# De-Confusing Pseudo-Labels in Source-Free Domain Adaptation

Idit Diamant<sup>1</sup><sup>®</sup>, Amir Rosenfeld<sup>1</sup><sup>®</sup>, Idan Achituve<sup>1</sup><sup>®</sup>, Jacob Goldberger<sup>2</sup><sup>®</sup>, and Arnon Netzer<sup>1</sup><sup>®</sup>

<sup>1</sup> Sony Semiconductor Israel

**Abstract.** Source-free domain adaptation aims to adapt a source-trained model to an unlabeled target domain without access to the source data. It has attracted growing attention in recent years, where existing approaches focus on self-training that usually includes pseudo-labeling techniques. In this paper, we introduce a novel noise-learning approach tailored to address noise distribution in domain adaptation settings and learn to **de-confuse the pseudo-labels**. More specifically, we learn a noise transition matrix of the pseudo-labels to capture the label corruption of each class and learn the underlying true label distribution. Estimating the noise transition matrix enables a better true class-posterior estimation, resulting in better prediction accuracy. We demonstrate the effectiveness of our approach when combined with several source-free domain adaptation methods: SHOT, SHOT++, and AaD. We obtain state-of-the-art results on three domain adaptation datasets: VisDA, DomainNet, and OfficeHome.

Keywords: Source-free domain adaptation · Noise learning

## 1 Introduction

Unsupervised domain adaptation (UDA) involves adjusting a predictive model trained on a labeled source domain to perform well on an unlabeled target domain despite possible domain shifts. Source-free domain adaption (SFDA) imposes an additional constraint that the source domain data are unavailable during adaptation to the target domain. Consequently, SFDA depends heavily on unsupervised learning and self-training techniques.

In SFDA, a key obstacle is to reduce the accumulation of errors caused by domain misalignment. Existing methods in SFDA focus on self-training using target pseudo-labels and entropy minimization techniques. In [28] the effective reduction of error accumulation was achieved by assigning pseudo-labels to the target data using class clusters in the penultimate layer of the model, where the pseudo-labels were determined based on distances from the centroids of the clusters. However, due to the domain shift, these generated pseudo-labels tend

to be noisy. For this reason, several methods have focused on the refinement of the target pseudo-labels during training (e.g., [5, 18, 61, 63]). One common solution is to update the pseudo-labels at each epoch, thus refining them to better match the target distribution [28, 63]. In [63], the authors handle this problem by pseudo-labeling using a strong pre-trained network and filtering out samples with low confidence.

In this work we do not refine the pseudo-labels during target training and do not make any explicit decision on the PLs correctness. Instead, we approach SFDA by adopting a learning with label-noise (LLN) perspective and suggest a novel approach tailored for domain adaptation settings to De-Confuse the **P**seudo-Labels. We call our method *DCPL*. It involves modifying the established technique of optimizing a noise transition matrix (or a confusion matrix) to adapt the classifier predictions to the pseudo-labels [45, 49]. Traditional learning-withlabel-noise (LLN) methods typically address label noise stemming from errors made by human annotators during labeling. In our specific scenario, the label noise is derived from pseudo-labels generated with a specified source model. This label noise emerges from domain shifts, where the characteristics of data in the target domain deviate from those in the source domain used for training. In this context, the outputs of the source model carry valuable information about the noise distribution. This information can be leveraged to improve the estimation of the transition matrix. We propose integrating this knowledge into the noise transition matrix learning method to achieve a more accurate prediction of the underlying true label distribution, consequently enhancing the accuracy of classifier predictions. Empirically, we demonstrate that our approach is easy to integrate with current source-free unsupervised domain adaptation (SFDA) techniques and yields improved prediction accuracy compared to baseline methods, hence achieving new state-of-the-art results. Our main contributions are as follows:

- We adopt an LLN perspective to approach SFDA using constant noisy labels that do not change during training.
- We introduce DCPL, a novel approach tailored for domain adaptation settings, in which a noise transition matrix is learned to capture the label corruption of the pseudo-labels, thus enabling a better true class-posterior estimation.
- We propose integrating knowledge derived from the source model into the noise transition learning to achieve a more accurate prediction of the underlying true label distribution.

We demonstrate the effectiveness of our approach by implementing it across several SFDA methods: SHOT, SHOT++, and AaD. Our results establish new state-of-the-art performance on three domain adaptation datasets: VisDA, DomainNet, and OfficeHome. Our code can be found at: https://github.com/ ssi-research/DCPL\_SFDA.

# 2 Related Work

Unsupervised Domain Adaptation. UDA attracted much attention in recent years [9, 10, 43, 51, 52]. One common approach for addressing UDA involves employing distributional matching (DM). This is usually done by directly minimizing the domain discrepancy statistics [35, 46] or by using some form of domain-adversarial learning [10, 51]. However, in [24] it was shown that DMbased approaches have limited generalization ability under what they term "realistic shifts". One such shift is when the source label distribution is balanced while the target label distribution is long-tailed. A possible remedy proposed in [24] is to use self-supervised methods which were found to be useful for UDA [1, 11, 47, 57].

Source-Free Domain Adaptation. Recently, in light of the growing emphasis on privacy and data protection, there has been a rising interest in source-free domain adaptation [16,25,28,40,41,54,59]. In SFDA, a model is first trained on the source data and then updated based on the target data, usually using unsupervised techniques and self-training due to the absence of labeled data [28]. Some studies modify the training procedure of the source model to ease the adaptation process by using self-supervision techniques [21], mixup augmentation [22], and/or uncertainty quantification [42]. Another line of research adds unsupervised loss terms to encourage discriminability and diversity in the output during the target adaptation process [6.28,58–60]. For instance, AaD [59] designs the loss such that it encourages the prediction of neighboring samples to be similar and the prediction of non-neighboring samples to be different. This results in simple discriminability and diversity terms. In SHOT [28] the authors use the conditional entropy and marginal entropy of the predictions as well as cross-entropy with target pseudo-labels. Using target pseudo-labels, the authors managed to reduce the error accumulation caused by domain misalignment. SHOT++ [29] builds on SHOT by adding a rotation-prediction head to allow self-supervised representation learning on the target domain. The samples are then divided into unlabeled and "labeled" based on the prediction confidence, and an additional step of semi-supervised learning is applied [4]. The authors of [61] noticed that the early-time training phenomenon is present in the SFDA settings. Hence, they added an auxiliary loss to prevent the model from deviating too much from early-time predictions.

