

“© 2019 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.”

Attract or Distract: Exploit the Margin of Open Set

Qianyu Feng¹, Guoliang Kang², Hehe Fan^{1,3*}, Yi Yang¹

¹ ReLER, University of Technology Sydney

² School of Computer Science, Carnegie Mellon University ³ Baidu Research

{qianyu.feng, hehe.fan}@student.uts.edu.au, yi.yang@uts.edu.au, kgl.prm1@gmail.com

Abstract

Open set domain adaptation aims to diminish the domain shift across domains, with partially shared classes. There exist unknown target samples out of the knowledge of source domain. Compared to the close set setting, how to separate the unknown (unshared) class from the known (shared) ones plays a key role. Whereas, previous methods did not emphasize the semantic structure of the open set data, which may introduce bias into the domain alignment and confuse the classifier around the decision boundary. In this paper, we exploit the semantic structure of open set data from two aspects: 1) **Semantic Categorical Alignment**, which aims to achieve good separability of target known classes by categorically aligning the centroid of target with the source. 2) **Semantic Contrastive Mapping**, which aims to push the unknown class away from the decision boundary. Empirically, we demonstrate that our method performs favourably against the state-of-the-art methods on representative benchmarks, e.g. Digit datasets and Office-31 datasets.

1. Introduction

Recent days have witnessed the advancement in many computer vision tasks [19, 37, 38, 13, 29, 21, 14, 43]. The success achieved can be largely attributed to the sufficient amount of labeled in-domain data. However, it is common that test data comes from a different distribution against the training data. Such so-called domain shift may degenerate the model performance heavily. Domain adaptation deals with this issue by diminishing the discrepancy across two domains. The widely considered *close set* setting assumes that both domains share the same set of underlying categories. However, in practice, it is common that some unshared (unknown) classes exist in the target. The methods

*Part of this work was done when Hehe Fan was an intern at Baidu Research. The code is publicly available at <https://github.com/qy-feng/margin-openset.git>

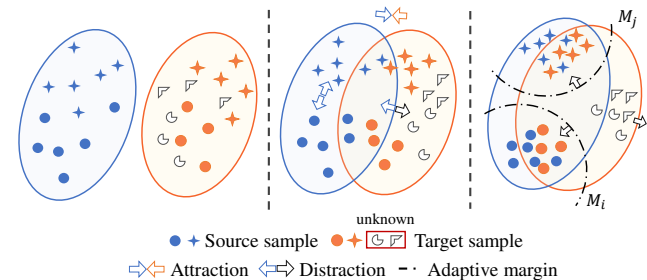


Figure 1. Visualization of data distribution with the proposed method. **Left:** Data before the adaptation with the existence of domain shift and unknown samples. **Middle:** Neighbors from the same class are pulled closer, while samples from unknown class are pushed away. **Right:** With the proposed method, representations become more discriminative. Samples from target domain can be better aligned within the corresponding neighborhood or distracted away from the known classes.

developed for close set domain adaptation may not be trivially transferred to such *open set* setting.

In this paper, we focus on the open set visual domain adaptation which aims to deal with the domain shift and the identification of unknown objects simultaneously, in the absence of target domain labels. Compared to the close set domain adaptation, how to separate the unknown class from the known ones plays a key role. Up to now, the open set recognition still remains as a pending issue. First raised by Busto *et al.* [27], they proposed to deal with open-set domain adaptation as an assignment task. [1] separated the unknown according to whether the sample can be reconstructed with the shared feature or not. While the above methods use part of labeled data from uninteresting classes as unknown samples, it is not possible to represent all the unknown categories in the wild. Another setting has been raised by Saito *et al.* [34] where unknown samples only exist in the target domain, which is closer to a realistic scenario. Saito *et al.* regarded the unknown samples as a separate class together with an adversarial loss to distinguish them. It is worth noting that the existence of unknown samples hinders the alignment across domain. In

the meanwhile, the disalignment inter-class across domain also makes it harder to distinguish the unknown samples.

Considering the aforementioned problem, we take the semantic structure of open set data into account to make the unknown class more separable and thus improve the model’s predictive ability on target domain data. Specifically, we focus on enlarging two kinds of margins: 1) the margin across the known classes and 2) the margin between the unknown samples and the known classes. For the first one, we aim to make the known classes more separable and for the second one, we expect to push the unknown class away from the decision boundary. As shown in Fig. 1, during training, samples coming from different domains but within the same class (e.g. the blue circle and the red circle) “attract” each other. For each domain, the margin between different known classes (e.g. the blue circle and the blue star) are enlarged. Moreover, the samples within unknown class (e.g. the irregular polygons) are “distracted” from the samples from known classes.

We propose using *semantic categorical alignment (SCA)* and *semantic contrastive mapping (SCM)* to achieve our goal. For semantic categorical alignment, due to the absence of target annotations, we indirectly promote the separability across target known classes by categorically aligning their centers with those in the source domain. For the source domain, although to an extent, the separability across known classes can be achieved through imposing the cross-entropy loss on the labeled data, we explicitly model and enhance such separability with the contrastive-center loss [28]. Empirically we demonstrate that our method leads to more discriminative features and benefit the semantic alignment across domains.

