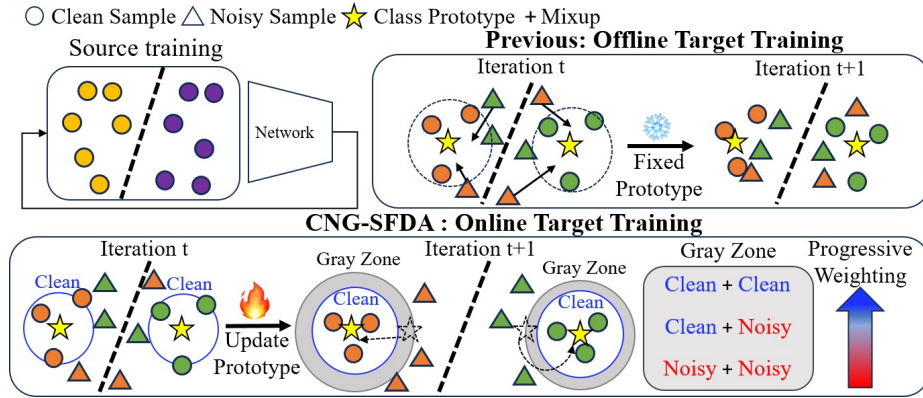


Fig. 1: The upper left presents the feature distribution of clean samples on the source domain. The upper right shows the previous method that utilizes fixed class prototypes in offline manner and learn nearest neighborhoods of class prototypes. The bottom depicts the process of CNG-SFDA that updates class prototypes in online manners and learns intermediate features in gray zone with progressive weight.

Recent advancements in SFDA primarily employ self-training methods [4, 25, 31, 33, 58, 64]. These techniques allow a source model to learn the target domain’s knowledge using pseudo-labels generated in the target domain and predominantly focus on generating reliable pseudo-labels from the neighborhood structure within the feature space. However, due to the domain shift, there are several ambiguous and noisy samples. Direct utilization of their noisy pseudo-labels in training can adversely impact the model; hence, distinguishing and appropriately managing clean (accurate) and noisy (inaccurate) pseudo-labels is highly critical.

Our preliminary studies, as well as other literature [7, 60, 62], have shown that samples with noisy pseudo-labels are mostly observed at cluster boundaries as shown in Fig. 1, underscoring the importance of considering clean and noisy pseudo-labels separately for each cluster. As illustrated in Fig. 1, several SFDA methods based self-training [8, 12, 14, 64] introduce class prototype using source classifier and utilize limited pseudo-labels near these prototypes for training. However, the prototypes in the previous approaches are fixed during target training, ignorantly using unreliable samples when target samples arrive online. To fully exploit the knowledge of the target domain in an SFDA scenario, it is crucial to update class prototypes based on target features that are added online, enabling effective learning from both reliable and unreliable pseudo-labels. Therefore, our key idea is to design a online-offline source-free domain adaptation method that effectively learns both clean and noisy pseudo-labels in an unseen target domain, thereby enabling the source model to better utilize the target domain.

In this work, we introduce Clean-and-Noisy region Guided SFDA, namely CNG-SFDA, a novel SFDA approach that focuses on addressing pseudo-labels of samples near cluster boundaries given the domain knowledge from a source model. We first partition Clean and Noisy regions within clusters formed in the target domain by leveraging online-updating cluster prototype (i.e., the centroid of each cluster) to measure the distance between the prototypes and the unlabeled target samples. Then, distinct training strategies are adopted for the clean and noisy regions due to their divergent local characteristics.

For the clean region, we conduct selective training using reliable pseudo-labels derived from predictions of the nearest neighborhoods. In the noisy region, we leverage mix-up inputs that represent intermediate features in gray zone bridging clean and noisy regions in Fig. 1, unlike other methods that directly employ or discard the unreliable pseudo-labels for training. To draw these mix-up inputs toward the clean region effectively, we assign a “clean probability” to each sample indicating the likelihood of having clean pseudo-labels and introduce a mixed-clean probability defined from the clean probability of all target samples with clean and noisy, thereby enabling effective training on samples in noisy regions.

To summarise, the main contributions of our work are:

- We propose CNG-SFDA, a novel SFDA approach that progressively update label prototypes in the target domain for better pseudo-supervision.
- We propose distinct training strategies for test samples with discriminate characteristics in their pseudo-label distributions.
- We validate CNG-SFDA on various benchmarks for online/offline SFDA in single/multi-source settings, which empirically demonstrate the superiority of our approach to SOTA by a large margin in the challenging scenarios.

2 Related Work

2.1 Domain Adaptation Scenarios

Table 1 introduces domain adaptation (DA) scenarios addressed by CNG-SFDA. **Domain Generalization (DG)**. DG [17] aims to build a model that generalizes well to unseen target domains. In this protocol, most approaches proposed to learn domain-invariant representations using multiple source domains. Recently, the more challenging Source-Free Domain Generalization (SFDG) [15] scenario has been handled which does not require access to the source domain. In the end, SFDG methods aim to build a general source model from models trained on each of the multiple source domains without access to the target domain.

Unsupervised Domain Adaptation (UDA). The main task for UDA [16,35] is to align the source distribution and target distribution from labeled source domain and unlabeled target domain. UDA leverages adversarial learning [5, 9, 16, 34, 49, 61] and generative approaches [3, 38, 40] to reduce the discrepancy between the distributions of source and target domains.

Our Approach. As shown in Table 1, CNG-SFDA addresses Source-Free Unsupervised Domain Adaptation (SFUDA) [31], Source-Free Multi-Source Unsupervised Domain Adaptation (SFMSUDA) [1], and Test-Time adaptation (TTA)

Settings	Source Domains		Target Domains Access	
	Source-Free	Multi-Source	Offline	Online
Domain Generalization (DG)	✗	✓	✗	✗
Source-Free Domain Generalization (SFDG)	✓	✓	✗	✗
Unsupervised Domain Adaptation (UDA)	✗	✗	✓	✗
Source-Free Unsupervised Domain Adaptation (SFUDA)	✓	✗	✓	✗
Source-Free Multi-Source Unsupervised Domain Adaptation (SFMSUDA)	✓	✓	✓	✗
Test-Time adaptation (TTA)	✓	✗	✗	✓

Table 1: Domain Adaptation Scenarios. We address the highlighted rows.

[51]. The goal of SFUDA is to adapt a trained model on the source domain to the target domain where the pre-trained source model and unlabeled target data are available. During adaptation, SFUDA has offline access to the target domain. In other words, the pre-trained source model learns the target domain’s knowledge by training on the target domain for several epochs. On the other hand, SFMSUDA aims to adapt a trained model on multiple source domains to an unseen target domain. In this protocol, as with SFUDA and SFMSUDA, access to the target domain is possible offline during adaptation. Finally, TTA is defined as a problem of accessing target domains online, unlike previous scenarios. Here, ‘online’ means that the pre-trained source model is accessible to the target domain during inference; hence we train only 1 epoch during adaptation’s phase.

