

Overcoming Concept Shift in Domain-Aware Settings through Consolidated Internal Distributions

Mohammad Rostami and Aram Galstyan

Information Sciences Institute, University of Southern California
{mrostami, galstyan}@isi.edu

Abstract

We develop an algorithm to improve the performance of a pre-trained model under *concept shift* without retraining the model from scratch when only unannotated samples of initial concepts are accessible. We model this problem as a domain adaptation problem, where the source domain data is inaccessible during model adaptation. The core idea is based on consolidating the intermediate internal distribution, learned to represent the source domain data, after adapting the model. We provide theoretical analysis and conduct extensive experiments to demonstrate that the proposed method is effective.

Introduction

When a trained neural network is used for prediction, the assumption is that the testing samples are drawn from the training distribution. However, concept shift is a natural phenomenon when a model is used in different domains (Vorburger and Bernstein 2006). Similar to other machine learning techniques, deep learning is vulnerable with respect to these distributional shifts during the model execution time. Distributional shifts would lead to model performance degradation which makes model retraining, i.e., *model adaptation*, inevitable in order to make the model generalizable again in new domains. Adapting deep neural networks is conditioned on availability of massive labeled datasets which may not always be feasible due to prohibitive costs of manual data annotation. Finetuning is a common strategy to improve model generalization in the presence of domain shift. Despite being effective in reducing the effect of domain shift, finetuning still requires annotated data in the target domain.

Unsupervised domain adaptation (UDA) is a similar learning setting in which only *unlabeled data* is accessible in the target domain (Glorot, Bordes, and Bengio 2011) along with labeled data in the source domain. The goal in UDA is to leverage from the source domain data to adapt the deep neural network to generalize well in the target domain. An effective approach for UDA is to align distributions of both domains by mapping data into a latent domain-invariant space (Daumé III 2007). As a result, a classifier that is trained using the source labeled data features in this space will generalize well in the target domain. UDA methods model this latent space as

the output of a shared deep encoder. The encoder network is trained such that the source and the target domains share a similar distribution in its output. This training procedure has been implemented using either adversarial learning (He et al. 2016; Sankaranarayanan et al. 2018; Pei et al. 2018; Zhang et al. 2019; Long et al. 2018) or by directly minimizing the distance between the two distributions in the embedding (Long et al. 2015; Ganin and Lempitsky 2015; Long et al. 2017; Kang et al. 2019; Rostami et al. 2019).

Most existing unsupervised domain adaptation (UDA) algorithms consider a joint learning setting, where the model is trained jointly on both the target domain unlabeled data and the source domain labeled data. As a result, these algorithms cannot be used for sequential model adaptation. Note that although a few source-free domain adaptation algorithms have been developed recently (Kundu et al. 2020; Li et al. 2020b; Liu, Zhang, and Wang 2021), these methods use adversarial learning to memorize the source domain to generate source domain pseudo-data points for model retraining which necessitates using additional networks, rather directly adapting the base classifier network, which makes the model and the optimization problem more complex. Our goal is to adapt the classifier model to generalize well in the target domain using solely the target domain unlabeled data to tackle domain shift. This setting can be considered as an improvement over using an off-the-shelf pre-trained model when unlabeled data in a source domain is available. Our approach also relaxes the necessity of sharing training data between the two domains.

Contributions: we develop a sequential model adaptation algorithm which is based on learning a parametric internal distribution for the source domain distribution in a shared embedding space. This internally learned distribution is used to align the source and the target distributions. In order to adapt the model to work well the target domain, we draw samples from the estimated internal distribution and enforce the target domain to share the same distribution in the embedding space by minimizing the distance between the two distributions. We provide a theoretical justification for the proposed algorithm, by establishing an upperbound for the expected risk in the target domain, and demonstrating that our algorithm minimizes this upperbound. We conduct experiments on five benchmarks and observe that our algorithm compares favorably against SOTA UDA methods.

Background and Related Work

UDA is closely related to our learning setting. Several discrepancy measures have been used in the literature to align two distributions to address UDA. A group of methods match the first-order and the second order statistics of the source and the target domains. This includes methods that use the Maximum Mean Discrepancy (MMD) (Long et al. 2015, 2017) and correlation alignment (Sun and Saenko 2016). A more effective approach is to use a probability distance metric that captures distributional differences in higher order statistics. The Wasserstein distance (WD) (Courty et al. 2016; Bhushan Damodaran et al. 2018) is such an example that is also a suitable metric for deep learning due to having non-vanishing gradients. This property is helpful because deep learning optimization problems are usually solved using the first-order optimization methods that rely on the objective function non-vanishing gradients. Damodaran et al. (Bhushan Damodaran et al. 2018) used the WD for domain alignment in a UDA setting which led to considerable performance improvement compared to the methods that rely on matching only lower-order statistical moments (Long et al. 2015; Sun and Saenko 2016). We rely on the sliced Wasserstein distance (SWD) variant of WD (Kolouri, Rohde, and Hoffman 2018) for domain alignment which can be computed more efficiently due to its closed-form formulation and has been used for UDA successfully (Lee et al. 2019; Stan and Rostami 2021). More recently, secondary mechanism have been used to improve upon the primary alignment mechanism (Mirza et al. 2022; Wang et al. 2022). To mitigate negative transfer, some methods separate between transferable and domain-specific knowledge (Dong et al. 2020, 2021).