Although the semantic categorical alignment helps make the decision boundary aligned across two domains, there may still exist confusing data of unknown class lying near the decision boundary. Thus we propose using semantic contrastive mapping to push the unknown class away from the boundary. In detail, we design the contrastive loss to make the margin between the unknown and known class larger than that between known classes.

As the target labels are not available, we use the predictions of the network at each iteration as the hypothesis of target labels to perform the semantic categorical alignment and the semantic contrastive mapping. We start our training from the source trained model to give a good initialization of target label hypothesis. Although the hypothesis of target labels may not be accurate, SCA and SCM in itself are robust to such noisy labels. Empirically we find that the estimated SCA/SCM loss works as a good proxy to improve the model’s performance on the target domain data.

In a nutshell, our contributions can be summarized as

- We propose using semantic categorical alignment to achieve good separability of target known classes and

semantic contrastive mapping to push the unknown class away from the decision boundary. Both benefits the adaptation performance noticeably.

- Our method performs favourably against the-state-of-the-art methods on two representative benchmarks, i.e. on Digits dataset, it achieves 84.3% accuracy on the average, 1.9% higher than the state-of-the-art; on Office-31 dataset, we achieve 89.7% with AlexNet and 89.1% with VGG.

2. Related work

Domain adaptation for visual recognition aims to bridge the knowledge gap between across different domains. Approaches for open set recognition attempt to figure out the unknown samples while identifying samples from known classes. In a real scenario, data not only come from diverse domains but also varies in a wide range of categories. This paper focuses on dealing with the overlap of these two problems.

Many methods [22, 23, 8, 39, 2, 36, 17, 42, 24] have been proposed for unsupervised domain adaptation (UDA), including the deep network. These work bring significant results focusing on the *closed set* in the following aspects. *Distribution-based learning.* Many approaches aim to learn features invariant to domain with a distance metric [22, 23, 40], e.g., KL-divergence, Maximum Mean Discrepancy (MMD), Wasserstein distance, but they neglected the alignment of conditional distribution. The categorical information is exploited to align domains at a fine-grained level together with the pseudo-label. The marginal distribution and conditional distribution can also be jointly aligned with a combined MMD proposed by [31]. [35, 12, 32, 33] pay attention to the discriminative property of the representation. This paper is also related to work [16, 10, 3, 25, 5, 6, 41] considering the categorical semantic compactness and separability as the same time. *Task-oriented learning.* Approaches [39, 2, 7] tend to align the domain discrepancy in an adversarial style. Ganin *et al.* [7] proposed to learn domain-invariant feature by using an adversarial loss which reverses the gradients during the back-propagation. Bousmalis *et al.* [2] enabled the network to separate the generated features into domain-specific subspace and domain-invariant subspace.

The aforementioned methods tackle with domain adaptation in the closed-set scenario. Inspired by recent work in open-set recognition, the problem of Open Set Domain Adaptation (OSDA) is raised by Busto *et al.* [27]. With several classes trained as unknown samples, Busto *et al.* proposed to solve this problem by learning a mapping across domains with an assignment problem of the target samples. As is aforementioned, it is not possible to cover all the unknown samples with selected categories. Another setting

domain. However, the feature generator G tries to minimize the difference between the source and the target. When an expert D fails to figure out which domain the sample comes from, the G learns the domain-invariant representation.

We use the cross-entropy loss together with the softmax function for the known source samples classification,

$$\begin{aligned}\mathcal{L}_{cls}(x_s, y_s) &= -\log(p(y = y_s | \mathbf{x}_s)), \\ &= -\log(D \circ G(\mathbf{x}_s))_{y_s}.\end{aligned}\quad (1)$$

Following [34], in an attempt to make a boundary for an unknown sample, we utilize a binary cross entropy loss,

$$\begin{aligned}\mathcal{L}_{adv}(\mathbf{x}_t) &= -\frac{1}{2} \log(p(y = N + 1 | \mathbf{x}_t)) \\ &\quad -\frac{1}{2} \log(1 - p(y = N + 1 | \mathbf{x}_t)).\end{aligned}\quad (2)$$

By the gradient reverse layer [7], we can flip the sign of the gradient during backward, which allows us to update the parameters of G and D simultaneously. Then, the objective of the ADA module can be formulated as

$$\begin{aligned}\mathcal{L}_{ADA} &= \min_G (\mathcal{L}_{cls}(\mathbf{x}_s, y_s) - \mathcal{L}_{adv}(\mathbf{x}_t)) + \\ &\quad \min_D (\mathcal{L}_{cls}(\mathbf{x}_s, y_s) + \mathcal{L}_{adv}(\mathbf{x}_t)).\end{aligned}\quad (3)$$

The ADA module only initially aligns samples in the target with source known samples and learns a rough boundary between the known and the unknown.

