UMAD: Universal Model Adaptation under Domain and Category Shift

Jian Liang [1,2]   Dapeng Hu[*3]   Jiashi Feng [4]   Ran He [1,2]

[1] Institute of Automation, Chinese Academy of Sciences (CASIA)   [2] University of Chinese Academy of Sciences (UCAS)   [3] National University of Singapore (NUS)   [4] Sea AI Lab (SAIL)

liangjian92@gmail.com   dapeng.hu@u.nus.edu   jshfeng@gmail.com   rhe@nlpr.ia.ac.cn

Abstract

Learning to reject unknown samples (not present in the source classes) in the target domain is fairly important for unsupervised domain adaptation (UDA). There exist two typical UDA scenarios, i.e., open-set, and open-partial-set, and the latter assumes that not all source classes appear in the target domain. However, most prior methods are designed for one UDA scenario and always perform badly on the other UDA scenario. Moreover, they also require the labeled source data during adaptation, limiting their usability in data privacy-sensitive applications. To address these issues, this paper proposes a Universal Model ADaptation (UMAD) framework which handles both UDA scenarios without access to the source data nor prior knowledge about the category shift between domains. Specifically, we aim to learn a source model with an elegantly designed two-head classifier and provide it to the target domain. During adaptation, we develop an informative consistency score to help distinguish unknown samples from known samples. To achieve bilateral adaptation in the target domain, we further maximize localized mutual information to align known samples with the source classifier and employ an entropic loss to push unknown samples far away from the source classification boundary, respectively. Experiments on open-set and open-partial-set UDA scenarios demonstrate that UMAD, as a unified approach without access to source data, exhibits comparable, if not superior, performance to state-of-the-art data-dependent methods.

1. Introduction

Benefiting from the increasing number of labeled data, deep learning has achieved great success, particularly in supervised learning. However, annotating data in a new situation is always time-cost and expensive *e.g.* semantic segmentation in autonomous driving [11]. Thus, researchers resort to an alternative solution by fully exploiting knowl-
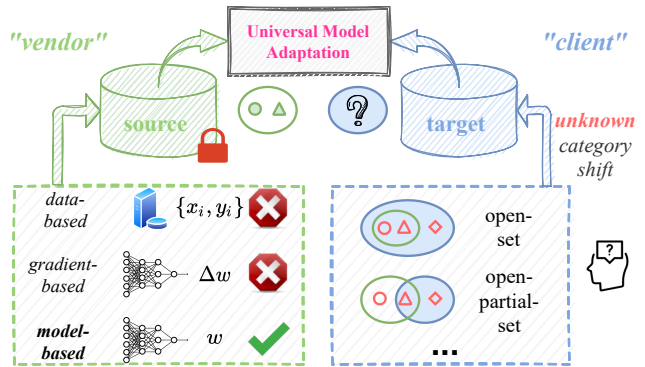

Figure 1. To pursue a broad range of applicability, we consider a challenging problem where the trained model instead of the data is provided from the source domain and meanwhile the detailed category shift (*i.e.*, open-set or open-partial-set) is unknown.

edge from data or models available in similar domains (referred to transfer learning (TL) [46]). As a particular case of TL, single-source unsupervised domain adaptation (UDA) [1] attracts increasing attention over the last decade, which has been successfully applied in many real-world applications, *e.g.*, image classification [15, 41], semantic segmentation [18, 62], and object detection [7, 22].

Typically, a UDA problem involves two related but not identical domains: source and target domains. The goal of UDA is to recognize the unlabeled samples in the target domain by leveraging knowledge from labeled data in the source domain. For a better illustration, we denote $\mathcal{C}_s$ and $\mathcal{C}_t$ as the label set of the source domain and the target domain, respectively. This well-studied identical label space assumption [15] ($\mathcal{C}_s = \mathcal{C}_t$, known as closed-set) is recently extended to various asymmetric settings, *e.g.*, partial-set [4] ($\mathcal{C}_s \supset \mathcal{C}_t$), open-set [59] ($\mathcal{C}_s \subset \mathcal{C}_t$), and open-partial-set [68] ($\mathcal{C}_s \cap \mathcal{C}_t \neq \varnothing, \mathcal{C}_s \not\supseteq \mathcal{C}_t, \mathcal{C}_s \not\subseteq \mathcal{C}_t$). Despite showing favorable results on one single setting with specific category overlap, previous UDA methods can not guarantee generalization to other settings. For example, a open-partial-set method [68] even under-performs the baseline source-only method for partial-set UDA [53]. A pio-

---

*The first two authors share co-first authorship.

1

neering work based on self-supervision is proposed in [53] to handle arbitrary settings across domains, however, it cannot well handle unknown samples in the target domain. Besides, with the increasing attention on data privacy, recent studies like [23, 33, 36] attempt to protect personal private data by only offering models trained in the source domain. For instance, the car company (*vendor*) in self-driving will only deploy the system but not share the training data. Assuming the presence of unknown samples in the target domain, this paper focuses on a new but realistic UDA problem illustrated in Fig. 1 where the trained model instead of the source data is provided and the detailed category shift (*i.e.*, open-set or open-partial-set) is unknown meanwhile.

Generally, there are three challenges in such a source data-free universal adaptation problem, *i.e.*, i) how to train a source model and adapt it to the unlabeled target domain; ii) how to guarantee the adaptation performance when some source classes may be absent in the target domain; and iii) how to guarantee the adaptation performance when some unknown classes exist in the target domain. To tackle them together, we propose a unified framework called Universal Model ADaptation (UMAD). Inspired by prior model adaptation methods [36, 37] which are mainly tailored to closed-set UDA, we also first train a self-defined source model and then provide it to the target domain for separate adaptation. Differently, we design a two-head classifier within the source model and introduce an informative consistency score to measure the uncertainty of target samples. With this score, we then propose a novel rejection mechanism to distinguish unknown samples from known samples in the target domain, where the key threshold is automatically determined by the averaged score of interpolated target samples via Mixup [72]. As for the bilateral model adaptation, we first reject unknown samples by pushing them far away from the source classification boundary. To align known samples with the corresponding source classifier, we then propose a new localized mutual information objective with a generalized diversity term.

Contributions. We study the problem of universal model adaptation, with unknown samples in the target domain and no access to the source data. We make the following contributions: 1). we propose a simple yet effective model adaptation framework that aligns known samples against domain shift and rejects unknown samples under unknown category shift; 2). we devise a new information consistency score and an automatic thresholding scheme to reject unknown samples. 3). we propose a localized mutual information objective to implicitly align known samples with the source model under various category shifts. 4). we validate the effectiveness of each component in UMAD via extensive experiments. Our source data-free UMAD even exhibits comparable performance to state-of-the-art data-dependent methods on open-set and open-partial-set UDA tasks.

2. Related Work

Unsupervised domain adaptation. Before the deep learning era, researchers utilize hand-crafted features and always resort to instance weighting [19, 60, 71] or feature alignment [16, 42, 61] to address covariate shift. In recent years, benefiting from representation learning, deep domain adaptation methods [15, 41, 63, 64] have almost dominated this field. To mitigate the gap across domains, there are two prevailing feature-level paradigms, *i.e.*, adversarial training [15] and discrepancy minimization [41]. By contrast, other studies [21, 36, 58] focus on the network outputs and expect the target data can be structurally close to the classification boundary. Some other works investigate the network components like batch normalization [34] and dropout [30] or study the properties of learned features [6, 66].

