

Unsupervised Model Adaptation for Continual Semantic Segmentation

Serban Stan,¹ Mohammad Rostami^{1,2}

¹ University of Southern California

² Information Sciences Institute

sstan@usc.edu, rostamim@usc.edu

Abstract

We develop an algorithm for adapting a semantic segmentation model that is trained using a labeled source domain to generalize well in an unlabeled target domain. A similar problem has been studied extensively in the unsupervised domain adaptation (UDA) literature, but existing UDA algorithms require access to both the source domain labeled data and the target domain unlabeled data for training a domain agnostic semantic segmentation model. Relaxing this constraint enables a user to adapt pretrained models to generalize in a target domain, without requiring access to source data. To this end, we learn a prototypical distribution for the source domain in an intermediate embedding space. This distribution encodes the abstract knowledge that is learned from the source domain. We then use this distribution for aligning the target domain distribution with the source domain distribution in the embedding space. We provide theoretical analysis and explain conditions under which our algorithm is effective. Experiments on benchmark adaptation tasks demonstrate our method achieves competitive performance even compared with joint UDA approaches.

Introduction

Image segmentation is an essential computer vision ability for delivering technologies such as autonomous driving (Zhang, Fidler, and Urtasun 2016) and automatic object tracking (Wang et al. 2019). Advances in deep learning have led to the development of image segmentation algorithms with close to human-level performance (Zeng, Wu, and Ji 2017). However, this success is conditioned on the availability of large and high-quality manually annotated datasets to satisfy the required sample complexity bounds for training deep neural networks. As a result, data annotation is a major bottleneck to address the problem of *domain shift*, where *domain discrepancy* exists between the distributions of training and testing domains (Luo et al. 2019) and the trained model needs to be adapted to generalize again after being fielded. This is particularly important in continual learning, where the goal is to enable a learning agent to learn new domains autonomously. Retraining the model from scratch is not a feasible solution for continual learning because manual data annotation is an expensive and time-consuming process for image segmentation, e.g., as much as 1.5 hours for a single

image of the current benchmark datasets (Cordts et al. 2016; Rostami, Huber, and Lu 2018). A practical alternative is to adapt the trained model using only unannotated data.

The problem of model adaptation for image segmentation has been studied extensively in the unsupervised domain adaptation (UDA) framework. The goal in UDA is to train a model for an unannotated target domain by transferring knowledge from a secondary related source domain in which annotated data is accessible or easier to generate, e.g., a synthetically generated domain. Knowledge transfer can be achieved by extracting domain-invariant features from the source and the target domains to address domain discrepancy. As a result, if we train a classifier using the source domain features as its input, the classifier will generalize on the target domain since the distributions of features are indistinguishable. Distributional alignment can be achieved by matching the distributions at different levels of abstraction, including appearance (Hoffman et al. 2018a; Sankaranarayanan et al. 2018), feature (Hoffman et al. 2018a; Murez et al. 2018), output (Zhang, David, and Gong 2017) levels.

A large group of the existing UDA algorithms for image segmentation use adversarial learning for extracting domain-invariant features (Luc et al. 2016; Bousmalis et al. 2017; Hoffman et al. 2018b; Murez et al. 2018; Saito et al. 2018; Sankaranarayanan et al. 2018; Dhoubib, Redko, and Lartzien 2020). Broadly speaking, a domain discriminator network can be trained to distinguish whether an input data point comes from the source or the target domain. This network is fooled by a feature generator network which is trained to make the domains similar at its output. Adversarial training (Goodfellow et al. 2014) of these two networks leads to learning a domain-agnostic embedding space. A second class of UDA algorithms directly minimize suitable loss functions that enforce domain alignment (Wu et al. 2018; Zhang, David, and Gong 2017; Gabourie et al. 2019; Zhang et al. 2019; Lee et al. 2019; Yang and Soatto 2020). Adversarial learning requires delicate optimization initialization, architecture engineering, and careful selection of hyper-parameters to be stable (Roth et al. 2017). In contrast, defining a suitable loss function for direct domain alignment may not be trivial.

A major limitation of existing UDA algorithms is that domain alignment can be performed only if the source and the target domain data are accessible concurrently. However, the source annotated data may not be necessarily accessible

during the model adaptation phase in a continual learning scenario. In this paper, we focus on a more challenging, yet more practical model adaptation scenario. We consider that a pretrained model is given and the goal is to adapt this model to generalize well in a target domain using solely unannotated target domain data. Our algorithm can be considered as an improvement over using an off-the-shelf pre-trained model naively by benefiting from the unannotated data in the target domain. This is a step towards lifelong image segmentation ability (Isele and Cosgun 2018; Rostami, Kolouri, and Pilly 2019).

Contributions: our main contribution is to relax the need for source domain annotated data for model adaptation. Our idea is to learn a prototypical distribution that encodes the abstract knowledge, learned for image segmentation using the source domain annotated data. The prototypical distribution is used for aligning the distributions across the two domains in an embedding space. We also provide theoretical analysis to justify the proposed model adaptation algorithm and determine the conditions under which our algorithm is effective. Finally, we provide experiments on the GTA5→Cityscapes and SYNTHIA→Cityscapes benchmark domain adaptation image segmentation tasks to demonstrate that our method is effective and leads to competitive performance, even when compared against existing UDA algorithms.

Related Work

We discuss related work on domain adaptation and semantic segmentation, focusing on direct distribution alignment.

Semantic Segmentation

Traditional semantic segmentation algorithms use hand-engineered extracted features which are fed into a classifier (Shotton, Johnson, and Cipolla 2008; Tighe and Lazebnik 2010), where the classifier is trained using supervised learning. In contrast, current state of the art approaches use convolutional neural networks (CNNs) for feature extraction (Long, Shelhamer, and Darrell 2015). A base CNN subnetwork is converted into a fully-convolutional network (FCN) for feature extraction and then is combined with a classifier subnetwork to form an end-to-end classifier. The whole pipeline is trained in an end-to-end deep supervised learning training scheme. Due to large size of learnable parameters of the resulting semantic segmentation network, a huge pixel-level manually annotated dataset is required for training.

We can use weakly supervised annotation such as using bounding boxes to reduce manual annotation cost (Pathak, Krahenbuhl, and Darrell 2015; Papandreou et al. 2015), but even obtaining weakly annotated data for semantic segmentation can be time-consuming. Additionally, a trained model using weakly annotated datasets may not generalize well during testing. Another approach for relaxing the need for manually annotated datasets is to use synthetic datasets which are generated using computer graphics (Ros et al. 2016; Cordts et al. 2016). These datasets can be annotated automatically. But a trained model using synthetic datasets, might not generalize well to real-world data due to the existence of domain shift problem (Sankaranarayanan et al. 2018). Unsupervised domain adaptation is developed to address this problem.