3.3. Semantic Categorical Alignment

We try to address the issues existing in ADA by further exploring the semantic structure of open-set data. To separate the unknown class from the known in the target domain, we should 1) make each known class more concentrate and the alignment between the source and the target more accurate and 2) push the unknown class away from the decision boundary. In this section, we aim to solve the first problem. We introduce the Semantic Categorical Alignment (SCA), which aims to compact the representation of known classes and distinguish each known class from others. There are two steps in SCA. First, the contrastive-center loss[28] is adopted to enhance the discriminative property of generated features of source samples. Second, each centroid of known classes from target will be aligned with the corresponding centroid of class in source domain. In this way, representations of source samples will finally become more discriminative, meanwhile, the known target centroids will be aligned more accurate.

To compact the source samples that belong to the same class in the feature space, we apply the following contrastive-center loss to the source samples,

$$\mathcal{L}_{cct} = \frac{1}{2} \sum_{i=1}^m \frac{\|\mathbf{x}_s^i - \mathbf{c}_s^{y_s^i}\|_2^2}{(\sum_{j=1, j \neq y_s^i}^N \|\mathbf{x}_s^i - \mathbf{c}_s^j\|_2^2) + \delta}, \quad (4)$$

where m denotes the number of samples in a mini-batch during training procedure, \mathbf{x}_s^i denotes the i -th training sample from the source domain. $\mathbf{c}_s^{y_s^i}$ denotes the centroid of class y_s^i in the source domain. δ is a constant used for preventing zero-denominator. In our experiments, $\delta = 1$ is set to be 10^{-6} by default.

To align the two centroids of a known class between the source and target, we try to minimize the distance between the pair of centroids $dist(\mathbf{c}_s^k, \mathbf{c}_t^k) = \|\mathbf{c}_s^k - \mathbf{c}_t^k\|^2$, where \mathbf{c}_s^k and \mathbf{c}_t^k represent the centroids of class k from the source and target domain, respectively.

Due to the randomness and deviation in each mini-batch, we align the global centroids (calculated from all samples) instead of the local centroids (calculated from a mini-batch). However, it is not easy to directly obtain the global centroids. We propose to partially update them with the local centroids at every iteration, according to their cosine similarities to the centroids in the source domain. Specifically, we first compute the initial global centroids based on the prediction of the pretrained model as follows,

$$\mathbf{c}_{(0)}^k = \frac{1}{n^k} \sum_{j=0}^{n^k} G(\mathbf{x}_i^k), \quad (5)$$

where n^k denotes the number of samples with label as k . The pretrained model is trained on the source domain within the supervised classification paradigm. For the target samples, we use the results of prediction as pseudo labels. In each iteration, we compute a set of local centroids $\mathbf{a}_{(I)}^k$ using the mini-batch samples, where I denotes the iteration. We compute the local centroids as the average of all samples in each iteration. Then, the source centroid \mathbf{c}_s^k and target centroid \mathbf{c}_t^k are updated with re-weighting as follows,

$$\begin{aligned}\rho_s &= \rho(\mathbf{a}_{s(I)}^k, \mathbf{c}_{s(I-1)}^k), \\ \mathbf{c}_{s(I)}^k &\leftarrow \rho_s \mathbf{a}_{s(I)}^k + (1 - \rho_s) \mathbf{c}_{s(I-1)}^k,\end{aligned}\quad (6)$$

$$\begin{aligned}\rho_t &= \rho(\mathbf{a}_{t(I)}^k, \mathbf{c}_{s(I-1)}^k), \\ \mathbf{c}_{t(I)}^k &\leftarrow \rho_t \mathbf{a}_{t(I)}^k + (1 - \rho_t) \mathbf{c}_{t(I-1)}^k,\end{aligned}\quad (7)$$

where $\rho(\cdot, \cdot)$ is defined as $\rho(x_i, x_j) = (\frac{x_i \cdot x_j}{\|x_i\| \times \|x_j\|} + 1)/2$. Finally, the categorical center alignment loss is formulated as follows,

$$\mathcal{L}_{cca} = \sum_{k=1}^N dist(\mathbf{c}_{s(I)}^k, \mathbf{c}_{t(I)}^k). \quad (8)$$

The benefits of SCA are intuitive: 1) The contrastive center loss, *i.e.*, Eq. (4), enhances the compactness of the representations which also enlarges the margin inter-class. 2) The categorical center alignment loss, Eq. (8), guarantees that the centroids of the same class are aligned between

the source domain and the target domain. 3) The dynamic update together ensures that the SCA aligns the global and up-to-date categorical distributions. Furthermore, the reweighting technique weakens the incorrect pseudo-labels and therefore can alleviate the accumulated error of the pseudo-labels.

3.4. Semantic Contrastive Mapping

SCA aligns the centroids of the same class between the source domain and the target domain. For the non-centroid samples in the target domain, we employ a contrastive loss function to encourage the known samples to move closer to their centroids and enforce the unknown samples to stay far away from all the centroids of known classes. By this way, we can align the non-centroid samples in the target domain. We refer to this process as the Semantic Contrastive Mapping (SCM).

