

Open Compound Domain Adaptation

Ziwei Liu^{1*} Zhongqi Miao^{2*} Xingang Pan¹ Xiaohang Zhan¹
Dahua Lin¹ Stella X. Yu² Boqing Gong³

¹ The Chinese University of Hong Kong ² UC Berkeley / ICSI ³ Google Inc.

<https://liuziwei7.github.io/projects/CompoundDomain.html>

Abstract

A typical domain adaptation approach is to adapt models trained on the annotated data in a source domain (e.g., sunny weather) for achieving high performance on the test data in a target domain (e.g., rainy weather). Whether the target contains a single homogeneous domain or multiple heterogeneous domains, existing works always assume that there exist clear distinctions between the domains, which is often not true in practice (e.g., changes in weather). We study an open compound domain adaptation (OCDA) problem, in which the target is a compound of multiple homogeneous domains without domain labels, reflecting realistic data collection from mixed and novel situations. We propose a new approach based on two technical insights into OCDA: 1) a curriculum domain adaptation strategy to bootstrap generalization across domains in a data-driven self-organizing fashion and 2) a memory module to increase the model’s agility towards novel domains. Our experiments on digit classification, facial expression recognition, semantic segmentation, and reinforcement learning demonstrate the effectiveness of our approach.

1. Introduction

Supervised learning can achieve competitive performance for a visual task when the test data is drawn from the same underlying distribution as the training data. This assumption, unfortunately, often does not hold in reality, e.g., the test data may contain the same class of objects as the training data but different backgrounds, poses, and appearances [39, 44].

The goal of domain adaptation is to adapt the model learned on the training data to the test data of a different distribution [39, 32, 11]. Such a distributional gap is often formulated as a shift between discrete concepts of well defined data domains, e.g., images collected in sunny weather versus those in rainy weather. Though domain generalization [22, 20] and latent domain adaptation [15,

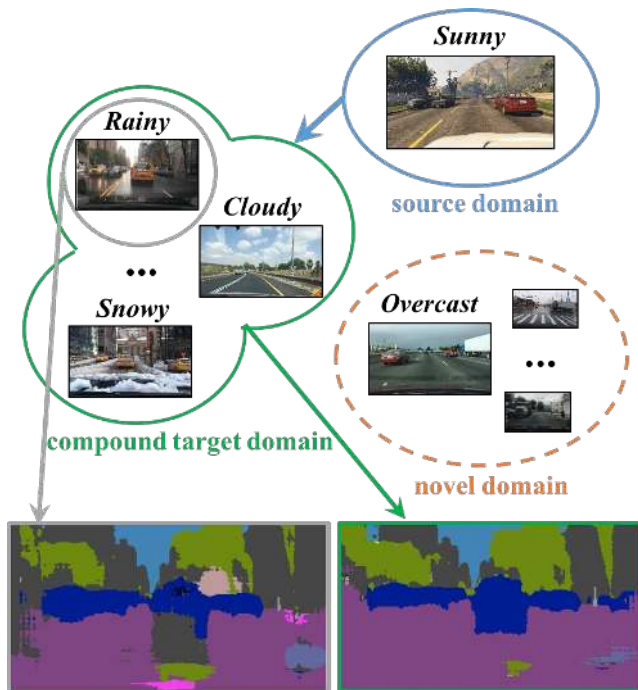


Figure 1: **Open compound domain adaptation.** Unlike existing domain adaptation which assumes clear distinctions between discrete domains (cf. the examples in gray frames), our compound target domain is a combination of multiple traditionally homogeneous domains without any domain labels. We also allow novel domains to show up at the inference time.

[10] have attempted to tackle complex target domains, most existing works usually assume that there is a known clear distinction between domains [11, 7, 46, 28, 40].

Such a known and clear distinction between domains is hard to define in practice, e.g., test images could be collected in mixed, continually varying, and sometimes never seen weather conditions. With numerous factors jointly contributing to data variance, it becomes implausible to separate data into discrete domains.

We propose to study *open compound domain adaptation* (OCDA), a continuous and more realistic setting for domain

*Equal contribution.

Table 1: **Comparison of domain adaptation settings.** Domain Labels tell to which domain each instance belongs. Open Classes refer to novel classes showing up during testing but not training. Open Domains are the domains of which no instances are seen during training.

Domain Adaptation Setting	# Target Domains	Domain Labels	Open Classes	Open Domains
Unsupervised Domain Adaptation	single	known	×	×
Multi-Target Domain Adaptation	multiple	known	×	×
Open/Partial Set Domain Adaptation	single	known	✓	×
Open Compound Domain Adaptation	multiple	unknown	×	✓

adaptation (cf. Figure 1 and Table 1). The task is to learn a model from labeled *source domain* data and adapt it to unlabeled *compound target domain* data which could differ from the source domain on various factors. Our target domain can be regarded as a combination of multiple traditionally homogeneous domains where each is distinctive on one or two major factors, and yet none of the domain labels are given. For example, the five well-known datasets on digits recognition (SVHN [31], MNIST [19], MNIST-M [6], USPS [18], and SynNum [6]) mainly differ from each other by the backgrounds and text fonts. It is not necessarily the best practice, and not feasible under some scenarios, to consider them as distinct domains. Instead, our compound target domain pools them together. Furthermore, at the inference stage, OCDA tests the model not only in the compound target domain but also in open domains that have previously unseen during training.

In our OCDA setting, the target domain no longer has a predominantly uni-modal distribution, posing challenges to existing domain adaptation methods. We propose a novel approach based on two technical insights into OCDA: 1) a curriculum domain adaptation strategy to bootstrap generalization across domain distinction in a data-driven self-organizing fashion and 2) a memory module to increase the model’s agility towards novel domains.

Unlike existing curriculum adaptation methods [54, 5, 27, 23, 56, 55] that rely on some holistic measure of instance difficulty, we schedule the learning of unlabeled instances in the compound target domain according to their *individual gaps* to the labeled source domain, so that we solve an incrementally harder domain adaptation problem till we cover the entire target domain.

