
Do We Really Need to Access the Source Data? Source Hypothesis Transfer for Unsupervised Domain Adaptation

Jian Liang¹ Dapeng Hu¹ Jiashi Feng¹

Abstract

Unsupervised domain adaptation (UDA) aims to leverage the knowledge learned from a labeled source dataset to solve similar tasks in a new unlabeled domain. Prior UDA methods typically require to access the source data when learning to adapt the model, making them risky and inefficient for decentralized private data. This work tackles a practical setting where only a trained source model is available and investigates how we can effectively utilize such a model without source data to solve UDA problems. We propose a simple yet generic representation learning framework, named *Source Hypothesis Transfer* (SHOT). SHOT freezes the classifier module (hypothesis) of the source model and learns the target-specific feature extraction module by exploiting both information maximization and self-supervised pseudo-labeling to implicitly align representations from the target domains to the source hypothesis. To verify its versatility, we evaluate SHOT in a variety of adaptation cases including closed-set, partial-set, and open-set domain adaptation. Experiments indicate that SHOT yields state-of-the-art results among multiple domain adaptation benchmarks.

1. Introduction

Deep neural networks have achieved remarkable success in a variety of applications across different fields but at the expense of laborious large-scale training data annotation. To avoid expensive data labeling, domain adaptation (DA) methods are developed to fully utilize previously labeled datasets and unlabeled data on hand in a transductive manner, which have obtained promising results in object recognition (Long et al., 2015; Tzeng et al., 2017), semantic

segmentation (Zhang et al., 2017; Hoffman et al., 2018), etc. Over the last decade, increasing efforts have been devoted to deep domain adaptation, especially under the vanilla unsupervised closed-set setting (Panareda Busto & Gall, 2017) where two domains share the same label space but the target data are not labeled. One prevailing paradigm is to mitigate the distribution divergence between domains by matching the distribution statistical moments at different orders (Sun et al., 2016; Zellinger et al., 2017; Peng et al., 2019a). For example, the most-favored Maximum Mean Discrepancy (MMD) (Gretton et al., 2007) measure minimizes a certain distance between weighted sums of all raw moments. Another popular paradigm leverages the idea of adversarial learning (Goodfellow et al., 2014) and introduces an additional domain classifier to minimize the Proxy A -distance (Ben-David et al., 2010) across domains.

Nowadays, data are distributed on different devices and usually contain private information, e.g., those on personal phones or from surveillance cameras. Existing DA methods need to access the source data during learning to adapt, which is not efficient for data transmission and may violate the data privacy policy. In this paper, we address an *interesting but challenging* unsupervised DA setting with only a trained source model provided as supervision. It differs from vanilla unsupervised DA in that the source model instead of the source data is provided to the unlabeled target domain. Such a setting helps protect the privacy in the source domain (e.g., individual hospital profiles), and transferring a light trained model is much efficient than heavy data transmission. For example, as shown in Table 1, the storage size of a trained model is much smaller than that of a compressed dataset.

Table 1. Storage size. The backbone networks of source models are LeNet (LeCun et al., 1998) and ResNet-101 (He et al., 2016), respectively. [†](Hoffman et al., 2018), [‡](Peng et al., 2017).

Storage size (MB)	Digits [†]	VisDA-C [‡]
Source Dataset	33.2	7884.8
Source Model	0.9	172.6

In terms of privacy protection, this setting seems somewhat similar to a recently proposed transfer learning setting in Hypothesis Transfer Learning (HTL) (Kuzborskij & Orabona,

¹Department of ECE, National University of Singapore (NUS). Correspondence to: Jian Liang <liangjian92@gmail.com>.

2013), where the learner does not have direct access to the source domain data and can only operate on the hypotheses induced from the source data. However, our method differs significantly from that work. Conventional HTL always requires labeled data in the target domain (Tommasi et al., 2013) or multiple hypotheses from different source domains (Mansour et al., 2009), which however cannot be applied to unsupervised DA.

To address such a challenging unsupervised DA setting, we propose a simple yet generic solution called Source HypOthesis Transfer (SHOT). Inspired by prior work (Tzeng et al., 2017), SHOT assumes that the same deep neural network model consists of a feature encoding module and a classifier module (hypothesis) across domains. It aims to learn a target-specific feature encoding module to generate target data representations that are well aligned with source data representations, without accessing source data and labels for target data. SHOT is designed based on the following observations. If we have learned source-like representations for target data, then the classification outputs from the source classifier (hypothesis) for target data should be similar to those of source data, *i.e.*, close to one-hot encodings. Thus, SHOT freezes the source hypothesis and fine-tunes the source encoding module by maximizing the mutual information between intermediate feature representations and outputs of the classifier, as information maximization (Hu et al., 2017) encourages the network to assign disparate one-hot encodings to the target feature representations.

Even though information maximization forces feature representations to fit the hypothesis well, it may still align target feature representations to the wrong source hypothesis. To avoid this, we propose a novel self-supervised pseudo-labeling method to augment the target representation learning. Considering pseudo labels generated by the source classifier may be inaccurate and noisy for target data, we propose to attain intermediate class-wise prototypes for the target domain itself and further obtain cleaner pseudo labels via supervision from these prototypes to guide the mapping module learning. Such self-supervised labeling fully exploits the global structure of the target domain and would learn feature representations that correctly fit the source hypothesis. Besides, we investigate label smoothing (Müller et al., 2019), weight normalization (Salimans & Kingma, 2016), and batch normalization (Ioffe & Szegedy, 2015), within the network architecture of the source model, to boost adaptation performance. Empirical evidence shows that both SHOT and its baseline method (source only) benefit from these techniques.

We apply SHOT to the vanilla closed-set (Saenko et al., 2010) DA scenario and a variety of other unsupervised DA tasks like the partial-set (Cao et al., 2018) and open-set (Panareda Busto & Gall, 2017) cases. Experiments show

that SHOT achieves state-of-the-art results for multiple and various domain adaptation tasks. For instance, on the Office-Home dataset, SHOT advances the best average accuracy from 67.6% (Wang et al., 2019) to **71.8%** for the closed-set setting and from 71.8% (Xu et al., 2019) to **79.3%** for the partial transfer case, and from 69.5% (Liu et al., 2019) to **72.8%** for the open-set scenario.

2. Related Work

Unsupervised Domain Adaptation. Lots of unsupervised DA methods have been developed and successfully used in cross-domain applications like object recognition (Liang et al., 2018; 2019b), object detection (Chen et al., 2018), semantic segmentation (Tsai et al., 2018), and sentiment classification (Glorot et al., 2011). Besides the two prevailing paradigms, moment matching and adversarial alignment, as aforementioned, there are several sample-based DA methods that estimate the importance weights of source samples (Jiang & Zhai, 2007; Huang et al., 2007) or sample-to-sample similarity via optimal transport (Bhushan Damodaran et al., 2018). Also, a few works (Bousmalis et al., 2016; Ghifary et al., 2016) impose data reconstruction as an auxiliary task to ensure feature invariance, and some studies consider batch normalization (Cariucci et al., 2017; Wang et al., 2019) and adversarial dropout (Saito et al., 2018a; Lee et al., 2019b) within the network architecture. However, all these methods assume the target user’s access to the source domain data, which is unsafe and sometimes unpractical since source data may be private and decentralized (stored on another device). In addition, our work is, as far as we know, the first DA framework to address almost all unsupervised DA scenarios including multi-source (Peng et al., 2019a) and multi-target (Peng et al., 2019b), open-set DA (Panareda Busto & Gall, 2017), and partial-set DA (Cao et al., 2018).