Domain Adaptation

Domain adaptation methods reduce domain discrepancy by aligning distributions using annotated data in a target domain and unannotated data in a target domain. A group of UDA methods use a shared cross-domain encoder to map data into a shared embedding space and train the encoder by minimizing a probability distance measure across the two domains at its output. The Wasserstein distance (WD) (Courtney et al. 2016; Bhushan Damodaran et al. 2018a) is an example of such measures which captures higher-order statistics. Damodaran et al. (Bhushan Damodaran et al. 2018a) demonstrate that using WD leads to performance improvement over methods that rely on matching lower-order statistics (Long et al. 2015; Sun and Saenko 2016). In this work, we rely on the sliced Wasserstein distance (SWD) variant of WD (Lee et al. 2019) for domain alignment. SWD has a closed form solution and can be computed more efficiently than WD.

Current UDA methods assume that the source and the target domain data are accessible concurrently during domain alignment. However, since usually a model is already pre-trained on the source domain, it is beneficial if we can adapt it using the target domain unannotated data. This model adaptation setting been explored for non-deep models (Dredze and Crammer 2008; Jain and Learned-Miller 2011; Wu 2016), but these works cannot be extended to semantic segmentation tasks. In this work, we benefit from prototypical distributions to align two distributions indirectly. The core idea is that the image pixels that belong to each semantic class form a data cluster in a shared embedding space (Rostami 2019). The centroid for this cluster is called class prototype. Recently, UDA has been addressed by aligning the prototype pairs across two domains (Pan et al. 2019; Chen et al. 2019a; Rostami and Galstyan 2020). Inspired by these works, we extend the notion of class prototypes to prototypical distributions in the embedding space. A prototypical distribution for image segmentation is a multimodal distribution that encodes the knowledge learned from the source domain. Our work is based on enforcing the two domains to share a similar prototypical distribution in the embedding.

Problem Formulation

Consider an image domain \mathbf{X}^s and a semantic segmentation model $f_{\theta}(\cdot) : \mathbf{X}^s \rightarrow \mathbf{Y}^s$ with learnable parameters θ which receives an input image $x^s \in \mathbf{X}^s$ and predicts pixel-wise category labels $y^s \in \mathbf{Y}^s$. The goal is to train the model such that the expected error, i.e. true risk, between the prediction and the ground truth is minimized, i.e., $\theta^* = \arg \min_{\theta} \{\mathbb{E}_{x^s \sim P_S(\mathbf{X})^s} (\mathcal{L}(f_{\theta}(x^s), y^s))\}$, where $P_S(\mathbf{X}^s)$ and $\mathcal{L}(\cdot)$ denote the input data distribution and a suitable loss function, respectively. In practice, we use empirical risk minimization (ERM) and the cross-entropy loss for solving for the optimal semantic segmentation model:

$$\begin{aligned} \hat{\theta} &= \arg \min_{\theta} \left\{ \frac{1}{N} \sum_{i=1}^N \mathcal{L}_{ce}(f_{\theta}(x^s), y^s) \right\} \\ \mathcal{L}_{ce} &= - \sum_{k=1}^K \sum_{h=1}^{H \times W} y_{ijk} \log(p_{ijk}), \end{aligned} \quad (1)$$

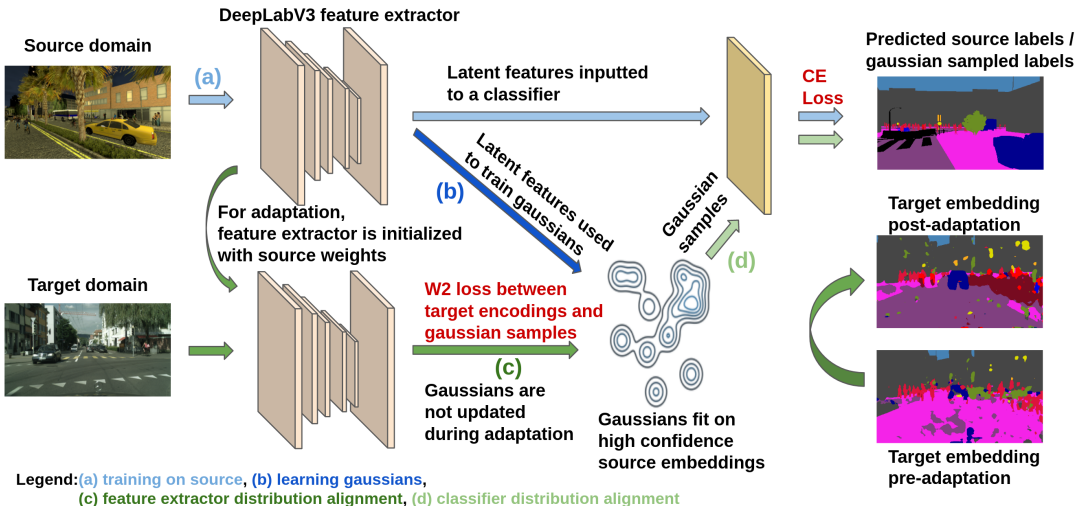


Figure 1: Diagram of the proposed model adaptation approach (best seen in color): (a) initial model training using the source domain labeled data, (b) estimating the prototypical distribution as a GMM distribution in the embedding space, (c) domain alignment is enforced by minimizing the distance between the prototypical distribution samples and the target unlabeled samples, (d) domain adaptation is enforced for the classifier module to fit correspondingly to the GMM distribution.

where N and K denote the training dataset size and the number of semantic categories. H and W denote the input image height and width, respectively. Also, $\mathbf{y}_{ij} = [y_{ijk}]_{k=1}^K$ is a one-hot vector that denotes the ground-truth semantic labels and $\mathbf{p}_{ij} = [p_{ijk}]_{k=1}^K$ is a probability vector of the predicted category probabilities by the model, i.e., a softmax layer is used as the last model layer. In practice, the model is fielded after training for testing and we do not store the source samples. If N is large enough and the base model is complex enough, the ERM-trained model will generalize well on unseen samples, drawn from the distribution $P_S(\mathbf{X}^s)$.