Specifically, we first train a neural network to 1) discriminate between classes in the labeled source domain and to 2) capture domain invariance from the easy target instances which differ the least from labeled source domain data. Once the network can no longer differentiate between the source domain and the easy target domain data, we feed the network harder target instances, which are further away from the source domain. The network learns to remain discriminative to the classification task and yet grow more robust to the entire compound target domain.

Technically, we must address the challenge of characterizing each instance’s gap to the source domain. We

first extract domain-specific feature representations from the data and then rank the target instances according to their distances to the source domain in that feature space, assuming that such features do not contribute to and even distract the network from learning discriminative features for classification. We use a class-confusion loss to distill the domain-specific factors and formulate it as a conventional cross-entropy loss with a randomized class label twist.

Our second technical insight is to prepare our model for open domains during inference with a memory module that effectively augments the representations of an input for classification. Intuitively, if the input is close enough to the source domain, the feature extracted from itself can most likely already result in accurate classification. Otherwise, the input-activated memory features can step in and play a more important role. Consequently, this memory-augmented network is more agile at handling open domains than its vanilla counterpart.

To summarize, we make the following contributions. **1)** We extend the traditional discrete domain adaptation to OCDA, a more realistic continuous domain adaptation setting. **2)** We develop an OCDA solution with two key technical insights: instance-specific curriculum domain adaptation for handling the target of mixed domains and memory augmented features for handling open domains. **3)** We design several benchmarks on classification, recognition, segmentation, and reinforcement learning, and conduct comprehensive experiments to evaluate our approach under the OCDA setting.

2. Related Works

We review literature according to Table 1.

Unsupervised Domain Adaptation. The goal is to retain recognition accuracies in new domains without ground truth annotations [39, 44, 47, 36]. Representative techniques include latent distribution alignment [11], back-propagation [6], gradient reversal [7], adversarial discrimination [46], joint maximum mean discrepancy [28], cycle consistency [16] and maximum classifier discrepancy [40]. While their results are promising, this traditional domain adaptation setting focuses on “one source domain, one target domain”, and cannot deal with more complicated scenarios where multiple target domains are present.

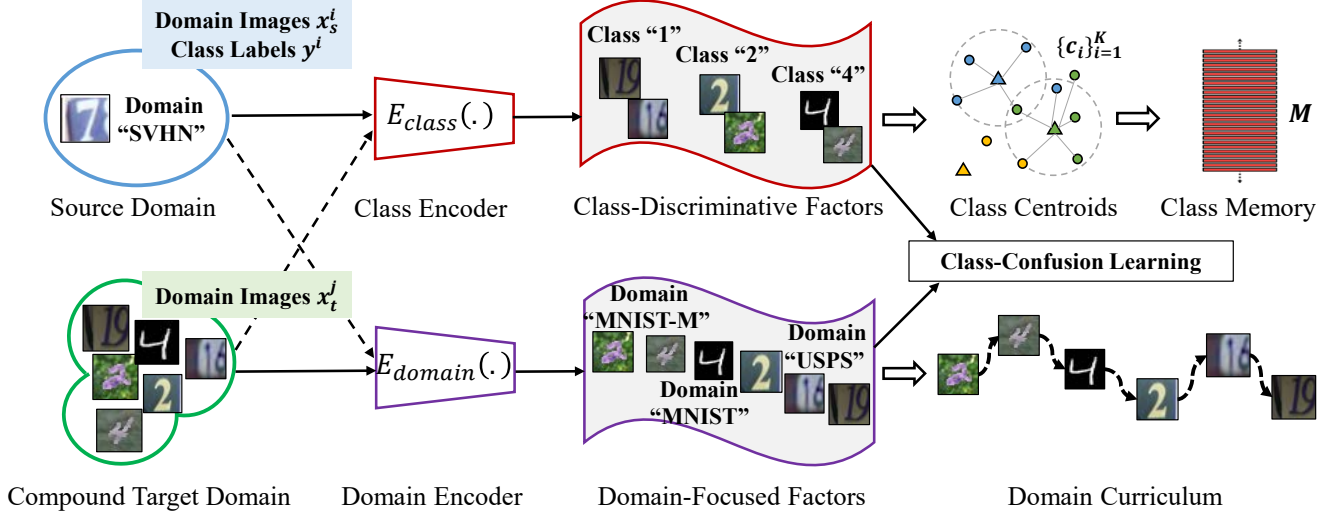


Figure 2: **Overview of disentangling domain characteristics and curriculum domain adaptation.** We separate characteristics specific to domains from those discriminative between classes. It is achieved by a class-confusion algorithm in an unsupervised manner. The teased out domain feature is used to construct a curriculum for domain-robust learning.

Latent & Multi-Target Domain Adaptation. The goal is to extend unsupervised domain adaptation to latent [15, 49, 30] or multiple [10, 8, 52] or continuous [2, 12, 29, 48] target domains, when only the source domain has class labels. These methods usually assume clear domain distinction or require domain labels (*e.g.* test instance i belongs to the target domain j), but this assumption rarely holds in the real-world scenario. Here we take one step further towards compound domain adaptation, where both category labels and domain labels in the test set are unavailable.

Open/Partial Set Domain Adaptation. Another route of research aims to tackle the category sharing/unsharing issues between source and target domain, namely open set [35, 41] and partial set [53, 3] domain adaptation. They assume that the target domain contains either 1) new categories that don't appear in source domain; or 2) only a subset of categories that appear in source domain. Both settings concern the "openness" of categories. Instead, here we investigate the "openness" of domains, *i.e.* there are novel domains existing that are absent in the training phase.

Domain Generalized/Agnostic Learning. Domain generalization [50, 21, 20] and domain agnostic learning [37, 4] aim to learn universal representations that can be applied in a domain-invariant manner. Since these methods focus on learning semantic representations that are invariant to the domain shift, they largely neglect the latent structures inside the target domains. In this work, we explicitly model the latent structures inside the compound target domain by leveraging the learned domain-focused factors for curriculum scheduling and dynamic adaptation.

3. Our Approach to OCDA

Figures 2 and 3 present our overall workflows. There are three major components: 1) disentangling domain characteristics with only class labels in the source domain, 2) scheduling data for curriculum domain adaptation, and 3) a memory module for handling new domains.