Hypothesis Transfer Learning (HTL). HTL (Kuzborskij & Orabona, 2013) can be seen as a generalization of parameter adaptation (Csurka, 2017), which assumes the optimal target hypothesis to be closely related to the source hypothesis. Like the famous fine-tuning strategy (Yosinski et al., 2014), HTL mostly acquires at least a small set of labeled target examples per class, limiting its applicability to the semi-supervised DA scenario (Yang et al., 2007; Nelakurthi et al., 2018). Besides, a seminal work (Mansour et al., 2009) infers the weights of different source hypotheses for each unlabeled target datum and utilizes a convex combination of source hypotheses. However, it only fits the multi-source scenario and needs an additive distribution estimation of each source domain. To our best knowledge, (Chidlovskii et al., 2016; Liang et al., 2019a) introduce similar settings as ours to the unsupervised DA problem, but they may not work well enough since they do not incorporate the end-to-

end feature learning module inside.

Pseudo Labeling (PL). Pseudo labeling (Lee, 2013) is originally proposed for semi-supervised learning and gains popularity in other transductive learning problems like DA. The main idea is to label unlabeled data with the maximum predicted probability and perform fine-tuning together with labeled data, which is quite efficient. For DA methods, (Zhang et al., 2018b; Choi et al., 2019) directly incorporate pseudo labeling as a regularization while (Long et al., 2017; 2018) leverage pseudo labels in the adaptation module to pursue discriminative (conditional) distribution alignment. (Zou et al., 2018) further designs an integrated framework to alternately solve target pseudo labels and perform model training. In the absence of labeled data, DeepCluster (Caron et al., 2018), one of the best self-supervised learning methods, generates pseudo labels via k-means clustering and utilizes them to re-train the current model. Considering the dataset shift, our framework combines advantages of both and develops a self-supervised pseudo labeling strategy to alleviate the harm from noisy pseudo labels.

Federated Learning (FL). FL (Bonawitz et al., 2019) is a distributed machine learning approach that trains a model across multiple decentralized edge devices without exchanging their data samples. A common usage of FL consumes a simple aggregate of updates from multiple local users for the learning algorithm at the server-side (Bonawitz et al., 2017; McMahan et al., 2018), offering a privacy-preserving mechanism. Recently, (Peng et al., 2020) introduces the first federated DA setting where knowledge is transferred from the decentralized nodes to a new node without any supervision itself and proposes an adversarial-based solution. Specifically, it trains one model per source node and updates the target model with the aggregation of source gradients to reduce domain shift. However, it may fail to address the vanilla DA setting with only one source domain available.

3. Method

We address the unsupervised DA task with only a pre-trained source model and without access to source data. In particular, we consider K -way classification. For a vanilla unsupervised DA task, we are given n_s labeled samples $\{x_s^i, y_s^i\}_{i=1}^{n_s}$ from the source domain \mathcal{D}_s where $x_s^i \in \mathcal{X}_s, y_s^i \in \mathcal{Y}_s$, and also n_t unlabeled samples $\{x_t^i\}_{i=1}^{n_t}$ from the target domain \mathcal{D}_t where $x_t^i \in \mathcal{X}_t$. The goal of DA is to predict the labels $\{y_t^i\}_{i=1}^{n_t}$ in the target domain, where $y_t^i \in \mathcal{Y}_t$, and the source task $\mathcal{X}_s \rightarrow \mathcal{Y}_s$ is assumed to be the same with the target task $\mathcal{X}_t \rightarrow \mathcal{Y}_t$. Here SHOT aims to learn a target function $f_t : \mathcal{X}_t \rightarrow \mathcal{Y}_t$ and infer $\{y_t^i\}_{i=1}^{n_t}$, with only $\{x_t^i\}_{i=1}^{n_t}$ and the source function $f_s : \mathcal{X}_s \rightarrow \mathcal{Y}_s$ available.

We address the above source model transfer task for UDA through three steps. First, we generate the source model

from source data. Secondly, we abandon source data and transfer the model (including source hypothesis) to the target domain without accessing source data. Thirdly, we further study how to design better network architectures for both models to improve adaptation performance.

3.1. Source Model Generation

We consider to develop a deep neural network and learn the source model $f_s : \mathcal{X}_s \rightarrow \mathcal{Y}_s$ by minimizing the following standard cross-entropy loss,

$$\mathcal{L}_{src}(f_s; \mathcal{X}_s, \mathcal{Y}_s) = -\mathbb{E}_{(x_s, y_s) \in \mathcal{X}_s \times \mathcal{Y}_s} \sum_{k=1}^K q_k \log \delta_k(f_s(x_s)), \quad (1)$$

where $\delta_k(a) = \frac{\exp(a_k)}{\sum_i \exp(a_i)}$ denotes the k -th element in the soft-max output of a K -dimensional vector a , and q is the one-of- K encoding of y_s where q_k is ‘1’ for the correct class and ‘0’ for the rest. To further increase the discriminability of the source model and facilitate the following target data alignment, we propose to adopt the label smoothing (LS) technique as it encourages examples to lie in tight evenly separated clusters (Müller et al., 2019). With LS, the objective function is changed to

$$\mathcal{L}_{src}^{ls}(f_s; \mathcal{X}_s, \mathcal{Y}_s) = -\mathbb{E}_{(x_s, y_s) \in \mathcal{X}_s \times \mathcal{Y}_s} \sum_{k=1}^K q_k^{ls} \log \delta_k(f_s(x_s)), \quad (2)$$

where $q_k^{ls} = (1 - \alpha)q_k + \alpha/K$ is the smoothed label and α is the smoothing parameter which is empirically set to 0.1.

3.2. Source Hypothesis Transfer with Information Maximization (SHOT-IM)

The source model parameterized by a DNN in Fig. 2 consists of two modules: the feature encoding module $g_s : \mathcal{X}_s \rightarrow \mathbb{R}^d$ and the classifier module $h_s : \mathbb{R}^d \rightarrow \mathbb{R}^K$, i.e., $f_s(x) = h_s(g_s(x))$. Here d is the dimension of the input feature. Previous DA methods align different domains by matching the data distributions in the feature space \mathbb{R}^d using maximum mean discrepancy (MMD) (Long et al., 2015) or domain adversarial alignment (Ganin & Lempitsky, 2015). However, both strategies assume the source and target domains share the same feature encoder and need to access the source data during adaptation. This is not applicable in the proposed DA setting. In contrast, ADDA (Tzeng et al., 2017) relaxes the parameter-sharing constraint and proposes a new adversarial framework for DA, which learns different mapping functions for each domain. Also, DIRT-T (Shu et al., 2018) first trains a parameter-sharing DA framework as initialization and then fine-tunes the whole network by minimizing the cluster assumption violation. Both methods indicate that learning a domain-specific feature encoding module is practicable and even works better than the parameter-sharing mechanism.

We follow this strategy and develop a Source HypOthesis Transfer (SHOT) framework by learning the domain-specific feature encoding module while fixing the source classifier module (hypothesis), as the source hypothesis encodes the distribution information of unseen source data. Specifically, SHOT uses the *same* classifier module for different domain-specific feature learning modules, namely, $h_t = h_s$. It aims to learn the optimal target feature learning module $g_t : \mathcal{X}_t \rightarrow \mathbb{R}^d$ such that the output target features can match the source feature distribution well and can be classified accurately by the source hypothesis directly. It is worth noting that SHOT merely utilizes the source data for just once to generate the source hypothesis, while ADDA and DIRT-T need to access the data from source and target simultaneously when learning the model.