**Pseudo Labels for SFDA.** Several SFDA studies have focused on correcting the pseudo-labels of the target samples under the assumption that they are noisy [2, 5, 18, 20, 31, 61]. In [18] only reliable samples from the target domain were chosen based on prediction statistics. However, this may lead to underrepresentation of the target domain data distribution. In [20] the authors used label refinement based on similarity statistics between classes across consecutive epochs. However, in cases of high fluctuations in the class assignment, the similarity matrix may be non-diagonally dominant resulting in uninformative

or even damaging labels. In [31] the labels were refined based on the predictions of neighboring examples. However, it requires storing a memory bank of examples which may impose a burden in low-resource environments. Recently, a co-learning strategy [63] was suggested to improve the quality of the generated target pseudo-labels using robust pre-trained networks such as Swin-B [34]. In addition, they refined the pseudo-labels and replaced uncertain ones with the current model output.

Learning with Label-Noise. In general, methods for LLN assume either instance-dependent noise or instance-independent noise [44]. In this study, we assume that the true label is conditionally independent of the features given the observed label. Common approaches for handling label noise with deep learning models include adding regularization [32,39,55], modifying the loss functions and values [3,12,33,36], sample selection [14,17,27,50,62], and architectural changes, such as adding a noise adaptation layer [13,45,56]. Our approach belongs to the last category since it is simple to incorporate into existing learning systems, easily scaled, and tends to work well. One obstacle to learning a confusion matrix is to prevent converging to degenerate solutions. Various methods have been proposed to address this issue. In [45,49] it was suggested to regularize the matrix trace. The authors of [13] use a joint learning of the network and the confusion matrix with a careful matrix initialization. In [64] it was suggested to maximize the total variation between the prediction vectors of samples. In [26] the authors proposed controlling the volume of the simplex formed by the columns of the confusion matrix by minimizing its log determinant. The authors of [30] advocated a bi-level optimization process of the network's parameters and the confusion matrix. In this study, we chose to minimize the trace as we found it to work well.

# 3 SFDA Problem Setup

This study was designed to address the challenge of source-free unsupervised domain adaptation (SFDA) in visual classification tasks. Throughout this work, we deal only with a single source domain and a closed set scenario, where both the source and target domains encompass an identical set of classes. We investigate SFDA within the context of a K-way image classification task. Typically, this problem is formulated as a two-stage training process.

In the initial stage, we train a predictive function  $h_s : \mathcal{X}_s \to \mathcal{Y}$  using  $n_s$ labeled samples  $\{x_{s,i}, y_{s,i}\}_{i=1}^{n_s}$  from a source domain  $\mathcal{D}_s$  where  $x_{s,i} \in \mathcal{X}_s$  and  $y_{s,i} \in \mathcal{Y}$  represent the samples and their associated labels respectively. Unless explicitly stated otherwise, we omit the sample indices. Generally,  $h_s(x_s) = \operatorname{argmax}(g_s(x_s))$  where  $g_s$  is the source model, comprising a feature extractor initialized using a pre-trained ImageNet model, and a task-specific classification head. In the second stage, we are provided with  $n_t$  unlabeled samples  $\{x_{t,i}\}_{i=1}^{n_t}$ from a target domain  $\mathcal{D}_t$  where  $x_t \in \mathcal{X}_t$ . We then train a target function  $h_t :$  $\mathcal{X}_t \to \mathcal{Y}$  to predict the labels  $\{y_{t,i}\}_{i=1}^{n_t}$  of the target samples, by only utilizing



**Fig. 1:** Schematic overview of our framework. DCPL learns both the target model and a transition matrix that adapts the model predictions to noisy pseudo-labels obtained by a general pre-trained network.

the source function  $h_s$  and the target domain data without access to the source data. We define  $h_t(x_t) = \operatorname{argmax}(g_t(x_t))$ , analogous to  $h_s$ .

# 4 De-Confusing Pseudo-Labels

Our method is based on the generation of target pseudo-labels, as is commonly done in SFDA approaches. Due to the domain shift between the source and target domains, these pseudo-labels inherently carry noise. Hence, we treat the target adaptation process with these pseudo-labels as a label-noise learning problem, introducing a novel modification to the well-established noise transition matrix LLN method [13, 49]. Conventional LLN methods, like the noise transition matrix approach, generally tackle label noise arising from errors made by human annotators during the labeling process, inconsistencies or disagreements among different annotators (inter-annotator variability), mislabeling of instances in the dataset, and/or the presence of ambiguous or subjective data points that are challenging to label accurately. However, in our scenario, label noise is introduced by domain shifts, where the characteristics of the data in the target domain differ from those in the source domain used for training. The label noise is derived from pseudo-labels generated with a specified source model. Consequently, the distribution of noise depends on the accuracy of the source model on the target dataset, which is associated with the domain gap between the source and target datasets, as well as the generalization capacity of the source model.

In this scenario, the outputs of the source model contain valuable information about the noise distribution that can be utilized to enhance the estimation of the transition matrix. Therefore, we propose integrating this knowledge into the

noise transition matrix learning method to achieve a more accurate prediction of the underlying true label distribution and hence enhance the accuracy of classifier predictions.

#### 4.1 Pseudo-label generation

At the beginning of the target adaptation process, we first apply a pseudolabeling function using a pre-trained network to produce high-quality pseudolabels similar to [63]. There are some advantages to using pre-trained models to produce pseudo-labels. One is that pre-training datasets, such as ImageNet, are large, diverse, and less source-biased and thus may better capture the target distribution. Another advantage is that modern state-of-the-art architectures are designed to learn robust and generalized features, enabling better transfer to target tasks [19]. However, since the pre-training and target label spaces are different, the final classification layer of the pre-trained model is not suitable for the target data. To address this issue, instead of using the pre-trained classifier, the classifier of the source model can be used.

The pseudo-label generation process includes using a feature extractor of a pre-trained robust network (Swin-B) instead of that of the source model. However, unlike Co-learning [63] (and also SHOT [28]) which updates the pseudolabels every epoch, we only invoke the pseudo-labeling function once, at the beginning of the adaptation process. The centroid  $C^{(k)}$  is thus calculated as the sum of the pre-trained features  $f_p$  weighted by estimated class probabilities of the source model  $g_s$ :

$$C^{(k)} = \frac{\sum_{x_t \in \mathcal{X}_t} \sigma^{(k)}(g_s(x_t)) f_p(x_t)}{\sum_{x_t \in \mathcal{X}_t} \sigma^{(k)}(g_s(x_t))},$$
(1)

where  $\sigma$  is a softmax function. Then, for each target instance, a pseudo-label  $\tilde{y}$  is generated based on its nearest centroid using the cosine distance. We denote the final pseudo-labeled target domain  $\{x_{t,i}, \tilde{y}_{t,i}\}_{i=1}^{n_t}$  by  $\tilde{\mathcal{D}}_t$ .

Utilizing robust pre-trained networks like Swin-B enhances the quality of the generated target pseudo-labels through the pseudo-labeling process. Nevertheless, the dependence of these generated pseudo-labels on the source model introduces potential noise, particularly in the presence of domain shifts.