Most aforementioned UDA methods focus on the closed-set scenario where the source label set is the same as the target label set. To be more practical, the open-set UDA scenario is first introduced in [47] where unknown samples exist in both domains. The following work [59] further assumes the presence of unknown samples only in the target domain, and it allows extracting features that separate the unknown target from known target samples via adversarial training. Mixing the properties of open-set and partial-set [4] together, [68] introduces a realistic scenario termed "universal" UDA. It exploits both the domain similarity and the prediction uncertainty of each sample to develop a weighting mechanism for discovering label sets shared by both domains and thus promote the common-class adaptation. In a narrow sense, [68] only deals well with open-partial-set UDA but worse with other special cases like open-set and partial-set UDA scenarios. A recent work [53] proposes a truly universal UDA method that works well on all four scenarios, yet it performs badly for open-set and open-partial-set UDA. Generally, different from the well-studied closed-set scenario, these new scenarios are more challenging due to the asymmetric label set across domains, which are prone to cause the negative transfer. We focus on developing a unified method for adaptation to a domain with unknown samples, *i.e.*, open-set and open-partial-set UDA.

Source data-free domain adaptation. Besides the attention on the varying label set in the target domain, researchers also consider domain adaptation in the absence of source data due to data privacy concerns and copyright laws. A pioneering work [8] comes up with several solutions using the source classifier, and [35] requires the statistics information like means and co-variances during feature alignment. However, they do not exploit representation learning but use fixed features instead. Inspired by the idea of model transfer [26, 43, 45], [36] proposes a new framework where only the source feature module is fine-tuned in the target domain to achieve domain alignment. Meanwhile, [33] proposes clustering regularization and weight regu-
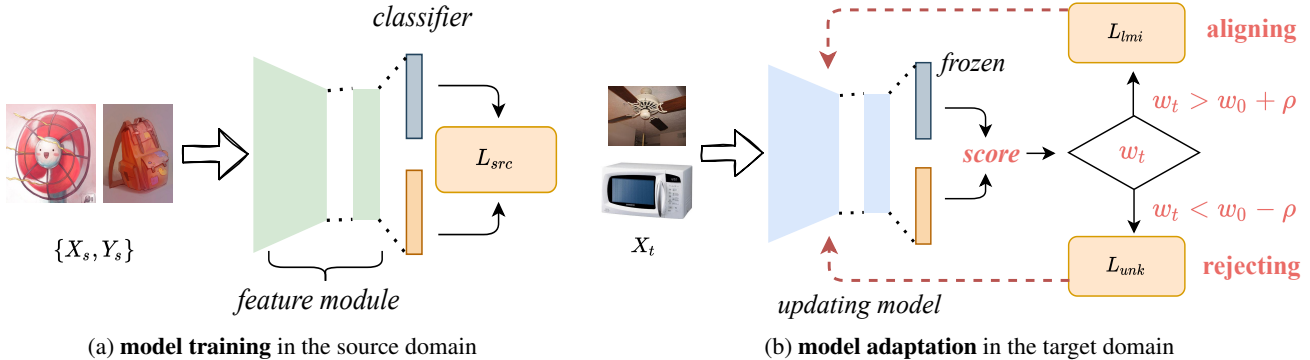
Figure 2. Overview of Universal Model ADaptation (UMAD). UMAD trains a deep neural network with a two-head classifier and provides the trained model to the target domain. During adaptation, UMAD freezes the two-head classifier and fine-tunes the source feature module from two different directions, *i.e.*, aligning known samples against domain shift and rejecting unknown samples against category shift.

larization that encourages similarity to the source model, and [23, 24] mainly focus on building a stronger pre-trained source model by generating unknown samples. Following works address this problem via variational inference [67], generative adversarial network [25], and mutual information [28], respectively. In addition, [50] exploits federated learning and transfers knowledge from the decentralized nodes to a new node with a different data domain. However, as evidenced in [73], sharing the gradient information is still risky. Besides, these methods above are designed for only one specific UDA scenario, *e.g.*, [33] for closed-set UDA, [24] for open-set UDA, and [23] for open-partial-set UDA. Even [36] provides results for UDA under three different scenarios, it still needs to know the UDA scenario in advance and subsequently adjusts the algorithm. By contrast, our method is a unified framework for both open-set and open-partial-set UDA with the same parameters.

Open-set learning. Open-set learning tries to handle "unknown" classes that are not contained in the training set, which is vital when deployed to real applications. [2] proposes an approach that builds per-class probabilistic models of the input not belonging to the known classes and combines these in the final estimate of each class probability including the unknown class. [17] shows that the maximum of the softmax outputs, or confidence, can be used to detect out-of-distribution (OOD) inputs. [39] exploits temperature scaling and perturbs an input in the direction of maximally increasing the max-softmax. Besides the aleatoric uncertainty, [14] employs Bayesian neural networks to consider epistemic uncertainty, which is computationally more efficient than an ensemble of networks [27]. [29] models the distribution of intermediate layer's activation by a Gaussian distribution for each class, and the OOD score is given by Mahalanobis distance and logistic regression, assuming the availability of OOD samples. Inspired by [58], the discrepancy between two classifiers is utilized to reject unknown samples in [70]. Our method differs from [70] in that we

propose a new information consistency score and an automatic thresholding scheme, which works under both domain shift and category shift.

3. Methodology

In a source data-free UDA task, there exist a source domain $\mathcal{D}_s = \{(x_i^s, y_i^s)\}_{i=1}^{N_s}$ consisting of N_s labeled samples and a target domain $\mathcal{D}_t = \{(x_i^t)\}_{i=1}^{N_t}$ consisting of N_t unlabeled samples, and $\mathcal{D}_s$ is not accessible during adaptation to the target domain. $\mathcal{C}_s \ni y_i^s$ and $\mathcal{C}_t \ni y_i^t$ denote the label sets of the source domain and the target domain, respectively. This paper mainly focuses on $\bar{\mathcal{C}}_t = \mathcal{C}_t \setminus \mathcal{C}_s \neq \varnothing$, which contains two open-set scenarios as special cases in Fig. 1. For a better illustration, $\bar{\mathcal{C}}_t$ denotes the target-only classes, and $\mathcal{C} = \mathcal{C}_s \cap \mathcal{C}_t$ denotes the classes existing in both domains. The ultimate goal is to recognize target samples belonging to $\mathcal{C}$ and reject target samples belonging to $\bar{\mathcal{C}}_t$ as "unknown" (*i.e.*, recognized as the $(K+1)$-th class, where $K = |\mathcal{C}_s|$ denotes the number of classes in the source domain).

3.1. Model training in the source domain

Although [50] offers a data-inaccessible solution via federated learning, the manner of gradient aggregation may be risky [73] and it still needs the presence of source data during adaptation. Instead, a trained source model could be provided to different target domains and would not leak the source data to our best knowledge. Then we mainly follow previous methods [33, 36] and adopt the model transfer strategy to address the data privacy issues. As shown in Fig. 2 (a), we adopt a similar network architecture as [36], and the only difference lies in the classifier layer. Denote by $g_s : \mathcal{X}_s \to \mathcal{R}^d$ the feature module within the network, and $h_v : \mathcal{R}^d \to \mathcal{R}^K, v = \{1, 2\}$ represent two different classifiers taking features from g_s, respectively. For simplicity, we denote $p^v(x_s) = h_v(g_s(x_s))$ as the K-dimensional softmax class probability from the v-th classifier for an input

x_s. The source model with two heads is learned by minimizing the following regularized co-training loss,

$$\mathcal{L}_{src} = -\mathbb{E}_{v,(x_s,y_s)\in\mathcal{D}_s} \sum_k q_k \log p_k^v(x_s) + \lambda\|W_1^T W_2\|, \quad (1)$$

where $p_k^v(x_s)$ denotes the k-th element of the soft-max output $p^v(x_s)$, and $q_k = (1-\alpha)\mathbb{1}[y_s = k] + \alpha/K$ is the smoothed one-of-K encoding of y_s, and α is the smoothing parameter empirically set to 0.1. Besides, λ is determined based on validation set, and the latter orthogonal constraint [5,57] (W_v denotes the fc weight in h_v, $v = 1, 2$) is employed to expect each head to learn different features.