Essentially, we expect to learn the optimal target encoder g_t so that the target data distribution $p(g_t(x_t))$ matches the source data distribution $p(g_s(x_s))$ well. However, feature-level alignment does not work at all since it is impossible to estimate the distribution of $p(g_s(x_s))$ without access to the source data. We view the challenging problem from another perspective: *if the domain gap is mitigated, what kind of outputs should unlabeled target data have?* We argue the ideal target outputs should be similar to one-hot encoding but differ from each other. For this purpose, we adopt the information maximization (IM) loss (Krause et al., 2010; Shi & Sha, 2012; Hu et al., 2017), to make the target outputs individually certain and globally diverse. In practice, we minimize the following \mathcal{L}_{ent} and \mathcal{L}_{div} that together constitute the IM loss:

$$\begin{aligned} \mathcal{L}_{ent}(f_t; \mathcal{X}_t) &= -\mathbb{E}_{x_t \in \mathcal{X}_t} \sum_{k=1}^K \delta_k(f_t(x_t)) \log \delta_k(f_t(x_t)), \\ \mathcal{L}_{div}(f_t; \mathcal{X}_t) &= \sum_{k=1}^K \hat{p}_k \log \hat{p}_k \\ &= D_{KL}(\hat{p}, \frac{1}{K} \mathbf{1}_K) - \log K, \end{aligned} \quad (3)$$

where $f_t(x) = h_t(g_t(x))$ is the K -dimensional output of each target sample, $\mathbf{1}_K$ is a K -dimensional vector with all ones, and $\hat{p} = \mathbb{E}_{x_t \in \mathcal{X}_t} [\delta(f_t^{(k)}(x_t))]$ is the mean output embedding of the whole target domain. IM would work better than conditional entropy minimization (Grandvalet & Bengio, 2005) widely used in prior DA works (Vu et al., 2019; Saito et al., 2019), since IM can circumvent the trivial solution where all unlabeled data have the same one-hot encoding via the fair diversity-promoting objective \mathcal{L}_{div} .

3.3. Source Hypothesis Transfer Augmented with Self-supervised Pseudo-labeling

As shown in Fig. 1, we show the t-SNE visualizations of features learned by SHOT-IM and the source model only. Intuitively, the target feature representations are distributed in a mess for the source model only, and using the IM loss

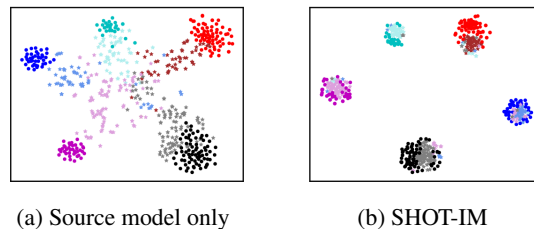


Figure 1. t-SNE visualizations for a 5-way classification task. Circles in dark colors denote unseen source data and stars in light denote target data. Different colors represent different classes.

helps align the target data with the unseen source data well. However, the target data may be matched to the wrong source hypothesis to some degree in Fig. 1(b).

We argue that the harmful effects result from the inaccurate network outputs. For instance, a target sample from the 2nd class with the normalized network output $[0.4, 0.3, 0.1, 0.1, 0.1]$ may be forced to have an expected output $[1.0, 0.0, 0.0, 0.0, 0.0]$. To alleviate such effects, we further apply pseudo-labeling (Lee, 2013) for each unlabeled data to better supervise the target data encoding training. However, pseudo labels that are conventionally generated by source hypotheses are still noisy due to domain shift. Inspired by DeepCluster (Caron et al., 2018), we propose a novel self-supervised pseudo-labeling strategy. First, we attain the centroid for each class in the target domain, similar to weighted k-means clustering,

$$c_k^{(0)} = \frac{\sum_{x_t \in \mathcal{X}_t} \delta_k(\hat{f}_t(x_t)) \hat{g}_t(x_t)}{\sum_{x_t \in \mathcal{X}_t} \delta_k(\hat{f}_t(x_t))}, \quad (4)$$

where $\hat{f}_t = \hat{g}_t \circ h_t$ denotes the previously learned target hypothesis. These centroids can robustly and more reliably characterize the distribution of different categories within the target domain. Then, we obtain the pseudo labels via the nearest centroid classifier:

$$\hat{y}_t = \arg \min_k D_f(\hat{g}_t(x_t), c_k^{(0)}), \quad (5)$$

where $D_f(a, b)$ measures the cosine distance between a and b . Finally, we compute the target centroids based on the new pseudo labels:

$$\begin{aligned} c_k^{(1)} &= \frac{\sum_{x_t \in \mathcal{X}_t} \mathbf{1}(\hat{y}_t = k) \hat{g}_t(x_t)}{\sum_{x_t \in \mathcal{X}_t} \mathbf{1}(\hat{y}_t = k)}, \\ \hat{y}_t &= \arg \min_k D_f(\hat{g}_t(x_t), c_k^{(1)}). \end{aligned} \quad (6)$$

We term \hat{y}_t as self-supervised pseudo labels since they are generated by the centroids obtained in an unsupervised manner. In practice, we update the centroids and labels in Eq. (6) for multiple rounds. However, experiments verify that updating for once gives sufficiently good pseudo labels.

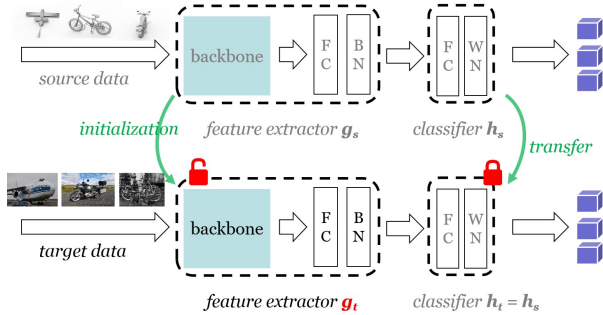


Figure 2. The pipeline of our SHOT framework. The source model consists of the feature learning module and the classifier module (hypothesis). SHOT keeps the hypothesis frozen and utilizes the feature learning module as initialization for target domain learning.

To summarize, given the source model $f_s = g_s \circ h_s$ and pseudo labels generated as above, SHOT freezes the hypothesis from source $h_t = h_s$ and learns the feature encoder g_t with the full objective as

$$\begin{aligned} \mathcal{L}(g_t) = & \mathcal{L}_{ent}(h_t \circ g_t; \mathcal{X}_t) + \mathcal{L}_{div}(h_t \circ g_t; \mathcal{X}_t) - \\ & \beta \mathbb{E}_{(x_t, \hat{y}_t) \in \mathcal{X}_t \times \hat{\mathcal{Y}}_t} \sum_{k=1}^K \mathbb{1}_{[k=\hat{y}_t]} \log \delta_k(h_t(g_t(x_t))), \end{aligned} \quad (7)$$

where $\beta > 0$ is a balancing hyper-parameter.

Remark. The proposed SHOT framework can also be easily extended to other UDA tasks like partial-set DA (Cao et al., 2018) and open-set DA (Panareda Busto & Gall, 2017). More details can be found in the appendix.

3.4. Network Architecture of Source Model

Now we study how to train a better source hypothesis for our own problem. We discuss some architecture choices for the neural network model to parameterize both the feature encoding module and hypothesis. First, we need to look back at the expected network outputs for cross-entropy loss in Eq. (1). If $y_s = k$, then maximizing $f_s^{(k)}(x_s) = \frac{\exp(w_k^T g_s(x_s))}{\sum_i \exp(w_i^T g_s(x_s))}$ means minimizing the distance between $g_s(x_s)$ and w_k , where w_k is the k -th weight vector in the last FC layer. Ideally, all the samples from the k -th class would have a feature embedding near to w_k . If unlabeled target samples are given the correct pseudo labels, it is easily understandable that source feature embeddings are similar to target ones via the pseudo-labeling term. The intuition behind is quite similar to previous studies (Long et al., 2013; Xie et al., 2018) where a simplified MMD is exploited for discriminative domain confusion.

Since the weight norm matters in the inner distance within the soft-max output, we adopt weight normalization (WN) (Salimans & Kingma, 2016) to keep the norm of each weight vector w_i the same in the FC classifier layer. Besides, as indicated in prior studies, batch normalization (BN) (Ioffe

& Szegedy, 2015) can reduce the internal dataset shift since different domains share the same mean (zero) and variance which can be considered as first-order and second-order moments. Based on these considerations, we form the complete framework of SHOT as shown in Fig. 2. In the experiments, we extensively study the effect of each aforementioned architecture design in Fig. 3.

4. Experiments

4.1. Setup

To testify its versatility, we evaluate SHOT in a variety of unsupervised DA scenarios, covering several popular visual benchmarks below. Full code is available at <https://github.com/tim-learn/SHOT>.