3.1. Disentangling Domain Characteristics

We separate characteristics specific to domains from those discriminative between classes. They allow us to construct a curriculum for increment domain adaptation.

We first train a neural network classifier using the labeled source domain data $\{x^i, y^i\}_i$. Let $E_{class}(\cdot)$ denote the encoder up to the second-to-the-last layer and $\Phi(E_{class}(\cdot))$ the classifier. The encoder captures primarily the class-discriminative representation of the data.

We assume that all the factors not covered by this class-discriminative encoder reflect domain characteristics. They can be extracted by another encoder $E_{domain}(\cdot)$ that satisfies two properties: **1) Completeness:** $Decoder(E_{class}(x), E_{domain}(x)) \approx x$, *i.e.*, the outputs of the two encoders shall provide sufficient information for a decoder to reconstruct the input, and **2) Orthogonality:** the domain encoder $E_{domain}(x)$ shall have little mutual information with the class encoder $E_{class}(x)$. We leave the algorithmic details for meeting the first property to the appendices as they are not our novelty.

For the orthogonality between $E_{domain}(x)$ and $E_{class}(x)$, we propose a **class-confusion algorithm**, which

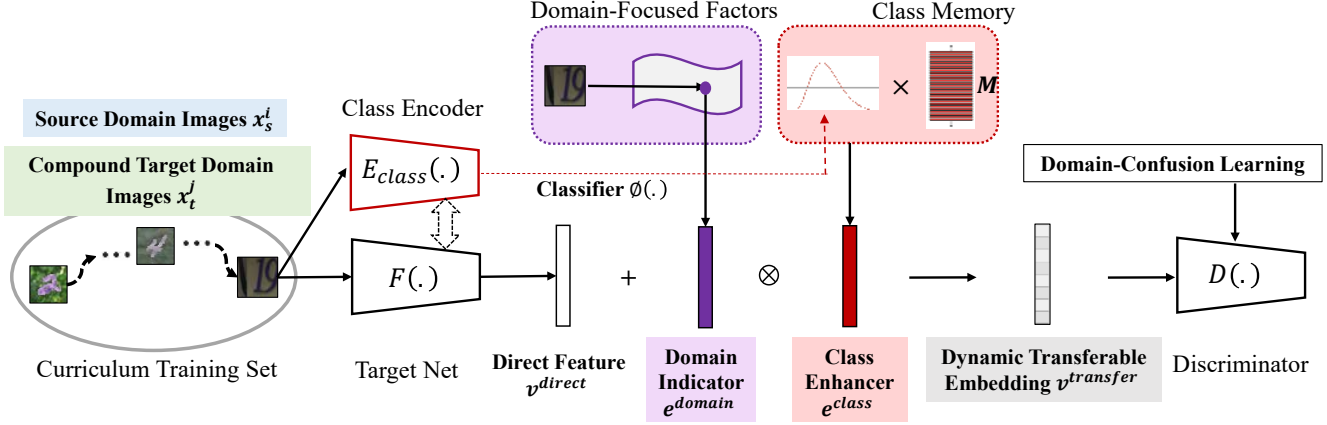


Figure 3: **Overview of the memory-enhanced deep neural network.** We enhance our network with a memory module that facilitates knowledge transfer from the source domain to target domain instances, so that the network can dynamically balance the input information and the memory-transferred knowledge for more agility towards previously unseen domains.

alternates between the two sub-problems below:

$$\min_{E_{domain}} - \sum_i z_{random}^i \log D(E_{domain}(x^i)), \quad (1)$$

$$\min_D - \sum_i y^i \log D(E_{domain}(x^i)), \quad (2)$$

where superscript i is the instance index, and $D(\cdot)$ is a discriminator the domain-encoder $E_{domain}(\cdot)$ tries to confuse. We first train the discriminator $D(\cdot)$ with the labeled data in the source domain. For the data in the target domain, we assign them pseudo-labels by the classifier $\Phi(E_{class}(\cdot))$ we have trained earlier. The learned domain encoder $E_{domain}(\cdot)$ is class-confusing due to z_{random}^i , a random label uniformly chosen in the label space. As the classifier $D(\cdot)$ is trained, the first sub-problem essentially learns the domain-encoder such that it classifies the input x^i into a random class z_{random}^i . Algorithm 1 details our domain disentanglement process.

Figure 4 (a) and (b) visualize the examples embedded by the class encoder $E_{class}(\cdot)$ and domain encoder $E_{domain}(\cdot)$, respectively. The class encoder places instances in the same class in a cluster, while the domain encoder places instances according to their common appearances, regardless of their classes.

3.2. Curriculum Domain Adaptation

We rank all the instances in the compound target domain according to their distances to the source domain, to be used for curriculum domain adaptation [54]. We compute the *domain gap* between a target instance x_t and the source domain $\{x_s^m\}$ as their mean distance in the domain feature space: $\text{mean}_m(\|E_{domain}(x_t) - E_{domain}(x_s^m)\|_2)$.

We train the network in stages, a few epochs at a time, gradually recruiting more instances that are increasingly far from the source domain. At each stage of the curriculum

Algorithm 1 Domain Disentanglement.

Input: The class encoder $E_{class}(\cdot)$ and classifier Φ have been trained using source-domain data, $Decoder(\cdot)$: the decoder, C : the number of classes, γ : a constant.

for k iterations **do**

 Sample mini-batch $\{x^i\}$.

 Compute pseudo labels $y_{pseudo}^i \leftarrow \Phi(E_{class}(x^i))$.

 Update the discriminator D .

 Prepare random labels $z_{random}^i \sim \text{uniform}\{0, 1, \dots, C-1\}$.

 Compute adversarial loss: $L_{adv} \leftarrow \sum_i -z_{random}^i \log(D(E_{domain}(x^i)))$.

 Compute reconstruction loss: $L_{rec} \leftarrow \sum_i \|Decoder(E_{class}(x^i), E_{domain}(x^i)) - x^i\|_2$.