3.2. Model adaptation in the target domain

Before we transfer the trained source model f_s in the target domain $\mathcal{D}_t$, we need to answer the following questions, (i) how to distinguish target samples belonging to $\mathcal{C}$ from those belonging to $\bar{\mathcal{C}}_t$, and (ii) how to mitigate the domain shift between $\mathcal{D}_s$ and $\mathcal{D}_t$ using samples from $\mathcal{C}$. For the first question, there are three common solutions to measure the uncertainty of target samples, $i.e.$, confidence, entropy, and consistency. Different from the mixture of three criteria in [13], we develop a new informative consistency score based on the two-head classifier, combining the advantages of both consistency and information entropy. It simply measures the inner product distance between the soft-max probabilities from h_1, h_2,

$$\mathbf{iscore}(x_t; g_t, h_1, h_2) = \mathbf{w}_t = <p^1(x_t),\ p^2(x_t)>, \quad (2)$$

where $<,>$ denotes the inner product. First, $\mathbf{w}_t$ always lies in the range of $[0,1]$ and gets rid of the non-trivial normalization. Second, when $p^1 = p^2$, $\mathbf{w}_t$ owns the same property as the collision entropy, which can measure the informative uncertainty. Third, the discrepancy between different classifiers has been proven effective for OOD detection [70]. Hereafter, we term $\mathbf{w}_t$ in Eq. (2) as "information consistency score". The larger $\mathbf{w}_t$ is, the more likely x_t belongs to $\mathcal{C}$ (known classes). We can confirm unknown samples have smaller values than known samples in Fig. 3. We prove in Fig. 3 that unknown samples generally have smaller **iscore** values than known ones.

As can be shown from Fig. 3, despite the domain shift, samples with $\mathcal{C}$ (known) are expected to have large scores and samples from $\bar{\mathcal{C}}_t$ (unknown) are expected to have small scores. Therefore, we aim to utilize a threshold w_0 to distinguish known samples and unknown samples in the target domain. To avoid the sensitivity to w_0, we further introduce a slack margin ρ and only consider the target samples beyond two thresholds $w_0 - \rho$ and $w_0 + \rho$ during adaptation. Since we are provided with the source model instead of the raw source data, it is hard to exploit prevailing paradigms like domain adversarial training and discrepancy minimization under the source data-free problem setting. Inspired
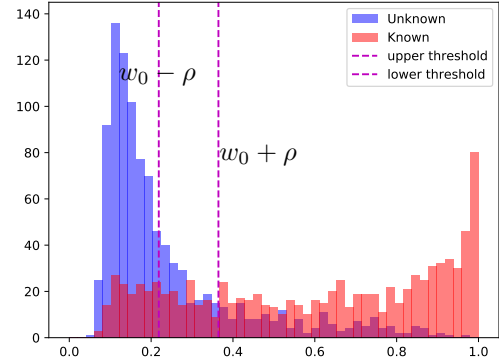

Figure 3. Histogram of the informative consistency score of all the target samples (**known** & **unknown**) through the source model.

by [36], we also freeze the classifier within the source model and try to merely learn the target-specific feature module by fine-tuning the source feature module.

Lemma 1 *For a linear multi-class classifier f_m, p is the soft-max output for the feature z. Given a well-trained data set $Z_s = \{z|\ \|f_m(z)\|_2 \to 1\}$ and the distance from a data point to a set $D(z_t, Z_s) = \min_{z\in Z_s} d_{cosine}(z_t, z)$, then we have $\lim_{\|f_m(z_t)\|_2\to 1} D(z_t, Z_s) \to 0$.*

Proof: For the sake of simplicity, we consider a linear classifier without the bias term, $i.e.$, $f_m(z) = \delta(W^T z)$, where $W \in \mathcal{R}^{d\times K}$, $\delta(\cdot)$ denotes the soft-max function, and d denotes the length of a feature vector z, and K denotes the size of classes. The non-negative cosine distance is defined as $d_{cosine} = 1 - \frac{<x,y>}{\|x\|\cdot\|y\|}$. Denote $W = [w_1, w_2, \cdots, w_K]$, then we can obtain the following term,

$$\|f_m(z)\|_2 = \|\delta(W^T z)\|_2 = \sum_j \left(\frac{\exp(w_j^T z)}{\sum_k \exp(w_k^T z)}\right)^2,$$

$$\|f_m(z)\|_2 \to 1 \Rightarrow \exists j \in [1, K],\ w_j^T z \gg w_k^T z,\ k \neq j,$$
$$\Rightarrow \exists j \in [1, K],\ d_{cosine}(z, w_j) \to 0. \quad (3)$$

Then the distance from z_t to Z_s can be computed as

$$\begin{aligned}
D(z_t, Z_s) &= \min_{z\in Z_s} d_{cosine}(z_t, z)\\
&\leq \min_{z\in Z_s} d_{cosine}(z_t, w_j) + d_{cosine}(z, w_j),\\
&\leq d_{cosine}(z_t, w_j) + \min_{z\in Z_s} d_{cosine}(z, w_j),
\end{aligned} \quad (4)$$

where $j = \arg\min_k d_{cosine}(z_t, w_k)$.

Since the source data set contains data from each class, it is easy to find the later term $\min_{z\in Z_s} d_{cosine}(z, w_j) \to 0$. In this manner, the conclusion $D(z_t, Z_s) \to 0$ is arrived. $\square$

4

According to Lemma 1, for "known" samples with consistency scores higher than $w_0 + \rho$, we devise a novel loss objective called localized mutual information $\mathcal{L}_{lmi}$ in the following, to encourage target features with diversified one-hot network outputs,

$$\mathcal{L}_{lmi} = \mathbb{E}_{v,x_t \in \mathcal{X}_t^+} p^v(x_t) \log p^v(x_t) - D_{kl}(Q_t^v \| \hat{\mathbf{Q}}_\mathbf{t}^\mathbf{v}),$$
$$\mathcal{X}_t^+ = \{x_t | x_t \in X_t, \mathbf{w}_t > w_0 + \rho\}, Q_t^v = \mathbb{E}_{x_t \in \mathcal{X}_t^+} p^v(x_t),$$
$$(5)$$

where D_{kl} is the Kullback-Leibler divergence function, Q_t^v denotes the batch-level mean probability of the v-th classifier, and $\hat{Q}_t^v$ is a K-dimensional uniform vector in [36, 37]. It is easy to find $\mathcal{L}_{lmi}$ considers the diversity and circumvents the trivial solution where all outputs collapse to a particular class. However, for partial-set or open-partial-set UDA, it is not reasonable since it would force target samples to be wrongly recognized as classes that only exist in the source domain. To address this problem in class-imbalanced domain adaptation, [31] proposes to use a moving average $\hat{Q}_t^v$ of $p(x_t)$. However, for closed-set UDA, it relies on a large batch size during training to estimate an accurate global label distribution. Besides, within a mini-batch, assuming the local label distribution is close to the global one is also not desirable. Instead, we propose a new operation called **flattening** as,

$$\text{Flatten}(p, T)_i = p_i^T / \sum_i p_i^T, \ 0 < T < 1, \qquad (6)$$

which is utilized to estimate the local label distribution. To avoid the influence of limited batch size and imbalanced label distribution, the batch-level mean probability Q_t^v is required to be close to its flattened one in the following,

$$\hat{\mathbf{Q}}_\mathbf{t}^\mathbf{v} = \text{Flatten}(Q_t^v, T). \qquad (7)$$

Note that $\mathcal{L}_{lmi}$ becomes the mutual information term when $T = 0$, and the simplified entropy minimization term for $T = 1$. In this paper, we empirically set this parameter to 0.1 for all open-set and open-partial-set UDA tasks.