Office (Saenko et al., 2010) is a standard DA benchmark which contains three domains (Amazon (**A**), DSLR (**D**), and Webcam (**W**)) and each domain consists of 31 object classes under the office environment. (Gong et al., 2012) further extracts 10 shared categories between **Office** and Caltech-256 (**C**), and forms a new benchmark named **Office-Caltech**. Both benchmarks are small-sized.

Office-Home (Venkateswara et al., 2017) is a challenging medium-sized benchmark, which consists of four distinct domains (Artistic images (**Ar**), Clip Art (**Cl**), Product images (**Pr**), and Real-World images (Rw)). There are totally 65 everyday objects categories in each domain.

VisDA-C (Peng et al., 2017) is a challenging large-scale benchmark that mainly focuses on the 12-class synthesis-to-real object recognition task. The source domain contains 152 thousand synthetic images generated by rendering 3D models while the target domain has 55 thousand real object images sampled from Microsoft COCO.

Digits is a standard DA benchmark that focuses on digit recognition. We follow the protocol of (Hoffman et al., 2018) and utilize three subsets: SVHN (**S**), MNIST (**M**), and USPS (**U**). Like (Long et al., 2018), we train our model using the training sets of each domain and report the recognition results on the standard test set of the target domain.

Baseline Methods. For vanilla unsupervised DA in digit recognition, we compare SHOT with ADDA (Tzeng et al., 2017), ADR (Saito et al., 2018a), CDAN (Long et al., 2018), CyCADA (Hoffman et al., 2018), CAT (Deng et al., 2019), and SWD (Lee et al., 2019a). For object recognition, we compare with DANN (Ganin & Lempitsky, 2015), DAN (Long et al., 2015), ADR, CDAN, CAT, SAFN (Xu et al., 2019), BSP (Chen et al., 2019), and TransNorm (Wang et al., 2019) and SWD. For specific partial-set and open-set DA tasks, we compare with IWAN (Zhang et al., 2018a), SAN (Cao et al., 2018), ETN (Cao et al., 2019), SAFN, and ATI- λ (Panareda Busto & Gall, 2017), OSBP (Saito et al.,

Table 2. Classification accuracies (%) on **Digits** dataset for *vanilla closed-set DA*. S: SVHN, M:MNIST, U: USPS.

Method (Source→Target)	S→M	U→M	M→U	Avg.
Source only (Hoffman et al., 2018)	67.1±0.6	69.6±3.8	82.2±0.8	73.0
ADDA (Tzeng et al., 2017)	76.0±1.8	90.1±0.8	89.4±0.2	85.2
ADR (Saito et al., 2018a)	95.0±1.9	93.1±1.3	93.2±2.5	93.8
CDAN+E (Long et al., 2018)	89.2	98.0	95.6	94.3
CyCADA (Hoffman et al., 2018)	90.4±0.4	96.5±0.1	95.6±0.4	94.2
rRevGrad+CAT (Deng et al., 2019)	98.8±0.0	96.0±0.9	94.0±0.7	96.3
SWD (Lee et al., 2019a)	98.9±0.1	97.1±0.1	98.1±0.1	98.0
Source model only	67.1±0.9	87.8±2.3	89.6±0.4	81.5
SHOT-IM (ours)	89.6±5.0	96.8±0.4	91.9±0.4	92.8
SHOT (full, ours)	98.9±0.0	98.4±0.6	98.0±0.2	98.4
Target-supervised (Oracle)	99.4±0.0	99.4±0.0	98.0±0.1	98.9

2018c), OpenMax (Bendale & Boult, 2016), STA (Liu et al., 2019), respectively. For multi-source and multi-target DA, we compare with DCTN (Xu et al., 2018), MCD (Saito et al., 2018b), M^3 SDA- β (Peng et al., 2019a), FADA (Peng et al., 2020), DANN, and DADA (Peng et al., 2019b), respectively. Note that results are directly cited from published papers if we follow the same setting. Source model only denotes using the entire model from source for target label prediction. SHOT-IM is a special case of SHOT, where the self-supervised pseudo-labeling is ignored by letting $\beta = 0$.

4.2. Implementation Details

Network architecture. For the digit recognition task, we use the same architectures with CDAN (Long et al., 2018), namely, the classical LeNet-5 (LeCun et al., 1998) network is utilized for USPS↔MNIST and a variant of LeNet is utilized for SVHN→MNIST. Detailed networks are provided in Appendix A. For the object recognition task, we employ the pre-trained ResNet-50 or ResNet-101 (He et al., 2016) models as the backbone module like (Long et al., 2018; Deng et al., 2019; Xu et al., 2019; Peng et al., 2019a). Following (Ganin & Lempitsky, 2015), we replace the original FC layer with a bottleneck layer (256 units) and a task-specific FC classifier layer in Fig. 2. A BN layer is put after FC inside the bottleneck layer and a weight normalization layer is utilized in the last FC layer.

Network hyper-parameters. We train the whole network through back-propagation, and the newly added layers are trained with learning rate 10 times that of the pre-trained layers (backbone in Fig. 2). Concretely, we adopt mini-batch SGD with momentum 0.9 and weight decay $1e^{-3}$ and learning rate $\eta_0 = 1e^{-2}$ for the new layers and layers learned from scratch for all experiments except $\eta_0 = 1e^{-3}$ for **VisDA-C**. We further adopt the same learning rate scheduler $\eta = \eta_0 \cdot (1 + 10 \cdot p)^{-0.75}$ as (Ganin & Lempitsky, 2015; Long et al., 2018), where p is the training progress changing from 0 to 1. Besides, we set the batch size to 64 for all the tasks. We utilize $\beta = 0.3$ for all experiments except $\beta = 0.1$ for **Digits**.

For **Digits**, we train the optimal source hypothesis using the test set of the source dataset as validation. For other datasets without train-validation splits, we randomly specify a 0.9/0.1 split in the source dataset and generate the optimal source hypothesis based on the validation split. The maximum number of epochs for Digits, Office, Office-Home, VisDA-C and Office-Caltech is empirically set as 30, 100, 50, 10 and 100, respectively. For learning in the target domain, we update the pseudo-labels epoch by epoch, and the maximum number of epochs is empirically set as 15. We randomly run our methods for three times with different random seeds {2019, 2020, 2021} via **PyTorch** and report the average accuracies. Note that we do not exploit any target augmentation such as the ten-crop ensemble (Long et al., 2018) for evaluation.

4.3. Results of Digit Recognition

For digit recognition, we evaluate SHOT on three closed-set adaptation tasks, SVHN→MNIST, USPS→MNIST, and MNIST→USPS. The classification accuracies of SHOT and prior work are reported in Table 2. Obviously, SHOT obtains the best mean accuracies for each task and outperforms prior work in terms of the average accuracy. Compared with source model only, SHOT-IM always achieves better results, and SHOT performs better than SHOT-IM due to the contribution of self-supervised pseudo-labeling. It is worth noting that SHOT is comparable to the target-supervised results in M→U. This may be because MNIST is much larger than USPS, which alleviates the domain shift well.

 Table 3. Classification accuracies (%) on small-sized **Office** dataset for *vanilla closed-set DA* (ResNet-50).

Method (Source→Target)	A→D	A→W	D→A	D→W	W→A	W→D	Avg.
ResNet-50 (He et al., 2016)	68.9	68.4	62.5	96.7	60.7	99.3	76.1
DANN (Ganin & Lempitsky, 2015)	79.7	82.0	68.2	96.9	67.4	99.1	82.2
DAN (Long et al., 2015)	78.6	80.5	63.6	97.1	62.8	99.6	80.4
CDAN+E (Long et al., 2018)	92.9	94.1	71.0	98.6	69.3	100.	87.7
rRevGrad+CAT (Deng et al., 2019)	90.8	94.4	72.2	98.0	70.2	100.	87.6
SAFN+ENT (Xu et al., 2019)	90.7	90.1	73.0	98.6	70.2	99.8	87.1
CDAN+BSP (Chen et al., 2019)	93.0	93.3	73.6	98.2	72.6	100.	88.5
CDAN+TransNorm (Wang et al., 2019)	94.0	95.7	73.4	98.7	74.2	100.	89.3
Source model only	80.8	76.9	60.3	95.3	63.6	98.7	79.3
SHOT-IM (ours)	90.6	91.2	72.5	98.3	71.4	99.9	87.3
SHOT (full, ours)	94.0	90.1	74.7	98.4	74.3	99.9	88.6

4.4. Results of Object Recognition (Vanilla Closed-set)

Next, we evaluate SHOT on a variety of object recognition benchmarks including **Office**, **Office-Home** and **VisDA-C** under the vanilla closed-set DA setting. As shown in Table 3, SHOT performs the best for two challenging tasks, D→A and W→A, and performs worse than previous state-of-the-art method (Wang et al., 2019) for other tasks. This may be because SHOT needs a relatively large target domain to learn the hypothesis f_t while D and W are small as the target domain. Generally, SHOT obtains competitive performance even with no direct access to the source domain data.