### 4.2 Adaptation as learning with noisy pseudo-labels

After obtaining the pseudo-labels of all the target inputs, we now want to apply a noisy-label learning approach to overcome the label noise. In this specific context, our focus revolves around learning techniques designed for handling noisy data by estimating the label noise matrix. These methods have proven highly successful, yielding cutting-edge results [26, 49]. Our domain adaptation approach relies on the estimated noise matrix to achieve a noise-robust adaptation. We propose using a modification of this LLN approach to capture the label noise corruption and improve the true label probability estimation. Since the noise in our setting originates from the source model, we adjust this approach to the domain adaptation setting by leveraging pseudo-label softmax information and using it to regularize the transition matrix learning. Fig. 1 provides a schematic overview of our overall framework.

Training methods for noise robustness that involve estimating the noise matrix are all grounded in a fundamental observation. Specifically, the clean class-posterior, denoted as p(y|x), can be inferred by utilizing the noisy class-posterior, represented as  $p(\tilde{y}|x)$ , and the class-dependent noise matrix  $T \in [0, 1]^{K \times K}$ , where  $t_{ij} = p(\tilde{y} = i|y = j)$ . This relationship is expressed as follows:

$$p(\tilde{y}|x) = T \cdot p(y|x). \tag{2}$$

Various works have incorporated changes to the architecture of classification networks to improve the representation of the label noise matrix in noisy datasets [13, 26, 45, 49]. These modifications include adding a noise adaptation layer on top of the softmax layer and developing a specialized architecture. The noise adaptation layer is designed to emulate the label transition behavior in learning a network. These adjustments have resulted in improved generalization by adjusting the network output based on the estimated label transition probability. In our implementation of training with noisy labels, we adhere to the approach outlined in [45, 49].

Denote the softmax label prediction of the adapted network for a sample  $x_t$  by  $p_{\theta}(y|x_t) = \sigma(g_t(x_t))$ , where  $\theta$  is the parameter set of the network, and the transition matrix estimator by  $\hat{T}$  which approximates T. The product of these two represents the estimated class probability of the pseudo-labels  $\hat{T} \cdot p_{\theta}(y|x_t)$ . During inference, we eliminate the noise adaptation layer defined by T since we want to predict the true (clean) label. Given the target training inputs  $\{x_{t,i}\}_{i=1}^{n_t}$  and the generated noisy pseudo-labels  $\{\tilde{y}_{t,i}\}_{i=1}^{n_t}$ , the standard cross-entropy loss function is:

$$\mathcal{L}_{ce} = -\mathbb{E}_{(x_t, \tilde{y}_t) \in \tilde{\mathcal{D}}_t} \text{CE} \ (p_\theta(y|x_t), \tilde{y}_t). \tag{3}$$

Following [45, 49], we optimize the network parameters alongside the estimated transition matrix  $\hat{T}$  and minimize the average cross-entropy between the pseudo-labels and the estimated noisy label distribution:

$$\mathcal{L}_{ce} = -\mathbb{E}_{(x_t, \tilde{y}_t) \in \tilde{\mathcal{D}}_t} \operatorname{CE} \left( \hat{T} \cdot p_\theta(y|x_t), \tilde{y}_t \right) + \lambda \operatorname{tr}(\hat{T}), \tag{4}$$

where  $tr(\hat{T})$  denotes the trace of matrix  $\hat{T}$  and  $\lambda$  is a balancing hyper-parameter.

By minimizing this cross-entropy loss, the prediction of the noisy label probabilities,  $\hat{T} \cdot p_{\theta}(y|x_t)$ , is pushed towards the distribution of the generated pseudolabels  $\tilde{y}(x_t)$ . However, convergence to the correct solution of  $\hat{T}$  and  $p_{\theta}(y|x_t)$  is not guaranteed, since there are infinite combinations that can construct a noisy label prediction with a perfect match to the pseudo-label distribution. To ensure that  $\hat{T}$  converges to the true noise transition matrix, namely  $\hat{T} \to T$ , a regularization term is added to the loss, thus minimizing the trace of the estimated noise transition matrix. With this regularization, for  $\hat{T}$  to converge to the true

noise transition matrix  $\hat{T} \to T$ , two assumptions must hold: 1)  $\hat{T}$  and T are diagonally dominant, and 2) the full model, namely, the composition of the noise model on the base classifier, aligns with the label distribution of the pseudo-label generator. Then, it can be shown that the diagonal elements of the estimated transition matrix upper bound the true matrix:

$$t_{ii} = \sum_{j} \hat{t}_{ij} \bar{p}_{ji} \le \sum_{j} \hat{t}_{ii} \bar{p}_{ji} = \hat{t}_{ii} (\sum_{j} \bar{p}_{ji}) = \hat{t}_{ii},$$
(5)

where  $\bar{p}_{ji} \triangleq \mathbb{E}_{p(x|y=i)}[p(\arg\max_{k}p_{\theta}^{(k)}(y|x) = j|y = i)]$ . Moreover, it was shown in [49] (Lemma 1) that obtaining  $\hat{T}$  with a minimal trace uniquely coincides with T. Note that assumption (2) holds by bringing  $\hat{T} \cdot p_{\theta}(y|x_t)$  closer to the true noisy distribution  $\tilde{y}$ , which is encouraged by minimizing the cross-entropy. As for assumption (1), using a large pre-trained network (Swin-B [34]) and a source model that is trained on a source domain with a reasonable domain gap from the target domain will probably generate a diagonally dominant T (see Section 5 for a transition matrix example).

**Domain adaptation regularization.** Minimizing the trace as above provides a structural regularization on the transition matrix. However, having a trained model as the source of the pseudo-labels allows us to impose an additional, more powerful prior, specific to the setting of domain adaptation. More specifically, we leverage the softmax values of the pseudo-labels, positing that they correlate well (on average) with the true confusion matrix of the pseudo-labeler. The softmax values of the pseudo-labels are determined by computing the cosine distance to the target centroids. Computing the mean softmax vector for each cluster provides insights into class similarities, which can serve as indicators of the pseudo-label noise. Stacking these vectors forms a prior confusion matrix. The aim is to leverage this prior to regularize the learning of the transition matrix  $\hat{T}$  which is initialized in a class-agnostic manner.