Since the pseudo labels of target samples are not totally correct, we select reliable samples whose classification probabilities are over a threshold. We set the threshold to $\frac{1}{N+1}$ in our method. SCM aims to reduce the distance between the reliable known samples and their centroids, while enlarge the distance between the reliable unknown samples and all centroids.

$$\mathcal{L}_{con}(x_t; G) = (1 - z)\mathcal{D}_k(x_t^k, c_s^k) - \frac{z}{N} \sum_{k=1}^N \mathcal{D}_u(x_t^k, c_s^k), \quad (9)$$

where z is equal to 0 if x_t is predicted as class $\in 1, 2, \dots, N$, otherwise, z equals to 1. \mathcal{D}_k denotes the distance between target known samples and the corresponding source centroid. \mathcal{D}_u denotes the distance between target unknown samples and all the source known classes. Inspired by the energy-base model in [11], functions are designed as follows,

$$\mathcal{D}_k(x_t^k, c_s^k) = (1 - \rho)^\omega \text{dist}(x_t^k, c_s^k)^2, \quad (10)$$

$$\mathcal{D}_u(x_t^{N+1}, c_s^k) = -\rho^\omega (\max\{0, M^k - \text{dist}(x_t^{N+1}, c_s^k)\})^2, \quad (11)$$

where ρ denotes the cosine similarity. To ensure an efficient and accurate measurement of the distances, we also use a hyper-parameter ω to re-weight distances calculated in the loss. M^k is a categorical adaptive margin to measure the radius of neighborhood of class k , defined as follows,

$$M^k = \frac{1}{N} \sum_{j=1, j \neq k}^N \text{dist}(c_t^j, c_s^k). \quad (12)$$

3.5. Objective

In the proposed method, considering the intra-class compactness and inter-class separability, we design the two modules SCA and SCM based on the adversarial learning

Algorithm 1 Exploit the Margin of Open Set, e denotes the training step, I denotes the iteration times.

Require: Labeled samples batches $X_s = \{(x_{s_i}, y_{s_i})\}_{i=1}^{n_s}$ from source domain, unlabeled samples batches $X_t = \{x_{t_j}\}_{j=1}^{n_t}$ from target domain. B_i^s and B_i^t denote the i_{th} mini-batch data in the training set.

Ensure: Parameters in the network θ_G, θ_D

- 1: **1st Stage**
 - 2: Pretrain G and D based on $\{X_s, Y_s\}$, update θ_G, θ_D
 - 3: **2nd Stage**
 - 4: $e = 0$
 - 5: **while** not converge **do**
 - 6: Calculate the current global centroids $c_{s(e)}^k$ and $c_{t(e)}^k$
 - 7: **for** $I = 1$ **to** max_iter **do**
 - 8: Update $c_{s(e)}^k$ and $c_{t(e)}^k$ by using Eq. 6 and Eq. 7
 - 9: Calculate pair distance between $c_{s(e)}^k$ and $c_{t(e)}^k$
 - 10: Select reliable target samples $\hat{X}_{t(I)}$
 - 11: Calculate pair distance between $c_{s(e)}^k, \hat{X}_{t(I)}$
 - 12: Train $model_m$ with B_i^s, B_i^t by optimizing loss in Eq. 13, update θ_G, θ_D
 - 13: $e = e + 1$
 - 14: **end for**
 - 15: **end while**
-

in ADA. Formally, the final objective is defined in Eq. 13.

$$\begin{aligned} \mathcal{L}_{total} &= \mathcal{L}_{ADA} + \mathcal{L}_{SCA} + \mathcal{L}_{SCM} \\ &= \mathcal{L}_{cls} + \mathcal{L}_{adv} + \lambda_s \mathcal{L}_{cct} + \lambda_c \mathcal{L}_{cca} + \lambda_t \mathcal{L}_{con}. \end{aligned} \quad (13)$$

In each iteration, the network updates the class centroids and network parameters simultaneously. The overall algorithm is shown in Algorithm 1. SCA attempts to enlarge the margins between known classes in source and categorically align the centroids across domain. SCM attempts to align all the known target samples to its source neighborhoods, while keeping the distance between unknown samples and the centroids of known classes around an adaptively determined margin. With SCA, the discriminator in ADA is access to more discriminative representation and well-aligned semantic features. On the other side, SCM aids to distinguish the unknown samples from the other known classes.

4. Experiments

4.1. Setup

In this section, we evaluate the proposed method on the open set domain adaptation task using two benchmarks, *i.e.*, Digit datasets and *Office-31* [30]. Considering the setting where unknown samples only exist in the target domain. We compare the performance of our method OSDA+BP [34] and other baselines including: Open-set SVM (OSVM) [15] and other methods combined with OSVM, *e.g.*, Maximum