Most existing UDA methods use a strong assumption. It is assumed that the source and the target domain data are accessible simultaneously and the model is trained jointly on both datasets. Sequential model adaptation can be considered as a more challenging setting, where the source domain data is inaccessible after an initial model training for various reasons, including privacy and security concerns. It is different from continual learning (Chen and Liu 2018; Rostami, Isele, and Eaton 2020) in that not all tasks are labelled and addressing forgetting effects (Kirkpatrick et al. 2017; Rostami et al. 2020) is the primary concern. By addressing the sequential model adaptation setting, we can also address UDA when the source domain data cannot be shared (Rostami et al. 2018). This learning setting for domain adaptation has been explored for non-deep models (Dredze and Crammer 2008; Jain and Learned-Miller 2011; Wu 2016). However, these works address sequential model adaptation when the input distribution can be estimated with a parametric distribution and the base models have a small number of parameters. Hence, it is not trivial to extend the above works for the end-to-end training procedure of deep neural networks when deep learning is necessary for decent performance. Recently, adversarial learning has been used to address source-free DA (Kundu et al. 2020; Li et al. 2020b). These methods memorize the source domain using additional networks to generate a source pseudo-dataset that can be used as a surrogate for the source dataset. While these methods relax the need for access to the source data, generating realistic data points for complex

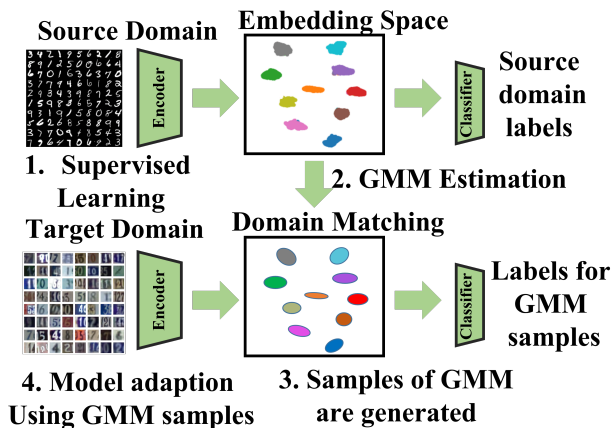


Figure 1: Architecture of the proposed method: (top) a neural network is trained on a source domain with annotated data. The internally learned distribution at the output space of the encoder is modeled via a GMM. (bottom) the model is updated to address concept shift in the target domain with annotated data by matching the target domain distribution with the GMM using a pseudo-dataset.

datasets is infeasible. Moreover, having an additional network makes optimization problem challenging. Quite differently, we rely on an internally learned distribution by a base model for the source domain to align the two source and target domain distributions indirectly in an embedding space to adapt the base classifier sequentially without using additional networks. We will improve upon a previous work that does not assume a source-free constraint (Rostami et al. 2019).

When a deep network is trained on a classification problem, the model would have decent performance only if the data points that belong to each class form a single cluster in an embedding space which is modeled by the network responses at higher layers. In other words, the source domain input distribution is transformed into a multimodal internal distribution, where each mode of the distribution encodes one of the classes. This internal distribution encodes the knowledge gained from supervised learning in the source domain. Domain shift occurs when changes in the input distribution lead to discrepancies between the transformed input distribution and the internally learned distribution. UDA can be addressed by aligning the target domain distribution with the learned source domain internal distribution. We estimate the internal distribution in the embedding via a parametric Gaussian mixture model (GMM). Previously, this idea has been used to address learning settings different from our setting (Oliveira, Minku, and Oliveira 2019; Wu, Hua, and Zhang 2005; Pfülb and Gepperth 2021) We develop an algorithm for addressing sequential model adaptation by enforcing the target domain to share the same internal distribution.