To tackle with "unknown" samples with consistency scores smaller than $w_0 - \rho$, we do not introduce an additional $(K + 1)$-th class but employ the following entropic open-set loss [9] to push the target samples far away from the source classification boundary,

$$\mathcal{L}_{unk} = -\mathbb{E}_{v,x_t \in \mathcal{X}_t^-} \sum_k \frac{1}{K} \log p_k^v(x_t), \qquad (8)$$

where $\mathcal{X}_t^- = \{x_t | x_t \in X_t, \mathbf{w}_t < w_0 - \rho\}$, and K denotes the size of source label set $\mathcal{C}_s$. The minimum of the loss $\mathcal{L}_{unk}$ for a sample x_t is achieved when each soft-max scores $p^v(x_t)$ equals to the uniform vector. In that case, the consistency score becomes $w_t = 1/K$ which is quite small.

Combining these two different objectives in Eq. (5) and Eq. (8) for known and unknown samples together, the final

objective of UMAD during target adaptation is

$$\mathcal{L}_{tgt} = \mathcal{L}_{unk} - \mathcal{L}_{lmi}. \qquad (9)$$

3.3. How to decide the threshold w_0?

As the threshold parameter, w_0 is crucial for open-set learning including our UMAD approach, we further provide an intuitive strategy on how to select it. Assuming that samples in known classes $\mathcal{C}$ have confident predictions and samples in unknown classes $\bar{\mathcal{C}}_t$ have confused predictions, we can use Eq. (2) to calculate the mean informative consistency score of unknown samples for w_0. But the target domain is completely unlabeled, it is impracticable to obtain such a score. Instead, we synthesize such "negative" samples via *mixup* [72]. Then the mean score of "negative" samples is obtained by

$$w_u = \mathbb{E}_{x_i, x_j \in \mathcal{X}_t} \textbf{iscore}(0.5x_i + 0.5x_j). \qquad (10)$$

Intuitively, for a pair of known samples with high confidences, the confidence after mixing will be halved. By contrast, for a pair of unknown samples, the confidence after mixing is believed to be similar. In other words, w_u is much smaller than $\mathbb{E}_{x \in \mathcal{X}_t} \textbf{iscore}(x)$ if the number of samples in $\bar{\mathcal{C}}_t$ is small, and w_u is relatively smaller with the number of samples in $\bar{\mathcal{C}}_t$ increases, indicating w_u could be a suitable choice for w_0. It corresponds with the assumption that "*mixing a known sample with another known sample gets an unknown sample while mixing an unknown sample with a known sample gets another unknown sample*". Besides, we empirically set the slack margin $\rho = 0.1 * w_0$, its sensitivity analysis could be found in the experiment.

Table 1. Details of class splits in each setting. $|\mathcal{C}_s \cap \mathcal{C}_t|$: # shared classes across domains, $|\mathcal{C}_s - \mathcal{C}_t|$: # private classes in the source domain, and $|\mathcal{C}_t - \mathcal{C}_s|$: # private classes in the target domain.

Tasks	Datasets	$|\mathcal{C}_s \cap \mathcal{C}_t|$	$|\mathcal{C}_s - \mathcal{C}_t|$	$|\mathcal{C}_t - \mathcal{C}_s|$
	Office [52]	10	0	11
OSDA [40]	Office-Home [65]	25	0	40
	VisDA-C [51]	6	0	6
	Office [52]	10	10	11
OPDA [68]	Office-Home [65]	10	5	50
	VisDA-C [51]	6	3	3
	DomainNet [49]	150	50	145

4. Experiments

4.1. Setup

Datasets. Office [52] is a widely-used DA benchmark that consists of three subsets, *i.e.*, Amazon (A), Dslr (D), and Webcam (W). Each subset contains 31 object classes under the office environment. **Office-Home** [65] is another

Table 2. HOS (%) averaged over three runs of each method on **Office** and **VisDA-C** for OSDA. (Best in red and second best in blue.)

Method	SDF	OSDA	OPDA	A → D	A → W	D → A	D → W	W → A	W → D	Avg.	VisDA-C
OSBP [59]	✗	✓	✗	78.0	75.9	67.0	76.3	69.0	78.9	74.2	46.9
UAN [68]	✗	✗	✓	54.2	57.4	73.7	75.2	59.8	67.6	64.6	50.8
CMU [13]	✗	✗	✓	71.6	70.5	80.2	81.2	70.8	70.8	74.2	24.1
ROS [3]	✗	✓	✗	65.8	71.7	87.2	94.8	82.0	98.2	83.3	50.1
DANCE [54]	✗	✓	✓	82.0	74.7	68.0	82.1	52.2	82.5	73.6	59.7
DCC† [32]	✗	✓	✓	58.3	54.8	67.2	89.4	85.3	80.9	72.6	70.7
OVANet [56]	✗	✓	✓	89.4	87.6	85.2	97.2	87.8	97.9	90.8	56.1
SHOT [36]	✓	✓	✗	80.2	71.6	64.3	93.1	64.0	91.8	77.5	28.1
Inheritune† [24]	✓	✓	✗	78.0	81.4	83.1	92.2	91.3	99.7	87.6	74.8
OSHT-SC† [12]	✓	✓	✗	91.3	92.4	90.8	95.2	89.6	96.0	92.5	78.6
UMAD	✓	✓	✓	88.5	84.4	86.8	95.0	88.2	95.9	89.8	80.2

popular benchmark that consists of four subsets, *i.e.*, Artistic images (Ar), Clip Art (Cl), Product images (Pr), and Real-world images (Re). Each subset contains 65 object classes. **VisDA-C** [51] is a large-scale testbed, consisting of 12 object classes in two subsets, *i.e.*, real images and synthesis data. Typically, only the synthesis-to-real transfer task is investigated. **DomainNet** [49] is the largest domain adaptation dataset with about 0.6 million images, and we follow [13] to use three subsets, *i.e.*, Painting (P), Real (R), and Sketch (S). Details of classes split is listed in Table 1.

Baseline methods. Using the code from each corresponding author, we implement several methods for comparisons with our method, including OSBP [59], UAN [68], CMU [13], ROS [3], DANCE [53], OVANet [56], DCC [32], and SHOT [36]. We adopt class-balanced source sampling on all these methods. We also compare our method with Inheritune [24], USFDA [23], and OSHT-SC [12] using available results. Note that, as for the traditional data-dependent methods, OSBP and ROS are proposed for open-set UDA (OSDA), UAN and CMU and DCC are proposed for open-partial-set UDA (OPDA), and DANCE and OVANet are recently proposed to solve both UDA scenarios. SHOT is a popular source data-free method for closed-set UDA but can also be used as a baseline for open-set UDA, Inheritune and OSHT-SC are methods for open-set UDA, and USFDA is proposed to solve open-partial-set UDA. For comprehensive comparisons, we classify the above methods according to their applicable UDA scenarios, *i.e.*, OSDA and OPDA, and whether they are source data-free (SDF). Methods with † denotes results copied from the original papers or OVANet [56]. Further, full results of DCC [32] are kindly provided by the authors.

Implementation details. All the experiments are conducted via **Pytorch** [48]. We adopt mini-batch SGD to learn the feature encoder by fine-tuning from the ImageNet pretrained **ResNet-50** model with the learning rate 0.001, and new layers (bottleneck layer and classification layer) from scratch with the learning rate of 0.01. Batch Normalization [20] is employed in the bottleneck layer. For training

in both domains, we set the maximum number of iteration as 3,000, and use the suggested training settings in [36], including learning rate scheduler, momentum (0.9), weight decay (1e-3), bottleneck size (256), and batch size (64).