In this work, we define SFUDA and SFMSUDA as offline source-free domain adaptation (offline SFDA) because they are takes place offline. Then, TTA is defined as online source-free domain adaptation (online SFDA).

2.2 Online-Offline Source-Free Domain Adaptation

Recently, several methods [4, 25, 31, 33, 51, 64] have been proposed for handling the domain shift in online-offline source-free domain adaptation (SFDA) settings. SHOT [31] introduces a class centroid used for pseudo-labeling for offline SFDA. Tent [51] optimizes entropy minimization for online SFDA. CoWA-JMDS [28] utilizes the joint model-data structure as sample-wise weights to represent target domain knowledge effectively. CRCo [64] proposes a probability-based similarity between target samples by embedding the source domain class relationship.

In addition, self-training methods [4, 12, 14, 25] have been proposed as well. In particular, ProxyMix and PS [12, 14] conceptualize the target domain as disjoint parts, such as proxy source and target domain. PS [14] gathers samples with low entropy and incorporates them with mix-up regularization for training. ProxyMix [12] defines the weights of the pretrained source model as class prototypes and assumes nearby target samples as the proxy source domain. Then, it applies inter-domain mix-up between the proxy source domain and the remaining target samples, as well as intra-domain mix-up among target samples, to support the linear behavior among training samples.

Limitation. The aforementioned methods selectively train only reliable samples or neglect to consider the characteristics of both reliable and unreliable samples during mix-up regularization. Additionally, since the class prototype is not up-

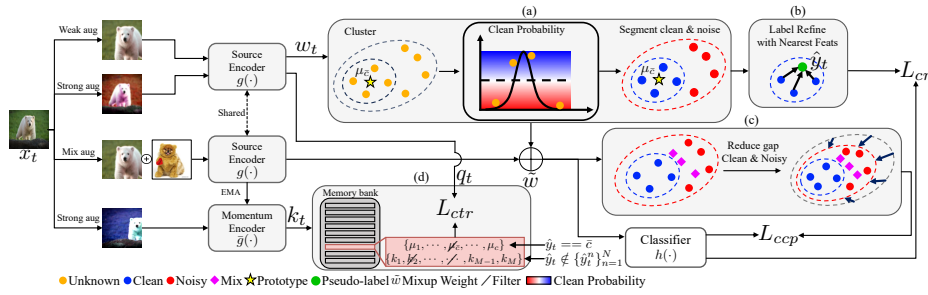


Fig 2: **Framework of CNG-SFDA** (a): partition the cluster into ‘Clean’ (close to the cluster prototype) and ‘Noisy’ (far from the cluster prototype) region based on the clean probability distribution. (b): training samples in clean regions with pseudo-labels generated from prediction of closest features. (c): reducing a gap between clean and noisy region with mix-up weight. (d): prototype/instance-aware contrastive learning.

dated during training, it may be challenging to apply in online environments where the domain changes.

Our approach. To address these limitations, we propose a method suitable for both offline and online Source-Free Domain Adaptation (SFDA) by defining class prototypes and clean-and-noisy regions that are updated online, and assigning the “clean probability” to each sample indicating the likelihood of having clean pseudo-labels. In particular, our method enhances the linear behavior for target samples by adjusting the model’s mix-up weight using the mixed-clean probability, defined from the clean probability of all target samples with clean and noisy pseudo-labels, during mix-up regularization.

3 The Proposed Method

In this work, we address Source-Free Domain Adaptation (SFDA) for the image classification task. In SFDA, we are given a source model trained on labeled source data $\mathcal{D}_s = \{x_s^i, y_s^i\}_{i=1}^{N_s}$ and unlabeled target data $\mathcal{D}_t = \{x_t^i\}_{i=1}^{N_t}$, where x_s^i, x_t^i represent the input images and N_s, N_t denote the numbers of data in the source and target domains, respectively. We focus on the closed-set SFDA where the target domain shares the same C classes as the source domain.

Given a pre-trained source model $f(\cdot)$ only, the objective of SFDA is to adapt the model to work on the unlabeled target data \mathcal{D}_t . We assume a general architecture of the source model $f(\cdot) = h(g(\cdot))$ comprising an encoder $g(\cdot)$ and a classifier $h(\cdot)$, and our method is introduced below.

Overview. In the domain adaptation phase, we employ MoCo framework following [4]. We initialize the momentum model $\bar{f}(\cdot) = \bar{h}(\bar{g}(\cdot))$ with the parameters from the source model $f(\cdot)$ at the beginning of adaptation, and $\bar{f}(\cdot)$ is updated by the exponential moving average of the source encoder $g(\cdot)$ and the source clas-

sifier $h(\cdot)$ at each mini-batch step to enhance the stability of the feature space. As illustrated in Fig. 2, the source encoder $g(\cdot)$ and the momentum encoder $\bar{g}(\cdot)$ take input images $\{\mathcal{T}(x_t^i)\}_{i=1}^{N_t}$ from the target domain $\{x_t^i\}_{i=1}^{N_t}$ with weak and two strong augmentations $\mathcal{T} = \{T_w, T_s, T_s\}$ and mixed images $\mathcal{D}_{mix} = \{\tilde{x}_t^i\}_{i=1}^{N_t}$ generated from $\{x_t^i\}_{i=1}^{N_t}$, respectively.

Then, for the purpose of categorizing the target data into individual clusters, we produce pseudo-labels $\{\hat{y}_t^i\}_{i=1}^{N_t}$ for the target data via a nearest neighbor soft voting [37] in a memory queue Q_w , representing the target feature space. As depicted in Fig 2 (a), to filter the target samples with noisy pseudo-labels \hat{y}_t in each cluster, we partition the clusters into clean (close to μ_k) and noisy (far from μ_k) regions using a clean probability derived via the distance from the cluster prototype μ_k (i.e., centroid for k -th cluster).

Next, to leverage different characteristics of the clean and noisy regions, we employ distinct training strategies for each region. For samples in the clean regions, we train $g(\cdot)$ and $h(\cdot)$ on them with the clean pseudo-labels \hat{y}_t generated through nearest neighbors features, as shown in Fig. 2 (b). In contrast, for those in the noisy regions, we employ a cluster compactness learning strategy that incorporates mix-up [59] between samples and the clean probability of each sample, as shown in (c) of Fig. 2.

Finally, as shown in Fig. 2 (d), to generate reliable \hat{y}_t from the nearest features, we adopt contrastive learning both between the prototype and sample embeddings, as well as among the samples themselves.

3.1 Distinguishing Clean and Noisy Regions in Clusters

In conventional Domain Adaptation (DA) scenarios, a significant domain discrepancy between the source and target domains often results in inaccurate pseudo-labels generated via the source model, negatively impacting the model’s performance on the target domain. In particular, SFDA is more challenging compared to traditional DA, as it offers no access to the source data, and thus efficiently leveraging the knowledge of the target domain is critical. In this work, we utilize the cluster structure and cluster prototype μ_k in the target as the knowledge of the target domain.