We define this prior  $T_c$  on the confusion matrix as follows. Let  $\mathcal{D}_t^k = \{x_t : \tilde{y}_t = k\}$  be the set of all samples pseudo-labeled as class k. We first compute the pseudo-label logits  $s(x_t)$  obtained via cosine-similarity with the centroids calculated with Eq. (1):

$$s(x_t)[k] = \frac{f_p(x_t) \cdot C^{(k)}}{\|f_p(x_t)\| \|C^{(k)}\|}.$$
(6)

We then compute the prior matrix  $T_c$  with the mean softmax vector for each class k as follows:

$$T_c[k,j] = \mathbb{E}_{x_t \in \mathcal{D}_t^k} \frac{\exp(s(x_t)[j]/\tau)}{\sum_{i=1}^K \exp(s(x_t)[i]/\tau)},\tag{7}$$

where  $\tau$  is a temperature (set to 0.01 unless stated otherwise) that is applied for sharpening the predictions, given that the logits are computed using cosine

### Algorithm 1 DCPL algorithm

**Input**: Source model  $g_s$ , Target dataset  $\mathcal{D}_t = \{x_{t,i}\}_{i=1}^{n_t}$ , pre-trained feature extractor  $f_p$ , N - number of epochs,  $\mathcal{L}_{SFDA}$  - additional loss terms,  $\lambda, \gamma$  - hyper-parameters weighing the regularization terms for the transition matrix

**Output**: target adaptation model  $g_t$ 

- 1: Initialize weights of adaptation model  $g_t$  with  $g_s$
- 2: Calculate target centroids  $C^{(k)}$  using Eq. 1
- 3: Calculate pseudo-labels  $\tilde{y}_t$  based on Cosine distance from  $C^{(k)}$
- 4: Compute  $T_c$  using Eq. 7
- 5: Initialize  $\hat{T}$  with identity matrix
- 6: for epoch i=1:N do
- 7: Compute loss:  $-\mathbb{E}_{(x_t, \tilde{y}_t) \in \tilde{\mathcal{D}}_t} \operatorname{CE} (\hat{T} \cdot p_{\theta}(y|x_t), \tilde{y}_t) + \lambda \operatorname{tr}(\hat{T}) + \gamma \|\hat{T} T_c\|_2^2 + \mathcal{L}_{SFDA}$
- 8: Update  $\hat{T}$  and  $g_t$  weights

9: end for

10: return adapted model  $g_t$ 

similarity constrained within the range [-1, 1]. Finally, we add a term to the loss to encourage the learned transition matrix to be similar to the estimated prior:

$$\mathcal{L}_{DCPL} = -\mathbb{E}_{(x_t, \tilde{y}_t) \in \tilde{\mathcal{D}}_t} \operatorname{CE}\left(\hat{T} \cdot p_\theta(y|x_t), \tilde{y}_t\right) + \lambda \operatorname{tr}(\hat{T}) + \gamma \|\hat{T} - T_c\|_2^2, \qquad (8)$$

where  $\gamma$  is a balancing hyper-parameter.

Equipped with our new cross-entropy loss, defined in Eq. 8, we can use it instead of the standard cross-entropy loss term in existing SFDA frameworks, and retain additional regularization terms specific to the SFDA methods, such as the information-maximization loss in SHOT [28] or the discriminability and diversity loss terms in AaD [59]. The final loss is then:

$$\mathcal{L} = \mathcal{L}_{DCPL} + \mathcal{L}_{SFDA}.$$
(9)

Algorithm 1 illustrates the DCPL approach.

#### 5 Experiments

In this section, we describe our experimental setup and present the results for three different domain adaptation benchmarks: OfficeHome, VisDA and Domain-Net.

### 5.1 Datasets

**VisDA** [38] is a large-scale benchmark for synthetic-to-real adaptation. It contains over 200K images from 12 classes. **DomainNet** [37] is a large UDA benchmark for image classification containing 345 classes in 6 domains. Since its full version suffers from labeling noise, following [63] we use a subset that contains 126 classes from 4 domains: Real (R), Clipart (C), Painting (P) and Sketch (S). **OfficeHome** [53] is an image classification benchmark containing 65 categories of objects. It consists of four distinct domains: Art (A), Clipart (C), Product (P) and RealWorld images (R).

**Table 1:** Performance comparison for VisDA with a Resnet-101 backbone. DCPL and co-learning results are obtained using SwinB-1K. \* indicates results obtained with our implementation.

$\overline{\rm Method}~({\rm Synthesis}{\rightarrow}{\rm Real})$	Plane	Bcycle	Bus	$\operatorname{Car}$	Horse	Knife	Mcycl	Person	Plant	Sktbrd	Train	Truck	AVG
Source only	60.9	21.6	50.9	67.6	65.8	6.3	82.2	23.2	57.3	30.6	84.6	8.0	46.6
A <sup>2</sup> Net [54]	94.0	87.8	85.6	66.8	93.7	95.1	85.8	81.2	91.6	88.2	86.5	56.0	84.3
G-SFDA [58]	96.1	88.3	85.5	74.1	97.1	95.4	89.5	79.4	95.4	92.9	89.1	42.6	85.4
SFDA-DE [8]	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 [23]	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
RCHC [7]	96.0	90.1	85.0	73.1	95.9	97.3	87.6	84.7	93.8	91.3	86.7	51.2	86.1
RCHC++ [7]	97.6	88.9	88.4	84.0	97.6	97.4	92.2	86.2	97.4	92.8	92.6	41.2	87.8
TPDS [48]	97.6	91.5	89.7	83.4	97.5	96.3	92.2	82.4	96.0	94.1	90.9	40.4	87.6
Co-learn [63]	99.0	90.0	84.2	81.0	98.1	97.9	94.9	80.1	94.8	95.9	94.4	48.1	88.2
AaD [59]*	96.8	89.9	83.9	82.7	96.7	95.7	89.5	79.8	95.5	92.6	88.8	56.5	87.4
AaD w/ Co-learn [63]*	97.4	90.6	84.5	83.7	97.1	97.1	91.5	84.9	96.6	94.2	90.4	57.5	88.8
AaD w/ DCPL (ours)	97.9	92.7	83.8	77.9	97.4	97.3	90.6	83.7	96.6	96.5	92.0	66.6	89.4
SHOT [28]	94.3	88.5	80.1	57.3	93.1	94.9	80.7	80.3	91.5	89.1	86.3	58.2	82.9
SHOT w/ ELR [61]	95.8	84.1	83.3	67.9	93.9	97.6	89.2	80.1	90.6	90.4	87.2	48.2	84.1
SHOT w/ ELR (w/ SwinB)*	96.2	87.5	80.9	72.1	94.9	97.5	90.3	81.6	93.4	92.9	87.9	55.0	85.8
SHOT w/ Co-learn [63]*	95.7	87.9	83.4	65.2	94.3	96.3	87.2	81.8	92.1	92.2	88.6	59.1	85.3
SHOT w/ DCPL (ours)	97.0	89.1	83.5	63.3	95.5	97.7	88.6	81.1	93.9	95.3	90.0	65.7	86.7
SHOT++ [29]	97.7	88.4	90.2	86.3	97.9	98.6	92.9	84.1	97.1	92.2	93.6	28.8	87.3
SHOT++ w/ co-learn $[63]^*$	97.8	89.6	91.2	85.4	97.9	98.1	93.6	85.7	97.2	94.7	93.3	37.2	88.5
SHOT++ w/ DCPL (ours)	98.0	91.3	88.5	77.9	98.0	98.7	93.3	84.6	97.3	94.9	93.8	54.7	89.3

**Table 2:** Performance comparison for DomainNet with a Resnet-50 backbone. DCPL and co-learning results are obtained using SwinB-1K. \* indicates results obtained with our implementation.