Problem Statement

Consider a source domain \mathcal{S} which consists of the distribution $p_{\mathcal{S}}(\mathbf{x})$ and the labeling function $f(\cdot) : \mathbb{R}^d \rightarrow \mathcal{Y} \subset \mathbb{R}^k$. Given a family of parametric functions f_{θ} , e.g., a deep neural network with learnable parameter θ , our

goal is to solve for an optimal model with minimum expected risk, i.e., $\mathbb{E}_{\mathbf{x}^s \sim p_S(\mathbf{x})}(\mathcal{L}(f(\mathbf{x}^s), f_\theta(\mathbf{x}^s)))$, where $\mathcal{L}(\cdot)$ is a proper loss function. To this end, we are given a labeled source dataset $\mathcal{D}_S = (\mathbf{X}_S, \mathbf{Y}_S)$, with $\mathbf{X}_S = [\mathbf{x}_1^s, \dots, \mathbf{x}_N^s] \in \mathcal{X} \subset \mathbb{R}^{d \times N}$ and $\mathbf{Y}_S = [\mathbf{y}_1^s, \dots, \mathbf{y}_N^s] \in \mathcal{Y} \subset \mathbb{R}^{k \times N}$, where the data points are drawn i.i.d. $\mathbf{x}_i^s \sim p_S(\mathbf{x})$. Given a sufficiently large dataset, we can solve for the optimal parameters using the standard empirical risk minimization (ERM): $\hat{\theta} = \arg \min_{\theta} \hat{e}_{\theta}(\mathbf{X}_S, \mathbf{Y}_S, \mathcal{L}) = \arg \min_{\theta} \sum_i \mathcal{L}(f_{\theta}(\mathbf{x}_i^s), \mathbf{y}_i^s)$. Consider that after training on the source domain, we encounter a target domain sequentially with an unlabeled dataset $\mathcal{D}_T = (\mathbf{X}_T)$, where $\mathbf{X}_T = [\mathbf{x}_1^t, \dots, \mathbf{x}_M^t] \in \mathbb{R}^{d \times M}$ and $\mathbf{x}_i^t \sim p_T(\mathbf{x})$. As a result, using ERM is not feasible in the target domain. We know a priori that the two domains share the same classes but also distributional discrepancy exists between the domains, i.e., $p_S \neq p_T$. This leads to poor generalization of $f_{\hat{\theta}}$ on the target domain due to domain shift. Our goal is to adapt the model using solely the dataset \mathcal{D}_T such that it generalizes well in the target domain in the absence of source dataset (see Figure 1).

In order to circumvent the challenge of distributional gap between the two domains, we can decompose the mapping $f_{\theta}(\cdot)$ into a deep encoder $\phi_v(\cdot) : \mathcal{X} \rightarrow \mathcal{Z} \subset \mathbb{R}^p$ and a classifier subnetwork $h_w(\cdot) : \mathcal{Z} \rightarrow \mathcal{Y}$ such that $f_{\theta} = h_w \circ \phi_v$, where $\theta = (\mathbf{w}, \mathbf{v})$. Here, \mathcal{Z} denotes an intermediate embedding space between the input space and the label space in which we assume that the classes have become geometrically separable. Given $\hat{\theta}$ and \mathcal{D}_T , if we adapt ϕ_v such that the discrepancy between the source and target distributions, i.e., the distance between $\phi(p_S(\mathbf{x}^s))$ and $\phi(p_T(\mathbf{x}^t))$, is minimized in the embedding space (making the embedding domain agnostic), then the classifier h_w will generalize well on the target domain, despite being trained using only the source labeled data points. Many prior classic UDA algorithms use this strategy but implement it by assuming that \mathcal{D}_S is always accessible. This makes computing the distance between the distributions $\phi(p_S(\mathbf{x}^s))$ and $\phi(p_T(\mathbf{x}^t))$ feasible and hence UDA reduces to selecting a proper probability metric and then solving a typical deep learning minimization problem (Long et al. 2015, 2017; Bhushan Damodaran et al. 2018; Courty et al. 2016; Bhushan Damodaran et al. 2018). The major challenge in the sequential model adaptation setting is that as opposed to the classical UDA framework, the term $\phi(p_S(\mathbf{x}^s))$ cannot be computed and we need a new approach to estimate this term using a suitable surrogate term.

Proposed Algorithmic Solution

We propose to solve the challenge of sequential model adaptation through aligning the source and the target distribution indirectly via an intermediate internally learned distribution in the embedding space. We set a softmax function at the output of the encoder just before passing data representations into the classifier subnetwork. As a result, the classifier can be assumed as a maximum *a posteriori* (MAP) estimator which assigns a membership probability distribution to any given data point. This assumption implies that when an optimal model is trained for the source domain, the encoder would transform the source distribution in the input space

into a multi-modal distribution $p_J(\mathbf{z})$ with k components in the embedding space, each corresponding to one class (see Figure 1, right). This internal distribution emerges because the classes should become separable in the embedding space as the result of learning for a generalizable model with a softmax layer. This internal distribution is a multimodal distribution in which each distribution mode would capture one of the classes. Performance degradation occurs when shifts in the input distribution violate boundaries between the classes in the embedding space. If we update the model such that the internal distribution remains stable in the target domain after model adaptation, i.e., the source and the target domain would share similar internal distributions in the embedding space, then the classifier subnetwork would generalize well in the target domain due to negligible domain gap.