 Update the domain encoder E_{domain} with: $\nabla_{\theta_{E_{domain}}}(L_{adv} + \gamma L_{rec})$.

end for

learning, we minimize two losses: One is the cross-entropy loss defined over the labeled source domain, and the other is the domain-confusion loss [46] computed between the source domain and the currently covered target instances. Figure 4 (c) illustrates a curriculum in our experiments.

3.3. Memory Module for Open Domains

Existing domain adaptation methods often use the features v_{direct} extracted directly from the input for adaptation. When the input comes from a new domain that significantly differs from the seen domains during training, this representation becomes inadequate and could fool the classifier. We propose a memory module to enhance our model; It allows knowledge transfer from the source domain so that the network can dynamically balance the input-conveyed information and the memory-transferred knowledge for more classification agility towards previously unseen domains.

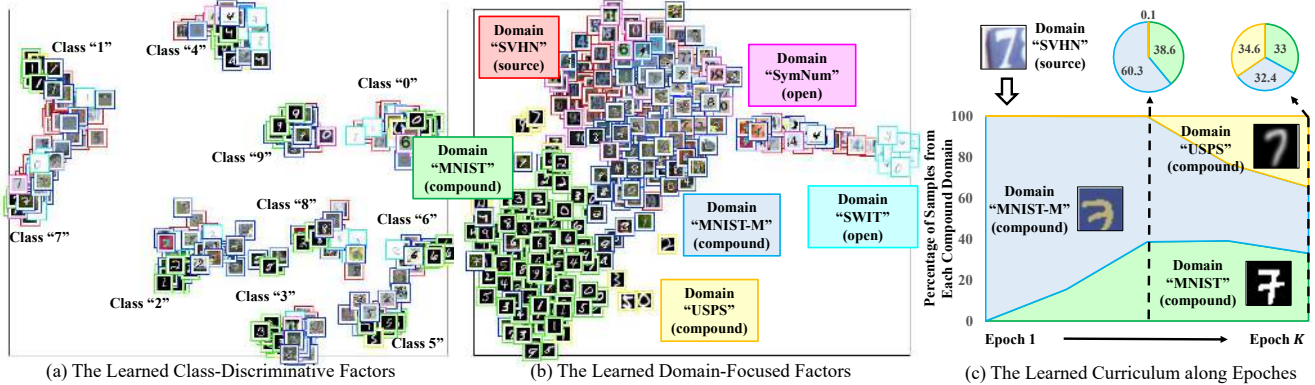


Figure 4: **t-SNE Visualization** of our (a) class-discriminative features, (b) domain features, and (c) curriculum. Our framework disentangles the mixed-domain data into class-discriminative factors and domain-focused factors. We use the domain-focused factors to construct a learning curriculum for domain adaptation.

Class Memory M . We design a memory module M to store the class information from the source domain. Inspired by [43, 34, 26] on prototype analysis, we also use class centroids $\{c_k\}_{k=1}^K$ to construct our memory M , where K is the number of object classes.

Enhancer $v_{enhance}$. For each input instance, we build an enhancer to augment its direct representation v_{direct} with knowledge in the memory about the source domain: $v_{enhance} = (\Psi(v_{direct}))^T M = \sum_{k=1}^K \psi_k c_k$, where $\Psi(\cdot)$ is a softmax function. We add this enhancer to the direct representation v_{direct} , weighted by a domain indicator.

Domain Indicator e_{domain} . With open domains, the network must dynamically calibrate how much knowledge to transfer from the source domain and how much to rely on the direct representation v_{direct} of the input. Intuitively, the larger domain gap between an input x and the source domain, the more weight on the memory feature. We design a domain indicator for such domain awareness: $e_{domain} = T(E_{domain}(x))$, where $T(\cdot)$ is a lightweight network with the \tanh activation functions and $E_{domain}(\cdot)$ is the domain encoder we have learned earlier.

Source-Enhanced Representation $v_{transfer}$. Our final representation of the input is a dynamically balanced version between the direct image feature and the memory enhanced feature:

$$v_{transfer} = v_{direct} + e_{domain} \otimes v_{enhance}, \quad (3)$$

which transfers class-discriminative knowledge from the labeled source domain to the input in a domain-aware manner. Operator \otimes is element-wise multiplication. Adopting cosine classifiers [25, 9], we ℓ_2 -normalize this representation before sending it to the softmax classification layer. All of these choices help cope with domain mismatch when the input is significantly different from the source domain.

4. Experiments

Datasets. To facilitate a comprehensive evaluation on various tasks (*i.e.*, classification, segmentation, and navigation), we carefully design four open compound domain adaptation (OCDA) benchmarks: C-Digits, C-Faces, C-Driving, and C-Mazes, respectively.

1. **C-Digits:** This benchmark aims to evaluate the classification adaptation ability under different appearances and backgrounds. It is built upon five classic digits datasets (SVHN [31], MNIST [19], MNIST-M [6], USPS [18] and SynNum [6]), where SVHN is used as the source domain, MNIST, MNIST-M, and USPS are mixed as the compound target domain, and SynNum is the open domain. We employ SWIT [1] as an additional open domain for further analysis.
2. **C-Faces:** This benchmark aims to evaluate the classification adaptation ability under different camera poses. It is built upon the Multi-PIE dataset [13], where C05 (frontal view) is used as source domain, C08-C14 (left side view) are combined as the compound target domain, and C19 (right side view) is kept out as the open domain.
3. **C-Driving:** This benchmark aims to evaluate the segmentation adaptation ability from simulation to different real driving scenarios. The GTA-5 [38] dataset is adopted as the source domain, while the BDD100K dataset [51] (with different scenarios including “rainy”, “snowy”, “cloudy”, and “overcast”) is taken for the compound and open domains.
4. **C-Mazes:** This benchmark aims to evaluate the navigation adaptation ability under different environmental appearances. It is built upon the GridWorld environment [17], where mazes with different colors are used