Method	$C \rightarrow I$	$P C \rightarrow P$	$R \ C \rightarrow S$	$P \rightarrow C$	$P \rightarrow R$	$P \rightarrow S$	$R \rightarrow C$	$R \rightarrow F$	$P R \rightarrow S$	$S \rightarrow C$	$S \rightarrow F$	$S \rightarrow R$	R AVG
Source only	49.0	62.4	51.1	48.1	72.2	43.9	58.5	64.0	49.0	60.4	53.4	61.1	56.1
AaD [59]*	61.1	77.4	62.2	63.3	81.5	61.3	71.0	70.0	59.8	74.8	67.2	78.8	69.0
AaD w/ Co-learn [63]*	66.1	81.2	64.0	66.1	83.5	62.8	72.4	71.4	62.3	75.8	70.2	82.2	71.5
AaD w/ DCPL (ours)	65.8	82.1	64.6	65.2	83.9	62.7	72.7	72.3	62.0	75.5	70.2	82.2	71.6
SHOT [28]	62.0	78.0	59.9	63.9	78.8	57.4	68.7	67.8	57.7	71.5	65.5	76.3	67.3
SHOT w/ Co-learn [63]*	62.7	80.8	58.7	61.1	82.0	56.2	67.1	70.1	52.9	70.9	67.8	80.9	67.6
SHOT w/ DCPL (ours)	62.9	80.6	61.7	63.3	82.1	60.7	70.1	71.0	59.8	73.4	68.7	80.7	69.6
SHOT++ [29]	63.0	80.4	62.5	66.9	80.0	60.0	71.2	68.7	61.2	72.8	66.6	78.1	69.3
SHOT++ w/ Co-learn [63]*	63.6	81.2	61.3	64.3	82.2	59.7	68.8	70.3	57.0	73.1	68.3	81.3	69.3
SHOT++ w/ DCPL (ours)	62.7	80.7	62.8	64.7	82.5	62.3	72.0	71.6	62.3	74.6	68.6	81.1	70.5

## 5.2 Implementation Details

In all our experiments we use PyTorch and train on NVIDIA V100 and A100 GPUs. We use pre-trained ResNet [15] architectures as backbones. Specifically, Resnet50 is used for OfficeHome and DomainNet, while Resnet-101 is used for VisDA. As in [28, 29], we replace the original fully connected layer with a bot-tleneck layer followed by batch normalization and a task-specific fully connected classifier with weight normalization. We follow the source model training procedure applied in [28, 29]. Specifically, we train the whole network through back-propagation, and the newly added layers are trained with a learning rate 10 times larger than that of the pre-trained layers. We use mini-batch SGD with a momentum of 0.9 and a weight decay of 1e-3.

**Table 3:** Performance comparison for OfficeHome with a Resnet-50 backbone. DCPL and co-learning results are obtained using SwinB-1K. \* indicates results obtained with our implementation.

Method	$A \rightarrow C$	$A \rightarrow P$	$A \rightarrow B$	$R \ C \rightarrow A$	$C \rightarrow F$	$P C \rightarrow R$	$P \rightarrow A$	$P \rightarrow C$	$P \rightarrow R$	$R \rightarrow A$	$R \rightarrow C$	$R \rightarrow P$	P AVG
Source only	44.6	67.3	74.8	52.7	62.7	64.8	53.0	40.6	73.2	65.3	45.4	78.0	60.2
A <sup>2</sup> Net [54]	58.4	79.0	82.4	67.5	79.3	78.9	68.0	56.2	82.9	74.1	60.5	85.0	72.8
SFDA-DE [8]	59.7	79.5	82.4	69.7	78.6	79.2	66.1	57.2	82.6	73.9	60.8	85.5	72.9
CoWA-JMDS [23]	56.9	78.4	81.0	69.1	80.0	79.9	67.7	57.2	82.4	72.8	60.5	84.5	72.5
TPDS [48]	59.3	80.3	82.1	70.6	79.4	80.9	69.8	56.8	82.1	74.5	61.2	85.3	73.5
Co-learn [63]*	58.1	84.4	87.6	75.9	85.5	86.9	74.6	55.1	87.9	78.0	58.2	88.4	76.7
AaD [59]	59.3	79.3	82.1	68.9	79.8	79.5	67.2	57.4	83.1	72.1	58.5	85.4	72.7
AaD w/ co-learn [63]*	61.6	83.6	85.5	74.6	84.5	84.2	72.3	60.7	85.7	76.8	63.0	88.0	76.7
AaD w/ DCPL (ours)	61.9	84.5	87.1	75.7	85.8	85.6	73.6	59.8	86.5	78.1	62.9	89.0	77.5
SHOT [28]	57.1	78.1	81.5	68.0	78.2	78.1	67.4	54.9	82.2	73.3	58.8	84.3	71.8
SHOT w/ ELR [61]	58.7	78.9	82.1	68.5	79.0	77.5	68.2	57.1	81.9	74.2	59.5	84.9	72.6
SHOT w/ co-learn [63]*	60.8	80.8	85.5	72.4	81.9	83.3	72.4	59.4	84.7	77.1	62.3	87.3	75.7
SHOT w/ DCPL (ours)	59.9	84.3	87.8	76.8	85.8	86.6	74.8	58.6	87.4	77.9	61.1	89.0	77.5
SHOT++ [29]	57.9	79.7	82.5	68.5	79.6	79.3	68.5	57.0	83.0	73.7	60.7	84.9	73.0
SHOT++ w/ co-learn [63]*	60.1	82.2	86.1	73.2	83.3	83.7	72.2	59.6	85.1	77.0	63.2	87.9	76.1
SHOT++ w/ DCPL (ours)	61.2	84.3	88.0	76.7	86.1	86.6	74.2	59.3	87.6	78.2	62.4	89.5	77.8

For the target adaptation process, we generate the pseudo-labels at the beginning of training using a pre-trained network Swin-Base [34], since Swin transformers [34] were shown to be robust under domain shifts [19]. In our experiments, we used two pre-trained SwinB networks, one of which was pre-trained on Imagenet-1K, and the other which was pre-trained on Imagenet-22K and then finetuned on Imagenet-1K (ImageNet-22K-1K). The learning rate is set to 1e-2 for OfficeHome, and 1e-3 for VisDA and DomainNet. The temperature  $\tau$ for calculating  $T_c$  is set to 0.01 when using the Imagenet-1K pretrained network and 0.05 for ImageNet-22K-1K.  $\lambda$  was set to 0.01 [49] for all datasets to avoid overfitting and maintain simplicity for the practitioners. For  $\gamma$  we chose 0.01 for OfficeHome and DomainNet which have a large number of classes and 1 for VisDA with a smaller amount of classes to account for the ratio in class amount (squared due to square in the loss term). We train the adapted networks for 50 epochs in all cases. All of our results report the mean accuracy over three runs with different random seeds (2019, 2020, and 2021) as was done in [28].