Method	SVHN-MNIST				USPS-MNIST				MNIST-USPS				Average			
	OS	OS*	ALL	UNK	OS	OS*	ALL	UNK	OS	OS*	ALL	UNK	OS	OS*	ALL	UNK
OSVM	54.3	63.1	37.4	10.5	43.1	32.3	63.5	97.5	79.8	77.9	84.2	89.0	59.1	57.7	61.7	65.7
MMD+OSVM	55.9	64.7	39.1	12.2	62.8	58.9	69.5	82.1	80.0	79.8	81.3	81.0	68.0	68.8	66.3	58.4
BP+OSVM	62.9	75.3	39.2	0.7	84.4	92.4	72.9	0.9	33.8	40.5	21.4	44.3	60.4	69.4	44.5	15.3
OSDA+BP[34]	63.0	59.1	71.0	82.3	92.3	91.2	94.4	97.6	92.1	94.9	88.1	78.0	82.4	81.7	84.5	85.9
Ours w/o SCA	65.6	61.6	73.9	85.4	93.6	95.4	87.9	82.8	86.5	86.1	88.1	88.5	81.9	81.0	83.3	85.6
Ours w/o SCM	65.5	61.0	74.8	87.8	92.5	93.8	87.1	81.1	84.6	84.0	86.0	87.7	80.9	79.6	82.6	85.5
Ours	68.6	65.5	75.3	84.3	93.1	95.2	92.8	91.7	91.3	92.0	90.7	87.8	84.3	84.2	86.3	87.9

Table 1. Accuracy (%) of experiments on Digit dataset.




Figure 3. A comparison between the existing method and the proposed method. **First row**: visualization of features of the state-of-the-art method OSDA+BP[34]. **Second row**: visualization of features generated by our method. The left two column show features of source and target in SVHN \rightarrow MNIST, the right two columns are features of source and target in MNIST \rightarrow USPS. The color red in the target represents the unknown class.

Mean Discrepancy(MMD) [9], BP [7], ATI- λ [27]. OSVM classifies test samples into unknown class with a threshold of probability when the predicted probability is lower than the threshold for other classes. OSVM also requires no unknown samples in the source domain during training. MMD+OSVM is a combination method with OSVM and MMD-based method for network in [22]. MMD is discrepancy measure metric used to match the distribution across domains. BP+OSVM combines OSVM with a domain classifier, BP [7], which is a representative of adversarial learning applied in unsupervised domain adaptation.

Digits We begin by exploring three Digit datasets, *i.e.* SVHN [26], MNIST [20] and USPS [20]. SVHN contains colored digit images of size 32×32 , where more than one digit may appear in a single image. MNIST includes 28×28 grey digit images and USPS consists of 16×16 grey digit images. We conduct 3 common adaptation scenarios including SVHN to MNIST, USPS to MNIST and MNIST to USPS.

Office-31 [30] is a standard benchmark for domain adaptation. There exist three distinct domains: Amazon (A) with 2817 images from the merchants, Webcam (W) with 795 images of low resolution and DSLR (D) with 498 images

of high resolution. Each domain shares 31 categories with the others. We examine the full transfer scenarios in our experiments.

Implementation For Digit datasets, we employ the same architecture with [34]. For *Office-31*, we employ two representative CNN architectures, AlexNet [19] and VGGNet [37], to extract the visual features. For both the generator and classifier, we use one-layer FC followed with Leaky-RELU and Batch-Normalization. For *Office-31*, we initialize the feature extractor from the ImageNet [4] pre-trained model. For both datasets, we first train our model with labeled source domain data. All networks are trained by Adam [18] optimizer with weight decay 10^{-6} . The initial learning rates for Digit and *Office-31* datasets is 2×10^{-4} and 10^{-3} respectively. Learning rate decreases following a cosine ramp-down schedule. We set the hyper-parameters $\lambda_s = 0.02$, $\lambda_c = 0.005$, and $\lambda_t = 10^{-4}$ in all the experiments. Following [27], we report the accuracy averaged over the classes in the OS and OS*. The average accuracy of all classes including the unknown one is denoted as OS. Accuracy measures only on the known classes of the target domain is denoted as OS*. All the results reported are the accuracy averaged over three independent running.

	Adaptation Scenario													
	A-D		A-W		D-A		D-W		W-A		W-D		AVG	
	OS	OS*	OS	OS*	OS	OS*	OS	OS*	OS	OS*	OS	OS*	OS	OS*
Method w/o unknown classes in source domain (AlexNet)														
OSVM	59.6	59.1	57.1	55.0	14.3	5.9	44.1	39.3	13.0	4.5	62.5	59.2	40.6	37.1
MMD + OSVM	47.8	44.3	41.5	36.2	9.9	0.9	34.4	28.4	11.5	2.7	62.0	58.5	34.5	28.5
BP+OSVM	40.8	35.6	31.0	24.3	10.4	1.5	33.6	27.3	11.5	2.7	49.7	44.8	29.5	22.7
ATI- λ [27] + OSVM	72.0	-	65.3	-	66.4	-	82.2	-	71.6	-	92.7	-	75.0	-
OSDA+BP[34]	76.6	76.4	70.1	69.1	62.5	62.3	94.4	94.6	82.3	82.2	96.8	96.9	80.4	80.2
Ours w/o SCA	87.8	89.0	85.6	87.1	74.2	73.8	97.2	98.1	74.9	73.9	98.5	99.0	86.5	87.0
Ours w/o SCM	89.8	91.2	88.0	90.6	77.8	77.9	97.6	98.6	75.1	75.0	98.0	99.3	87.7	88.8
Ours	91.0	92.7	89.5	89.6	81.8	83.0	97.8	98.8	78.7	81.4	98.5	99.7	89.7	90.7
Method w/o unknown classes in source domain (VGGNet)														
OSVM	82.1	83.9	75.9	75.8	38.0	33.1	57.8	54.4	54.5	50.7	83.6	83.3	65.3	63.5
MMD + OSVM	84.4	85.8	75.6	75.7	41.3	35.9	61.9	58.7	50.1	45.6	84.3	83.4	66.3	64.2
BP+OSVM	83.1	84.7	76.3	76.1	41.6	36.5	61.1	57.7	53.7	49.9	82.9	82.0	66.4	64.5
OSDA+BP[34]	85.8	85.8	76.9	76.6	89.4	91.5	96.0	96.6	83.4	83.1	97.1	97.3	88.0	88.5
Ours	90.1	92.0	86.4	87.7	81.6	88.4	97.9	99.8	80.3	82.6	98.2	99.3	89.1	91.6