Table 4. Classification accuracies (%) on medium-sized **Office-Home** dataset for *vanilla closed-set DA* (ResNet-50).

Method (Source→Target)	Ar→Cl	Ar→Pr	Ar→Re	Cl→Ar	Cl→Pr	Cl→Re	Pr→Ar	Pr→Cl	Pr→Re	Re→Ar	Re→Cl	Re→Pr	Avg.
ResNet-50 (He et al., 2016)	34.9	50.0	58.0	37.4	41.9	46.2	38.5	31.2	60.4	53.9	41.2	59.9	46.1
DANN (Ganin & Lempitsky, 2015)	45.6	59.3	70.1	47.0	58.5	60.9	46.1	43.7	68.5	63.2	51.8	76.8	57.6
DAN (Long et al., 2015)	43.6	57.0	67.9	45.8	56.5	60.4	44.0	43.6	67.7	63.1	51.5	74.3	56.3
CDAN+E (Long et al., 2018)	50.7	70.6	76.0	57.6	70.0	70.0	57.4	50.9	77.3	70.9	56.7	81.6	65.8
CDAN+BSP (Chen et al., 2019)	52.0	68.6	76.1	58.0	70.3	70.2	58.6	50.2	77.6	72.2	59.3	81.9	66.3
SAFN (Xu et al., 2019)	52.0	71.7	76.3	64.2	69.9	71.9	63.7	51.4	77.1	70.9	57.1	81.5	67.3
CDAN+TransNorm (Wang et al., 2019)	50.2	71.4	77.4	59.3	72.7	73.1	61.0	53.1	79.5	71.9	59.0	82.9	67.6
Source model 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
SHOT-IM (ours)	55.4	76.6	80.4	66.9	74.3	75.4	65.6	54.8	80.7	73.7	58.4	83.4	70.5
SHOT (full, ours)	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

Table 5. Classification accuracies (%) on large-scale **VisDA-C** dataset for *vanilla closed-set DA* (ResNet-101).

Method (Synthesis → Real)	plane	bcycl	bus	car	horse	knife	meycl	person	plant	sktbrd	train	truck	Per-class
ResNet-101 (He et al., 2016)	55.1	53.3	61.9	59.1	80.6	17.9	79.7	31.2	81.0	26.5	73.5	8.5	52.4
DANN (Ganin & Lempitsky, 2015)	81.9	77.7	82.8	44.3	81.2	29.5	65.1	28.6	51.9	54.6	82.8	7.8	57.4
DAN (Long et al., 2015)	87.1	63.0	76.5	42.0	90.3	42.9	85.9	53.1	49.7	36.3	85.8	20.7	61.1
ADR (Saito et al., 2018a)	94.2	48.5	84.0	72.9	90.1	74.2	92.6	72.5	80.8	61.8	82.2	28.8	73.5
CDAN (Long et al., 2018)	85.2	66.9	83.0	50.8	84.2	74.9	88.1	74.5	83.4	76.0	81.9	38.0	73.9
CDAN+BSP (Chen et al., 2019)	92.4	61.0	81.0	57.5	89.0	80.6	90.1	77.0	84.2	77.9	82.1	38.4	75.9
SAFN (Xu et al., 2019)	93.6	61.3	84.1	70.6	94.1	79.0	91.8	79.6	89.9	55.6	89.0	24.4	76.1
SWD (Lee et al., 2019a)	90.8	82.5	81.7	70.5	91.7	69.5	86.3	77.5	87.4	63.6	85.6	29.2	76.4
Source model 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
SHOT-IM (ours)	93.7	86.4	78.7	50.7	91.0	93.5	79.0	78.3	89.2	85.4	87.9	51.1	80.4
SHOT (full, ours)	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

As expected, on the medium-sized **Office-Home** dataset, SHOT significantly outperforms previously published state-of-the-art approaches, advancing the average accuracy from 67.6% (Wang et al., 2019) to 71.8% in Table 4. Besides, SHOT performs the best among 10 out of 12 separate tasks. For a large-scale synthesis-to-real **VisDA-C** dataset in Table 5, SHOT still achieves the best per-class accuracy and performs the best among 7 out of 12 tasks. Carefully comparing SHOT with prior work, we find that SHOT performs well for the most challenging class ‘truck’. Again, the proposed self-supervised pseudo-labeling strategy works well as SHOT always performs better than SHOT-IM.

Table 6. Average accuracies on three closed-set UDA datasets.

Methods / Datasets	Office	Office-Home	VisDA-C
Source model only	79.3	60.2	46.6
naive pseudo-labeling (PL) (Lee, 2013)	83.0	64.1	76.6
Self-supervised PL (ours)	87.6	68.9	80.7
\mathcal{L}_{ent}	83.5	55.5	63.3
$\mathcal{L}_{ent} + \mathcal{L}_{div}$	87.3	70.5	80.4
$\mathcal{L}_{ent} + \mathcal{L}_{div} +$ naive PL (Lee, 2013)	87.5	70.3	82.9
$\mathcal{L}_{ent} + \mathcal{L}_{div} +$ Self-supervised PL	88.6	71.8	82.9

Ablation Study. We study the advantages of the self-supervised pseudo-labeling (PL) strategy over naive one (Lee, 2013) in Table 6. It is clearly shown that both PL strategies work well and our self-supervised PL is always much better than the naive one. Surprisingly, merely using \mathcal{L}_{ent} along with self-supervised PL even produces lower results. But using both \mathcal{L}_{ent} and \mathcal{L}_{div} , the results become better than the self-supervised PL baseline, which indicates the importance of the diversity-promoting objective \mathcal{L}_{div} .

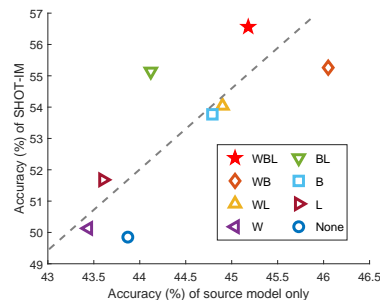


Figure 3. Accuracies (%) on the Ar→Cl task for *closed-set UDA*. [Weight normalization/ Batch normalization/ Label smoothing]

We then evaluate the contribution of each component within the network architecture, and show the results in Fig. 3. It seems that the accuracies of SHOT-IM partially depend on the quality of source model only. The results of source model only are especially boosted with the help of BN that is also adopted in prior studies (Xu et al., 2019; Wang et al., 2019). Using WN and LS may decrease the accuracies of source model only, but we still obtain better results via SHOT-IM. All these components complement to each other, and using each two of them together helps source model only and SHOT-IM achieve the best performance. Finally, SHOT-IM achieves the best performance when all these components are used in the network.

4.5. Results of Object Recognition (Beyond Vanilla Closed-set)

We further extend SHOT to two other DA scenarios, multi-source (Peng et al., 2019a) and multi-target (Peng et al.,

Table 7. (OS) Classification accuracies (%) on **Office-Home** dataset for *partial-set and open-set DA* (ResNet-50).