Evaluation metrics. Generally, we report the mean value over three random trials. For both open-set and open-partial-set UDA tasks, we report the HOS score [3, 13], which is a harmonic mean of accuracy over known samples and accuracy over unknown samples, in the main text. Denote acc_{kn} the average per-class accuracy over the known classes and acc_{ukn} the accuracy of the unknown class. HOS score can be easily obtained by

$$HOS = 2 \times \frac{\text{acc}_{kn} \times \text{acc}_{ukn}}{\text{acc}_{kn} + \text{acc}_{ukn}}. \tag{11}$$

The HOS score is considered fair to balance the rejection of unknown samples and the recognition of known samples.

4.2. Results

Open-set UDA. As for open-set UDA (OSDA), we compare our method with source data-free OSDA methods including SHOT, Inheritune, and OSHT-SC and make comparisons with data-dependent methods including universal methods like DANCE and OVANet, OSDA methods like OSBP and ROS, and OPDA methods like UAN and CMU. We report the results on **Office** and **VisDA-C** in Table 2 and the results on **Office-Home** in Table 3. As shown in Table 2 and Table 3, on OSDA tasks, methods designed for OPDA, *i.e.*, UAN and CMU, generally underperform methods taking into account the OSDA scenario like ROS and OVANet. This observation demonstrates the motivation of our work, *i.e.*, the scenario-tailored UDA method tends to be suffering under other DA scenarios and it is desirable to devise universal methods for DA with unknown target samples. Making comparisons among all methods, we find that our method achieves the best on two benchmarks, *i.e.*, 66.3% on the challenging 12-task benchmark **Office-Home** and 80.2% on the large-scale benchmark **VisDA-C**, and only underperforms the data-dependent method OVANet

Table 3. HOS (%) averaged over three runs of each method on **Office-Home** for OSDA.

Method	SDF	OSDA	OPDA	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.
OSBP [59]	×	✓	×	53.4	62.6	63.0	56.3	60.1	66.5	55.8	51.7	61.9	57.4	48.6	59.7	58.1
UAN [68]	×	×	✓	34.7	22.4	9.4	38.9	22.9	21.8	47.4	39.7	30.9	34.4	35.8	22.0	30.0
CMU [13]	×	×	✓	55.0	57.0	59.0	59.3	58.2	60.6	59.2	51.3	61.2	61.9	53.5	55.3	57.6
ROS [3]	×	✓	×	53.4	66.3	73.9	54.9	62.3	66.9	58.7	50.4	72.6	65.1	52.6	72.0	62.4
DANCE [54]	×	✓	✓	6.5	9.0	9.9	20.4	10.4	9.2	28.4	12.8	12.6	14.2	7.9	13.2	12.9
DCC† [32]	×	✓	✓	66.7	67.9	66.7	48.6	66.5	63.7	54.9	53.7	70.5	62.1	58.2	72.4	61.9
OVANet [56]	×	✓	✓	58.4	66.3	69.3	60.3	65.1	67.2	58.8	52.4	68.7	67.6	58.6	66.6	63.3
SHOT [36]	✓	✓	×	37.7	41.8	48.4	56.4	39.8	40.9	60.0	41.5	49.7	61.8	41.4	43.6	46.9
OSHT-SC† [12]	✓	✓	×	40.9	32.3	40.8	30.6	23.8	24.2	49.8	31.8	40.2	31.3	46.8	46.1	36.6
UMAD	✓	✓	✓	59.2	71.8	76.6	63.5	69.0	71.9	62.5	54.6	72.8	66.5	57.9	70.7	66.4

Table 4. HOS (%) averaged over three runs of each method on **Office**, **VisDA-C**, and **DomainNet** for OPDA.

Method	SDF	OSDA	OPDA	A→D	A→W	D→A	D→W	W→A	W→D	Avg.	VisDA-C	P→R	P→S	R→P	R→S	S→P	S→R	Avg.
OSBP [59]	×	✓	×	61.7	59.5	62.5	62.4	60.6	65.1	62.0	37.7	52.2	35.0	46.5	35.8	38.6	52.1	43.4
UAN [68]	×	×	✓	52.0	53.2	70.8	76.0	64.8	64.1	63.5	34.8	0.1	28.5	37.7	31.4	31.4	12.2	23.6
CMU [13]	×	×	✓	76.9	74.1	82.8	88.5	79.2	85.2	81.1	32.9	50.5	43.9	49.4	43.5	44.0	48.1	46.6
ROS [3]	×	✓	×	29.8	26.8	86.4	86.6	83.9	96.0	68.3	30.3	20.5	30.0	36.9	28.7	19.9	23.2	26.5
DANCE [54]	×	✓	✓	78.8	72.0	78.8	91.6	72.3	91.7	80.9	3.9	38.8	43.8	48.1	43.8	39.4	20.9	39.1
DCC† [32]	×	✓	✓	88.5	78.5	70.2	79.3	75.9	88.6	80.2	43.0	56.9	43.7	50.3	43.3	44.9	56.1	49.2
OVANet [56]	×	✓	✓	84.0	79.4	75.1	95.8	82.1	95.8	85.4	44.7	55.8	45.5	51.0	43.2	46.7	56.7	49.8
USFDA† [23]	✓	×	✓	85.5	79.8	83.2	90.6	88.7	81.2	84.8	-	-	-	-	-	-	-	-
SHOT [36]	✓	✓	×	73.5	67.2	59.3	88.3	77.1	84.4	74.9	44.0	35.0	30.8	37.2	28.3	31.9	32.2	32.6
UMAD	✓	✓	✓	79.1	77.4	87.4	90.7	90.4	97.2	87.0	58.3	59.0	44.3	50.1	42.1	32.0	55.3	47.1

and the source data-free OSDA method OSHT-SC on the standard benchmark **Office**, which verifies the effectiveness of our method. Furthermore, our method is the only one that consistently achieves the leading performance on all three benchmarks, unlike other state-of-the-art methods. Specifically, the best data-dependent universal DA method OVANet only achieves a HOS score of 56.1% on **VisDA-C** and the state-of-the-art source data-free OSDA method OSHT-SC only obtains a HOS score of 36.6% on **Office-Home**, both of which are far worse than the corresponding HOS scores of our method.

Open-partial-set UDA. As for open-partial-set UDA (OPDA), similarly, we compare our methods with various scenario-specific or universal UDA methods. We report the results on **Office**, **VisDA-C**, and **DomainNet** in Table 4 and the results on **Office-Home** in Table 5. As shown in Table 4 and Table 5, we have a similar observation that OPDA methods tend to outperform OSDA methods on various OPDA benchmarks, *e.g.*, CMU significantly beats the state-of-the-art OSDA method ROS on 3 out of all 4 benchmarks. Our method still achieves consistently leading performance on all benchmarks. To be specific, our method achieves the best on two benchmarks, *i.e.*, 87.0% on **Office** and 58.3% on **VisDA-C**, the third best HOS score of 48.0% on the largest DA benchmark **DomainNet**, and the second best HOS score of 69.4% on **Office-Home**. Besides, although slightly falling behind the state-of-the-art data-dependent universal UDA method OVANet on **Do-**

mainNet and **Office-Home**, our method significantly outperforms OVANet on **VisDA-C** without accessing source data.

Summary. Comparing the results for OSDA and OPDA, we find that universal UDA methods like OVANet and UMAD are able to handle UDA scenarios with unknown category shifts, which makes them stand out among various scenario-specific UDA methods. Besides the stable and superior performance, UMAD is more practical, in terms of better privacy-protection and lighter data transmission, than data-dependent methods with access to raw source data.

4.3. Analysis

Varying unknown classes. We compare UMAD with other methods with varying unknown classes $|\mathcal{C}_t - \mathcal{C}_s|$, and report results on **Office-Home** in Fig. 4. The compared methods include universal UDA methods like DANCE and OVANet, OPDA method like CMU, and the source data-free OSDA method SHOT. With an increasing number of unknown classes, UMAD stably outperforms most, if not all, other methods on both UDA scenarios, benefited from the proposed rejection mechanism.