Table 2. Accuracy (%) of each method on scores of OS and OS*. A, D and W correspond to Amazon, DSLR and Webcam respectively. The ablation versions of our method w/o SCA and w/o SCM are also reported.

4.2. Results on Digit Dataset

In the three Digit datasets, numbers 0~4 are chosen as known classes. Samples from the known classes make up the source samples. In the target samples, numbers 5~9 are regarded as one unknown class. Besides the scores of OS and OS*, we also report the total accuracy for samples in target and the accuracy of unknown class, which are denoted as *ALL* and *UNK*, respectively.

As shown in Table 1, our method produces competitive results compared to other methods. Results of our method outperform the other methods in SVHN \rightarrow MNIST, MNIST \rightarrow USPS and the average scores. It is shown that our method is better at recognizing the unknown samples(5~9) while maintaining the performance of identifying the known classes. For SVHN \rightarrow MNIST, the semantic gap between them is large, as there may exist several digits in the images of SVHN. Thus the accuracies of SVHN \rightarrow MNIST are lower than the other two scenarios. Our method outperforms existing methods on the average scores. For instance, our approach achieves 87.9% in the average score of unknown class. This is 2.0% higher than OSDA+BP [34]. Learned features obtained by the trained model of the three scenarios are visualized in Fig 3. We observe that the distribution of unknown samples is more decentralized in SVHN \rightarrow MNIST because of the large divergence of the two domains. Compared with OSDA+BP [34], the proposed method could better centralize samples of the same known class and distinguish unknown samples from known ones.

4.3. Results on Office-31 Dataset

We compare our method with other works on *Office-31* dataset following [27]. There are 31 classes in this dataset,

and the first 10 classes in alphabetical order are selected as known classes. In [34], 21-31 classes are selected as unknown samples in the target domain, which only exist in the target domain. We evaluate the experiment on all the 6 scenario tasks: A \rightarrow D, A \rightarrow W, D \rightarrow A, D \rightarrow W, W \rightarrow A, W \rightarrow D. The improvement on some hard transfer tasks is encouraging to prove the effectiveness and value of the proposed method.

Results of *Office-31* are shown in Table 2. With features extracted with AlexNet, our method significantly outperforms the state-of-the-art methods except W \rightarrow A. Our approach achieves 89.7% in OS score based on AlexNet, overpassing the state-of-the-art by 2%. For OS*, our approach improves the results of state-of-the-art from 80.2% to 90.7%. Based on VGG, our method reaches 89.1% on OS score and 91.6 on OS* score, outperforming the state-of-the-art method by 1.1% and 3.1%, respectively. Especially, under the harder scenarios of A \rightarrow D, A \rightarrow W, D \rightarrow A, our method brings large improvement.

4.4. Ablation Study

For a straightforward understanding of the proposed method, we further evaluate each module via ablation experiments on Digit datasets. We alternately remove the SCA and SCM from our model. Results are reported in Table 1 above our final results. A decrease in performance is observed when removing SCA or SCM. Particularly, when the discriminative learning (w/o SCA) or contrastive learning (w/o SCM) is ablated, the accuracy of OS* or UNK or both of them will decrease significantly. We reconfirm the effect of each module based on AlexNet in the experiments of *Office-31*. Results in Table 2 also indicate the importance of learning the discriminative representation and contrastive