Partial-set DA (Source→Target)	Ar→Cl	Ar→Pr	Ar→Re	Cl→Ar	Cl→Pr	Cl→Re	Pr→Ar	Pr→Cl	Pr→Re	Re→Ar	Re→Cl	Re→Pr	Avg.
ResNet-50 (He et al., 2016)	46.3	67.5	75.9	59.1	59.9	62.7	58.2	41.8	74.9	67.4	48.2	74.2	61.3
IWAN (Zhang et al., 2018a)	53.9	54.5	78.1	61.3	48.0	63.3	54.2	52.0	81.3	76.5	56.8	82.9	63.6
SAN (Cao et al., 2018)	44.4	68.7	74.6	67.5	65.0	77.8	59.8	44.7	80.1	72.2	50.2	78.7	65.3
ETN (Cao et al., 2019)	59.2	77.0	79.5	62.9	65.7	75.0	68.3	55.4	84.4	75.7	57.7	84.5	70.5
SAFN (Xu et al., 2019)	58.9	76.3	81.4	70.4	73.0	77.8	72.4	55.3	80.4	75.8	60.4	79.9	71.8
Source model only	45.2	70.4	81.0	56.2	60.8	66.2	60.9	40.1	76.2	70.8	48.5	77.3	62.8
SHOT-IM (ours)	57.9	83.6	88.8	72.4	74.0	79.0	76.1	60.6	90.1	81.9	68.3	88.5	76.8
SHOT (full, ours)	64.8	85.2	92.7	76.3	77.6	88.8	79.7	64.3	89.5	80.6	66.4	85.8	79.3
Open-set DA (Source→Target)	Ar→Cl	Ar→Pr	Ar→Re	Cl→Ar	Cl→Pr	Cl→Re	Pr→Ar	Pr→Cl	Pr→Re	Re→Ar	Re→Cl	Re→Pr	Avg.
ResNet (He et al., 2016)	53.4	52.7	51.9	69.3	61.8	74.1	61.4	64.0	70.0	78.7	71.0	74.9	65.3
ATI-λ (Panareda Busto & Gall, 2017)	55.2	52.6	53.5	69.1	63.5	74.1	61.7	64.5	70.7	79.2	72.9	75.8	66.1
OSBP (Saito et al., 2018c)	56.7	51.5	49.2	67.5	65.5	74.0	62.5	64.8	69.3	80.6	74.7	71.5	65.7
OpenMax (Bendale & Boulton, 2016)	56.5	52.9	53.7	69.1	64.8	74.5	64.1	64.0	71.2	80.3	73.0	76.9	66.7
STA (Liu et al., 2019)	58.1	53.1	54.4	71.6	69.3	81.9	63.4	65.2	74.9	85.0	75.8	80.8	69.5
Source model only	36.3	54.8	69.1	33.8	44.4	49.2	36.8	29.2	56.8	51.4	35.1	62.3	46.6
SHOT-IM (ours)	62.5	77.8	83.9	60.9	73.4	79.4	64.7	58.7	83.1	69.1	62.0	82.1	71.5
SHOT (full, ours)	64.5	80.4	84.7	63.1	75.4	81.2	65.3	59.3	83.3	69.6	64.6	82.3	72.8

2019b). For the multi-source setting, we first learn multiple source hypotheses from different source domains and directly transfer them to the target domain one by one. Then the prediction scores of different optimized target hypotheses via SHOT are added up to find the label which has the maximum value. For the multi-target setting, we naively combine these target domains and treat it as the new target domain for SHOT. The results of SHOT and previously published state-of-the-arts are shown in Table 8. It is easy to find that SHOT achieves the best results for both settings. Note that FADA (Peng et al., 2020) does not access to the source domain but utilizes the gradient update information like federated learning, but it performs much worse than our SHOT for the multi-source setting.

Table 8. Classification accuracies (%) on **Office-Caltech** dataset for *multi-source and multi-target DA* (ResNet-101). [*] denotes the **rest** three domains except the single source / target.]

Multi-source (R →)	R → A	R → C	R → D	R → W	Avg.
ResNet-101 (He et al., 2016)	88.7	85.4	98.2	99.1	92.9
DAN (Long et al., 2015)	91.6	89.2	99.1	99.5	94.8
DCTN (Xu et al., 2018)	92.7	90.2	99.0	99.4	95.3
MCD (Saito et al., 2018b)	92.1	91.5	99.1	99.5	95.6
M ³ SDA-β (Peng et al., 2019a)	94.5	92.2	99.2	99.5	96.4
FADA (Peng et al., 2020)	84.2	88.7	87.1	88.1	87.1
Source model only	95.4	93.7	98.9	98.3	96.6
SHOT-IM (ours)	96.2	96.1	98.5	99.7	97.6
SHOT (full, ours)	96.4	96.2	98.5	99.7	97.7
Multi-target (→ R)	A → R	C → R	D → R	W → R	Avg.
ResNet-101 (He et al., 2016)	90.5	94.3	88.7	82.5	89.0
SE (French et al., 2018)	90.3	94.7	88.5	85.3	89.7
MCD (Saito et al., 2018b)	91.7	95.3	89.5	84.3	90.2
DANN (Ganin & Lempitsky, 2015)	91.5	94.3	90.5	86.3	90.7
DADA (Peng et al., 2019b)	92.0	95.1	91.3	93.1	92.9
Source model only	90.7	96.1	90.2	90.9	92.0
SHOT-IM (ours)	95.7	97.2	96.3	96.1	96.3
SHOT (full, ours)	96.2	97.3	96.3	96.2	96.5

Finally, we evaluate the generalization of SHOT by extending it to two challenging DA tasks, PDA and ODA, and follow the protocols on **Office-Home** in (Cao et al., 2019)

and (Liu et al., 2019), respectively. For PDA, there are totally 25 classes (the first 25 in alphabetical order) in the target domain and 65 classes in the source domain, while for ODA, the source domain consists of the same 25 classes but the target domain contains 65 classes including unknown samples. More details about how SHOT works for these two scenarios are provided in *Appendix B and C*, respectively. Results in Table 7 validate the effectiveness of SHOT for these two challenging asymmetric DA tasks.

Special Case. One may wonder whether SHOT works if we cannot train the source model by ourselves. To find the answer, we utilize the most popular off-the-shelf pre-trained ImageNet models ResNet-50 (He et al., 2016) and consider a PDA task (**ImageNet** → **Caltech**) to evaluate the effectiveness of SHOT with the same basic setting as (Cao et al., 2019). Obviously, in Table 9, SHOT still achieves slightly higher accuracy than the state-of-the-art ETN (Cao et al., 2019) even without accessing the source data.

Table 9. Results of a PDA task (**ImageNet** → **Caltech**). † utilizes the training set of ImageNet besides pre-trained ResNet-50 model.

Methods	ResNet-50	ETN†	SHOT-IM (ours)	SHOT (full, ours)
Accuracy	69.7±0.0	83.2±0.2	81.7±0.5	83.3±0.1

5. Conclusion

In this paper, we address a practical unsupervised DA setting with a simple yet generic representation learning framework named SHOT. SHOT merely needs the well-trained source model and offers the feasibility of unsupervised DA without access to the source data which may be private and decentralized. Specifically, SHOT learns the optimal target-specific feature learning module to fit the source hypothesis by exploiting the information maximization and self-supervised pseudo-labeling. Experiments for both digit and object recognition verify that SHOT achieves competitive and even state-of-the-art performance.

Acknowledgements: Jiashi Feng was partially supported by AISG-100E-2019-035 and MOE2017-T2-2-151. The authors would like to acknowledge the (partial) support from the Open Project Program (No. KBDat1505) of Jiangsu Key Laboratory of Big Data Analysis Technology, Nanjing University of Information Science & Technology. The authors also thank Quanhong Fu and Weihao Yu for their help to improve the technical writing aspect of this paper.

Algorithm 1 SHOT algorithm for closed-set UDA task.