The empirical version of the internal source distribution is encoded by the source data samples $\{(\phi_v(\mathbf{x}_i^s), \mathbf{y}_i^s)\}_{i=1}^N$. We consider that $p_J(\mathbf{z})$ is a GMM with k components:

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

where α_j denote mixture weights, i.e., prior probability for each class, $\boldsymbol{\mu}_j$ and $\boldsymbol{\Sigma}_j$ denote the mean and co-variance for each component. Since we have labeled data points, we can compute the GMM parameters using MAP estimates (Rostami 2021). Let \mathcal{S}_j denote the support set for class j in the training dataset, i.e., $\mathcal{S}_j = \{(\mathbf{x}_i^s, \mathbf{y}_i^s) \in \mathcal{D}_S | \arg \max \mathbf{y}_i^s = j\}$. Then, the MAP estimate for the GMM parameters is:

$$\begin{aligned} \hat{\alpha}_j &= \frac{|\mathcal{S}_j|}{N}, \quad \hat{\boldsymbol{\mu}}_j = \sum_{(\mathbf{x}_i^s, \mathbf{y}_i^s) \in \mathcal{S}_j} \frac{1}{|\mathcal{S}_j|} \phi_v(\mathbf{x}_i^s), \\ \hat{\boldsymbol{\Sigma}}_j &= \sum_{(\mathbf{x}_i^s, \mathbf{y}_i^s) \in \mathcal{S}_j} \frac{1}{|\mathcal{S}_j|} (\phi_v(\mathbf{x}_i^s) - \hat{\boldsymbol{\mu}}_j)^\top (\phi_v(\mathbf{x}_i^s) - \hat{\boldsymbol{\mu}}_j). \end{aligned} \quad (2)$$

Our major idea is to use this internal distributional estimate to circumvent the major challenge of sequential model adaptation. In order to adapt the model to work well for the target domain, we update the model such that the encoder subnetwork matches the target domain distribution with the internal distribution for the source domain in the embedding space, making the embedding domain-invariant. To this end, we can draw random samples from the internal distributional estimate and generate a labeled pseudo-dataset: $\hat{\mathcal{D}} = (\mathbf{Z}_P, \mathbf{Y}_P)$, where $\mathbf{Z}_P = [\mathbf{z}_1^p, \dots, \mathbf{z}_{N_p}^p] \in \mathbb{R}^{p \times N_p}$, $\mathbf{Y}_P = [\mathbf{y}_1^p, \dots, \mathbf{y}_{N_p}^p] \in \mathbb{R}^{k \times N_p}$, $\mathbf{z}_i^p \sim \hat{p}_J(\mathbf{z})$, and the labels are ascribed according to the classifier subnetwork prediction. To generate a clean pseudo-dataset, we also set a threshold τ and include only those generated samples for which the classifier prediction confidence is greater than τ . The sequential model adaptation problem then reduces to solving the alignment problem as the following:

$$\min_{\mathbf{v}, \mathbf{w}} \sum_{i=1}^N \mathcal{L}(h_w(\mathbf{z}_i^p), \mathbf{y}_i^p) + \lambda D(\phi_v(p_T(\mathbf{X}_T)), \hat{p}_J(\mathbf{Z}_P)), \quad (3)$$

where $D(\cdot, \cdot)$ denotes a probability distribution metric to measure the distributional discrepancy, and λ is a trade-off parameter between the two terms (see Figure 1, left).

Algorithm 1: SDAUP (λ, ITR)

```

1: Initial Training:
2: Input: source dataset  $\mathcal{D}_S = (\mathbf{X}_S, \mathbf{Y}_S)$ ,
3: Training on Source Domain:
4:  $\hat{\theta}_0 = (\hat{w}_0, \hat{v}_0) = \arg \min_{\theta} \sum_i \mathcal{L}(f_{\theta}(x_i^s), y_i^s)$ 
5: Internal Distribution Estimation:
6: Use Eq. (2) and estimate  $\alpha_j, \mu_j$ , and  $\Sigma_j$ 
7: Model Adaptation:
8: Input: target dataset  $\mathcal{D}_T = (\mathbf{X}_T)$ 
9: Pseudo-Dataset Generation:
10:  $\hat{\mathcal{D}}_P = (\mathbf{Z}_P, \mathbf{Y}_P) =$ 
11:  $([z_1^p, \dots, z_N^p], [y_1^p, \dots, y_N^p])$ , where:
12:  $z_i^p \sim \hat{p}_j(z), 1 \leq i \leq N_p$ 
13:  $y_i^p = \arg \max_j \{h_{\hat{w}_0}(z_i^p)\}$ 
14: for  $itr = 1, \dots, ITR$  do
15: draw data batches from  $\hat{\mathcal{D}}_T$  and  $\hat{\mathcal{D}}_P$ 
16: Update the model by solving Eq. (3)
17: end for