#### 5.3 Main Results

In this section, we report the results of DCPL on all of the datasets. Tables 1-3 present performance comparisons of our method to state-of-the-art methods. Since DCPL only adjusts the cross-entropy loss, it can be combined with other common SFDA methods. Here we investigate its combination with SHOT [28], SHOT++ [29] and AaD [59]. For SHOT and SHOT++ which contain cross-entropy components, we replace this component with our DCPL cross-entropy. For AaD, which originally did not include a pseudo-labeling process or a cross-entropy term loss, we follow co-learning [63] and add a cross-entropy term in our comparisons. Finally, we also report the accuracy of "source only" which is the source model (trained with the source dataset alone) on the target dataset.

Method	Prior loss	OfficeHome	DomainNet	VisDA
SHOT w/ co-learn	NR	75.7	67.6	85.3
SHOT w/ DCPL	-	77.4	69.3	85.9
SHOT w/ DCPL	+	77.5	69.6	86.7

Table 4: Prior loss (Eq. 8) contribution to DCPL performance

Table 5: Influence of PL network (SwinB-1K versus SwinB-22K-1K) on performance

Method	PL net	OfficeHome	DomainNet	VisDA
SHOT w/ DCPL SHOT w/ DCPL	SwinB-1K SwinB-22K-1K	77.5 83.8	69.6 73.7	86.7 88.0
$\overline{ \begin{array}{c} {\rm SHOT}{\rm ++ \ w/ \ DCPL} \\ {\rm SHOT}{\rm ++ \ w/ \ DCPL} \end{array} }$	SwinB-1K SwinB-22K-1K	77.8 83.8	$70.5 \\ 74.2$	89.2 89.9

The tables show the following: (1) DCPL outperforms all previous methods on all three datasets, and (2) DCPL significantly outperforms Co-learning irrespective of the SFDA method it was applied to.

#### 5.4 Ablation Study

In this section, we analyze the contribution of the noise transition matrix and the influence of our new prior loss regularization on performance.

**Prior loss regularization ablation.** Table 4 shows the contribution of our prior loss to the transition matrix learning when applied with SHOT [28] implementation for all three datasets. In all cases, the performance increased when the prior was added with the largest contribution in VisDA.

**Pseudo-labeling network ablation.** The strength of the pseudo-labeler has a direct effect on the quality of the adapted network, as a stronger network will reduce label noise. To test this, we report results using the SwinB version trained on either ImageNet1K or trained on ImageNet22K and finetuned on ImageNet1K. Table 5 depicts the performance of both SHOT [28] and SHOT++ [29] with DCPL.

**Oracle comparison.** Table 6 presents a performance comparison of DCPL to an oracle implementation on the DomainNet dataset. The oracle version was implemented using the real noise transition matrix based on the target labels (which we do not have access to in UDA). The comparison to the oracle allows us to analyze the performance increase potential of DCPL when the estimated transition matrix is made accurate. In the table, we also compare DCPL to a naive baseline implementation in which the transition matrix is set to a fixed identity matrix ( $\hat{T} = I$ ); i.e., if we assume noiseless pseudo-labels and do not learn the TM. Comparing this naive baseline to DCPL shows an increase of 4.3%

**Table 6:** Performance comparison to Oracle and Baseline (without using a transition matrix) for DomainNet.

Method	PL net	$C \rightarrow I$	$P \ C \rightarrow R$	$C \rightarrow S$	$P \rightarrow C$	$P \rightarrow R$	$P \rightarrow S$	$R \rightarrow C$	$R \rightarrow P$	$R \rightarrow S$	$S \rightarrow C$	$T S \rightarrow P$	$P \to I$	₹ AVG
SHOT [28]		62.0	78.0	59.9	63.9	78.8	57.4	68.7	67.8	57.7	71.5	65.5	76.3	67.3
w/ DCPL (w/ $\hat{T} = I$ )	SwinB-1K	63.3	79.9	53.9	57.7	81.5	52.1	61.9	68.9	49.0	66.7	67.5	80.9	65.3
w/ DCPL (ours)	SwinB-1K	62.9	80.6	61.7	63.3	82.1	60.7	70.1	71.0	59.8	73.4	68.7	80.7	69.6
w/ DCPL Oracle	SwinB-1K	64.7	83.5	64.2	67.6	83.8	64.4	73.7	72.9	65.4	76.3	70.8	83.2	72.5
w/ DCPL (w/ $\hat{T} = I$ )	swinB-22K-1K	65.8	83.5	62.1	67.4	84.8	62.1	71.1	72.9	61.2	72.2	70.1	84.8	71.5
w/ DCPL (ours)	swinB-22K-1K	69.2	84.6	65.3	68.9	86.2	64.9	74.1	75.0	63.3	75.5	72.9	84.7	73.7
w/ DCPL Oracle	swinB-22K-1K	71.7	86.8	68.6	74.6	87.7	68.8	76.1	76.3	68.4	78.7	75.5	87.8	76.8



Fig. 2: Real and learned noise transition matrices for VisDA.

using SwinB-1K and 2.2% using SwinB-22K-1K. This demonstrates the contribution of TM over SwinB. Comparing this naive baseline to the oracle version of DCPL shows an increase of 7.2% to 72.5% and 5.3% to 76.8%, using SwinB-1K and SwinB-22K-1K respectively. As expected, DCPL achieves an intermediate accuracy of 69.6% and 73.7%. These results show that better estimations of the transition matrix can generate significant improvement gains.

Noise transition matrix. Fig. 2 shows an example of real as compared to learned confusion matrices for the VisDA dataset. It is clear that the real and learned matrices are diagonally dominant, which is compatible with our assumption. Additionally, the comparison of the matrices shows that the main confusions are learned successfully, e.g., Truck samples classified as Cars, Bus and Car samples classified as Truck, and Person samples classified as Skateboard. On the other hand, there are also estimation errors in the transition matrix, e.g., the Knife category seems to be estimated as a corrupted class, although it has high accuracy estimation (above 90%) with almost no misclassifications.