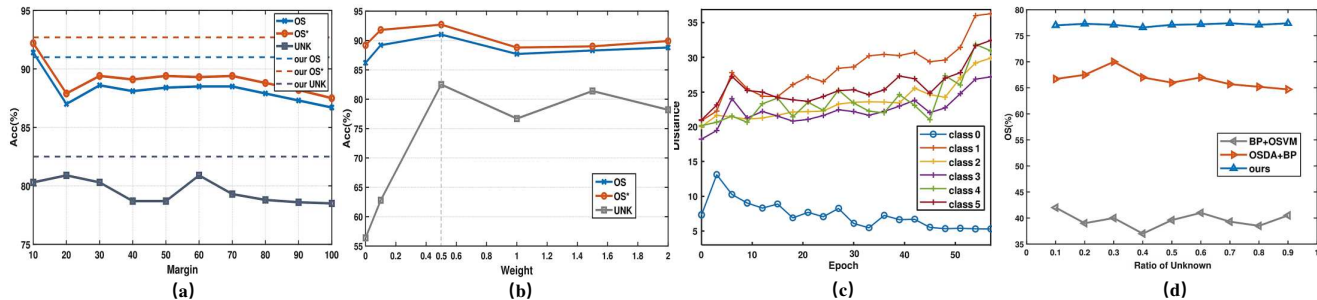


Figure 4. (a): A comparison between static margin and our adaptive margin. (b): A comparison of the behavior of our method with different re-weight ω in the contrastive loss. (c): Distances between centroid of class “backpack” labeled “0” in the target with centroids of class 0~5 in source . (d): Performance of the proposed method under different ratios of unknown samples.

mapping simultaneously. It indicates that the two modules take effect jointly. The discriminative representation helps to push away the unknown samples while the distraction of unknown samples assists the alignment of known categories.

Effect of adaptive margin. As the margin is designed for the contrastive loss as an adaptive one, we also observe the behavior of the model with static margin on $A \rightarrow D$. We choose constant value of margin $m_s \in \{10, 20, 30, \dots, 100\}$ for comparison. When m_s is equal to 0, the contrastive term in the objective is only to align all the target samples predicted as known with the corresponding centroid in source. When m_s is assigned with a large value, the model tends to penalize all the target samples predicted as unknown samples with large loss. According to results in Fig. 4(a), the accuracies of OS and OS* are trending downward when using a constant margin. The accuracy of UNK raises when the margin is 20 and 60. To further investigate the changes of distances between categories during the training of model, we visualize the centroid of known class “Backpack” in the target with the centroids of known classes in the source domain. The “Backpack” is labeled as 0. Results are shown in Fig. 4(c). With the training of model, the distance between the centroids of class 0 in the source and target domains is declining. This indicates that the distribution between the source and target domains are aligned for class 0. In the meanwhile, distances between centroids of target class 0 and the other classes in source domain improve with the increase of training epoch. In spite of the reduce of discrepancy between the two domains, the distances with the other classes in source are increasing. This demonstrates that it is improper to use a static margin to measure the energy for pulling apart unknown samples. With the alignment across domain and the separation between different class, the radius of the neighborhood of each class would change. This also explains the reason why the adaptive margin produces higher scores than the static margin.

Effect of re-weighting the contrastive loss. There is another hyper-parameter ω which re-weights the dis-

tances in \mathcal{L}_{con} . We conduct experiments with $\omega \in \{0, 0.1, 0.5, 1, 1.5, 2\}$. As shown in Fig. 4(b), our approach achieves the best results when ω is equal to 0.5. This parameter smooths the weight of calculated with cosine similarity. When ω is 0, the re-weighting term is equal to 1, which means the contrastive loss is calculated without the re-weighting. It also reveals that the effectiveness of the re-weighting term could help the model to better measure the distances between unknown samples and centroids in source.

Effect of the ratio of unknown samples. In this section, we aim to investigate the robustness of our method under different ratios of unknown samples in the target data. In Fig. 4(d), we compare our results with method BP+OSVM and OSDA+BP under ratio $\in (0, 1)$. Unknown samples are randomly sampled according to the ratio. It can be seen that the accuracy of our model fluctuates little and is always above the baseline methods, which implies the robustness of our method.

5. Conclusion

In this paper, we focus on the open set domain adaptation setting where only source domain annotations are available and the unknown class exists in the target. To better identify the unknown and meanwhile diminish the domain shift, we take the semantic margin of open set data into account through semantic categorical alignment and semantic contrastive mapping, aiming to make the known classes more separable and push the unknown class away from the decision boundary, respectively. Empirically, we demonstrate that our method is comparable to or even better than the state-of-the-art methods on representative open set benchmarks, *i.e.* Digits and Office-31. The effectiveness of each component of our method is also verified. Our method implies that explicitly taking the semantic margin of open set data into account is beneficial. And it is promising to make more exploration in this direction in the future.