To address this, we propose an approach that guides the model to learn the cluster structure in the feature space with respect to the prototypes of each cluster in the target domain. Initially, to enable efficient nearest-neighbor search for generating the pseudo-labels \hat{y}_t from the cluster structures, we store the features $w_t = g(T_w(x_t))$ and their class-wise probability $p_t = \sigma(f(T_w(x_t)))$ in a memory queue $Q_w = \{w_t^j, p_t^j\}_{j=1}^M$ of the length of M , where σ denotes the softmax function. Then, in order to assign the target images x_t into each cluster, we generate its pseudo-labels \hat{y}_t by performing a soft voting [37] among the K closest neighbors of x_t in Q_w . The soft voting is done by averaging their probability outputs as:

$$\hat{p}_t^{(i,c)} = \frac{1}{K} \sum_{j=1}^K p_t^{(j,c)}, \quad (1)$$

and \hat{y}_t^i is categorized into cluster c by using $\operatorname{argmax}_c \hat{p}_t^{(i,c)}$.

To determine whether a given target sample x_t^i is clean or noisy within the cluster, as shown (a) in Fig. 2, we introduce a Gaussian Mixture Model (GMM) [44]. Given the observation x_t^i , a probability of the latent variable z_i that x_t^i belongs to k -th cluster is defined as:

$$\gamma_{ik} = p(z_i = k | x_t^i) = \frac{\pi_k \mathcal{N}(x_t^i | \mu_k, \sigma_k)}{\sum_k \pi_k \mathcal{N}(x_t^i | \mu_k, \sigma_k)} \quad (2)$$

where z_i represents the assignment of each cluster, μ_k and σ_k indicate the mean and variance of the Gaussian distribution $\mathcal{N}(x_t^i | \mu_k, \sigma_k)$ for the k -th cluster, and π_k is the prior probability of cluster k , i.e., $p(z_i = k)$. Here, assuming that z_i has uniform distribution, the Eq. (2) is expressed as:

$$\gamma_{ik} = p(z_i = k | x_t^i) = \frac{\mathcal{N}(x_t^i | \mu_k, \sigma_k)}{\sum_k \mathcal{N}(x_t^i | \mu_k, \sigma_k)}, \quad (3)$$

whose μ_k , σ_k , and z_i can be solved via EM algorithm [44].

In this work, we hypothesize that x_t^i has a clean label if \hat{y}_t^i from our model and z_i estimated via the GMM are identical. To connect \hat{y}_t^i and z_i in GMM, we adopt the μ_k , σ_k , and z_i computed using $f(\cdot)$. Specifically, we represent our predictions $p(\hat{y}_t^i = k | x_t^i, \theta)$ as the posterior probability, along with the output feature $g(x_t^i)$ as the given observation. Thus, following [44], μ_k , σ_k are represented as:

$$\mu_k = \operatorname{norm}\left(\frac{\sum_i p(\hat{y}_t^i = k | x_t^i, \theta) g(x_t^i)}{\sum_i p(\hat{y}_t^i = k | x_t^i, \theta)}\right), \quad (4)$$

$$\sigma_k = \frac{\sum_i p(\hat{y}_t^i = k | x_t^i, \theta) (g(x_t^i) - \mu_k)^T (g(x_t^i) - \mu_k)}{\sum_i p(\hat{y}_t^i = k | x_t^i, \theta)} \quad (5)$$

where $\operatorname{norm}(\cdot)$ is l2-normalization such that $\|\mu_k\|_2=1$. With μ_k, σ_k estimated from $f(\cdot)$, the Eq. (3) is defined by the probability density function:

$$\begin{aligned} \gamma_{ik} &= \frac{\exp(-(g(x_t^i) - \mu_k)^T (g(x_t^i) - \mu_k) / 2\sigma_k)}{\sum_k \exp(-(g(x_t^i) - \mu_k)^T (g(x_t^i) - \mu_k) / 2\sigma_k)} \\ &= \frac{\exp(g(x_t^i)^T \mu_k / \sigma_k)}{\sum_k \exp(g(x_t^i)^T \mu_k / \sigma_k)}. \end{aligned} \quad (6)$$

Thus, the probability of the sample being clean is defined as:

$$\gamma_{\hat{y}_t^i = z_i} = p(\hat{y}_t^i = z_i | x_t^i) = \frac{\exp(g(x_t^i)^T \mu_{z_i} / \sigma_{z_i})}{\sum_k \exp(g(x_t^i)^T \mu_k / \sigma_k)} \quad (7)$$

by using $g(x_t^i)^T \mu_k$ as the distance metric between the prototype and the sample embedding. From this probability distribution, our method enables the partition of each cluster into closer (Clean) and farther (Noisy) regions relative to the prototype through the probability threshold α . Thus, the samples in the clean and noisy region are represented as:

$$\{x_t^i\}_{i=1}^{N_t} = \begin{cases} \{x_t^{cr}\}_{cr=1}^{N_{cr}}, & \text{if } p(\hat{y}_t^i = z_i | x_t^i) \geq \alpha \\ \{x_t^{nr}\}_{nr=1}^{N_{nr}}, & \text{otherwise,} \end{cases} \quad (8)$$

where N_{cr} and N_{nr} denote the number of samples in clean and noisy regions.

3.2 Training Strategy for Clean and Noisy Regions

In preliminary study, we observed that clean and noisy regions within a cluster have divergent local characteristics. In the clean region, the samples, as well as their closest features, have reliable pseudo-labels \hat{y}_t generated by Eq. 1, as they are also in the clean regions. On the other hand, the noisy regions contain samples with incorrect \hat{y}_t , as they are located near the cluster boundaries. To deal with this, we employ distinct training strategies for $\{x_t^{cr}\}_{cr=1}^{N_{cr}}$ and $\{x_t^{nr}\}_{nr=1}^{N_{nr}}$.

Learning from clean samples. In the clean region, we exclusively train on $\{x_t^{cr}\}_{cr=1}^{N_{cr}}$ with the pseudo-labels \hat{y}_t , as they are more reliable than x_t^{nr} from the noisy regions. The loss for the samples in the clean region is given by:

$$L_{cr} = -\mathbb{E}_{x_t^{cr} \in \mathcal{X}_t} \sum_{c=1}^C \hat{y}_t^c \log p_t^c. \quad (9)$$

Learning from noisy samples. For the noisy region, we avoid directly using \hat{y}_t for $\{x_t^{nr}\}_{nr=1}^{N_{nr}}$ for training but rather leverage *mix-up* input \tilde{x}_t for the effective noisy training, as the mix-up samples lie in between the stable and noisy samples [39]. We consider $g(\tilde{x}_t)$ as the features to bridge the gap between the clean and noisy regions. Here, \tilde{x}_t^i and its pseudo-label \tilde{y}_t^i are represented as $\lambda x_t^i + (1-\lambda)x_t^j$ and $\lambda \hat{y}_t^i + (1-\lambda)\hat{y}_t^j$, where $\{x_t^i, x_t^j\} \in D_t$ and mix-up ratio $\lambda \in [0, 1]$. To effectively attract \tilde{x}_t to the clean region, we introduce the mixed-clean probability $p(\tilde{y}_t^i = z_i|\tilde{x}_t^i)$, assigning greater significance as \tilde{x}_t^i approaches the clean region. We define $p(\tilde{y}_t^i = z_i|\tilde{x}_t^i)$ as a weighted sum of the clean probabilities of x_t^i and x_t^j :

$$p(\tilde{y}_t^i = z_i|\tilde{x}_t^i) = \lambda p(z_i|x_t^i) + (1-\lambda)p(z_j|x_t^j). \quad (10)$$