```

The first term in Eq. (3) is to ensure that the classifier continues to perform well on the internal distribution (note that the pseudo-dataset approximates this distribution). The second term is the domain alignment matching loss which enforces the target domain to share the internal distribution in the embedding space to make the embedding domain-invariant. The major remaining question is selection of the distribution metric. We choose SWD metric (Courty et al. 2016) to compute $D(\cdot, \cdot)$ due to its suitability for deep learning due to possessing non-vanishing gradients when two distributions have non-overlapping supports (Bonnotte 2013; Lee et al. 2019). SWD inherits this property from WD, yet the advantage of using SWD over WD is that SWD can be computed efficiently using a closed form solution. Additionally, empirical version of SWD can be computed using the samples that are drawn from the corresponding two distributions, as it is the case in Eq. (3). Hence, Eq. (3) can be solved using first-order optimization techniques (see the Appendix for more details on properties of SWD). Our proposed solution, named Sequential Model Adaptation Using Internal distribution (SMAUI), is presented and visualized in Algorithm 1. Figure 1 presents the high-level description of SMAUI.

Theoretical Analysis

We demonstrate that our algorithm optimizes an upperbound for the target domain expected risk. We treat the embedding space as the hypothesis space within the standard PAC-learning in our analysis. We consider the hypothesis class for PAC-learning to be the set of classifier subnetworks parameterized by w , $\mathcal{H} = \{h_w(\cdot) | h_w(\cdot) : \mathcal{Z} \rightarrow \mathbb{R}^k, v \in \mathbb{R}^V\}$. We denote the observed risk on the source and the target domains by e_S and e_T , respectively. Also, let $\hat{\mu}_S = \frac{1}{N} \sum_{n=1}^N \delta(\phi_v(x_n^s))$ and $\hat{\mu}_T = \frac{1}{M} \sum_{m=1}^M \delta(\phi_v(x_m^t))$ denote the empirical source and the empirical target distributions in the embedding space. These distributions are computed using the representations of the training data in the embedding space. Similarly, let $\hat{\mu}_P = \frac{1}{N_p} \sum_{q=1}^{N_p} \delta(z_n^q)$ de-

note the empirical internal distribution built using the pseudo-dataset. Moreover, let h_{w^*} denote the optimal model that minimizes the combined source and target risks $e_C(w^*)$, i.e. $w^* = \arg \min_w e_C(w) = \arg \min_w \{e_S + e_T\}$. In the presence of enough labeled target domain data, this is the best model that can be learned jointly on both domains. Finally, since we generate a clean pseudo-dataset, we can let $\tau = \mathbb{E}_{z \sim p_j(\hat{z})} (\mathcal{L}(h(z), h_{\hat{w}_0}(z)))$ as the expected risk of the optimal model for the source domain data on the generated pseudo-dataset using the GMM. We conclude:

Theorem 1: Consider that we generate a pseudo-dataset using the estimate for the internally learned distribution and update the model sequentially using the algorithm 1. Then, the following holds:

$$\begin{aligned}
e_T \leq & e_S + W(\hat{\mu}_S, \hat{\mu}_P) + W(\hat{\mu}_T, \hat{\mu}_P) + (1 - \tau) + e_{C'}(w^*) \\
& + \sqrt{(2 \log(\frac{1}{\xi})/\zeta)} \left(\sqrt{\frac{1}{N}} + \sqrt{\frac{1}{M}} + 2\sqrt{\frac{1}{N_p}} \right),
\end{aligned} \tag{4}$$

where $W(\cdot, \cdot)$ denotes the WD distance and ξ is a constant.

Proof: The complete proof is included in the Appendix.

We can use Theorem 1 to explain why our algorithm can adapt the model that is trained using the source domain to generalize well on the target domain. We can see that SMAUI algorithm minimizes the upperbound of the target domain risk as given in Eq. (4). The first three terms are minimized directly in our optimization. The source risk e_S is minimized through the initial training on the source domain prior to the model adaptation step. We minimize the second term in the upperbound of Eq. (4) by intentionally fitting a GMM distribution on the source domain distribution in the embedding. We note that this term is small if the source domain distribution can be fit well with a GMM distribution in reality. Hence, this is a constraint for applicability of our method. Note, however, all the parametric learning methods face a similar limitation and this is not particular to our algorithm. The third term is minimized directly because it is the second term in the objective function of Eq. (3). The fourth and fifth terms are not minimized by our algorithm, rather state conditions under which the algorithm would work. The term $(1 - \tau)$ is small if we set $\tau \approx 1$ to generate a clean pseudo-dataset and cancel out the outliers. We note that if τ is chosen very close to 1, then the pseudo-dataset samples would all be close to the mean and hence pseudo-dataset will not capture the higher-order moments of the internal distribution. The term $e_{C'}(w^*)$ is a constant term. This term will be small if the domains are related, i.e., share the same classes with the same consistent label-space encoding, and in presence of the target labeled data, the base model can be trained to work well on both domains. This means that aligning the distributions in the embedding must be a possibility for our algorithm to work. Finally, the last term in Eq. (4) is a constant term that merely states that in order to train a good model, we need sufficiently large source and target datasets and also we need to generate a large pseudo-dataset. We conclude assuming that the domains are related and applicability of the GMM estimation for the source distribution, SMAUI minimizes all the remaining terms in the upperbound of Eq. (4).