Input: source model $f_s = g_s \circ h_s$, target data $\{x_i^t\}_{i=1}^{n_t}$, maximum number of epochs T_m , trade-off parameter β .
Initialization: Freeze the final classifier layer $h_t = h_s$, and copy the parameters from g_s to g_t as initialization.
for epoch = 1 **to** T_m **do**
 Obtain self-supervised pseudo labels via Eq. (6)
 for iter = 1 **to** n_b **do**
 # min-batch optimization
 Sample a batch from target data and get the corresponding pseudo labels.
 Update the parameters in g_t via $\mathcal{L}(g_t)$ in Eq. (7).
 end for
end for

6. Appendix

A. UDA scenarios and Network Architecture

Concerning the DA scenarios, the closet-set setting (Saenko et al., 2010) is the most favored, where the source and the target domain share the same label space. Later, the original one-to-one adaptation is extended to the multi-source (Peng et al., 2019a) and the multi-target (Peng et al., 2019b) setting, respectively. To be more realistic, open-set DA (Panareda Busto & Gall, 2017) assumes the target domain has some unseen classes, while partial-set DA (Cao et al., 2018) considers a setting where the label space of the source domain subsumes that of the target domain.

For the digit experiments, we follow (Long et al., 2018) and show the detailed networks in Fig. 4.

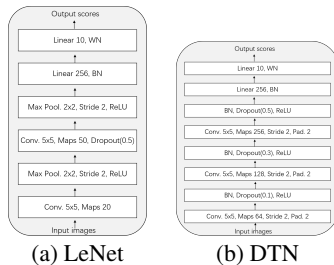


Figure 4. Network architectures used for digit experiments. (a) for MNIST \leftrightarrow USPS and (b) for SVHN \rightarrow MNIST.

B. Details of SHOT for Partial-set DA (PDA)

Looking at the diversity-promoting term \mathcal{L}_{div} in Eq. (3), it encourages the target domain to own a uniform target label distribution. It sounds reasonable for closed-set DA, but not

suitable for partial-set DA. In reality, the target domain only contains some classes of the whole classes in the source domain, making the label distribution sparse. Hence, we drop the second term \mathcal{L}_{div} for PDA.

Besides, within the self-supervised pseudo-labeling strategy, we usually need to obtain K centroids. However, for the PDA task, there are some tiny centroids which should be considered as empty like k-means clustering. Particularly, SHOT discards tiny centroids whose size is smaller than T_c in Eq. (6) for PDA. Next, we show the average accuracy on **Office-Home** with regards to the choice of T_c in Fig. 5.

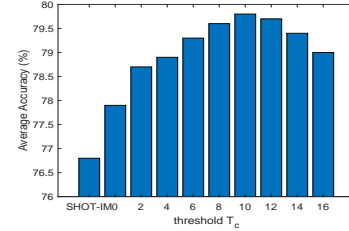


Figure 5. Average accuracies (%) of SHOT-IM and SHOT for all partial-set DA tasks on **Office-Home**. (ResNet-50)

C. Details of SHOT for Open-set DA (ODA)

Conversely, the target domain in an open-set DA problem contains some unknown classes besides the classes seen in the source domain. However, it is hard to change the final classifier layer for SHOT and learn it totally in the target domain. Thus we resort to a confidence thresholding strategy to reject unknown samples in the target domain during the learning of SHOT. In effect, we utilize the entropy of network output to compute the uncertainty, and normalize it in the range of $[0,1]$ by dividing $\log K$. For each epoch, we first compute all the uncertainty values and perform a 2-class k-means clustering. The cluster with larger mean uncertainty would be treated as the unknown class and not be considered to update the target centroids and compute the objective in \mathcal{L}_{ent} . Next, we show the changes in OS score, OS* score, and unknown accuracy during the process of SHOT in Fig. 6. Specifically, the OS score includes the unknown class and measures the per-class accuracy, i.e., $OS = \frac{1}{K+1} \sum_{k=1}^{K+1} Acc_k$, where K indicates the number of known classes and $(K+1)$ -th class is an unknown class. OS* score only measure the per-class accuracy on the known classes, i.e., $OS^* = \frac{1}{K} \sum_{k=1}^K Acc_k$.

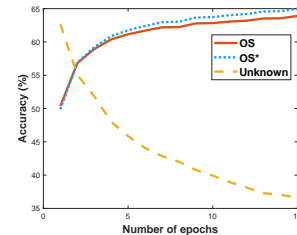


Figure 6. Different scores of SHOT w.r.t. the number of epochs for the open-set DA task Ar \rightarrow Cl on **Office-Home**. (ResNet-50)

References

- Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., and Vaughan, J. W. A theory of learning from different domains. *Springer MLJ*, 79(1-2):151–175, 2010.
- Bendale, A. and Boulton, T. E. Towards open set deep networks. In *CVPR*, 2016.
- Bhushan Damodaran, B., Kellenberger, B., Flamary, R., Tuia, D., and Courty, N. Deepjdot: Deep joint distribution optimal transport for unsupervised domain adaptation. In *ECCV*, 2018.
- Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., and Seth, K. Practical secure aggregation for privacy-preserving machine learning. In *ACM CCS*, 2017.
- Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., Kiddon, C., Konecny, J., Mazzocchi, S., McMahan, H. B., et al. Towards federated learning at scale: System design. *arXiv preprint arXiv:1902.01046*, 2019.
- Bousmalis, K., Trigeorgis, G., Silberman, N., Krishnan, D., and Erhan, D. Domain separation networks. In *NeurIPS*, 2016.
- Cao, Z., Long, M., Wang, J., and Jordan, M. I. Partial transfer learning with selective adversarial networks. In *CVPR*, 2018.
- Cao, Z., You, K., Long, M., Wang, J., and Yang, Q. Learning to transfer examples for partial domain adaptation. In *CVPR*, 2019.
- Cariucci, F. M., Porzi, L., Caputo, B., Ricci, E., and Bulò, S. R. Autodial: Automatic domain alignment layers. In *ICCV*, 2017.
- Caron, M., Bojanowski, P., Joulin, A., and Douze, M. Deep clustering for unsupervised learning of visual features. In *ECCV*, 2018.
- Chen, X., Wang, S., Long, M., and Wang, J. Transferability vs. discriminability: Batch spectral penalization for adversarial domain adaptation. In *ICML*, 2019.
- Chen, Y., Li, W., Sakaridis, C., Dai, D., and Van Gool, L. Domain adaptive faster r-cnn for object detection in the wild. In *CVPR*, 2018.
- Chidlovskii, B., Clinchant, S., and Csurka, G. Domain adaptation in the absence of source domain data. In *KDD*, 2016.
- Choi, J., Jeong, M., Kim, T., and Kim, C. Pseudo-labeling curriculum for unsupervised domain adaptation. In *BMVC*, 2019.
- Csurka, G. A comprehensive survey on domain adaptation for visual applications. In *Domain Adaptation in Computer Vision Applications*, pp. 1–35. Springer, 2017.
- Deng, Z., Luo, Y., and Zhu, J. Cluster alignment with a teacher for unsupervised domain adaptation. In *ICCV*, 2019.
- French, G., Mackiewicz, M., and Fisher, M. Self-ensembling for visual domain adaptation. In *ICLR*, 2018.
- Ganin, Y. and Lempitsky, V. Unsupervised domain adaptation by backpropagation. In *ICML*, 2015.
- Ghifary, M., Kleijn, W. B., Zhang, M., Balduzzi, D., and Li, W. Deep reconstruction-classification networks for unsupervised domain adaptation. In *ECCV*, 2016.
- Glorot, X., Bordes, A., and Bengio, Y. Domain adaptation for large-scale sentiment classification: a deep learning approach. In *ICML*, 2011.
- Gong, B., Shi, Y., Sha, F., and Grauman, K. Geodesic flow kernel for unsupervised domain adaptation. In *CVPR*, 2012.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *NeurIPS*, 2014.
- Grandvalet, Y. and Bengio, Y. Semi-supervised learning by entropy minimization. In *NeurIPS*, 2005.
- Gretton, A., Borgwardt, K., Rasch, M., Schölkopf, B., and Smola, A. J. A kernel method for the two-sample problem. In *NeurIPS*, 2007.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *CVPR*, 2016.
- Hoffman, J., Tzeng, E., Park, T., Zhu, J.-Y., Isola, P., Saenko, K., Efros, A., and Darrell, T. Cycada: Cycle-consistent adversarial domain adaptation. In *ICML*, 2018.
- Hu, W., Miyato, T., Tokui, S., Matsumoto, E., and Sugiyama, M. Learning discrete representations via information maximizing self-augmented training. In *ICML*, 2017.
- Huang, J., Gretton, A., Borgwardt, K., Schölkopf, B., and Smola, A. J. Correcting sample selection bias by unlabeled data. In *NeurIPS*, 2007.
- Ioffe, S. and Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015.