We then represent the mix-up weight, denoted as \tilde{w} , using $p(\tilde{y}_t^i = z_i|\tilde{x}_t^i)$ with an exponential function:

$$\tilde{w} = \exp(p(\tilde{y}_t^i = z_i|\tilde{x}_t^i)) \quad (11)$$

which varies depending on the regions where x_t^i and x_t^j are located by Eq. 10. If both x_t^i and x_t^j are in the noisy region, \tilde{w} has a low value and vice versa for the clean region. We argue that using \tilde{x}_t with \tilde{w} enhances the compactness of the cluster since the model can learn intermediate features between the clean and noisy region, as illustrated in (c) in Fig. 2. Together with the mix-up samples, the Cluster ComPactness (CCP) loss, denoted as L_{ccp} , is given by:

$$L_{ccp} = -\tilde{w} \cdot \mathbb{E}_{\tilde{x}_t \in \tilde{\mathcal{X}}_t} \sum_{c=1}^C \tilde{y}_t^c \log \tilde{p}_t^c, \quad (12)$$

where $\tilde{p}_t = \sigma(f(\tilde{x}_t))$ are the predicted probabilities for \tilde{x}_t .

Plus, we add a regularization promoting diverse predictions similar to [4]:

$$L_{div} = \mathbb{E}_{x_t \in \mathcal{X}_t} \sum_{c=1}^C \bar{p}_t^c \log \bar{p}_t^c \quad (13)$$

where $\bar{p}_t = \mathbb{E}_{x_t \in \mathcal{X}_t} \sigma(f(T_s(x_t)))$ and L_{div} prevents the model from trusting false pseudo-labels.

3.3 Instance and Prototype-aware Contrastive Learning

In our method, we generate the pseudo-labels of the target samples from the nearest features, as discussed in Eq. 1. However, the pseudo-labels might be incorrect if the predictions of their nearest features are not correlated with the class of the target. To better align the nearest features, we introduce the contrast learning framework. Prior approaches [21, 26, 41, 65] consider each pair of instances as attractive and dispersing samples to enable instance discrimination. However, the limitation of these approaches is that it is difficult to attract different instances with similar semantic features in the same cluster, thereby posing challenges in encoding class-semantic information between instances.

To address these challenges, we contrast the prototype μ_k obtained in Section 3.1 and queries $q_t = g(T_s(x_t))$ for prototype-aware contrastive learning. As described in Fig. 2 (d), we pair q_t and $\mu_{k=\bar{c}}$ (i.e., the prototype feature of corresponding class $\bar{c} = \hat{y}_t$ of q_t), as the positive pair. In contrast, the negative pair is given by q_t and $\{\mu_{k \neq \bar{c}}\}_{k=1}^C$, which are the prototype features not corresponding to class \bar{c} . Our prototype-aware contrastive loss is defined as:

$$L_{prt} = -\log \frac{\exp q_t \cdot \mu_{k=\bar{c}}/\tau}{\sum_{k \neq \bar{c}} q_t \cdot \mu_k/\tau}, \quad (14)$$

and τ is a temperature controlling the scale of predictions.

Moreover, to learn instance-discriminate features, we leverage queries q_t and keys $k_t = \bar{g}(T_s(x_t))$ as the positive pair as in [18]. To build the negative pairs \mathcal{N}_p of q_t , we store k_t to the memory queue $Q_k = \{k_t^i\}_{i=1}^M$ with the length M , but Q_k may have features which have the same class as the corresponding class \hat{y}_t of q_t . In our work, to eliminate negative pairs with the same class, we utilize the nearest feature information in Q_k . We exclude the negative pairs when $\{\hat{y}_t^n\}_{n=1}^N$ of their N closest features are at least one the same as \hat{y}_t , as shown in Fig. 2 (d). Our instance-aware contrastive loss reflecting the behavior above is given by:

$$L_{inst} = -\log \frac{\exp q_t \cdot k_t/\tau}{\sum_{j \in \mathcal{N}_p} q_t \cdot k_t^j/\tau} \quad (15)$$

where \mathcal{N}_p is defined as $\{j | 1 \leq j \leq M, \hat{y}_t^j \notin \{\hat{y}_t^n\}_{n=1}^N\}$, and the overall contrastive loss combining L_{prt} and L_{inst} is expressed as:

$$L_{ctr} = L_{inst} + L_{prt}. \quad (16)$$

In the end, by summarizing all the losses from (9), (12), (13) and (16), we get the global loss function as below:

$$L_{overall} = \gamma_1 L_{cr} + \gamma_2 L_{ccp} + \gamma_3 L_{div} + \gamma_4 L_{ctr}, \quad (17)$$

and we set $\gamma_1 = \gamma_2 = \gamma_3 = \gamma_4 = 1.0$ during training.

4 Experiments

4.1 Experimental Setup

Datasets. **VisDA-C** [43] is a large-scale dataset consisting of 12 classes for Synthetic-to-Real object classification. The source domain has 152k synthetic images, while the target domain has 55k real-world images. We compare the per-class top-1 accuracies and their class-wise averages. **DomainNet-126** is a subset of DomainNet [42], comprising 4 domains: Real (R), Sketch (S), Clipart (C), and Painting (P) with 126 classes, following the setup in [45]. **PACS** [29] has 4 domains: Art-Painting (A), Cartoon (C), Photo (P), and Sketch (S) with 7 classes. For DomainNet-126 and PACS, we compare the accuracy for each domain shift and the average of all domain shifts. We will perform experiments that deal with the domain shifts within each dataset.

Architecture. For a fair comparison with prior methods, we adopt the same architectures and training strategies for the source model. Specifically, we use ResNet [19] as our backbones: ResNet-18 for PACS, ResNet-50 for DomainNet-126, and ResNet-101 for the VisDA-C. We add a 256-dimensional bottleneck layer, which is a fully-connected layer, followed by BatchNorm [22] after the backbone, and apply WeightNorm [47] to the classifier, as done in [4, 31, 33].

Implementation details. *For the source training*, we initialize the model with ImageNet-1K [10] pre-trained weights. We train the pre-trained model on the source data, the same as in [4]. *For offline SFDA*, we opt for the SGD optimizer with a momentum of 0.9 and a learning rate of $2e-4$ for all datasets. We set the threshold α of the clean probability to 0.5, the temperature τ to 0.07, the memory bank size M to 16384, and the momentum value for the EMA update to 0.999. *For online SFDA*, we turn on/off a soft-voting for pseudo labeling when 1024 features-probability pairs are accumulated in the memory queue. Other hyper-parameters are the same as in offline SFDA. More details in Appendix.