Experimental Validation

We compare our algorithm against several UDA algorithms and UDA tasks due to closeness of the UDA learning setting to the sequential model adaptation setting. Our code is provided at <https://github.com/rostami-m/SDAUP>.

Datasets: We validate our method on five standard UDA benchmarks and adapted them for sequential task learning: **Digit recognition tasks, Office-31 Dataset, ImageCLEF-DA Dataset, Office-Caltech Dataset, and VisDA-2017**. Details about these datasets are included in the Appendix.

Empirical Evaluation: Since sequential model adaptation is not a well-explored problem, we follow the UDA literature for evaluation due to the topic proximity. For this reason, we use the metrics and the features used in the UDA literature for fair comparison. We use the VGG16 network as the base model for the digit recognition tasks. The network is initialized with random weights. We use the Decaf6 features for the Office-Caltech tasks. For the Office-31 and ImageCLEF-DA datasets, we use the ResNet-50 backbone which is pre-trained on the ImageNet. For the VisDA-2017 dataset, we use the ResNet-101 backbone pre-trained on the ImageNet. A point of strength for our algorithm is that there are only two major algorithm-specific hyper-parameters and tuning them is not challenging. We set $\tau = 0.99$ and $\lambda = 10^{-3}$.

We report the average classification rate on the target domain and the standard deviation based on ten runs for each UDA task. We train the base model using the source labeled data. We report the performance of the model before adaptation as a baseline which also demonstrates the effect of domain shift. Then we adapt the model using the target unlabeled data using SMAUI algorithm and report the target domain performance. In our Tables, bold font denotes the best performance. The baseline performance before model adaptation is reported in the first row, then the UDA methods based on adversarial learning, then the UDA methods based on direct matching which are separated by a line in the middle, followed by our result in the last rows of the tables.

We compare our method against methods that are based on adversarial learning: GtA (Sankaranarayanan et al. 2018), DANN (Ganin et al. 2016), ADDA (Tzeng et al. 2017), MADA (Pei et al. 2018), SymNets (Zhang et al. 2019), CDAN (Long et al. 2018), MMAN (Ma, Zhang, and Xu 2019), and DANCE (Saito et al. 2020). We have also included methods based on direct distribution matching: DAN (Long et al. 2015), DRCN (Ghifary et al. 2016), CORAL (Sun, Feng, and Saenko 2016), RevGrad (Ganin and Lempitsky 2015), CAN (Kang et al. 2019), JAN (Long et al. 2017), WD-GRL (Shen et al. 2018), JDDA (Chen et al. 2019), ETD (Li et al. 2020a), and SRADA (Wang and Zhang 2020). Finally, we also compared against UDAwSD (Li et al. 2020b) and SHOT (Liang, Hu, and Feng 2020) which are source-free UDA methods. For each dataset, we include results of the works for which the original paper has used that dataset and reported the corresponding performance. For more details on the experimental setup, please refer to the Appendix.

Results

Results for the digit recognition tasks are reported in Table 1. Despite the sequential training regime, we observe

SMAUI outperforms the other methods in one of these tasks and its performance is quite competitive in the other two tasks. SMAUI leads to strong results compared to the joint UDA methods because stabilizing the internal distribution for both domain would mitigate domain shift. We also observe that performance of the methods based on direct probability matching has improved recently which might be because of using secondary mechanisms to improve generalization.

Table 2 summarizes the results for Office-31 dataset. We see in two of the tasks SMAUI leads to the best results and for the rest of the tasks is still competitive. Note, however, it seems there is no clear winner algorithm across all the tasks of this dataset. This maybe because it is known that some labels in this dataset are noisy and some images contain objects that belong to other classes (Bousmalis et al. 2016). The approach we use to align the two distributions is vulnerable with respect to this label pollution. We conclude that label pollution can degrade the performance of SMAUI.

Results for UDA tasks of the ImageCLEF-DA dataset are reported in Table 3. We see that although SMAUI does not use the source samples during model adaptation, it leads to a significant performance boost over the prior methods for this dataset. This may be because the Caltech-256, the ILSVRC 2012, and the Pascal VOC 2012 datasets have the equal size and the domains are balanced across the classes. As a result, matching the source and the target distributions to the same internal multi-modal distribution is more straightforward for the ImageCLEF-DA dataset. Since we rely on empirical distributions for alignment, balanced datasets represent true distributions better. We have reported results for the Office-Caltech dataset in Table 4. We note that SMAUI leads to a relatively competitive performance on the tasks and in terms of average performance outperforms the listed UDA methods in this dataset. We conclude that addition of the Caltech-256 dataset might have helped mitigating the effect of label pollution which has improved our performance.