**t-SNE visualisations.** Fig. 3 shows t-SNE plots of target sample embeddings using our method for VisDA and DomainNet (Clipart to Sketch and Real to Sketch). These plots show the influence of using the noise transition matrix on



**Fig. 3:** t-SNE visualization for (a) Clipart to Sketch adaptation from the DomainNet dataset, (b) Real to Sketch adaptation from the DomainNet dataset, (c) synthetic to real adaptation for the VisDA dataset: (1) at the beginning of the target adaptation process, which demonstrates the source model's ability to separate the different classes, (2) at the end of target adaptation using DCPL with  $\hat{T} = I$ , (3) at the end of target adaptation using DCPL.

the target cluster arrangement. In each figure, three t-SNE plots of the target embeddings are presented: (1) at the beginning of the target adaptation process, which demonstrates the ability of the source model to separate the different classes, (2) at the end of target adaptation using DCPL with  $\hat{T} = I$  (using SHOT), (3) at the end of target adaptation with DCPL (using SHOT). These plots clearly show the contribution of the transition matrix to the discriminability of the different classes.

# 6 Conclusion

In this work, we introduce DCPL, a new SFDA approach that de-confuses the pseudo-labels by learning a noise transition matrix. Inspired by LLN methods and modified for the setting of domain adaptation using information inherent in the source of the pseudo-labels, the transition matrix captures the label corruption, which in turn enables a better true class-posterior estimation. We demonstrate the effectiveness of our approach when applied to several SFDA methods and provide state-of-the-art results on three domain adaptation datasets, showing that it is complementary to each of the methods we experimented with, as well as applicable to a wide range of scenarios.

# References

- Achituve, I., Maron, H., Chechik, G.: Self-supervised learning for domain adaptation on point clouds. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision. pp. 123–133 (2021)
- 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)
- Arazo, E., Ortego, D., Albert, P., O'Connor, N., McGuinness, K.: Unsupervised label noise modeling and loss correction. In: International Conference on Machine Learning (ICML) (2019)
- Berthelot, D., Carlini, N., Goodfellow, I., Oliver, A., Papernot, N., Raffel, C.: Mixmatch: A holistic approach to semi-supervised learning. In: Advances in Neural Information Processing Systems (NeurIPs) (2019)
- Chen, W., Lin, L., Yang, S., Xie, D., Pu, S., Zhuang, Y.: Self-supervised noisy label learning for source-free unsupervised domain adaptation. In: 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 10185– 10192. IEEE (2022)
- Cui, S., Wang, S., Zhuo, J., Li, L., Huang, Q., Tian, Q.: Fast batch nuclearnorm maximization and minimization for robust domain adaptation. arXiv preprint arXiv:2107.06154 (2021)
- 7. Diamant, I., Jennings, R.H., Dror, O., Habi, H.V., Netzer, A.: Reconciling a centroid-hypothesis conflict in source-free domain adaptation (2022)
- Ding, N., Xu, Y., Tang, Y., Wang, Y., Tao, D.: Source-free domain adaptation via distribution estimation. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
- Ganin, Y., Lempitsky, V.: Unsupervised domain adaptation by backpropagation. In: International Conference on Machine Learning (ICML) (2015)
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., Lempitsky, V.: Domain-adversarial training of neural networks. Journal of Machine Learning Research 17(1), 2096–2030 (2016)
- Ge, Y., Chen, D., Li, H.: Mutual mean-teaching: Pseudo label refinery for unsupervised domain adaptation on person re-identification. In: International Conference on Learning Representations (ICLR) (2019)
- Ghosh, A., Kumar, H., Sastry, P.S.: Robust loss functions under label noise for deep neural networks. In: Proceedings of the AAAI conference on artificial intelligence. vol. 31 (2017)
- Goldberger, J., Ben-Reuven, E.: Training deep neural-networks using a noise adaptation layer. In: International Conference on Learning Representations (ICLR) (2017)
- Han, B., Yao, Q., Yu, X., Niu, G., Xu, M., Hu, W., Tsang, I., Sugiyama, M.: Co-teaching: Robust training of deep neural networks with extremely noisy labels. Advances in Neural Information Processing Systems (NeurIPs) (2018)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
- Hwang, U., Lee, J., Shin, J., Yoon, S.: SF(DA)\$^2\$: Source-free domain adaptation through the lens of data augmentation. In: International Conference on Learning Representations (ICLR) (2024)

- 16 I. Diamant et al.
- Jiang, L., Zhou, Z., Leung, T., Li, L.J., Fei-Fei, L.: Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In: International Conference on Machine Learning (ICML) (2018)
- Karim, N., Mithun, N., Rajvanshi, A., Chiu, H., Samarasekera, S., Rahnavard, N.: C-sfda: A curriculum learning aided self-training framework for efficient source free domain adaptation. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
- Kim, D., Wang, K., Sclaroff, S., Saenko, K.: A broad study of pre-training for domain generalization and adaptation. In: Proc. of the European Conference on Computer Vision (ECCV) (2022)
- Kumar, V., Lal, R., Patil, H., Chakraborty, A.: Conmix for source-free single and multi-target domain adaptation. In: 2023 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) (2023)
- Kundu, J.N., Bhambri, S., Kulkarni, A., Sarkar, H., Jampani, V., Babu, R.V.: Concurrent subsidiary supervision for unsupervised source-free domain adaptation. In: Proc. of the European Conference on Computer Vision (ECCV) (2022)
- Kundu, J.N., Kulkarni, A.R., Bhambri, S., Mehta, D., Kulkarni, S.A., Jampani, V., Radhakrishnan, V.B.: Balancing discriminability and transferability for sourcefree domain adaptation. In: International Conference on Machine Learning (ICML) (2022)
- Lee, J., Jung, D., Yim, J., Yoon, S.: Confidence score for source-free unsupervised domain adaptation (2022)
- 24. Li, B., Wang, Y., Che, T., Zhang, S., Zhao, S., Xu, P., Zhou, W., Bengio, Y., Keutzer, K.: Rethinking distributional matching based domain adaptation. arXiv preprint arXiv:2006.13352 (2020)
- Li, R., Jiao, Q., Cao, W., Wong, H.S., Wu, S.: Model adaptation: Unsupervised domain adaptation without source data. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
- Li, X., Liu, T., Han, B., Niu, G., Sugiyama, M.: Provably end-to-end label-noise learning without anchor points. In: International Conference on Machine Learning (ICML) (2021)
- Li, Y., Han, H., Shan, S., Chen, X.: Disc: Learning from noisy labels via dynamic instance-specific selection and correction. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 24070–24079 (2023)
- 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 (ICML) (2020)
- 29. 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 (TPAMI) (2021), in Press
- 30. Lin, Y., Pi, R., Zhang, W., Xia, X., Gao, J., Zhou, X., Liu, T., Han, B.: A holistic view of label noise transition matrix in deep learning and beyond. In: International Conference on Learning Representations (ICLR) (2023)
- Litrico, M., Bue, A.D., Morerio, P.: Guiding pseudo-labels with uncertainty estimation for source-free unsupervised domain adaptation. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
- Liu, S., Niles-Weed, J., Razavian, N., Fernandez-Granda, C.: Early-learning regularization prevents memorization of noisy labels. Advances in Neural Information Processing Systems (NeurIPs) (2020)
- Liu, T., Tao, D.: Classification with noisy labels by importance reweighting. IEEE Transactions on Pattern Analysis and Machine Intelligence 38(3), 447–461 (2015)