4.2 Experimental Results

Table. 2, 3, and 4 show the performance of CNG-SFDA and the compared methods across 3 different datasets in the single-source setting. Table. 5 shows results in the multi-source setting. In each table, † indicates the results that we reproduced from the source code, and ‘SF’ and ‘On’ denote ‘source-free’ and ‘online’, respectively. All reported results are average accuracies from three seeds.

Single-Source DA. Table. 2, 3, and 4 show the classification accuracy of CNG-SFDA and previous methods for UDA/SFDA in an offline setting. In Table. 2, CNG-SFDA outperforms UDA methods such as CAN [24] without access to the source data. In the more challenging SFDA, we surpass SOTA [33] by a 1.3% margin on per-class average accuracy in VisDA-C (Table. 2) and 0.8% margin on DomainNet-126 (Table. 3). Moreover, in PACS (Table. 4), we achieve competitive performance compared to SOTA. Unlike self-training methods (AdaContrast [4], GU-SFDA [33]) using nearest features for pseudo-labels, CNG-SFDA leverages clean probabilities within clusters, proving more effective.

Method	SF	On	plane	bcycl	bus	car	horse	knife	mcycl	person	plant	sktbrd	train	truck	Avg.
MCC(ECCV'20) [23]	✗	✗	88.7	80.3	80.5	71.5	90.1	93.2	85.0	71.6	89.4	73.8	85.0	36.9	78.8
RWOT(CVPR'20) [56]	✗	✗	95.1	80.3	83.7	90.0	92.4	68.0	92.5	82.2	87.9	78.4	90.4	68.2	84.0
Source Only	-	-	57.2	11.1	42.4	66.9	55.0	4.4	81.1	27.3	57.9	29.4	86.7	5.8	43.8
SHOT (ICML'20) [31]	✓	✗	95.3	87.5	78.7	55.6	94.1	94.2	81.4	80.0	91.8	90.7	86.5	59.8	83.0
NRC (NeurIPS'21) [58]	✓	✗	96.8	91.3	82.4	62.4	96.2	95.9	86.1	80.6	94.8	94.1	90.4	59.7	85.9
PS (AAAI'22) [14]	✓	✗	95.3	86.2	82.3	61.6	93.3	95.7	86.7	80.4	91.6	90.9	86.0	59.5	84.1
ProxyMix (NN'22) [12]	✓	✗	95.4	81.7	87.2	79.9	95.6	96.8	92.1	85.1	93.4	90.3	89.1	42.2	85.7
SFDA-DE (CVPR'22) [11]	✓	✗	95.3	91.2	77.5	72.1	95.7	97.8	85.5	86.1	95.5	93.0	86.3	61.6	86.5
CoWA-JMDS (ICML'22) [28]	✓	✗	96.2	89.7	83.9	73.8	96.4	97.4	89.3	86.8	94.6	92.1	88.7	53.8	86.9
AaD-SFDA [†] (NeurIPS'22) [57]	✓	✗	97.4	90.5	80.8	76.2	97.3	96.1	89.8	82.9	95.5	93.0	92.0	64.7	88.0
AdaContrast [†] (CVPR'22) [4]	✓	✗	97.2	83.6	84.0	77.2	96.9	94.4	91.6	84.9	94.5	93.3	94.1	45.9	86.5
NRC++ (TPAMI'23) [58]	✓	✗	97.4	91.9	88.2	83.2	97.3	96.2	90.2	81.1	96.3	94.3	91.4	49.6	88.1
C-SFDA [†] (CVPR'23) [25]	✓	✗	97.6	88.8	86.1	72.2	97.2	94.4	92.1	84.7	93.0	90.7	93.1	63.5	87.8
GU-SFDA [†] (CVPR'23) [33]	✓	✗	97.1	91.2	87.6	72.1	96.9	96.4	93.8	86.7	96.3	94.2	91.3	67.0	89.2
CNG-SFDA	✓	✗	98.0	90.2	91.3	88.9	98.1	96.2	94.7	88.6	97.7	97.3	95.3	67.0	90.5
Tent [†] (ICLR'21) [51]	✓	✓	86.9	57.7	77.4	56.8	87.3	62.4	86.6	62.9	71.2	39.9	84.8	24.7	66.5
SHOT [†] (ICML'21) [31]	✓	✓	90.5	77.0	76.2	47.5	87.9	62.1	75.9	74.4	83.3	47.0	84.2	41.6	70.6
AdaContrast [†] (CVPR'22) [4]	✓	✓	95.0	68.0	82.7	69.6	94.3	80.8	90.3	79.6	90.6	69.7	87.6	36.0	78.7
C-SFDA [†] (CVPR'23) [25]	✓	✓	95.9	75.6	88.4	68.1	95.4	86.1	94.5	82.0	89.2	80.2	87.3	43.8	82.1
TeSLA [†] (CVPR'23) [48]	✓	✓	95.4	87.4	83.8	70.1	95.1	90.0	84.8	83.2	93.6	67.9	85.4	49.3	82.2
CNG-SFDA	✓	✓	96.7	82.2	89.3	82.7	96.9	94.8	94.2	86.6	94.9	91.4	89.8	42.3	86.8

Table 2: Classification Accuracy (%) on **VisDA-C** for Single-Source Domain Adaptation (ResNet-101 Backbone).

Method	SF	On	R → C	R → P	P → C	C → S	S → P	P → R	S → R	Avg
MCC [23]	✗	✗	44.8	65.7	41.9	34.9	47.3	35.3	72.4	48.9
Source Only	-	-	55.5	62.7	53.0	46.9	50.1	46.3	75.0	55.6
SHOT [31]	✓	✗	67.7	68.4	66.9	60.1	66.1	59.9	80.8	67.1
AdaContrast [†] [4]	✓	✗	69.7	69.0	68.6	58.4	66.6	60.5	80.2	67.6
C-SFDA [†] [25]	✓	✗	70.8	71.1	68.5	62.1	67.4	62.7	80.4	69.0
GU-SFDA [†] [33]	✓	✗	74.2	70.4	68.8	64.0	67.5	65.7	76.5	69.6
CNG-SFDA	✓	✗	73.7	72.2	71.7	60.6	67.8	64.7	83.2	70.4
Tent [†] [51]	✓	✓	58.5	65.7	57.9	48.5	52.4	54.0	67.0	57.7
AdaContrast [†] [4]	✓	✓	61.1	66.9	60.8	53.4	62.7	64.5	78.9	62.6
C-SFDA [†] [25]	✓	✓	61.6	67.4	61.3	55.1	63.2	54.8	78.5	63.1
CNG-SFDA	✓	✓	67.1	69.0	65.6	59.1	66.4	59.9	81.5	66.9

Table 3: Classification Accuracy (%) on **DomainNet-126** for Single-Source Domain Adaptation (ResNet-50 Backbone).