Results for VisDA task is presented in Table 5. We observe a significant boost in performance for VisDA task. From inspecting Tables 1–5, we conclude that there is no single UDA method with the best performance on all the tasks. This is natural because these datasets are diverse in terms of difficulty, cross-domain gap, dataset size, label-pollution, etc, and any particular method may be more vulnerable in special cases. However, we note that although these UDA methods should serve an upperbound for SMAUI, SMAUI works reasonably well on all the domain adaptation tasks. These results demonstrate that although our motivation was to address sequential model adaptation, SMAUI can also be used as a standard joint training UDA algorithm with competitive results. Moreover, it can also be used as a source-free domain adaptation algorithm for preserving privacy (Stan and Rostami 2022a) and multi-source UDA (Stan and Rostami 2022b).

Analysis and Ablative Studies

To demonstrate that SMAUI algorithm implements what we anticipated and to provide a better intuition about its effectiveness, we have used the UMAP (McInnes et al. 2018) visualization tool to reduce the dimension of the data representations in the embedding space to two for 2D visualiza-

Method	$\mathcal{M} \rightarrow \mathcal{U}$	$\mathcal{U} \rightarrow \mathcal{M}$	$\mathcal{S} \rightarrow \mathcal{M}$	Method	$\mathcal{M} \rightarrow \mathcal{U}$	$\mathcal{U} \rightarrow \mathcal{M}$	$\mathcal{S} \rightarrow \mathcal{M}$
GtA	92.8 ± 0.9	90.8 ± 1.3	92.4 ± 0.9	CDAN	93.9	96.9	88.5
ADDA	89.4 ± 0.2	90.1 ± 0.8	76.0 ± 1.8	CyCADA	95.6 ± 0.2	96.5 ± 0.1	90.4 ± 0.4
SRADA	94.1	98.0	-	-	-	-	-
RevGrad	77.1 ± 1.8	73.0 ± 2.0	73.9	JDDA	-	97.0 ± 0.2	93.1 ± 0.2
DRCN	91.8 ± 0.1	73.7 ± 0.4	82.0 ± 0.2	OPDA	70.0	60.2	-
ETD	96.4 ± 0.3	96.3 ± 0.1	97.9 ± 0.4	MML	77.9	60.5	62.9
Source Only	90.1 ± 2.6	80.2 ± 5.7	67.3 ± 2.6	SMAUI	92.2 ± 0.5	98.2 ± 0.2	92.6 ± 1.0

Table 1: Classification accuracy for MINIST, USPS, and SVHN digit recognition datasets.

Method	$\mathcal{A} \rightarrow \mathcal{W}$	$\mathcal{D} \rightarrow \mathcal{W}$	$\mathcal{W} \rightarrow \mathcal{D}$	$\mathcal{A} \rightarrow \mathcal{D}$	$\mathcal{D} \rightarrow \mathcal{A}$	$\mathcal{W} \rightarrow \mathcal{A}$	Average
Source Only	68.4 ± 0.2	96.7 ± 0.1	99.3 ± 0.1	68.9 ± 0.2	62.5 ± 0.3	60.7 ± 0.3	75.6
GtA	89.5 ± 0.5	97.9 ± 0.3	99.8 ± 0.4	87.7 ± 0.5	72.8 ± 0.3	71.4 ± 0.4	86.5
DANN	82.0 ± 0.4	96.9 ± 0.2	99.1 ± 0.1	79.7 ± 0.4	68.2 ± 0.4	67.4 ± 0.5	82.2
ADDA	86.2 ± 0.5	96.2 ± 0.3	98.4 ± 0.3	77.8 ± 0.3	69.5 ± 0.4	68.9 ± 0.5	82.8
SymNets	90.8 ± 0.1	98.8 ± 0.3	100.0 ± 0	93.9 ± 0.5	74.6 ± 0.6	72.5 ± 0.5	88.4
MADA	82.0 ± 0.4	96.9 ± 0.2	99.1 ± 0.1	79.7 ± 0.4	68.2 ± 0.4	67.4 ± 0.5	82.2
CDAN	93.1 ± 0.2	98.2 ± 0.2	100.0 ± 0.0	89.8 ± 0.3	70.1 ± 0.4	68.0 ± 0.4	86.6
SRADA	95.2	98.6	100.0	91.7	74.5	73.7	89.0
UDAwSD	93.7 ± 0.2	98.5 ± 0.1	99.8 ± 0.2	92.7 ± 0.4	75.3 ± 0.5	77.8 ± 0.1	89.6
DAN	80.5 ± 0.4	97.1 ± 0.2	99.6 ± 0.1	78.6 ± 0.2	63.6 ± 0.3	62.8 ± 0.2	80.4
DRCN	72.6 ± 0.3	96.4 ± 0.1	99.2 ± 0.3	67.1 ± 0.3	56.0 ± 0.5	72.6 ± 0.3	77.7
RevGrad	82.0 ± 0.4	96.9 ± 0.2	99.1 ± 0.1	79.7 ± 0.4	68.2 ± 0.4	67.4 ± 0.5	82.2
CAN	94.5 ± 0.3	99.1 ± 0.2	99.8 ± 0.2	95.0 ± 0.3	78.0 ± 0.3	77.0 ± 0.3	90.6
JAN	85.4 ± 0.3	97.4 ± 0.2	99.8 ± 0.2	84.7 ± 0.3	68.6 ± 0.3	70.0 ± 0.4	85.8
JDDA	82.6 ± 0.4	95.2 ± 0.2	99.7 ± 0.0	79.8 ± 0.1	57.4 ± 0.0	66.7 ± 0.2	80.2
ETD	92.1	100.0	100.0	88.0	71.0	67.8	86.2
DANCE	88.6	97.5	100.0	89.4	69.5	68.2	85.5
SHOT	91.2	98.3	99.9	90.6	72.5	71.4	87.3
SMAUI	97.8 ± 2.1	95.6 ± 0.5	99.1 ± 0.3	97.8 ± 1.7	68.2 ± 4.5	71.7 ± 3.6	88.4