Now, consider that after training, we want to employ the source-trained model in a target domain \mathbf{X}^t with the distribution $P_T(\mathbf{X}^t)$, where $P_T(\mathbf{X}^t) \neq P_S(\mathbf{X}^s)$. Note that this situation emerges naturally in continual learning (Shin et al. 2017; Rostami et al. 2020) when data distribution changes over time. Within domain adaptation learning setting, this means that the target domain is encountered sequentially after learning the source domain. Due to existing distributional discrepancy, the source-trained model $f_{\hat{\theta}}$ will have poor generalization capability on the target domain. Our goal is to improve the model generalization by adapting the model such that the source and the target domains share a similar distribution in an embedding space. Following the state-of-the-art semantic segmentation models (Long, Shelhamer, and Darrell 2015), we consider a deep network as the base model $f_{\theta}(\cdot)$. This network is decomposed into a deep CNN encoder $\phi_u(\cdot) : \mathbb{R}^{H \times W} \rightarrow \mathbb{R}^P$, a category-wise CNN decoder $\psi_v(\cdot) : \mathbb{R}^P \rightarrow \mathbb{R}^{H \times W \times K}$, and a pixel-level classifier subnetwork $h_w(\cdot) : \mathcal{Z} \subset \mathbb{R}^K \rightarrow \mathbb{R}^K$ such that $f_{\theta} = h_w \circ \psi_v \circ \phi_u$, where $\theta = (w, v, u)$. In this decomposition, \mathcal{Z} denotes the shared cross-domain embedding space. If we adapt the base trained model $f_{\hat{\theta}}$ such that the domain discrepancy is minimized, i.e., the distance between the distributions $\psi(\phi(p_S(\mathbf{X}^s)))$ and $\psi(\phi(p_T(\mathbf{X}^t)))$ is minimized in

the embedding space, then the source-trained classifier $h_{\hat{w}}$ will generalize on both domains. Most UDA algorithms benefit from this approach to address annotated data scarcity in the target domain but all assume that the source samples are accessible for model adaptation. Since this makes computing the distance between the distributions $\psi(\phi(p_S(\mathbf{X}^s)))$ and $\psi(\phi(p_T(\mathbf{X}^t)))$ feasible, solving UDA reduces to aligning these distributions. Note, however, $\psi(\phi(p_T(\mathbf{X}^t)))$ cannot be computed directly in our learning framework due to absence of source samples and we need to estimate this distribution.

Proposed Algorithm

Figure 1 presents a high-level visual description of our approach. Our solution is based on aligning the source and the target distributions via an intermediate prototypical distribution in the embedding space. Since the last layer of the classifier is a softmax layer, we can treat the classifier as a maximum *a posteriori* (MAP) estimator. This composition implies that if after training, the model can generalize well in the target domain, it must transform the source input distribution into a multimodal distribution $p_J(z)$ with K separable components in the embedding space (see Figure 1 (a)). Each mode of this distribution represents one of the K semantic classes. This prototypical distribution emerges as a result of model training because the classes should become separable in the embedding space for a generalizable softmax classifier. Recently, this property has been used for UDA (Pan et al. 2019; Chen et al. 2019a), where the means for distribution modes are considered as the class prototype. The idea for UDA is to align the domain-specific prototypes for each class to enforce distributional alignment across the domains. Our idea is to adapt the trained model using the target unlabeled data such that in addition to the prototypes, the source-learned prototypical distribution does not change after adaptation. As a result, the classifier subnetwork will still generalize in the target

domain because its input distribution has been consolidated.

We model the prototypical distribution $p_J(\mathbf{z})$ as a Gaussian mixture model (GMM) with k components:

$$p_J(\mathbf{z}) = \sum_{j=1}^k \alpha_j \mathcal{N}(\mathbf{z} | \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j), \quad (2)$$

where α_j denote mixture weights, i.e., prior probability for each semantic class. For each component, $\boldsymbol{\mu}_j$ and $\boldsymbol{\Sigma}_j$ denote the mean and co-variance of the Gaussian (see Figure 1 (b)).

The empirical version of the prototypical distribution is accessible by the source domain samples $\{(\psi_v(\phi_v(\mathbf{x}_i^s)), \mathbf{y}_i^s)\}_{i=1}^N$ which we use for estimating the parameters. Note that since the labels are accessible in the source domain, we can estimate the parameters of each component independently via MAP estimation. Additionally, since p_{ijk} denotes the confidence of the classifier for the estimated semantic label for a given pixel. Hence, we can choose a threshold τ and compute the parameters using samples for which $p_{ijk} > \tau$ to cancel the effect of misclassified samples that would act as outliers. Let \mathbf{S}_j denote the support set for class j in the training dataset for which $p_{ijk} > \tau$, i.e., $\mathbf{S}_j = \{(\mathbf{x}_i^s, \mathbf{y}_i^s) \in \mathcal{D}_S | \arg \max \hat{\mathbf{y}}_i^s = j, p_{ijk} > \tau\}$. Then, the MAP estimates for the distribution parameters would be:

$$\begin{aligned} \hat{\alpha}_j &= \frac{|\mathbf{S}_j|}{\sum_{j=1}^k |\mathbf{S}_j|}, \quad \hat{\boldsymbol{\mu}}_j = \sum_{(\mathbf{x}_i^s, \mathbf{y}_i^s) \in \mathbf{S}_j} \frac{1}{|\mathbf{S}_j|} \psi_u(\phi_v(\mathbf{x}_i^s)), \\ \hat{\boldsymbol{\Sigma}}_j &= \frac{1}{|\mathbf{S}_j|} \sum_{(\mathbf{x}_i^s, \mathbf{y}_i^s) \in \mathbf{S}_j} (\psi_u(\phi_v(\mathbf{x}_i^s)) - \hat{\boldsymbol{\mu}}_j)^\top (\phi_v(\phi_v(\mathbf{x}_i^s)) - \hat{\boldsymbol{\mu}}_j). \end{aligned} \quad (3)$$

We take advantage of the prototypical distributional estimate in Eq.(3) as a surrogate for the source domain distribution to align the source and the target domain distribution in the absence of the source samples. We can adapt the model such that the encoder transforms the target domain distribution into the prototypical distributional in the embedding space. We use the prototypical distributional estimate and draw random samples to generate a labeled pseudo-dataset: $\mathcal{D}_P = (\mathbf{Z}_P, \mathbf{Y}_P)$, where $\mathbf{Z}_P = [\mathbf{z}_1^p, \dots, \mathbf{z}_{N_p}^p] \in \mathbb{R}^{K \times N_p}$, $\mathbf{Y}_P = [\mathbf{y}_1^p, \dots, \mathbf{y}_{N_p}^p] \in \mathbb{R}^{K \times N_p}$, $\mathbf{z}_i^p \sim \hat{p}_J(\mathbf{z})$. To improve the quality of the pseudo-dataset, we use the classifier sub-network prediction on drawn samples \mathbf{z}^p to select samples with $h_w(\mathbf{z}^p) > \tau$. After generating the pseudo-dataset, we solve the following optimization problem to align the source and the target distributions indirectly in the embedding:

$$\begin{aligned} \arg \min \mathbf{u}, \mathbf{v}, \mathbf{w} \{ & \frac{1}{N_p} \sum_{i=1}^{N_p} \mathcal{L}_{ce}(h_w(\mathbf{z}_i^{(p)}), \mathbf{y}_i^{(p)}) + \\ & + \lambda D(\psi_v(\phi_v(p_{\mathcal{T}}(\mathbf{X}_{\mathcal{T}}))), \hat{p}_J(\mathbf{Z}_P)) \}, \end{aligned} \quad (4)$$

where $D(\cdot, \cdot)$ denotes a probability distribution metric to enforce alignment of the target domain distribution with the prototypical distribution in embedding space and λ is a trade-off parameter between the two terms (see Figure 1 (c)).