Method	SF	On	P → A	P → C	P → S	A → P	A → C	A → S	Avg
Source Only	-	-	58.0	21.1	27.8	96.1	48.6	39.7	48.6
NEL [2]	✓	✗	82.6	80.5	32.3	98.4	84.3	56.1	72.4
GU-SFDA [†] [33]	✓	✗	88.6	82.2	69.0	96.3	84.6	73.9	82.4
CNG-SFDA	✓	✓	92.2	69.1	66.2	98.8	83.0	84.8	82.4
Tent [†] [51]	✓	✓	74.3	62.1	58.8	96.3	72.6	71.1	72.5
SHOT [†] [31]	✓	✓	76.1	59.1	60.8	96.9	73.1	73.7	73.2
CNG-SFDA	✓	✓	76.3	64.4	60.7	97.1	74.9	73.3	74.4

Table 4: Classification Accuracy (%) on **PACS** for Single-Source Domain Adaptation (ResNet-18 Backbone).

Multi-Source DA. Table 5 shows the performance of Multi-Source Unsupervised Domain Adaptation (MSUDA) on PACS. In MSUDA, each domain is treated as the target, while the source model is trained by aggregating all other domains without domain labels. In source-free MSUDA, CNG-SFDA achieves the highest accuracy by training pseudo-labels for the target with different strategies depending on whether the data is in clean or noisy regions, whereas SHOT++ [32] produces pseudo-labels using the centroid of the nearest cluster.

Online Source-Free DA. In Table 2, 3, 4, and 5, ‘online’ results show the classification accuracy of CNG-SFDA and the compared methods for online SFDA. In Table 2 and 3, CNG-SFDA significantly outperforms SOTA (4.6% and 3.8% improvements in VisDA-C and DomainNet-126). C-SFDA [25] may fail to capture all the knowledge from the target domain, particularly from the noisy data,

Method	SF	On	A	C	P	S	Avg	PL	L_{cr}	L_{ccp}	L_{div}	L_{inst}	L_{prt}	VisDA-C	DomainNet	PACS
SIB [20]	✗	✗	88.9	89.0	98.3	82.2	89.6	✓	✗	✗	✗	✗	✗	84.9	58.5	75.7
T-SVDNET [30]	✗	✗	90.4	90.6	98.5	85.4	91.2	✓	✓	✗	✗	✗	✗	88.0	65.8	78.8
iMSDA [27]	✗	✗	93.7	92.4	98.4	89.2	93.4	✓	✓	✓	✗	✗	✗	89.6	67.4	79.7
Source Only	-	-	77.9	72.8	95.7	63.5	77.5	✓	✓	✓	✓	✗	✗	89.0	67.7	80.6
SHOT [†] [31]	✓	✓	90.7	88.1	98.5	75.4	88.2	✓	✓	✓	✓	✓	✗	90.1	69.7	81.1
SHOT ^{++†} [32]	✓	✓	92.3	89.7	98.8	75.5	89.1	✓	✓	✓	✓	✓	✓	90.5	70.4	82.4
CNG-TTA	✓	✓	93.7	93.3	98.7	87.8	93.4									

Table 5: Classification Accuracy (%) on PACS for Multi-Source Domain Adaptation (ResNet-18 Backbone).
 Table 6: The effectiveness of each component across all datasets is validated by classification accuracy (%).




Fig. 3: Feature distribution of clean and noisy samples in target domain. Here, ‘clean’ and ‘noisy’ refer to cases where the Ground-Truth (GT) and the Pseudo-Label (PL) match and mismatch. Circle and triangular denote clean and noisy samples predicted by CNG-SFDA. For clean samples (circle), each color represent the GT class. For noisy samples (triangular), pink and black triangular represent cases where CNG-SFDA correctly (i.e., Prediction: Noisy, GT: Noisy) and incorrectly (i.e., Prediction: Noisy, GT: Clean) identifies noisy data.

due to its selective approach based on curriculum learning. However, CNG-SFDA utilizes the entire target domain’s knowledge by training effectively all samples. Plus, Table. 4, 5 show the effectiveness of CNG-SFDA on PACS for online SFDA.

4.3 Analysis and Discussion

Ablation study. Table. 6 shows the effectiveness of each component on all datasets. The ‘PL’ method as a baseline indicates the method generating the pseudo-labels from nearest features in Eq. 1. By applying ‘ L_{cr} ’ in Sec. 3.2, we boost the performance by 3.1%, 7.3%, and 3.1% on each domain. These results confirm that ‘ L_{cr} ’ learning only the samples with clean labels is more effective than ‘PL’ using all data for training. Plus, incorporating the cluster compactness loss ‘ L_{ccp} ’ in Sec. 3.2 and contrastive losses ‘ L_{inst} ’, ‘ L_{prt} ’ in Sec. 3.3 consistently contributes to performance improvement on all datasets.

Effect on clean and noisy region. Fig. 3 shows the effectiveness of distinguishing clean and noisy data within clusters. From the gray boxes in (a) and (c) of Fig. 3, we observed that applying CNG-SFDA resulted in a reduction of

Method	real \rightarrow	clipart \rightarrow	painting \rightarrow	sketch \rightarrow	Avg.	BackBone			
						ResNet18	ResNet50	ViT-B/16	
Source only	45.3	49.3	41.7	44.8	45.3	ERM [50]	82.1	84.6	87.1
AdaContrast [†] (CVPR'22)	38.3	38.8	33.7	32.7	35.9	DGCM (ICML'21) [36]	85.5	87.5	-
CoTTA (CVPR'22)	43.4	43.0	36.4	36.3	39.8	FACT (CVPR'21) [55]	84.5	88.2	-
SAR [†] (ICLR'23)	42.8	44.3	37.7	37.7	40.6	T3A (NeurIPS'21) [55]	81.7	84.5	86.0
RoTTA [†] (CVPR'23)	43.2	44.0	38.6	37.6	40.8	AdaNPC (ICML'23) [63]	83.1	85.7	88.7
RMT [†] (CVPR'23)	37.6	38.6	33.3	33.3	35.7	TSD [†] (CVPR'23) [53]	87.8	89.4	90.2
CNG-SFDA	37.7	39.2	31.5	33.7	35.5	CNG-SFDA (Ours)	86.5	90.4	91.7

Table 7: Classification Error Rate (%) on DomainNet-126 for Continual TTA.

Table 8: Classification Accuracy (%) on PACS with different backbone architecture for Multi-Source Unsupervised Domain Adaptation.



Fig. 4: Qualitative evaluation of samples in predicted clean and noisy regions by CNG-SFDA at each target domain (each row). PL and CP are Pseudo-Label and Clean Probability estimated by CNG-SFDA, respectively.

noisy samples and clustering around clean samples. Additionally, as shown in (a), CNG-SFDA correctly identified noisy samples early in training, effectively utilizing them with clean samples. Interestingly, as seen in (d), samples correctly predicted to have noisy pseudo-labels were indeed difficult samples to classify.