Choice of the uncertainty score. We compare the performance of UMAD with different uncertainty scores in Fig. 5 (a), including the L_2 distance, the cosine distance, the mean entropy, and our proposed inner product distance. Results further demonstrate that only mean entropy and the inner product distance are effective and inner product con-

Table 5. HOS (%) averaged over three runs of each method on **Office-Home** for OPDA.

Method	SDF	OSDA	OPDA	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.
OSBP [59]	×	✓	×	51.0	55.2	71.6	51.4	48.3	61.2	55.5	50.7	63.4	52.6	49.7	54.3	55.4
UAN [68]	×	×	✓	38.7	32.1	30.6	45.2	32.8	35.1	54.9	42.0	47.0	51.5	50.8	43.9	42.1
CMU [13]	×	×	✓	57.3	59.9	66.9	64.4	59.0	64.0	65.6	56.5	68.6	67.2	60.2	65.1	62.9
ROS [3]	×	✓	×	54.0	77.7	85.3	62.1	71.0	76.4	68.8	52.4	83.2	71.6	57.8	79.2	70.0
DANCE [54]	×	✓	✓	34.1	23.9	38.3	46.7	21.6	35.4	58.2	47.5	39.4	32.8	38.3	43.1	38.3
DCC† [32]	×	✓	✓	80.8	73.8	80.8	63.7	71.7	69.3	73.5	53.6	80.2	74.4	57.1	76.8	69.8
OVANet [56]	×	✓	✓	59.7	76.9	80.0	68.8	69.1	76.2	69.6	56.9	81.0	75.5	62.0	78.6	71.2
SHOT [36]	✓	✓	×	32.9	29.5	39.6	56.8	30.1	41.1	54.9	35.4	42.3	58.5	33.5	33.3	40.7
UMAD	✓	✓	✓	61.1	76.3	82.7	70.7	67.7	75.7	64.4	55.7	76.3	73.2	60.4	77.2	70.1

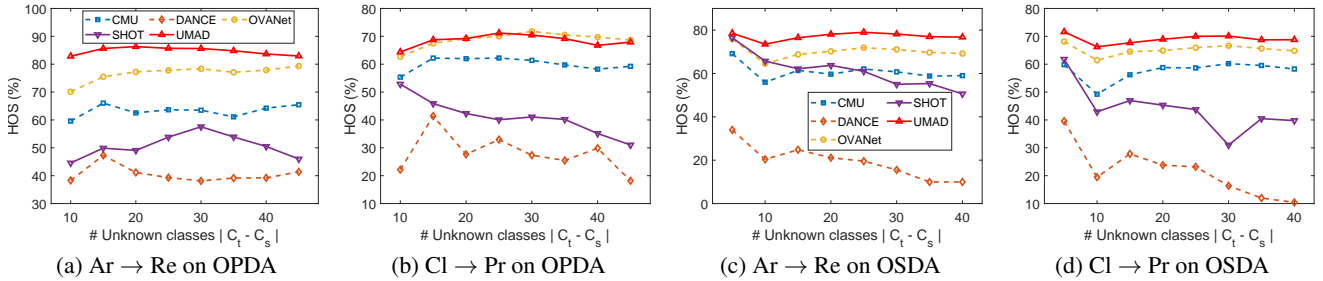(a) Ar → Re on OPDA (b) Cl → Pr on OPDA (c) Ar → Re on OSDA (d) Cl → Pr on OSDA

Figure 4. **HOS (%) of open-partial domain adaptation (OPDA) and open-set domain adaptation (OSDA)**. We vary the number of unknown classes (*i.e.*, $\|\mathcal{C}_t - \mathcal{C}_s\|$) of four different transfer tasks on the **Office-Home** dataset.

(a) uncertainty score (b) ablation on losses (c) sensitivity to ρ (d) sensitivity to T
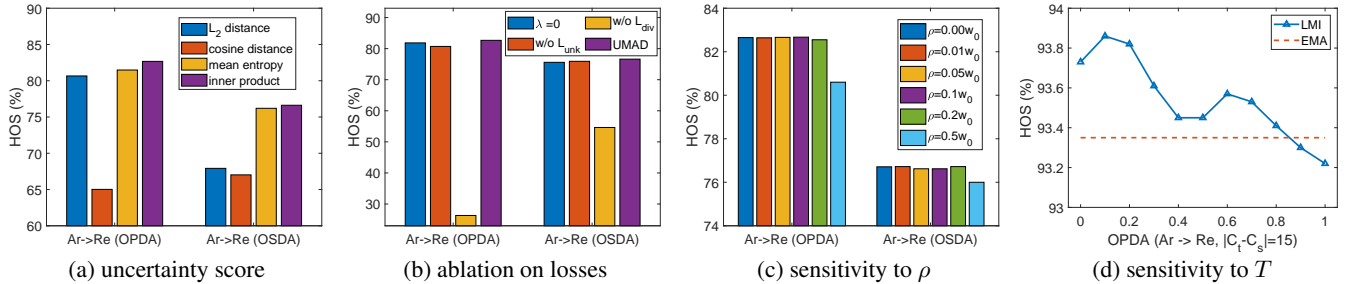
Figure 5. **Ablation study on different components of UMAD.** (a) shows results of other scores besides Eq. (2), (b) shows the contribution of each component in the final objective, and (c-d) plot the sensitivity to two different parameters.

sistently outperforms mean entropy on both UDA scenarios.

Ablation on losses. We provide the ablation on main loss objectives in UMAD, including the two-classifier orthogonal constraint, the novel diversity term, and the entropic loss. Results in Fig. 5 (b) demonstrate the importance of each loss. A significant decrease in the HOS score is expected with the removal of entropic loss. Because low accuracy of unknown samples will directly lead to low HOS score, according to Eq. (11).

Parameter sensitivity. We study the parameter sensitivity of ρ and T in Fig. 5 (c-d). We study the ratio of the slack margin in a wide range, *i.e.*, [0.00, 0.01, 0.05, 0.1, 0.2, 0.5], as shown in Fig. 5 (c). Results show that the performance of both UDA scenarios around the chosen ρ is stable and may be better. In Fig. 5 (d), we study the flattening fac-

tor T in the range of [0.0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0] and make comparisons with the vanilla EMA strategy [31]. The smaller T, the larger degree of flattening. Results show that within the value of 0.2, the performance of UMAD is not sensitive to T.

Extension to partial-set UDA. Besides the two UDA scenarios with unknown target samples investigated above, we ignore the rejection mechanism and apply UMAD to partial-set UDA [21,38] on **Office-Home** and report the average accuracy across 12 tasks over three runs in Table 6. We find our UMAD outperforms other data-dependent UDA methods including one of the state-of-the-art PDA methods—BA^{3}US [38]. Comparisons between different diversity terms show that our proposed LMI is the best and significantly outperforms the vanilla mutual informa-

Table 6. Avg. accuracy (%) on **Office-Home** for partial-set UDA.

Method	Avg.
MCC† [21]	75.1
JUMBOT† [10]	75.5
BA^{3}US† [38]	76.0
DCC† [32]	73.0
DANCE [54]	71.1
No Adaptation	64.4
UMAD ($w_0 = 0$, MI [36])	72.9
UMAD ($w_0 = 0$, EMA [31])	**78.0**
UMAD ($w_0 = 0$, LMI)	**78.4**

(a) No adaptation

(b) UMAD
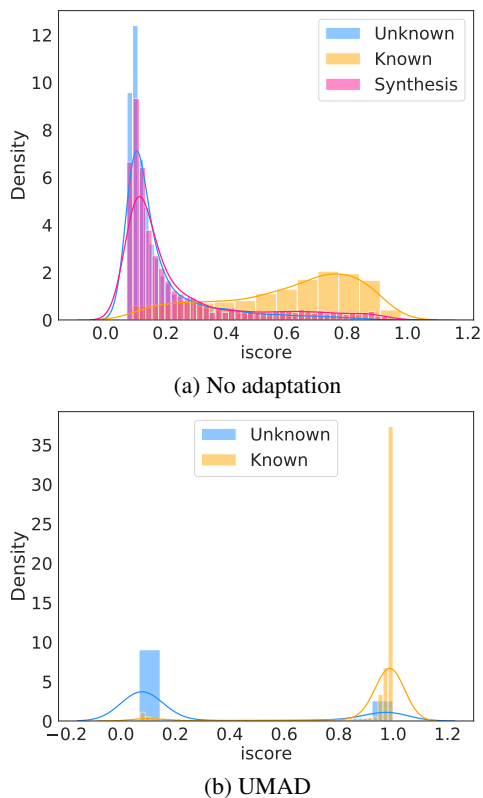
Figure 6. Histogram before/ after UMAD for Ar→Re (OPDA).

tion term [36] on partial-set UDA tasks.