- 34. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proc. of the IEEE International Conference on Computer Vision (ICCV) (2021)
- Morerio, P., Cavazza, J., Murino, V.: Minimal-entropy correlation alignment for unsupervised deep domain adaptation. In: International Conference on Learning Representations (ICLR) (2018)
- 36. Patrini, G., Rozza, A., Krishna Menon, A., Nock, R., Qu, L.: Making deep neural networks robust to label noise: A loss correction approach. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
- Peng, X., Bai, Q., Xia, X., Huang, Z., Saenko, K., Wang, B.: Moment matching for multi-source domain adaptation. In: Proc. of the IEEE International Conference on Computer Vision (ICCV) (2019)
- Peng, X., Usman, B., Kaushik, N., Hoffman, J., Wang, D., Saenko, K.: Visda: The visual domain adaptation challenge. arXiv preprint arXiv:1710.06924 (2017)
- 39. Pereyra, G., Tucker, G., Chorowski, J., Kaiser, L., Hinton, G.E.: Regularizing neural networks by penalizing confident output distributions. In: International Conference on Learning Representations, Workshop Track Proceedings (2017)
- 40. Qiu, Z., Zhang, Y., Lin, H., Niu, S., Liu, Y., Du, Q., Tan, M.: Source-free domain adaptation via avatar prototype generation and adaptation. In: International Joint Conference on Artificial Intelligence (2021)
- 41. Qu, S., Chen, G., Zhang, J., Li, Z., He, W., Tao, D.: Bmd: A general class-balanced multicentric dynamic prototype strategy for source-free domain adaptation. In: Proc. of the European Conference on Computer Vision (ECCV) (2022)
- 42. Roy, S., Trapp, M., Pilzer, A., Kannala, J., Sebe, N., Ricci, E., Solin, A.: Uncertainty-guided source-free domain adaptation. In: Proc. of the European Conference on Computer Vision (ECCV) (2022)
- Saito, K., Watanabe, K., Ushiku, Y., Harada, T.: Maximum classifier discrepancy for unsupervised domain adaptation. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
- Song, H., Kim, M., Park, D., Shin, Y., Lee, J.G.: Learning from noisy labels with deep neural networks: A survey. IEEE Transactions on Neural Networks and Learning Systems 34(11), 8135–8153 (2023)
- Sukhbaatar, S., Bruna, J., Paluri, M., Bourdev, L., Fergus, R.: Training convolutional networks with noisy labels. In: International conference on learning representations workshop (2015)
- 46. Sun, B., Saenko, K.: Deep coral: Correlation alignment for deep domain adaptation. In: Proc. of the European Conference on Computer Vision (ECCV) (2016)
- 47. Sun, Y., Tzeng, E., Darrell, T., Efros, A.A.: Unsupervised domain adaptation through self-supervision. arXiv preprint arXiv:1909.11825 (2019)
- Tang, S., Chang, A., Zhang, F., Zhu, X., Ye, M., Zhang, C.: Source-free domain adaptation via target prediction distribution searching. International Journal of Computer Vision (IJCV) 132(3), 654–672 (2023)
- Tanno, R., Saeedi, A., Sankaranarayanan, S., Alexander, D.C., Silberman, N.: Learning from noisy labels by regularized estimation of annotator confusion. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
- Tu, Y., Zhang, B., Li, Y., Liu, L., Li, J., Zhang, J., Wang, Y., Wang, C., Zhao, C.R.: Learning from noisy labels with decoupled meta label purifier. Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) pp. 19934–19943 (2023)

- 18 I. Diamant et al.
- Tzeng, E., Hoffman, J., Saenko, K., Darrell, T.: Adversarial discriminative domain adaptation. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
- 52. Tzeng, E., Hoffman, J., Zhang, N., Saenko, K., Darrell, T.: Deep domain confusion: Maximizing for domain invariance. arXiv preprint arXiv:1412.3474 (2014)
- 53. Venkateswara, H., Eusebio, J., Chakraborty, S., Panchanathan, S.: Deep hashing network for unsupervised domain adaptation. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
- Xia, H., Zhao, H., Ding, Z.: Adaptive adversarial network for source-free domain adaptation. In: Proc. of the IEEE International Conference on Computer Vision (ICCV) (2021)
- Xia, X., Liu, T., Han, B., Gong, C., Wang, N., Ge, Z., Chang, Y.: Robust earlylearning: Hindering the memorization of noisy labels. In: International Conference on Learning Representations (ICLR) (2020)
- Xiao, T., Xia, T., Yang, Y., Huang, C., Wang, X.: Learning from massive noisy labeled data for image classification. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2015)
- Xu, J., Xiao, L., López, A.M.: Self-supervised domain adaptation for computer vision tasks. IEEE Access 7, 156694–156706 (2019)
- Yang, S., Wang, Y., van de Weijer, J., Herranz, L., Jui, S.: Generalized source-free domain adaptation. In: Proc. of the IEEE International Conference on Computer Vision (ICCV) (2021)
- 59. Yang, S., Wang, Y., Wang, K., Jui, S., van de Weijer, J.: Attracting and dispersing: A simple approach for source-free domain adaptation (2022)
- 60. 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 (NeurIPs) (2021)
- Yi, L., Xu, G., Xu, P., Li, J., Pu, R., Ling, C., McLeod, A.I., Wang, B.: When source-free domain adaptation meets learning with noisy labels. In: International Conference on Learning Representations (ICLR) (2023)
- Yu, X., Han, B., Yao, J., Niu, G., Tsang, I., Sugiyama, M.: How does disagreement help generalization against label corruption? In: International Conference on Machine Learning (ICML). pp. 7164–7173 (2019)
- Zhang, W., Shen, L., Foo, C.S.: Rethinking the role of pre-trained networks in source-free domain adaptation. In: Proc. of the IEEE International Conference on Computer Vision (ICCV) (2023)
- Zhang, Y., Niu, G., Sugiyama, M.: Learning noise transition matrix from only noisy labels via total variation regularization. In: International Conference on Machine Learning (ICML) (2021)