Capability of Continual Test-Time Adaptation. We also conducted experiments on DomainNet-126 in continual Test-time adaptation settings [52]. Continual TTA is particularly useful for scenarios where the target data distribution is not known in advance or is constantly changing. Consequently, the source model is dynamically adapted to a sequence of test domains in an online manner. For DomainNet-126, we randomly generated four different domain se-




Fig. 5: Robustness analysis of different hyper-parameters.

Method	VisDa-C	DomainNet	PACS
w/o weighting	89.7	69.2	80.7
w/ linear weighting	90.0	70.2	82.1
w/ exponential weighting	90.5	70.4	82.4

Table 9: Classification Accuracy (%) on **all dataset** comparing linear vs exponential weighting in Eq. 11.

quences, ensuring that each domain serves as the source domain once. Table. 7 reveals that CNG-SFDA surpasses SOTA [13] for the continual TTA, proving that its robustness in dynamic scenarios where target distributions are changing.

Qualitative evaluation on distinguishing clean and noisy regions. Fig. 4 showcases samples with clean and noisy pseudo-labels predicted by CNG-SFDA, along with the PL and CP of each sample. We observed that noisy images tend to have lower CP values, indicating difficulty in classification. These noisy images often contain objects outside the class or are challenging to perceptually classify. On the other hand, clean images mostly comprise easy samples containing only objects corresponding to the ground truth classes. This qualitative analysis confirms the reliability of CNG-SFDA in identifying clean and noisy images.

Scalability on different models. In Table. 8, we validate our method using different backbones to ensure it is the model-agnostic method. CNG-SFDA outperforms better than the previous methods on PACS for MSUDA. Through these results, we confirm that our method works well with different backbones.

Robustness of hyper-parameters. Fig. 5 shows the robustness of CNG-SFDA versus the clean probability threshold α and K in Eq. 14. In addition, we analyze the sensitivity of the memory queue length, which is the criterion for turning on-off soft voting for pseudo-labeling for online SFDA. These results demonstrate the robustness of CNG-SFDA to the hyper-parameters.

Linear vs Exponential weight. Table. 9 presents the effectiveness of the exponential weighting in Eq. 11, comparing linear one and without one. CNG-SFDA outperforms previous SFDA methods, even with linear weighting.

5 Conclusion

We introduced CNG-SFDA, a novel online-offline SFDA approach in image classification. To ensure the comprehensive transfer of target domain knowledge to the source model, effective learning of clean and noisy pseudo-labels from the target domain is crucial. To address this, we propose to distinguish clean and noisy regions at the cluster level. Then, we apply different training strategies to each region by considering their respective characteristics. Our experiments and analysis prove the effectiveness of CNG-SFDA compared to other approaches.

Acknowledgments. This research was conducted with resources and endless support from VUNO Inc, and Won Hwa Kim was supported by Graduate School of AI at POSTECH (IITP-2019-0-01906).

References

1. Ahmed, S.M., Raychaudhuri, D.S., Paul, S., Oymak, S., Roy-Chowdhury, A.K.: Unsupervised multi-source domain adaptation without access to source data. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10103–10112 (2021)
2. Ahmed, W., Morerio, P., Murino, V.: Cleaning noisy labels by negative ensemble learning for source-free unsupervised domain adaptation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 1616–1625 (2022)
3. Bousmalis, K., Silberman, N., Dohan, D., Erhan, D., Krishnan, D.: Unsupervised pixel-level domain adaptation with generative adversarial networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3722–3731 (2017)
4. Chen, D., Wang, D., Darrell, T., Ebrahimi, S.: Contrastive test-time adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 295–305 (2022)
5. Chen, L., Chen, H., Wei, Z., Jin, X., Tan, X., Jin, Y., Chen, E.: Reusing the task-specific classifier as a discriminator: Discriminator-free adversarial domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7181–7190 (2022)
6. Cho, H., Nishimura, K., Watanabe, K., Bise, R.: Effective pseudo-labeling based on heatmap for unsupervised domain adaptation in cell detection. *Medical Image Analysis* **79**, 102436 (2022)
7. Choi, J., Jeong, M., Kim, T., Kim, C.: Pseudo-labeling curriculum for unsupervised domain adaptation. arXiv preprint arXiv:1908.00262 (2019)
8. Choi, S., Yang, S., Choi, S., Yun, S.: Improving test-time adaptation via shift-agnostic weight regularization and nearest source prototypes. In: European Conference on Computer Vision. pp. 440–458. Springer (2022)
9. Cui, S., Wang, S., Zhuo, J., Su, C., Huang, Q., Tian, Q.: Gradually vanishing bridge for adversarial domain adaptation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 12455–12464 (2020)
10. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
11. Ding, N., Xu, Y., Tang, Y., Xu, C., Wang, Y., Tao, D.: Source-free domain adaptation via distribution estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7212–7222 (2022)
12. Ding, Y., Sheng, L., Liang, J., Zheng, A., He, R.: Proxymix: Proxy-based mixup training with label refinery for source-free domain adaptation. *Neural Networks* **167**, 92–103 (2023)
13. Döbler, M., Marsden, R.A., Yang, B.: Robust mean teacher for continual and gradual test-time adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7704–7714 (2023)
14. Du, Y., Yang, H., Chen, M., Luo, H., Jiang, J., Xin, Y., Wang, C.: Generation, augmentation, and alignment: A pseudo-source domain based method for source-free domain adaptation. *Machine Learning* pp. 1–21 (2023)
15. Frikha, A., Chen, H., Krompaß, D., Runkler, T., Tresp, V.: Towards data-free domain generalization. In: Asian Conference on Machine Learning. pp. 327–342. PMLR (2023)