The first term in Eq. (4) is to update the classifier such that it keeps its generalization power on the prototypical

Algorithm 1 MAS³ (λ, τ)

```

1: Initial Training:
2:   Input: source domain dataset  $\mathcal{D}_S = (\mathbf{X}_S, \mathbf{Y}_S)$ ,
3:   Training on Source Domain:
4:      $\hat{\theta}_0 = (\hat{w}_0, \hat{v}_0 \hat{u}_0) = \arg \min_{\theta} \sum_i \mathcal{L}(f_{\theta}(\mathbf{x}_i^s), \mathbf{y}_i^s)$ 
5:   Prototypical Distribution Estimation:
6:     Use Eq. (3) and estimate  $\alpha_j, \boldsymbol{\mu}_j$ , and  $\boldsymbol{\Sigma}_j$ 
7:   Model Adaptation:
8:     Input: target dataset  $\mathcal{D}_T = (\mathbf{X}_T)$ 
9:     Pseudo-Dataset Generation:
10:     $\mathcal{D}_P = (\mathbf{Z}_P, \mathbf{Y}_P) =$ 
11:       $([\mathbf{z}_1^p, \dots, \mathbf{z}_{N_p}^p], [\mathbf{y}_1^p, \dots, \mathbf{y}_{N_p}^p])$ , where:
12:       $\mathbf{z}_i^p \sim \hat{p}_J(\mathbf{z}), 1 \leq i \leq N_p$ 
13:       $\mathbf{y}_i^p = \arg \max_j \{h_{w_0}(\mathbf{z}_i^p)\}, p_{ip} > \tau$ 
14:    for  $itr = 1, \dots, ITR$  do
15:      draw random batches from  $\mathcal{D}_T$  and  $\mathcal{D}_P$ 
16:      Update the model by solving Eq. (4)
17:    end for
```

distribution. The second term is a matching loss term used to update the model such that the target domain distribution is matched to the prototypical distribution in the embedding space. Given a suitable probability metric, Eq. (4) can be solved using standard deep learning optimization techniques.

The major remaining question is selecting a proper probability metric to compute $D(\cdot, \cdot)$. Note that the original target distribution is not accessible and hence we should select a metric that can be used to compute the domain discrepancy via the observed target domain data samples and the drawn samples from the prototypical distribution. Additionally, the metric should be smooth and easy to compute to make it suitable for gradient-based optimization that is normally used to solve Eq. (4). In this work, we use Sliced Wasserstein Distance (SWD) (Rabin et al. 2011). Wasserstein Distance (WD) has been used successfully for domain alignment in the UDA literature (Bhushan Damodaran et al. 2018b; Rostami et al. 2019; Xu et al. 2020; Li et al. 2020; Rostami and Galstyan 2020). SWD is a variant of WD that can be computed more efficiently (Lee et al. 2019). SWD benefits from the idea of slicing by projecting high-dimensional probability distributions into their marginal one-dimensional distributions. Since one-dimensional WD has a closed-form solution, WD between these marginal distributions can be computed fast. SWD approximates WD as a summation of WD between a number of random one-dimensional projections:

$$D(\hat{p}_J, p_{\mathcal{T}}) \approx \frac{1}{L} \sum_{l=1}^L \sum_{i=1}^M |\langle \gamma_l, \mathbf{z}_{p_l[i]}^p \rangle - \langle \gamma_l, \psi(\phi(\mathbf{x}_{t_l[i]}^t)) \rangle|^2 \quad (5)$$

where $\gamma_l \in \mathbb{S}^{f-1}$ is uniformly drawn random sample from the unit f -dimensional ball \mathbb{S}^{f-1} , and $p_l[i]$ and $t_l[i]$ are the sorted indices for the prototypical and the target domain samples, respectively. We utilize Eq. (5) to solve Eq. (4).

Our solution for source-free model adaptation, named Model Adaptation for Source-Free Semantic Segmentation (MAS³), is described conceptually in Figure 1 and the corresponding algorithmic solution is given in Algorithm 1.

Theoretical Analysis

We analyze our algorithm within standard PAC-learning and prove that Algorithm 1 optimizes an upper-bound of the expected error for the target domain under certain conditions.

Consider that the hypothesis space within PAC-learning is the set of classifier sub-networks $\mathcal{H} = \{h_w(\cdot) | h_w(\cdot) : \mathcal{Z} \rightarrow \mathbb{R}^k, w \in \mathbb{R}^W\}$. Let e_S and e_T denote the true expected error of the optimal domain-specific model from this space on the source and target domain respectively. We denote the joint-optimal model with h_{w^*} . This model has the minimal combined source and target expected error $e_C(w^*)$, i.e. $w^* = \arg \min_w e_C(w) = \arg \min_w \{e_S + e_T\}$. In other words, it is a model with the best performance for both domains.

Since we process the observed data points from these domains, let $\hat{\mu}_S = \frac{1}{N} \sum_{n=1}^N \delta(\psi(\phi_v(x_n^s)))$ and $\hat{\mu}_T = \frac{1}{M} \sum_{m=1}^M \delta(\psi(\phi_v(x_m^t)))$ denote the empirical source and the empirical target distributions in the embedding space that are built using the observed data points. Similarly, let $\hat{\mu}_P = \frac{1}{N_p} \sum_{q=1}^{N_p} \delta(z_q^p)$ denote the empirical prototypical distribution which is built using the generated pseudo-dataset.

Finally, note that when fitting the GMM and when we generate the pseudo-dataset, we only included those data points and pseudo-data points for which the model is confident about their predicted labels. For this reason, we can conclude that: $\tau = \mathbb{E}_{z \sim p_J(z)} (\mathcal{L}(h(z), h_{w_0}(z)))$.

Theorem 1: Consider that we generate a pseudo-dataset using the prototypical distribution and update the model for sequential UDA using algorithm 1. Then, the following holds:

$$e_T \leq e_S + W(\hat{\mu}_S, \hat{\mu}_P) + W(\hat{\mu}_T, \hat{\mu}_P) + (1 - \tau) + e_{C'}(w^*) + \sqrt{(2 \log(\frac{1}{\xi})/\zeta)} (\sqrt{\frac{1}{N}} + \sqrt{\frac{1}{M}} + 2\sqrt{\frac{1}{N_p}}), \quad (6)$$

where $W(\cdot, \cdot)$ denotes the WD distance and ξ is a constant which depends on the loss function $\mathcal{L}(\cdot)$.