Visualization of iscore. Fig. 6 (a) confirms that **iscore** in Eq. (2) is meaningful and the devised "synthesis" solution in Eq. (10) is reasonable. Fig. 6 (b) validates UMAD enlarges the discrepancy between known and unknown samples.

Limitations. So far, it is hard to conduct closed-set UDA validation methods [44, 55, 69] for UDA with unknown target samples. Besides, UMAD works well for both open-set and open-partial-set scenarios, but it still needs adjustment when adapting to a closed-set or partial-set scenario.

5. Conclusion

This paper investigated a more realistic open-set UDA setting where neither source data nor the prior about the label set overlap across domains is utilized during the target adaptation process. Using the trained model from the source domain, we proposed Universal Model ADaptation (UMAD). UMAD tackled category shift by devising an effective consistency score and an automatic thresholding scheme to detect samples from unknown classes. Besides, UMAD exploited a new localized mutual information objective to encourage target samples from known classes to fit the source model against domain shift. As a unified approach, UMAD has been validated to be effective across both open-set and open-partial-set UDA tasks on multiple datasets. As future work, UMAD can be extended to tackle complex tasks like segmentation and detection.

6. Notes

The early version of our manuscript (dubbed BATMAN) has been finished at Jan 2021. Code to reproduce experimental results will be released at `https://github.com/tim-learn/UMAD`, questions are welcome.

References

[1] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine Learning*, 79(1):151–175, 2010. 1

[2] Abhijit Bendale and Terrance E Boult. Towards open set deep networks. In *Proc. CVPR*, pages 1563–1572, 2016. 3

[3] Silvia Bucci, Mohammad Reza Loghmani, and Tatiana Tommasi. On the effectiveness of image rotation for open set domain adaptation. In *Proc. ECCV*, pages 422–438, 2020. 6, 7, 8

[4] Zhangjie Cao, Lijia Ma, Mingsheng Long, and Jianmin Wang. Partial adversarial domain adaptation. In *Proc. ECCV*, pages 135–150, 2018. 1, 2

[5] Minmin Chen, Kilian Q Weinberger, and John Blitzer. Co-training for domain adaptation. In *Proc. NeurIPS*, pages 2456–2464, 2011. 4

[6] Xinyang Chen, Sinan Wang, Mingsheng Long, and Jianmin Wang. Transferability vs. discriminability: Batch spectral penalization for adversarial domain adaptation. In *Proc. ICML*, pages 1081–1090, 2019. 2

[7] Yuhua Chen, Wen Li, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Domain adaptive faster r-cnn for object detection in the wild. In *Proc. CVPR*, pages 3339–3348, 2018. 1

[8] Boris Chidlovskii, Stéphane Clinchant, and Gabriela Csurka. Domain adaptation in the absence of source domain data. In *Proc. KDD*, pages 451–460, 2016. 2

[9] Akshay Raj Dhamija, Manuel Günther, and Terrance Boult. Reducing network agnostophobia. In *Proc. NeurIPS*, pages 9157–9168, 2018. 5

9

[10] Kilian Fatras, Thibault Séjourné, Rémi Flamary, and Nicolas Courty. Unbalanced minibatch optimal transport; applications to domain adaptation. In *Proc. ICML*, pages 3186–3197, 2021. 9

[11] Di Feng, Christian Haase-Schütz, Lars Rosenbaum, Heinz Hertlein, Claudius Glaeser, Fabian Timm, Werner Wiesbeck, and Klaus Dietmayer. Deep multi-modal object detection and semantic segmentation for autonomous driving: Datasets, methods, and challenges. *IEEE Transactions on Intelligent Transportation Systems*, 22(3):1341–1360, 2020. 1

[12] Zeyu Feng, Chang Xu, and Dacheng Tao. Open-set hypothesis transfer with semantic consistency. *IEEE Transactions on Image Processing*, 30:6473–6484, 2021. 6, 7

[13] Bo Fu, Zhangjie Cao, Mingsheng Long, and Jianmin Wang. Learning to detect open classes for universal domain adaptation. In *Proc. ECCV*, pages 567–583, 2020. 4, 6, 7, 8

[14] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *Proc. ICML*, pages 1050–1059, 2016. 3

[15] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *Proc. ICML*, pages 1180–1189, 2015. 1, 2

[16] Boqing Gong, Yuan Shi, Fei Sha, and Kristen Grauman. Geodesic flow kernel for unsupervised domain adaptation. In *Proc. CVPR*, pages 2066–2073, 2012. 2

[17] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *Proc. ICLR*, 2017. 3

[18] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *Proc. ICML*, pages 1989–1998, 2018. 1

[19] Jiayuan Huang, Arthur Gretton, Karsten Borgwardt, Bernhard Schölkopf, and Alex Smola. Correcting sample selection bias by unlabeled data. *Proc. NeurIPS*, pages 601–608, 2006. 2

[20] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proc. ICML*, pages 448–456, 2015. 6

[21] Ying Jin, Ximei Wang, Mingsheng Long, and Jianmin Wang. Minimum class confusion for versatile domain adaptation. In *Proc. ECCV*, pages 464–480, 2020. 2, 8, 9

[22] Taekyung Kim, Minki Jeong, Seunghyeon Kim, Seokeon Choi, and Changick Kim. Diversify and match: A domain adaptive representation learning paradigm for object detection. In *Proc. CVPR*, pages 12456–12465, 2019. 1

[23] Jogendra Nath Kundu, Naveen Venkat, R Venkatesh Babu, et al. Universal source-free domain adaptation. In *Proc. CVPR*, pages 4544–4553, 2020. 2, 3, 6, 7

[24] Jogendra Nath Kundu, Naveen Venkat, Ambareesh Revanur, R Venkatesh Babu, et al. Towards inheritable models for open-set domain adaptation. In *Proc. CVPR*, pages 12376–12385, 2020. 3, 6

[25] Vinod K Kurmi, Venkatesh K Subramanian, and Vinay P Namboodiri. Domain impression: A source data free domain adaptation method. In *Proc. WACV*, pages 615–625, 2021. 3

[26] Ilja Kuzborskij and Francesco Orabona. Stability and hypothesis transfer learning. In *Proc. ICML*, pages 942–950, 2013. 2

[27] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Proc. NeurIPS*, pages 6405–6416, 2017. 3

[28] Qicheng Lao, Xiang Jiang, and Mohammad Havaei. Hypothesis disparity regularized mutual information maximization. In *Proc. AAAI*, pages 8243–8251, 2021. 3

[29] Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. In *Proc. NeurIPS*, pages 7167–7177, 2018. 3

[30] Seungmin Lee, Dongwan Kim, Namil Kim, and Seong-Gyun Jeong. Drop to adapt: Learning discriminative features for unsupervised domain adaptation. In *Proc. ICCV*, pages 91–100, 2019. 2