- Jiang, J. and Zhai, C. Instance weighting for domain adaptation in nlp. In *ACL*, 2007.
- Krause, A., Perona, P., and Gomes, R. G. Discriminative clustering by regularized information maximization. In *NeurIPS*, 2010.
- Kuzborskij, I. and Orabona, F. Stability and hypothesis transfer learning. In *ICML*, 2013.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Lee, C.-Y., Batra, T., Baig, M. H., and Ulbricht, D. Sliced wasserstein discrepancy for unsupervised domain adaptation. In *CVPR*, 2019a.
- Lee, D.-H. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, 2013.
- Lee, S., Kim, D., Kim, N., and Jeong, S.-G. Drop to adapt: Learning discriminative features for unsupervised domain adaptation. In *ICCV*, 2019b.
- Liang, J., He, R., Sun, Z., and Tan, T. Aggregating randomized clustering-promoting invariant projections for domain adaptation. *IEEE TPAMI*, 41(5):1027–1042, 2018.
- Liang, J., He, R., Sun, Z., and Tan, T. Distant supervised centroid shift: A simple and efficient approach to visual domain adaptation. In *CVPR*, 2019a.
- Liang, J., He, R., Sun, Z., and Tan, T. Exploring uncertainty in pseudo-label guided unsupervised domain adaptation. *Elsevier PRJ*, 96:106996, 2019b.
- Liu, H., Cao, Z., Long, M., Wang, J., and Yang, Q. Separate to adapt: Open set domain adaptation via progressive separation. In *CVPR*, 2019.
- Long, M., Wang, J., Ding, G., Sun, J., and Yu, P. S. Transfer feature learning with joint distribution adaptation. In *ICCV*, 2013.
- Long, M., Cao, Y., Wang, J., and Jordan, M. Learning transferable features with deep adaptation networks. In *ICML*, 2015.
- Long, M., Zhu, H., Wang, J., and Jordan, M. I. Deep transfer learning with joint adaptation networks. In *ICML*, 2017.
- Long, M., Cao, Z., Wang, J., and Jordan, M. I. Conditional adversarial domain adaptation. In *NeurIPS*, 2018.
- Mansour, Y., Mohri, M., and Rostamizadeh, A. Domain adaptation with multiple sources. In *NeurIPS*, 2009.
- McMahan, H. B., Ramage, D., Talwar, K., and Zhang, L. Learning differentially private recurrent language models. In *ICLR*, 2018.
- Müller, R., Kornblith, S., and Hinton, G. E. When does label smoothing help? In *NeurIPS*, 2019.
- Nelakurthi, A. R., Maciejewski, R., and He, J. Source free domain adaptation using an off-the-shelf classifier. In *IEEE BigData*, 2018.
- Panareda Busto, P. and Gall, J. Open set domain adaptation. In *ICCV*, 2017.
- Peng, X., Usman, B., Kaushik, N., Hoffman, J., Wang, D., and Saenko, K. Visda: The visual domain adaptation challenge. *arXiv preprint arXiv:1710.06924*, 2017.
- Peng, X., Bai, Q., Xia, X., Huang, Z., Saenko, K., and Wang, B. Moment matching for multi-source domain adaptation. In *ICCV*, 2019a.
- Peng, X., Huang, Z., Sun, X., and Saenko, K. Domain agnostic learning with disentangled representations. In *ICML*, 2019b.
- Peng, X., Huang, Z., Zhu, Y., and Saenko, K. Federated adversarial domain adaptation. In *ICLR*, 2020.
- Saenko, K., Kulis, B., Fritz, M., and Darrell, T. Adapting visual category models to new domains. In *ECCV*, 2010.
- Saito, K., Ushiku, Y., Harada, T., and Saenko, K. Adversarial dropout regularization. In *ICLR*, 2018a.
- Saito, K., Watanabe, K., Ushiku, Y., and Harada, T. Maximum classifier discrepancy for unsupervised domain adaptation. In *CVPR*, 2018b.
- Saito, K., Yamamoto, S., Ushiku, Y., and Harada, T. Open set domain adaptation by backpropagation. In *ECCV*, 2018c.
- Saito, K., Kim, D., Sclaroff, S., Darrell, T., and Saenko, K. Semi-supervised domain adaptation via minimax entropy. In *ICCV*, 2019.
- Salimans, T. and Kingma, D. P. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. In *NeurIPS*, 2016.
- Shi, Y. and Sha, F. Information-theoretical learning of discriminative clusters for unsupervised domain adaptation. In *ICML*, 2012.
- Shu, R., Bui, H. H., Narui, H., and Ermon, S. A dirt-t approach to unsupervised domain adaptation. In *ICLR*, 2018.

- Sun, B., Feng, J., and Saenko, K. Return of frustratingly easy domain adaptation. In *AAAI*, 2016.
- Tommasi, T., Orabona, F., and Caputo, B. Learning categories from few examples with multi model knowledge transfer. *IEEE TPAMI*, 36(5):928–941, 2013.
- Tsai, Y.-H., Hung, W.-C., Schuler, S., Sohn, K., Yang, M.-H., and Chandraker, M. Learning to adapt structured output space for semantic segmentation. In *CVPR*, 2018.
- Tzeng, E., Hoffman, J., Saenko, K., and Darrell, T. Adversarial discriminative domain adaptation. In *CVPR*, 2017.
- Venkateswara, H., Eusebio, J., Chakraborty, S., and Panchanathan, S. Deep hashing network for unsupervised domain adaptation. In *CVPR*, 2017.
- Vu, T.-H., Jain, H., Bucher, M., Cord, M., and Pérez, P. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *CVPR*, 2019.
- Wang, X., Jin, Y., Long, M., Wang, J., and Jordan, M. I. Transferable normalization: Towards improving transferability of deep neural networks. In *NeurIPS*, 2019.
- Xie, S., Zheng, Z., Chen, L., and Chen, C. Learning semantic representations for unsupervised domain adaptation. In *ICML*, 2018.
- Xu, R., Chen, Z., Zuo, W., Yan, J., and Lin, L. Deep cocktail network: Multi-source unsupervised domain adaptation with category shift. In *CVPR*, 2018.
- Xu, R., Li, G., Yang, J., and Lin, L. Larger norm more transferable: An adaptive feature norm approach for unsupervised domain adaptation. In *ICCV*, 2019.
- Yang, J., Yan, R., and Hauptmann, A. G. Cross-domain video concept detection using adaptive svms. In *ACM-MM*, 2007.
- Yosinski, J., Clune, J., Bengio, Y., and Lipson, H. How transferable are features in deep neural networks? In *NeurIPS*, 2014.
- Zellinger, W., Grubinger, T., Lughofer, E., Natschläger, T., and Saminger-Platz, S. Central moment discrepancy (cmd) for domain-invariant representation learning. In *ICLR*, 2017.
- Zhang, J., Ding, Z., Li, W., and Ogunbona, P. Importance weighted adversarial nets for partial domain adaptation. In *CVPR*, 2018a.
- Zhang, W., Ouyang, W., Li, W., and Xu, D. Collaborative and adversarial network for unsupervised domain adaptation. In *CVPR*, 2018b.
- Zhang, Y., David, P., and Gong, B. Curriculum domain adaptation for semantic segmentation of urban scenes. In *ICCV*, 2017.
- Zou, Y., Yu, Z., Vijaya Kumar, B., and Wang, J. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *ECCV*, 2018.