Proof: the complete proof is included in the Appendix.

Theorem 1 justifies effectiveness of our algorithm. We observe that MAS^3 algorithm minimizes the upperbound expressed in Eq. (6). The source expected risk is minimized through the initial training on the source domain. The second term in Eq. (6) is minimized because we deliberately fit a GMM distribution on the source domain distribution in the embedding space. Note that minimizing this term is conditionally possible when the source domain distribution can be approximated well with a GMM distribution. However, similar constraint exist for all the parametric methods in statistics. Additionally, since we use a softmax in the last layer, this would likely to happen because the classes should become separable for a generalizable model to be trained. The third term in the Eq. (6) upperbound is minimized as the second term in Eq. (4). The fourth term is a constant term depending on the threshold we use and can be small if $\tau \approx 1$. Note that when selecting the source distribution samples for fitting the Gaussian distribution, if we set τ too close to 1, we may not have sufficient samples for accurate estimation of GMM and hence the second term may increase. Hence, there is a trade-off between minimizing the second and the fourth term

in Eq. (6). The term $e_{C'}(w^*)$ will be small if the domains are related, i.e., share the same classes and the base model can generalize well in both domains, when trained in the presence of sufficient labeled data from both domains. In other words, aligning the distributions in the embedding must be a possibility for our algorithm to work. This is a condition for all UDA algorithms to work. Finally, the last term in Eq. (6) is a constant term similar to most PAC-learnability bounds and can be negligible if sufficiently large source and target datasets are accessible and we generate a large pseudo-dataset.

Experimental Validation

We validate our algorithm using two benchmark domain adaptation tasks, and provide our code at <https://github.com/serbanstan/mas3-continual>.

Experimental Setup

Datasets and Evaluation Metrics: We validate MAS^3 on the standard GTA5 (Richter et al. 2016)→Cityscapes (Cordts et al. 2016) and the SYNTHIA (Ros et al. 2016)→Cityscapes benchmark UDA tasks for semantic segmentation.

GTA5 consists of 24,966 1914×1052 image instances.

SYNTHIA consists of 9,400 1280×760 image instances.

Cityscapes is a real-world dataset consisting of a training set with 2,957 instances and a validation set, used as testing set, with 500 instances of images with size 2040×1016 .

The GTA5 and SYNTHIA datasets are used as source domains and after training the model, we adapt it to generalize in the Cityscapes dataset as the target domain. To use a shared cross-domain encoder, we resize all images to 1024×512 size. This can be considered as a preprocessing step using a domain-specific mapping. Due to space limits, implementation details are included in the Appendix.

Evaluation: Following the literature, we report the results on the Cityscapes validation set and use the category-wise and the mean intersection over union (IoU) to measure segmentation performance (Hoffman et al. 2016). Note that while GTA5 has the same 19 category annotations as Cityscapes, SYNTHIA has 16 common category annotations. For this reason and following the literature, we report the results on the shared cross-domain categories for each task.

Comparison With State-of-the-art Methods: To the best of our knowledge, there is no prior source-free model adaptation algorithm for performance comparison. For this reason, we compare MAS^3 against UDA algorithms based on joint training due to proximity of these works to our learning setting. In our comparison, we have included both pioneer and recent UDA image segmentation method to be representative of the literature. We have compared our performance against the adversarial learning-based UDA methods: GIO-Ada (Chen et al. 2019b), ADVENT (Vu et al. 2019), AdaSegNet (Tsai et al. 2018), TGCF-DA+SE (Choi, Kim, and Kim 2019), PCEDA (Yang et al. 2020), and CyCADA (Hoffman et al. 2018a). We have also included methods that are based on direct distributional matching which are more similar to MAS^3 : FCNs in the Wild (Hoffman et al. 2016), CDA (Zhang, David, and Gong 2017), DCAN (Wu et al. 2018), SWD (Lee et al. 2019), Cross-City (Chen et al. 2017b).

Method	Adv.	road	sdwlk	bldng	light	sign	vgtn	sky	person	rider	car	bus	mcycl	bcycl	mIoU
Source Only (VGG16)	N	6.4	17.7	29.7	0.0	7.2	30.3	66.8	51.1	1.5	47.3	3.9	0.1	0.0	20.2
FCNs in the Wild	N	11.5	19.6	30.8	0.1	11.7	42.3	68.7	51.2	3.8	54.0	3.2	0.2	0.6	22.9
CDA	N	65.2	26.1	74.9	3.7	3.0	76.1	70.6	47.1	8.2	43.2	20.7	0.7	13.1	34.8
DCAN	N	9.9	30.4	70.8	6.70	23.0	76.9	73.9	41.9	16.7	61.7	11.5	10.3	38.6	36.4
SWD	N	83.3	35.4	82.1	12.2	12.6	83.8	76.5	47.4	12.0	71.5	17.9	1.6	29.7	43.5
Cross-City	Y	62.7	25.6	78.3	1.2	5.4	81.3	81.0	37.4	6.4	63.5	16.1	1.2	4.6	35.7
GIO-Ada	Y	78.3	29.2	76.9	10.8	17.2	81.7	81.9	45.8	15.4	68.0	15.9	7.5	30.4	43.0
ADVENT	Y	67.9	29.4	71.9	0.6	2.6	74.9	74.9	35.4	9.6	67.8	21.4	4.1	15.5	36.6
AdaSegNet	Y	78.9	29.2	75.5	0.1	4.8	72.6	76.7	43.4	8.8	71.1	16.0	3.6	8.4	37.6
TGCF-DA+SE	Y	90.1	48.6	80.7	3.2	14.3	82.1	78.4	54.4	16.4	82.5	12.3	1.7	21.8	46.6
PCEDA	Y	79.7	35.2	78.7	10.0	28.9	79.6	81.2	51.2	25.1	72.2	24.1	16.7	50.4	48.7
MAS ³ (Ours)	N	75.1	49.6	70.9	14.1	25.3	72.7	76.7	48.5	19.9	65.3	17.6	6.8	39.0	44.7

Table 1: Model adaptation comparison results for the SYNTHIA→Cityscapes task. We have used DeepLabV3 (Chen et al. 2017a) as the feature extractor with a VGG16 (Simonyan and Zisserman 2014) backbone. The first row presents the source-trained model performance prior to adaptation to demonstrate effect of initial knowledge transfer from the source domain.