References

- [1] Mahsa Baktashmotlagh, Masoud Faraki, Tom Drummond, and Mathieu Salzmann. Learning factorized representations for open-set domain adaptation. In *ICLR*, 2019. 1, 3
- [2] Konstantinos Bousmalis, Nathan Silberman, David Dohan, Dumitru Erhan, and Dilip Krishnan. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *CVPR*, 2017. 2
- [3] Binghui Chen, Weihong Deng, and Haifeng Shen. Virtual class enhanced discriminative embedding learning. In *Advances in Neural Information Processing Systems*, pages 1942–1952, 2018. 2
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 6
- [5] Hehe Fan, Xiaojun Chang, De Cheng, Yi Yang, Dong Xu, and Alexander G. Hauptmann. Complex event detection by identifying reliable shots from untrimmed videos. In *ICCV*, 2017. 2
- [6] Hehe Fan, Liang Zheng, Chenggang Yan, and Yi Yang. Unsupervised person re-identification: Clustering and fine-tuning. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 14(4):83, 2018. 2
- [7] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, 2015. 2, 4, 6
- [8] Muhammad Ghifary, W Bastiaan Kleijn, Mengjie Zhang, David Balduzzi, and Wen Li. Deep reconstruction-classification networks for unsupervised domain adaptation. In *ECCV*, 2016. 2
- [9] Arthur Gretton, Karsten Borgwardt, Malte Rasch, Bernhard Schölkopf, and Alex J Smola. A kernel method for the two-sample-problem. In *NIPS*, pages 513–520, 2007. 6
- [10] Qingji Guan and Yaping Huang. Multi-label chest x-ray image classification via category-wise residual attention learning. *Pattern Recognition Letters*, 2018. 2
- [11] Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *CVPR*, 2006. 5
- [12] Philip Haeusser, Thomas Frerix, Alexander Mordvintsev, and Daniel Cremers. Associative domain adaptation. In *ICCV*, 2017. 2
- [13] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017. 1
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 1
- [15] Lalit P Jain, Walter J Scheirer, and Terrance E Boult. Multi-class open set recognition using probability of inclusion. In *ECCV*, 2014. 5
- [16] Guoliang Kang, Lu Jiang, Yi Yang, and Alexander G Hauptmann. Contrastive adaptation network for unsupervised domain adaptation. In *CVPR*, 2019. 2
- [17] Guoliang Kang, Liang Zheng, Yan Yan, and Yi Yang. Deep adversarial attention alignment for unsupervised domain adaptation: the benefit of target expectation maximization. In *ECCV*, 2018. 2
- [18] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 6
- [19] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1097–1105, 2012. 1, 6
- [20] Yann LeCun, Léon Bottou, Yoshua Bengio, Patrick Haffner, et al. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 6
- [21] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. 1
- [22] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael I Jordan. Learning transferable features with deep adaptation networks. In *ICML*, 2015. 2, 6
- [23] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Unsupervised domain adaptation with residual transfer networks. In *Advances in Neural Information Processing Systems*, pages 136–144, 2016. 2
- [24] Yawei Luo, Ping Liu, Tao Guan, Junqing Yu, and Yi Yang. Significance-aware information bottleneck for domain adaptive semantic segmentation. In *ICCV*, 2019. 2
- [25] Yawei Luo, Liang Zheng, Tao Guan, Junqing Yu, and Yi Yang. Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In *CVPR*, 2019. 2
- [26] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bisacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. In *NIPS*, 2011. 6
- [27] Pau Panareda Busto and Juergen Gall. Open set domain adaptation. In *ICCV*, 2017. 1, 2, 6, 7
- [28] Ce Qi and Fei Su. Contrastive-center loss for deep neural networks. In *2017 IEEE International Conference on Image Processing (ICIP)*, pages 2851–2855, 2017. 2, 4
- [29] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, 2016. 1
- [30] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *ECCV*, 2010. 5, 6
- [31] Kuniaki Saito, Yoshitaka Ushiku, and Tatsuya Harada. Asymmetric tri-training for unsupervised domain adaptation. In *ICML*, 2017. 2
- [32] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. Adversarial dropout regularization. In *ICLR*, 2018. 2
- [33] Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In *CVPR*, 2018. 2
- [34] Kuniaki Saito, Shohei Yamamoto, Yoshitaka Ushiku, and Tatsuya Harada. Open set domain adaptation by backpropagation. In *ECCV*, 2018. 1, 3, 4, 5, 6, 7
- [35] Ozan Sener, Hyun Oh Song, Ashutosh Saxena, and Silvio Savarese. Learning transferrable representations for unsupervised domain adaptation. In *Advances in Neural Information Processing Systems*, pages 2110–2118, 2016. 2

- [36] Rui Shu, Hung H Bui, Hirokazu Narui, and Stefano Ermon. A dirt-t approach to unsupervised domain adaptation. *ICLR*, 2018. [2](#)
- [37] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. [1](#), [6](#)
- [38] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *CVPR*, 2015. [1](#)
- [39] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *CVPR*, 2017. [2](#)
- [40] 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](#)
- [41] Zhedong Zheng, Liang Zheng, Michael Garrett, Yi Yang, and Yi-Dong Shen. Dual-path convolutional image-text embedding with instance loss. *arXiv preprint arXiv:1711.05535*, 2017. [2](#)
- [42] Zhun Zhong, Liang Zheng, Zhiming Luo, Shaozi Li, and Yi Yang. Invariance matters: Exemplar memory for domain adaptive person re-identification. In *CVPR*, 2019. [2](#)
- [43] Zhun Zhong, Liang Zheng, Zhedong Zheng, Shaozi Li, and Yi Yang. Camstyle: A novel data augmentation method for person re-identification. *IEEE Transactions on Image Processing*, 28:1176–1190, 2019. [1](#)