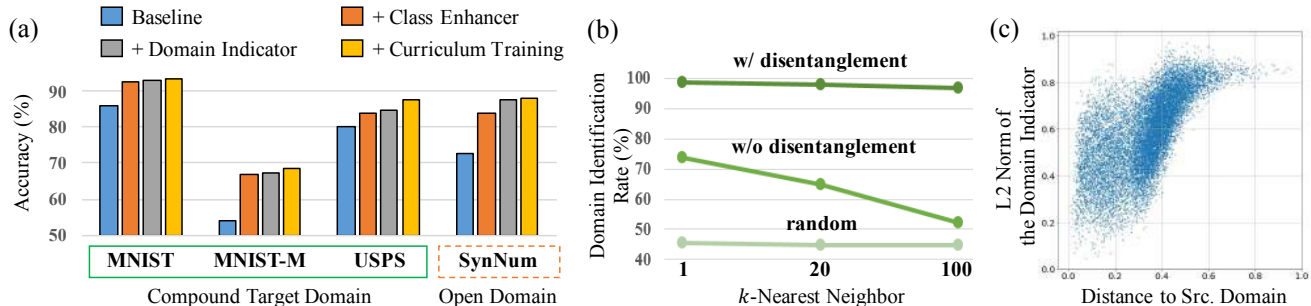



Figure 5: **Results of ablation studies** about (a) the memory-enhanced embeddings and curriculum domain adaptation, (b) the domain-focused factors disentanglement, and (c) the memory-induced domain indicator vs. gaps to the source.

Table 2: **Performance on the C-Digits benchmark.** The methods in gray are especially designed for multi-target domain adaptation. [†]MTDA uses domain labels, while [‡]BTDA and DADA use the open domain images during training.



Src. Domain SVHN →	Compound Domains (C)			Open (O)	Avg.	
	MNIST	MNIST-M	USPS	SynNum	C	C+O
ADDA [46]	80.1±0.4	56.8±0.7	64.8±0.3	72.5±1.2	67.2±0.5	68.6±0.7
JAN [28]	65.1±0.1	43.0±0.1	63.5±0.2	85.6±0.0	57.2±0.1	64.3±0.1
MCD [40]	69.6±1.4	48.6±0.5	70.6±0.2	89.8±2.9	62.9±1.0	69.9±1.3
MTDA [†] [8]	84.6±0.3	65.3±0.2	70.0±0.2	-	73.3±0.2	-
BTDA [‡] [4]	85.2±1.6	65.7±1.3	74.3±0.9	84.4±2.2	75.1±1.3	77.4±1.5
DADA [‡] [37]	-	-	-	-	-	80.1±0.4
Ours	90.9±0.2	65.7±0.5	83.4±0.3	88.2±0.8	80.0±0.3	82.1±0.5

as the source and open domains. Since reinforcement learning often assumes no prior access to the environments, there are no compound target domains here.

Network Architectures. To make a fair comparison with previous works [46, 8, 37], the modified LeNet-5 [19] and ResNet-18 [14] are used as the backbone networks for C-Digits and C-Faces, respectively. Following [45, 56, 33], a pre-trained VGG-16 [42] is the backbone network for C-Driving. We additionally test our approach on reinforcement learning using ResNet-18 following [17].

Evaluation Metrics. The C-digits performance is measured by the digit classification accuracy, and the C-Faces performance is measured by the facial expression classification accuracy. The C-Driving performance is measured by the standard mIOU, and the C-Mazes performance is measured by the average successful rate in 300 steps. We evaluate the performance of each method with five runs and report both the mean and standard deviation. Moreover, we report both results of individual domains and the averaged results for a comprehensive analysis.

Comparison Methods. For classification tasks, we choose for comparison state-of-the-art methods in both conventional unsupervised domain adaptation (ADDA [46], JAN [28], MCD [40]) and the recent multi-target domain adaptation methods (MTDA [8], BTDA [4], DADA [37]). Since MTDA [8], BTDA [4] and DADA [37] are the most related to our work, we directly contrast our results to


the numbers reported in their papers. For the segmentation task, we compare with three state-of-the-art methods, AdaptSeg [45], CBST [56], IBN-Net [33] and PyCDA [24]. For the reinforcement learning task, we benchmark with MTL, MLP [17] and SynPo [17], a representative work for adaptation across environments. We apply these methods to the same backbone networks as ours for a fair comparison.

4.1. Ablation Study

Effectiveness of the Domain-Focused Factors Disentanglement. Here we verify that the domain-focused factors disentanglement helps discover the latent structures in the compound target domain. It is probed by the domain identification rate within the k -nearest neighbors found by different encodings. Figure 5 (b) shows that features produced by our disentanglement have a much higher identification rate ($\sim 95\%$) than the counterparts without disentanglement ($\sim 65\%$).

Effectiveness of the Curriculum Domain Adaptation. Figure 5 (a) also reveals that, in the compound domain, the curriculum training contributes to the performance on USPS more than MNIST and MNIST-M. On the other hand, we can observe from Figure 4 and Table 2 that USPS is the furthest target domain from the source domain SVHN. It implies that curriculum domain adaptation makes it easy to adapt to the distant target domains through an easy-to-hard adaptation schedule.

Table 3: **Performance on the C-Faces benchmark.** The methods in gray are especially designed for multi-target domain adaptation. [†]MTDA uses domain labels during training.



Src. Domain C05 →	Compound Domains (C)				Open (O)	Avg.	
	C08	C09	C13	C14	C19	C	C+O
ADDA [46]	46.9±0.2	36.4±0.5	39.1±0.3	65.4±0.4	71.8±0.8	47.0±0.4	51.9±0.4
JAN [28]	63.5±0.3	40.6±1.0	83.5±0.4	92.0±0.8	52.5±1.5	69.7±0.6	66.2±0.8
MCD [40]	50.4±0.5	45.8±0.2	77.8±0.1	88.0±0.1	60.4±0.9	65.7±0.2	64.6±0.4
MTDA [†] [8]	49.0±0.2	48.2±0.1	53.1±0.2	84.3±0.1	-	58.7±0.2	-
Ours	73.3±0.2	55.1±0.4	84.1±0.1	88.9±0.3	72.7±0.6	75.4±0.3	74.8±0.3

Table 4: **Performance on the C-Driving (left) and C-Mazes benchmarks (right).** “SynPo+Aug.” indicates that we equip SynPo with proper color augmentation/randomization during training. Visual illustrations of both datasets are in Figure 6.