Results

Quantitative Performance Comparison:

SYNTHIA→Cityscapes: We report the quantitative results in table 1. We note that despite addressing a more challenge learning setting, *MAS*³ outperforms most of the UDA methods. Recently developed UDA methods based on adversarial learning outperform our method but we note that these methods benefit from a secondary type of regularization in addition to probability matching. Overall, *MAS*³ performs reasonably well even compared with these UDA methods that need source samples. Additionally, *MAS*³ has the best performance for some important categories, e.g., traffic light.

GTA5→Cityscapes: Quantitative results for this task are reported in Tables 2, 3. We observe a more competitive performance for this task but the performance comparison trend is similar. These results demonstrate that although the motivation in this work is source-free model adaptation, *MAS*³ can also be used as a joint-training UDA algorithm.

Qualitative Performance Validation:

In Figure 4, we have visualized exemplar frames for the Cityscapes dataset for the GTA5→Cityscapes task which are segmented using the model prior and after adaptation along with the ground-truth (GT) manual annotation for each image. Visual observation demonstrates that our method is able to significantly improve image segmentation from the source-only segmentation to the post-adaptation segmentation, noticeably on sidewalk, road, and car semantic classes for the GTA5-trained model. Examples of segmented frames for the SYNTHIA→Cityscapes task are included in the Appendix and provide similar observation.

Effect of Alignment in the Embedding Space:

To demonstrate that our solution implements what we anticipated, we have used UMAP (McInnes et al. 2018) visualization tool to reduce the dimension of the data representations in the embedding space to two for 2D visualization. Figure 2 represents the samples of the prototypical distribution along with the target domain data prior and after adaptation in the embedding space for the GTA5→Cityscapes task. Each point in Figure 2 denotes a single data point and each color

denotes a semantic class cluster. Comparing Figure 2b and Figure 2c with Figure 2a, we can see that the semantic classes in the target domain have become much more well-separated and more similar to the prototypical distribution after model adaptation. This means that domain discrepancy has been reduced using *MAS*³ and the source and the target domain distributions are aligned indirectly as anticipated using the intermediate prototypical distribution in the embedding space.

Ablation Study

A major advantage of our algorithm over methods based on adversarial learning is its simplicity in depending on a few hyper-parameters. We note that the major algorithm-specific hyper-parameters are λ and τ . We observed in our experiments that *MAS*³ performance is stable with respect to the trade-off parameter λ value. This is expected because in Eq. (4), the \mathcal{L}_{ce} loss term is small from the beginning due to prior training on the source domain. We investigated the impact of the confidence hyper-parameter τ value. Figure 3 presents the fitted GMM on the source prototypical distribution for three different values of τ . As it can be seen, when $\tau = 0$, the fitted GMM clusters are cluttered. As we increase the threshold τ and use samples for which the classifier is confident, the fitted GMM represents well-separated semantic classes which increases knowledge transfer from the source domain. This experiments also empirically validates what we deduced about importance of τ using Theorem 1.

Conclusions

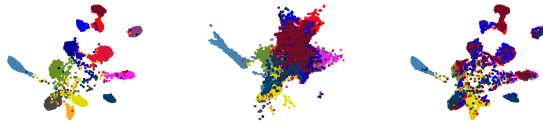
We developed an algorithm for adapting an image segmentation model to generalize in new domains after training using solely unlabeled data. Our algorithm is based on using an intermediate multi-modal prototypical distribution to minimize the the distributional cross-domain discrepancy in a shared embedding space. We estimate the prototypical distribution as a parametric GMM distribution. Experiments on benchmark tasks demonstrate our algorithm is effective and leads to competitive performance, even when compared to UDA algorithms that are based on joint-domain model training.

Method	Adv.	road	sdwk	bldng	wall	fence	pole	light	sign	vgtn	trm	...	mIoU
Source (VGG16)	N	25.9	10.9	50.5	3.3	12.2	25.4	28.6	13.0	78.3	7.3	...	24.9
FCNs Wld.	N	70.4	32.4	62.1	14.9	5.4	10.9	14.2	2.7	79.2	21.3	...	27.1
CDA	N	74.9	22.0	71.7	6.0	11.9	8.4	16.3	11.1	75.7	13.3	...	28.9
DCAN	N	82.3	26.7	77.4	23.7	20.5	20.4	30.3	15.9	80.9	25.4	...	36.2
SWD	N	91.0	35.7	78.0	21.6	21.7	31.8	30.2	25.2	80.2	23.9	...	39.9
CyCADA	Y	85.2	37.2	76.5	21.8	15.0	23.8	22.9	21.5	80.5	31.3	...	35.4
ADVENT	Y	86.9	28.7	78.7	28.5	25.2	17.1	20.3	10.9	80.0	26.4	...	36.1
AdaSegNet	Y	86.5	36.0	79.9	23.4	23.3	23.9	35.2	14.8	83.4	33.3	...	42.4
TGCF-DA+SE	Y	90.2	51.5	81.1	15.0	10.7	37.5	35.2	28.9	84.1	32.7	...	42.5
PCEDA	Y	90.2	44.7	82.0	28.4	28.4	24.4	33.7	35.6	83.7	40.5	...	44.6
MAS ³ (Ours)	N	75.5	53.7	72.2	20.5	24.1	30.5	28.7	37.8	79.6	36.9	...	43.9

Table 2: Domain adaptation results for different methods for the GTA5→Cityscapes task. Reporting the first 10 classes.

Method	Adv.	...	sky	person	rider	car	truck	bus	train	mcycl	beycl	mIoU
Source (VGG16)	N	...	63.9	52.1	7.9	66.3	5.2	7.8	0.9	13.7	0.7	24.9
FCNs Wld.	N	...	64.6	44.1	4.2	70.4	8.0	7.3	0.0	3.5	0.0	27.1
CDA	N	...	66.5	38.0	9.3	55.2	18.8	18.9	0.0	16.8	14.6	28.9
DCAN	N	...	69.5	52.6	11.1	79.6	24.9	21.2	1.30	17.0	6.70	36.2
SWD	N	...	74.1	53.1	15.8	79.3	22.1	26.5	1.5	17.2	30.4	39.9
CyCADA	Y	...	60.7	50.5	9.0	76.9	17.1	28.2	4.5	9.8	0.0	35.4
ADVENT	Y	...	70.2	47.1	8.4	81.5	26.0	17.2	18.9	11.7	1.6	36.1
AdaSegNet	Y	...	75.6	58.5	27.6	73.7	32.5	35.4	3.9	30.1	28.1	42.4
TGCF-DA+SE	Y	...	75.9	62.7	19.9	82.6	22.9	28.3	0.0	23.0	25.4	42.5
PCEDA	Y	...	75.1	54.4	28.2	80.3	23.8	39.4	0.0	22.8	30.8	44.6
MAS ³ (Ours)	N	...	78.7	49.6	16.5	77.4	26.0	42.6	18.8	15.3	49.9	43.9

Table 3: Domain adaptation results for different methods for the GTA5→Cityscapes task. Reporting the last 9 classes.