16. Ganin, Y., Lempitsky, V.: Unsupervised domain adaptation by backpropagation. In: International conference on machine learning. pp. 1180–1189. PMLR (2015)
17. Gulrajani, I., Lopez-Paz, D.: In search of lost domain generalization. arXiv preprint arXiv:2007.01434 (2020)
18. 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)
19. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. In: Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14. pp. 630–645. Springer (2016)
20. Hu, S.X., Moreno, P.G., Xiao, Y., Shen, X., Obozinski, G., Lawrence, N.D., Dami-anou, A.: Empirical bayes transductive meta-learning with synthetic gradients. arXiv preprint arXiv:2004.12696 (2020)
21. Huynh, T., Kornblith, S., Walter, M.R., Maire, M., Khademi, M.: Boosting contrastive self-supervised learning with false negative cancellation. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision. pp. 2785–2795 (2022)
22. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: International conference on machine learning. pp. 448–456. pmlr (2015)
23. Jin, Y., Wang, X., Long, M., Wang, J.: Minimum class confusion for versatile domain adaptation. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16. pp. 464–480. Springer (2020)
24. Kang, G., Jiang, L., Yang, Y., Hauptmann, A.G.: Contrastive adaptation network for unsupervised domain adaptation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4893–4902 (2019)
25. Karim, N., Mithun, N.C., Rajvanshi, A., Chiu, H.p., Samarasekera, S., Rahnavard, N.: C-sfda: A curriculum learning aided self-training framework for efficient source free domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 24120–24131 (2023)
26. Kim, S., Lee, G., Bae, S., Yun, S.Y.: Mixco: Mix-up contrastive learning for visual representation. arXiv preprint arXiv:2010.06300 (2020)
27. Kong, L., Xie, S., Yao, W., Zheng, Y., Chen, G., Stojanov, P., Akinwande, V., Zhang, K.: Partial disentanglement for domain adaptation. In: International Conference on Machine Learning. pp. 11455–11472. PMLR (2022)
28. Lee, J., Jung, D., Yim, J., Yoon, S.: Confidence score for source-free unsupervised domain adaptation. In: International Conference on Machine Learning. pp. 12365–12377. PMLR (2022)
29. Li, D., Yang, Y., Song, Y.Z., Hospedales, T.M.: Deeper, broader and artier domain generalization. In: Proceedings of the IEEE international conference on computer vision. pp. 5542–5550 (2017)
30. Li, R., Jia, X., He, J., Chen, S., Hu, Q.: T-svdnet: Exploring high-order prototypical correlations for multi-source domain adaptation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9991–10000 (2021)
31. Liang, J., Hu, D., Feng, J.: Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In: International conference on machine learning. pp. 6028–6039. PMLR (2020)

32. Liang, J., Hu, D., Wang, Y., He, R., Feng, J.: Source data-absent unsupervised domain adaptation through hypothesis transfer and labeling transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **44**(11), 8602–8617 (2021)
33. Litrico, M., Del Bue, A., Morerio, P.: Guiding pseudo-labels with uncertainty estimation for source-free unsupervised domain adaptation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 7640–7650 (2023)
34. Long, M., Cao, Z., Wang, J., Jordan, M.I.: Conditional adversarial domain adaptation. *Advances in neural information processing systems* **31** (2018)
35. Long, M., Zhu, H., Wang, J., Jordan, M.I.: Unsupervised domain adaptation with residual transfer networks. *Advances in neural information processing systems* **29** (2016)
36. Mahajan, D., Tople, S., Sharma, A.: Domain generalization using causal matching. In: *International Conference on Machine Learning*. pp. 7313–7324. PMLR (2021)
37. Mitchell, H.B., Schaefer, P.A.: A “soft” k-nearest neighbor voting scheme. *International journal of intelligent systems* **16**(4), 459–468 (2001)
38. Murez, Z., Kolouri, S., Kriegman, D., Ramamoorthi, R., Kim, K.: Image to image translation for domain adaptation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 4500–4509 (2018)
39. Na, J., Jung, H., Chang, H.J., Hwang, W.: Fixbi: Bridging domain spaces for unsupervised domain adaptation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 1094–1103 (2021)
40. Nam, H., Lee, H., Park, J., Yoon, W., Yoo, D.: Reducing domain gap by reducing style bias. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8690–8699 (2021)
41. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748* (2018)
42. 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)
43. Peng, X., Usman, B., Kaushik, N., Hoffman, J., Wang, D., Saenko, K.: Visda: The visual domain adaptation challenge. *arXiv preprint arXiv:1710.06924* (2017)
44. Reynolds, D.A., et al.: Gaussian mixture models. *Encyclopedia of biometrics* **741**(659-663) (2009)
45. Saito, K., Kim, D., Sclaroff, S., Darrell, T., Saenko, K.: Semi-supervised domain adaptation via minimax entropy. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 8050–8058 (2019)
46. Saito, K., Watanabe, K., Ushiku, Y., Harada, T.: Maximum classifier discrepancy for unsupervised domain adaptation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 3723–3732 (2018)
47. Salimans, T., Kingma, D.P.: Weight normalization: A simple reparameterization to accelerate training of deep neural networks. *Advances in neural information processing systems* **29** (2016)
48. Tomar, D., Vray, G., Bozorgtabar, B., Thiran, J.P.: Tesla: Test-time self-learning with automatic adversarial augmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 20341–20350 (2023)
49. Tzeng, E., Hoffman, J., Saenko, K., Darrell, T.: Adversarial discriminative domain adaptation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 7167–7176 (2017)
50. Vapnik, V.N.: *Statistical Learning Theory*. Wiley-Interscience (1998)

51. Wang, D., Shelhamer, E., Liu, S., Olshausen, B., Darrell, T.: Tent: Fully test-time adaptation by entropy minimization. arXiv preprint arXiv:2006.10726 (2020)
52. Wang, Q., Fink, O., Van Gool, L., Dai, D.: Continual test-time domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7201–7211 (2022)
53. Wang, S., Zhang, D., Yan, Z., Zhang, J., Li, R.: Feature alignment and uniformity for test time adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20050–20060 (2023)
54. Wilson, G., Cook, D.J.: A survey of unsupervised deep domain adaptation. ACM Transactions on Intelligent Systems and Technology (TIST) **11**(5), 1–46 (2020)
55. Xu, Q., Zhang, R., Zhang, Y., Wang, Y., Tian, Q.: A fourier-based framework for domain generalization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14383–14392 (2021)
56. Xu, R., Liu, P., Wang, L., Chen, C., Wang, J.: Reliable weighted optimal transport for unsupervised domain adaptation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4394–4403 (2020)
57. Yang, S., Jui, S., van de Weijer, J., et al.: Attracting and dispersing: A simple approach for source-free domain adaptation. Advances in Neural Information Processing Systems **35**, 5802–5815 (2022)
58. Yang, S., van de Weijer, J., Herranz, L., Jui, S., et al.: Exploiting the intrinsic neighborhood structure for source-free domain adaptation. Advances in neural information processing systems **34**, 29393–29405 (2021)
59. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. arXiv preprint arXiv:1710.09412 (2017)
60. 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)
61. Zhang, W., Ouyang, W., Li, W., Xu, D.: Collaborative and adversarial network for unsupervised domain adaptation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3801–3809 (2018)
62. Zhang, W., Zhu, L., Hallinan, J., Zhang, S., Makmur, A., Cai, Q., Ooi, B.C.: Boostmis: Boosting medical image semi-supervised learning with adaptive pseudo labeling and informative active annotation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20666–20676 (2022)
63. Zhang, Y., Wang, X., Jin, K., Yuan, K., Zhang, Z., Wang, L., Jin, R., Tan, T.: Adanpc: Exploring non-parametric classifier for test-time adaptation. In: International Conference on Machine Learning. pp. 41647–41676. PMLR (2023)
64. Zhang, Y., Wang, Z., He, W.: Class relationship embedded learning for source-free unsupervised domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7619–7629 (2023)
65. Zhang, Z., Chen, W., Cheng, H., Li, Z., Li, S., Lin, L., Li, G.: Divide and contrast: Source-free domain adaptation via adaptive contrastive learning. Advances in Neural Information Processing Systems **35**, 5137–5149 (2022)