Source GTA-5 →	Compound (C)			Open (O)	Avg.		Source M0 →	Open(O)				Avg.
	Rainy	Snowy	Cloudy	Overcast	C	C+O		M1	M2	M3	M4	O
Source Only	16.2	18.0	20.9	21.2	18.9	19.1	Source Only	0±0	0±0	0±0	0±0	0±0
AdaptSeg [45]	20.2	21.2	23.8	25.1	22.1	22.5	MTL	0±0	30±5	75±0	65±5	42.5±2.5
CBST [56]	21.3	20.6	23.9	24.7	22.2	22.6	MLP [17]	5±5	45±10	75±5	80±10	51.2±7.5
IBN-Net [33]	20.6	21.9	26.1	25.5	22.8	23.5	SynPo [17]	5±5	30±20	80±5	30±5	36.3±8.8
PyCDA [24]	21.7	22.3	25.9	25.4	23.3	23.8	SynPo+Aug.	0±5	40±10	95±5	45±5	45.0±6.3
Ours	22.0	22.9	27.0	27.9	24.5	25.0	Ours	80±2.5	75±10	85±5	90±5	82.5±5.6

Effectiveness of Memory-Enhanced Representations.

Recall that the memory-enhanced representations consist of two main components: the enhancer coming from the memory and the domain indicator. From Figure 5 (a), we observe that the class enhancer leads to large improvements on all target domains. It is because the enhancer from the memory transfers useful semantic concepts to the input of any domain. Another observation is that the domain indicator is the most effective on the open domain (“SynNum”), because it helps dynamically calibrate the representations by leveraging domain relations (Figure 5 (c)).

4.2. Comparison Results

C-Digits. Table 2 shows the comparison performances of different methods. We have the following observations. Firstly, ADDA [46] and JAN [28] boost the performance on the compound domain by enforcing global distribution alignment. However, they also sacrifice the performance on the open domain since there is no built-in mechanism for handling any new domains, “overfitting” the model to the seen domains. Secondly, MCD [40] improves the results on the open domain, but its accuracy degrades on the compound target domain. Maximizing the classifier discrepancy increases the robustness to the open domain; however, it also fails to capture the fine-grained latent structure in the compound target domain. Lastly, compared to other multi-target domain adaptation methods (MTDA [8] and DADA [37]), our approach discovers domain structures and performs domain-aware knowledge transfer, achieving substantial advantages on all the test domains.

C-Faces. Similar observations can be made on the C-Faces benchmark as shown in Table 3. Since face representations are inherently hierarchical, JAN [28] demonstrates com-

petitive results on C14 due to its layer-wise transferring strategy. Under the domain shift with different camera poses, our approach still consistently outperforms other alternatives for both the compound and open domains.

C-Driving. We compare with the state-of-the-art semantic segmentation adaptation methods such as AdaptSeg [45], CBST [56], and IBN-Net [33]. All methods are tested under real-world driving scenarios in the BDD100K dataset [51]. We can see that our approach has clear advantages on both the compound domain (1.1% gains) and the open domain (2.4% gains) as shown in Table 4 (left). We show detailed per-class accuracies in the appendices. The qualitative comparisons are shown in Figure 6 (a).

C-Mazes. To directly compare with SynPo [17], we also evaluate on the GridWorld environments they provided. The task in this benchmark is to learn navigation policies that can successfully collect all the treasures in the given mazes. Existing reinforcement learning methods suffer from environmental changes, which we simulate as the appearances of the mazes here. The final results are listed in Table 4 (right). Our approach transfers visual knowledge among navigation experiences and achieves more than 30% improvements over the prior arts.

4.3. Further Analysis

Robustness to the Complexity of the Compound Target Domain.

We control the complexity of the compound target domain by varying the number of traditional target domains / datasets in it. Here we gradually increase constituting domains from a single target domain (*i.e.*, MNIST) to two, and eventually three (*i.e.*, MNIST + MNIST-M + USPS). From Figure 7 (a), we observe that as the number of datasets increase, our approach only undergoes a moderate