(a) GMM samples (b) Pre-adaptation (c) Post-adaptation

Figure 2: Indirect distribution matching in the embedding space: (a) drawn samples from the GMM trained on the SYNTHIA distribution, (b) representations of the Cityscapes validation samples prior to model adaptation (c) representation of the Cityscapes validation samples after domain alignment.



(a) $\tau = 0$ mIoU=41.8 (b) $\tau = 0.8$ mIoU=42.7 (c) $\tau = 0.97$ mIoU=44.7

Figure 3: Ablation experiment to study effect of τ on the GMM learnt in the embedding space: (a) all samples are used; adaptation mIoU=41.8, (b) a portion of samples is used; adaptation mIoU=42.7, (c) samples with high model-confidence are used; adaptation mIoU=44.7

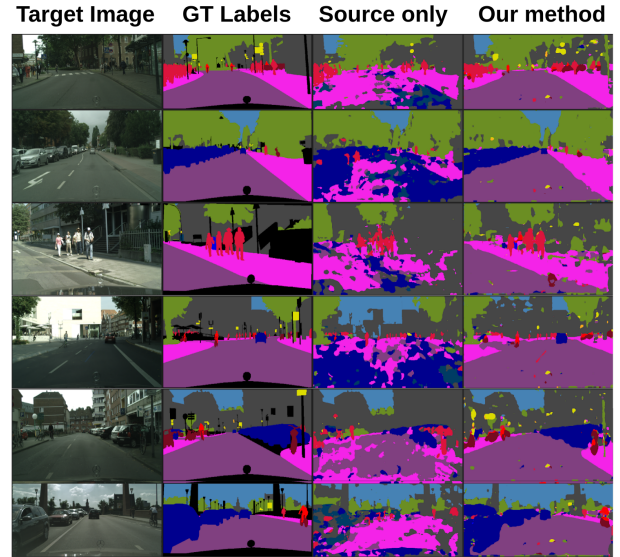


Figure 4: Qualitative performance: examples of the segmented frames for SYNTHIA→Cityscapes using the MAS³ method. Left to right: real images, manually annotated images, source-trained model predictions, predictions based on our method.

Acknowledgements

This material is based on research sponsored by Air Force Research Laboratory (AFRL) under agreement number FA8750-19-1-1000. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation therein. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of Air Force Laboratory, DARPA or the U.S. Government.

References

- Bhushan Damodaran, B.; Kellenberger, B.; Flamary, R.; Tuia, D.; and Courty, N. 2018a. Deepjdot: Deep joint distribution optimal transport for unsupervised domain adaptation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 447–463.
- Bhushan Damodaran, B.; Kellenberger, B.; Flamary, R.; Tuia, D.; and Courty, N. 2018b. Deepjdot: Deep joint distribution optimal transport for unsupervised domain adaptation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 447–463.
- Bousmalis, K.; Silberman, N.; Dohan, D.; Erhan, D.; and Krishnan, D. 2017. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3722–3731.
- Chen, L.-C.; Papandreou, G.; Schroff, F.; and Adam, H. 2017a. Rethinking atrous convolution for semantic image segmentation. In *arXiv:1706.05587 [cs.CV]*.
- Chen, Y.-H.; Chen, W.-Y.; Chen, Y.-T.; Tsai, B.-C.; Wang, Y.-C. F.; and Sun, M. 2017b. No more discrimination: Cross city adaptation of road scene segmenters. In *arXiv:1704.08509 [cs.CV]*.
- Chen, C.; Xie, W.; Huang, W.; Rong, Y.; Ding, X.; Huang, Y.; Xu, T.; and Huang, J. 2019a. Progressive feature alignment for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 627–636.
- Chen, Y.; Li, W.; Chen, X.; and Gool, L. V. 2019b. Learning semantic segmentation from synthetic data: A geometrically guided input-output adaptation approach. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1841–1850.
- Choi, J.; Kim, T.; and Kim, C. 2019. Self-ensembling with gan-based data augmentation for domain adaptation in semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6830–6840.
- Cordts, M.; Omran, M.; Ramos, S.; Rehfeld, T.; Enzweiler, M.; Benenson, R.; Franke, U.; Roth, S.; and Schiele, B. 2016. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3213–3223.
- Courty, N.; Flamary, R.; Tuia, D.; and Rakotomamonjy, A. 2016. Optimal transport for domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39(9):1853–1865.
- Dhouib, S.; Redko, I.; and Lartizien, C. 2020. Margin-aware adversarial domain adaptation with optimal transport. In *Thirty-seventh International Conference on Machine Learning*.
- Dredze, M., and Crammer, K. 2008. Online methods for multi-domain learning and adaptation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 689–697. Association for Computational Linguistics.
- Gabourie, A. J.; Rostami, M.; Pope, P. E.; Kolouri, S.; and Kim, K. 2019. Learning a domain-invariant embedding for unsupervised domain adaptation using class-conditioned distribution alignment. In *2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 352–359.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, 2672–2680.
- Hoffman, J.; Wang, D.; Yu, F.; and Darrell, T. 2016. Fcns in the wild: Pixel-level adversarial and constraint-based adaptation. In *arXiv:1612.02649 [cs.CV]*.
- Hoffman, J.; Tzeng, E.; Park, T.; Zhu, J.-Y.; Isola, P.; Saenko, K.; Efros, A.; and Darrell, T. 2018a. CyCADA: Cycle-consistent adversarial domain adaptation. In *International Conference on Machine Learning*, 1989–1998.
- Hoffman, J.; Tzeng, E.; Park, T.; Zhu, J.-Y.; Isola, P.; Saenko, K.; Efros, A.; and Darrell, T. 2018b. Cycada: Cycle-consistent adversarial domain adaptation. In *International conference on machine learning*, 1989–1998. PMLR.
- Isele, D., and Cosgun, A. 2018. Selective experience replay for lifelong learning. In *AAAI Conference on Artificial Intelligence 2018*, 3302–3309. Association for the Advancement of Artificial Intelligence (AAAI).
- Jain, V., and Learned-Miller, E. 2011. Online domain adaptation of a pre-trained cascade of classifiers. In *Proceedings of the 2011 IEEE Conference on Computer Vision and Pattern Recognition*, 577–584.
- Lee, C.-Y.; Batra, T.; Baig, M. H.; and Ulbricht, D. 2019. Sliced wasserstein discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 10285–10295.
- Li, M.; Zhai, Y.-M.; Luo, Y.-W.; Ge, P.-F.; and Ren, C.-X. 2020. Enhanced transport distance for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13936–13944.
- Long, M.; Cao, Y.; Wang, J.; and Jordan, M. 2015. Learning transferable features with deep adaptation networks. In *Proceedings of International Conference on Machine Learning*, 97–105.
- Long, J.; Shelhamer, E.; and Darrell, T. 2015. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3431–3440.
- Luc, P.; Couprie, C.; Chintala, S.; and Verbeek, J. 2016. Semantic segmentation using adversarial networks. In *NIPS Workshop on Adversarial Training*.
- Luo, Y.; Zheng, L.; Guan, T.; Yu, J.; and Yang, Y. 2019. Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2507–2516.
- McInnes, L.; Healy, J.; Saul, N.; and Großberger, L. 2018. UMAP: Uniform manifold approximation and projection. *Journal of Open Source Software* 3(29):861.
- Murez, Z.; Kolouri, S.; Kriegman, D.; Ramamoorthi, R.; and Kim, K. 2018. Image to image translation for domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4500–4509.
- Pan, Y.; Yao, T.; Li, Y.; Wang, Y.; Ngo, C.-W.; and Mei, T. 2019. Transferrable prototypical networks for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2239–2247.