[31] Bo Li, Yezhen Wang, Tong Che, Shanghang Zhang, Sicheng Zhao, Pengfei Xu, Wei Zhou, Yoshua Bengio, and Kurt Keutzer. Rethinking distributional matching based domain adaptation. *arXiv preprint arXiv:2006.13352*, 2020. 5, 8, 9

[32] Guangrui Li, Guoliang Kang, Yi Zhu, Yunchao Wei, and Yi Yang. Domain consensus clustering for universal domain adaptation. In *Proc. CVPR*, pages 9757–9766, 2021. 6, 7, 8, 9

[33] Rui Li, Qianfen Jiao, Wenming Cao, Hau-San Wong, and Si Wu. Model adaptation: Unsupervised domain adaptation without source data. In *Proc. CVPR*, pages 9641–9650, 2020. 2, 3

[34] Yanghao Li, Naiyan Wang, Jianping Shi, Xiaodi Hou, and Jiaying Liu. Adaptive batch normalization for practical domain adaptation. *Pattern Recognition*, 80:109–117, 2018. 2

[35] Jian Liang, Ran He, Zhenan Sun, and Tieniu Tan. Distant supervised centroid shift: A simple and efficient approach to visual domain adaptation. In *Proc. CVPR*, pages 2975–2984, 2019. 2

[36] Jian Liang, Dapeng Hu, and Jiashi Feng. Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In *Proc. ICML*, pages 6028–6039, 2020. 2, 3, 4, 5, 6, 7, 8, 9

[37] Jian Liang, Dapeng Hu, Yunbo Wang, Ran He, and Jiashi Feng. Source data-absent unsupervised domain adaptation through hypothesis transfer and labeling transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 2, 5

[38] Jian Liang, Yunbo Wang, Dapeng Hu, Ran He, and Jiashi Feng. A balanced and uncertainty-aware approach for partial domain adaptation. In *Proc. ECCV*, pages 123–140, 2020. 8, 9

[39] Shiyu Liang, Yixuan Li, and R Srikant. Enhancing the reliability of out-of-distribution image detection in neural networks. In *Proc. ICLR*, 2018. 3

[40] Hong Liu, Zhangjie Cao, Mingsheng Long, Jianmin Wang, and Qiang Yang. Separate to adapt: Open set domain adaptation via progressive separation. In *Proc. CVPR*, pages 2927–2936, 2019. 5

[41] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *Proc. ICML*, pages 97–105, 2015. 1, 2

[42] Mingsheng Long, Jianmin Wang, Guiguang Ding, Jiaguang Sun, and Philip S Yu. Transfer feature learning with joint distribution adaptation. In *Proc. ICCV*, pages 2200–2207, 2013. 2

[43] Yishay Mansour, Mehryar Mohri, and Afshin Rostamizadeh. Domain adaptation with multiple sources. In *Proc. NeurIPS*, pages 1041–1048, 2008. 2

[44] Pietro Morerio, Jacopo Cavazza, and Vittorio Murino. Minimal-entropy correlation alignment for unsupervised deep domain adaptation. In *Proc. ICLR*, 2018. 9

[45] Arun Reddy Nelakurthi, Ross Maciejewski, and Jingrui He. Source free domain adaptation using an off-the-shelf classifier. In *Proc. IEEE BigData*, pages 140–145, 2018. 2

[46] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359, 2009. 1

[47] Pau Panareda Busto and Juergen Gall. Open set domain adaptation. In *Proc. ICCV*, pages 754–763, 2017. 2

[48] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library, 2019. 6

[49] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proc. ICCV*, pages 1406–1415, 2019. 5, 6

[50] Xingchao Peng, Zijun Huang, Yizhe Zhu, and Kate Saenko. Federated adversarial domain adaptation. In *Proc. ICLR*, 2020. 3

[51] Xingchao Peng, Ben Usman, Neela Kaushik, Judy Hoffman, Dequan Wang, and Kate Saenko. Visda: The visual domain adaptation challenge. *arXiv preprint arXiv:1710.06924*, 2017. 5, 6

[52] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *Proc. ECCV*, pages 213–226, 2010. 5

[53] Kuniaki Saito, Donghyun Kim, Stan Sclaroff, and Kate Saenko. Universal domain adaptation through self supervision. In *Proc. NeurIPS*, 2020. 1, 2, 6

[54] Kuniaki Saito, Donghyun Kim, Stan Sclaroff, and Kate Saenko. Universal domain adaptation through self supervision. In *Proc. NeurIPS*, 2021. 6, 7, 8, 9

[55] Kuniaki Saito, Donghyun Kim, Piotr Teterwak, Stan Sclaroff, Trevor Darrell, and Kate Saenko. Tune it the right way: Unsupervised validation of domain adaptation via soft neighborhood density. In *Proc. ICCV*, pages 9184–9193, 2021. 9

[56] Kuniaki Saito and Kate Saenko. Ovanet: One-vs-all network for universal domain adaptation. In *Proc. ICCV*, pages 9000–9009, 2021. 6, 7, 8

[57] Kuniaki Saito, Yoshitaka Ushiku, and Tatsuya Harada. Asymmetric tri-training for unsupervised domain adaptation. In *Proc. ICML*, pages 2988–2997, 2017. 4

[58] Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In *Proc. CVPR*, pages 3723–3732, 2018. 2, 3

[59] Kuniaki Saito, Shohei Yamamoto, Yoshitaka Ushiku, and Tatsuya Harada. Open set domain adaptation by backpropagation. In *Proc. ECCV*, pages 153–168, 2018. 1, 2, 6, 7, 8

[60] Masashi Sugiyama, Shinichi Nakajima, Hisashi Kashima, Paul Von Buenau, and Motoaki Kawanabe. Direct importance estimation with model selection and its application to covariate shift adaptation. In *Proc. NeurIPS*, pages 1433–1440, 2007. 2

[61] Baochen Sun, Jiashi Feng, and Kate Saenko. Return of frustratingly easy domain adaptation. In *Proc. AAAI*, pages 2058–2065, 2016. 2

[62] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schulter, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *Proc. CVPR*, pages 7472–7481, 2018. 1

[63] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *Proc. CVPR*, pages 7167–7176, 2017. 2

[64] Eric Tzeng, Judy Hoffman, Ning Zhang, Kate Saenko, and Trevor Darrell. Deep domain confusion: Maximizing for domain invariance. *arXiv preprint arXiv:1412.3474*, 2014. 2

[65] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *Proc. CVPR*, pages 5018–5027, 2017. 5

[66] Ruijia Xu, Guanbin Li, Jihan Yang, and Liang Lin. Larger norm more transferable: An adaptive feature norm approach for unsupervised domain adaptation. In *Proc. ICCV*, pages 1426–1435, 2019. 2

[67] Hao-Wei Yeh, Baoyao Yang, Pong C Yuen, and Tatsuya Harada. Sofa: Source-data-free feature alignment for unsupervised domain adaptation. In *Proc. WACV*, pages 474–483, 2021. 3

[68] Kaichao You, Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Universal domain adaptation. In *Proc. CVPR*, pages 2720–2729, 2019. 1, 2, 5, 6, 7, 8

[69] Kaichao You, Ximei Wang, Mingsheng Long, and Michael Jordan. Towards accurate model selection in deep unsupervised domain adaptation. In *Proc. ICML*, pages 7124–7133, 2019. 9

[70] Qing Yu and Kiyoharu Aizawa. Unsupervised out-of-distribution detection by maximum classifier discrepancy. In *Proc. ICCV*, pages 9518–9526, 2019. 3, 4

[71] Bianca Zadrozny. Learning and evaluating classifiers under sample selection bias. In *Proc. ICML*, pages 903–910, 2004. 2

[72] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *Proc. ICLR*, 2018. 2, 5

[73] Ligeng Zhu, Zhijian Liu, and Song Han. Deep leakage from gradients. In *Proc. NeurIPS*, pages 14747–14756, 2019. 3