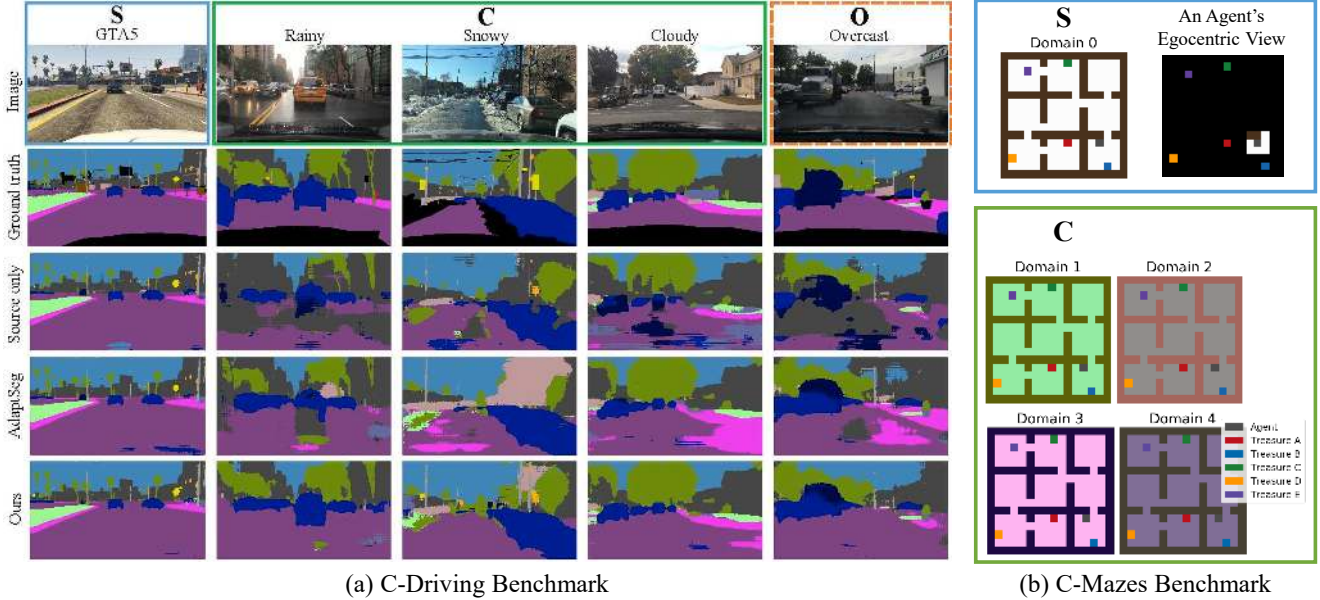


Figure 6: (a) **Qualitative results comparison** of semantic segmentation on the source domain (S), the compound target domain (C), and the open domain (O). (b) **Illustrations** of the 5 different domains in the C-Mazes benchmark. Our approach consistently outperforms existing domain adaptation methods across all compound and open target domains.

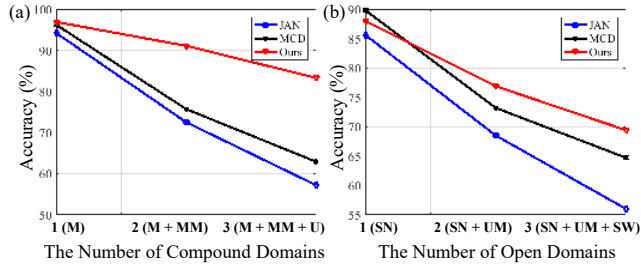


Figure 7: **Further analysis** on the (a) robustness to the complexity of the compound target domain and (b) robustness to the number of open domains. “M”, “MM” and “U” stand for MNIST, MNIST-M, and USPS, respectively, while “SN”, “UM” and “SW” stand for SynNum, USPS-M, and SWIT, respectively.

performance drop. The learned curriculum enables gradual knowledge transfer that is capable of coping with complex structures in the compound target domain.

Robustness to the Number of Open Domains. The performance change w.r.t. the number of open domains is demonstrated in Figure 7 (b). Here we include two new digits datasets, USPS-M (crafted in a similar way as MNIST-M) and SWIT [1], as the additional open domains. Compared to JAN [28] and MCD [40], our approach is more resilient to the various numbers of open domains. The domain indicator module in our framework helps dynamically calibrate the embedding, thus enhancing the robustness to open domains. Figure 8 presents the t-SNE visualization comparison between the obtained embeddings of JAN [28], MCD [40], and our approach.

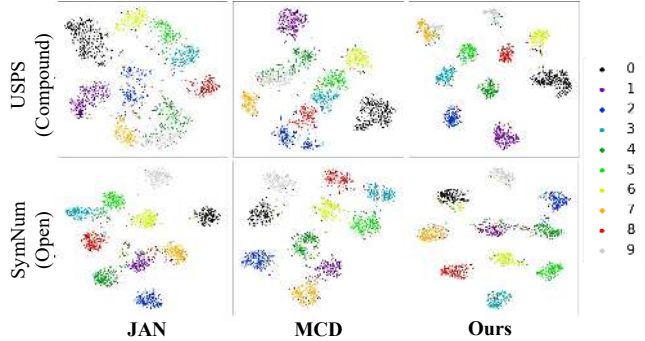


Figure 8: **t-SNE visualization** of the obtained embeddings. Compared to other methods, our approach is capable of producing class-discriminative features on both compound and open target domains.

5. Summary

We formalize a more realistic topic called open compound domain adaptation for domain-robust learning. We propose a novel model which includes a self-organizing curriculum domain adaptation to bootstrap generalization and a memory enhanced feature representation to build agility towards open domains. We develop several benchmarks on classification, recognition, segmentation, and reinforcement learning and demonstrate the effectiveness of our model.

Acknowledgements. This research was supported, in part, by the General Research Fund (GRF) of Hong Kong (No. 14236516 & No. 14203518), Berkeley Deep Drive, DARPA, NSF 1835539, and US Government fund through Etegent Technologies on Low-Shot Detection in Remote Sensing Imagery.