- Papandreou, G.; Chen, L.-C.; Murphy, K. P.; and Yuille, A. L. 2015. Weakly-and semi-supervised learning of a deep convolutional network for semantic image segmentation. In *Proceedings of the IEEE international conference on computer vision*, 1742–1750.
- Pathak, D.; Krahenbuhl, P.; and Darrell, T. 2015. Constrained convolutional neural networks for weakly supervised segmentation. In *Proceedings of the IEEE international conference on computer vision*, 1796–1804.
- Rabin, J.; Peyré, G.; Delon, J.; and Bernot, M. 2011. Wasserstein barycenter and its application to texture mixing. In *International Conference on Scale Space and Variational Methods in Computer Vision*, 435–446. Springer.
- Richter, S. R.; Vineet, V.; Roth, S.; and Koltun, V. 2016. Playing for data: Ground truth from computer games. In *European conference on computer vision*, 102–118. Springer.
- Ros, G.; Sellart, L.; Materzynska, J.; Vazquez, D.; and Lopez, A. M. 2016. The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3234–3243.
- Rostami, M., and Galstyan, A. 2020. Sequential unsupervised domain adaptation through prototypical distributions. *arXiv preprint arXiv:2007.00197*.
- Rostami, M.; Kolouri, S.; Eaton, E.; and Kim, K. 2019. Deep transfer learning for few-shot sar image classification. *Remote Sensing* 11(11):1374.
- Rostami, M.; Kolouri, S.; Pilly, P. K.; and McClelland, J. 2020. Generative continual concept learning. In *AAAI*, 5545–5552.
- Rostami, M.; Huber, D.; and Lu, T.-C. 2018. A crowdsourcing triage algorithm for geopolitical event forecasting. In *Proceedings of the 12th ACM Conference on Recommender Systems*, 377–381.
- Rostami, M.; Kolouri, S.; and Pilly, P. K. 2019. Complementary learning for overcoming catastrophic forgetting using experience replay. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, 3339–3345. AAAI Press.
- Rostami, M. 2019. *Learning Transferable Knowledge Through Embedding Spaces*. Ph.D. Dissertation, University of Pennsylvania.
- Roth, K.; Lucchi, A.; Nowozin, S.; and Hofmann, T. 2017. Stabilizing training of generative adversarial networks through regularization. In *Advances in Neural Information Processing Systems*, 2018–2028.
- Saito, K.; Watanabe, K.; Ushiku, Y.; and Harada, T. 2018. Maximum classifier discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3723–3732.
- Sankaranarayanan, S.; Balaji, Y.; Jain, A.; Nam Lim, S.; and Chellappa, R. 2018. Learning from synthetic data: Addressing domain shift for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3752–3761.
- Shin, H.; Lee, J. K.; Kim, J.; and Kim, J. 2017. Continual learning with deep generative replay. In *Advances in Neural Information Processing Systems*, 2990–2999.
- Shotton, J.; Johnson, M.; and Cipolla, R. 2008. Semantic texton forests for image categorization and segmentation. In *2008 IEEE conference on computer vision and pattern recognition*, 1–8. IEEE.
- Simonyan, K., and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Sun, B., and Saenko, K. 2016. Deep coral: Correlation alignment for deep domain adaptation. In *European conference on computer vision*, 443–450. Springer.
- Tighe, J., and Lazebnik, S. 2010. Superparsing: scalable nonparametric image parsing with superpixels. In *European conference on computer vision*, 352–365. Springer.
- Tsai, Y.-H.; Hung, W.-C.; Schuster, S.; Sohn, K.; Yang, M.-H.; and Chandraker, M. 2018. Learning to adapt structured output space for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7472–7481.
- Vu, T.-H.; Jain, H.; Bucher, M.; Cord, M.; and Pérez, P. 2019. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2517–2526.
- Wang, Q.; Zhang, L.; Bertinetto, L.; Hu, W.; and Torr, P. H. 2019. Fast online object tracking and segmentation: A unifying approach. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1328–1338.
- Wu, Z.; Han, X.; Lin, Y.-L.; Gokhan Uzunbas, M.; Goldstein, T.; Nam Lim, S.; and Davis, L. S. 2018. Dcan: Dual channel-wise alignment networks for unsupervised scene adaptation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 518–534.
- Wu, D. 2016. Online and offline domain adaptation for reducing bci calibration effort. *IEEE Transactions on Human-Machine Systems* 47(4):550–563.
- Xu, R.; Liu, P.; Wang, L.; Chen, C.; and Wang, J. 2020. Reliable weighted optimal transport for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4394–4403.
- Yang, Y., and Soatto, S. 2020. Fda: Fourier domain adaptation for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4085–4095.
- Yang, Y.; Lao, D.; Sundaramoorthi, G.; and Soatto, S. 2020. Phase consistent ecological domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9011–9020.
- Zeng, T.; Wu, B.; and Ji, S. 2017. Deepem3d: approaching human-level performance on 3d anisotropic em image segmentation. *Bioinformatics* 33(16):2555–2562.
- Zhang, Q.; Zhang, J.; Liu, W.; and Tao, D. 2019. Category anchor-guided unsupervised domain adaptation for semantic segmentation. In *Advances in Neural Information Processing Systems*, 435–445.
- Zhang, Y.; David, P.; and Gong, B. 2017. Curriculum domain adaptation for semantic segmentation of urban scenes. In *Proceedings of the IEEE International Conference on Computer Vision*, 2020–2030.
- Zhang, Z.; Fidler, S.; and Urtasun, R. 2016. Instance-level segmentation for autonomous driving with deep densely connected mrfs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 669–677.