References

- [1] Switzerland handwritten digits dataset. <https://github.com/kensanata/numbers>. Accessed: 2019-03-15. 5, 8
- [2] Andreea Bobu, Eric Tzeng, Judy Hoffman, and Trevor Darrell. Adapting to continuously shifting domains. *ICLR Workshop*, 2018. 3
- [3] Zhangjie Cao, Lijia Ma, Mingsheng Long, and Jianmin Wang. Partial adversarial domain adaptation. In *ECCV*, 2018. 3
- [4] Ziliang Chen, Jingyu Zhuang, Xiaodan Liang, and Liang Lin. Blending-target domain adaptation by adversarial meta-adaptation networks. In *CVPR*, 2019. 3, 6
- [5] Dengxin Dai, Christos Sakaridis, Simon Hecker, and Luc Van Gool. Model adaptation with synthetic and real data for semantic dense foggy scene understanding. *IJCV*, 2019. 2
- [6] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, 2015. 2, 5
- [7] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *JMLR*, 2016. 1, 2
- [8] Behnam Gholami, Pritish Sahu, Ognjen Rudovic, Konstantinos Bousmalis, and Vladimir Pavlovic. Unsupervised multi-target domain adaptation: An information theoretic approach. *arXiv preprint arXiv:1810.11547*, 2018. 3, 6, 7
- [9] Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In *CVPR*, 2018. 5
- [10] Boqing Gong, Kristen Grauman, and Fei Sha. Reshaping visual datasets for domain adaptation. In *NIPS*, 2013. 1, 3
- [11] Boqing Gong, Yuan Shi, Fei Sha, and Kristen Grauman. Geodesic flow kernel for unsupervised domain adaptation. In *CVPR*, 2012. 1, 2
- [12] Rui Gong, Wen Li, Yuhua Chen, and Luc Van Gool. Dlow: Domain flow for adaptation and generalization. In *CVPR*, 2019. 3
- [13] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker. Multi-pie. In *2008 8th IEEE International Conference on Automatic Face Gesture Recognition*, 2008. 5
- [14] Kaiying He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 6
- [15] Judy Hoffman, Brian Kulis, Trevor Darrell, and Kate Saenko. Discovering latent domains for multisource domain adaptation. In *ECCV*, 2012. 1, 3
- [16] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei A Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *ICML*, 2018. 2
- [17] Hexiang Hu, Liyu Chen, Boqing Gong, and Fei Sha. Synthesized policies for transfer and adaptation across tasks and environments. In *NIPS*, 2018. 5, 6, 7
- [18] J. J. Hull. A database for handwritten text recognition research. *TPAMI*, 1994. 2, 5
- [19] Yann LeCun, Léon Bottou, Yoshua Bengio, Patrick Haffner, et al. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 1998. 2, 5, 6
- [20] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Learning to generalize: Meta-learning for domain generalization. In *AAAI*, 2018. 1, 3
- [21] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *CVPR*, 2018. 3
- [22] Wen Li, Zheng Xu, Dong Xu, Dengxin Dai, and Luc Van Gool. Domain generalization and adaptation using low rank exemplar svms. *TPAMI*, 2017. 1
- [23] Xiaoxiao Li, Ziwei Liu, Ping Luo, Chen Change Loy, and Xiaoou Tang. Not all pixels are equal: Difficulty-aware semantic segmentation via deep layer cascade. In *CVPR*, 2017. 2
- [24] Qing Lian, Fengmao Lv, Lixin Duan, and Boqing Gong. Constructing self-motivated pyramid curriculums for cross-domain semantic segmentation: A non-adversarial approach. In *ICCV*, 2019. 6, 7
- [25] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Spheredface: Deep hypersphere embedding for face recognition. In *CVPR*, 2017. 5
- [26] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. Large-scale long-tailed recognition in an open world. In *CVPR*, 2019. 5
- [27] Ziwei Liu, Sijie Yan, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Fashion landmark detection in the wild. In *ECCV*, 2016. 2
- [28] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Deep transfer learning with joint adaptation networks. In *ICML*, 2017. 1, 2, 6, 7, 8
- [29] Massimiliano Mancini, Samuel Rota Bulo, Barbara Caputo, and Elisa Ricci. Adagraph: Unifying predictive and continuous domain adaptation through graphs. In *CVPR*, 2019. 3
- [30] Massimiliano Mancini, Lorenzo Porzi, Samuel Rota Bulo, Barbara Caputo, and Elisa Ricci. Inferring latent domains for unsupervised deep domain adaptation. *TPAMI*, 2019. 3
- [31] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bisacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. *NIPS*, 2011. 2, 5
- [32] Sinno Jialin Pan, Ivor W Tsang, James T Kwok, and Qiang Yang. Domain adaptation via transfer component analysis. *IEEE Transactions on Neural Networks*, 2010. 1
- [33] Xingang Pan, Ping Luo, Jianping Shi, and Xiaoou Tang. Two at once: Enhancing learning and generalization capacities via ibn-net. In *ECCV*, 2018. 6, 7
- [34] Yingwei Pan, Ting Yao, Yehao Li, Yu Wang, Chong-Wah Ngo, and Tao Mei. Transferrable prototypical networks for unsupervised domain adaptation. In *CVPR*, 2019. 5
- [35] Pau Panareda Busto and Juergen Gall. Open set domain adaptation. In *ICCV*, 2017. 3
- [36] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *ICCV*, 2019. 2
- [37] Xingchao Peng, Zijun Huang, Ximeng Sun, and Kate Saenko. Domain agnostic learning with disentangled representations. In *ICML*, 2019. 3, 6, 7

- [38] Stephan R. Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *ECCV*, 2016. 5
- [39] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *ECCV*, 2010. 1, 2
- [40] Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In *CVPR*, 2018. 1, 2, 6, 7, 8
- [41] Kuniaki Saito, Shohei Yamamoto, Yoshitaka Ushiku, and Tatsuya Harada. Open set domain adaptation by backpropagation. In *ECCV*, 2018. 3
- [42] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 6
- [43] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *NIPS*, 2017. 5
- [44] Antonio Torralba and Alexei A Efros. Unbiased look at dataset bias. In *CVPR*, 2011. 1, 2
- [45] 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 *CVPR*, 2018. 6, 7
- [46] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *CVPR*, 2017. 1, 2, 4, 6, 7
- [47] Hemant Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *CVPR*, 2017. 2
- [48] Zuxuan Wu, Xin Wang, Joseph E Gonzalez, Tom Goldstein, and Larry S Davis. Ace: Adapting to changing environments for semantic segmentation. In *ICCV*, 2019. 3
- [49] Caiming Xiong, Scott McCloskey, Shao-Hang Hsieh, and Jason J Corso. Latent domains modeling for visual domain adaptation. In *AAAI*, 2014. 3
- [50] Zheng Xu, Wen Li, Li Niu, and Dong Xu. Exploiting low-rank structure from latent domains for domain generalization. In *ECCV*, 2014. 3
- [51] Fisher Yu, Wenqi Xian, Yingying Chen, Fangchen Liu, Mike Liao, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving video database with scalable annotation tooling. *arXiv preprint arXiv:1805.04687*, 2018. 5, 7
- [52] Huanhuan Yu, Menglei Hu, and Songcan Chen. Multi-target unsupervised domain adaptation without exactly shared categories. *arXiv preprint arXiv:1809.00852*, 2018. 3
- [53] Jing Zhang, Zewei Ding, Wanqing Li, and Philip Ogunbona. Importance weighted adversarial nets for partial domain adaptation. In *CVPR*, 2018. 3
- [54] Yang Zhang, Philip David, Hassan Foroosh, and Boqing Gong. A curriculum domain adaptation approach to the semantic segmentation of urban scenes. *TPAMI*, 2019. 2, 4
- [55] Yang Zou, Zhiding Yu, Xiaofeng Liu, BVK Kumar, and Jinsong Wang. Confidence regularized self-training. In *ICCV*, 2019. 2
- [56] Yang Zou, Zhiding Yu, BVK Vijaya Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *ECCV*, 2018. 2, 6, 7