Table 2: Classification accuracy for UDA tasks for Office-31 dataset.

Method	$\mathcal{I} \rightarrow \mathcal{P}$	$\mathcal{P} \rightarrow \mathcal{I}$	$\mathcal{I} \rightarrow \mathcal{C}$	$\mathcal{C} \rightarrow \mathcal{I}$	$\mathcal{C} \rightarrow \mathcal{P}$	$\mathcal{P} \rightarrow \mathcal{C}$	Average
Source Only	74.8 ± 0.3	83.9 ± 0.1	91.5 ± 0.3	78.0 ± 0.2	65.5 ± 0.3	91.2 ± 0.3	80.8
DANN	82.0 ± 0.4	96.9 ± 0.2	99.1 ± 0.1	79.7 ± 0.4	68.2 ± 0.4	67.4 ± 0.5	82.2
SymNets	80.2 ± 0.3	93.6 ± 0.2	97.0 ± 0.3	93.4 ± 0.3	78.7 ± 0.3	96.4 ± 0.1	89.9
MADA	75.0 ± 0.3	87.9 ± 0.2	96.0 ± 0.3	88.8 ± 0.3	75.2 ± 0.2	92.2 ± 0.3	85.9
CDAN	76.7 ± 0.3	90.6 ± 0.3	97.0 ± 0.4	90.5 ± 0.4	74.5 ± 0.3	93.5 ± 0.4	87.1
SRADA	78.3	91.3	96.7	90.5	78.1	96.2	88.5
DAN	74.5 ± 0.4	82.2 ± 0.2	92.8 ± 0.2	86.3 ± 0.4	69.2 ± 0.4	89.8 ± 0.4	82.4
RevGrad	75.0 ± 0.6	86.0 ± 0.3	96.2 ± 0.4	87.0 ± 0.5	74.3 ± 0.5	91.5 ± 0.6	85.0
JAN	76.8 ± 0.4	88.0 ± 0.2	94.7 ± 0.2	89.5 ± 0.3	74.2 ± 0.3	91.7 ± 0.3	85.7
ETD	81.0	91.7	97.9	93.3	79.5	95.0	89.7
SMAUI	88.7 ± 1.2	99.5 ± 0.2	100 ± 0.0	94.9 ± 0.3	88.8 ± 0.9	99.8 ± 0.0	95.3

Table 3: Classification accuracy for UDA tasks for ImageCLEF-DA dataset.

Method	A→C	A→D	A→W	W→A	W→D	W→C	D→A	D→W	D→C	C→A	C→W	C→D	Average
Source Only	84.6	81.1	75.6	79.8	98.3	79.6	84.6	96.8	80.5	92.4	84.2	87.7	85.4
DANN	87.8	82.5	77.8	83.0	100	81.3	84.7	99.0	82.1	93.3	89.5	91.2	87.7
MMAN	88.7	97.5	96.6	94.2	100	89.4	94.3	99.3	87.9	93.7	98.3	98.1	94.6
RevGrad	85.7	89.2	90.8	93.8	98.7	86.9	90.6	98.3	83.7	92.8	88.1	87.9	88.9
DAN	84.1	91.7	91.8	92.1	100	81.2	90.0	98.5	80.3	92.0	90.6	89.3	90.1
CORAL	86.2	91.2	90.5	88.4	100	88.6	85.8	97.9	85.4	93.0	92.6	89.5	90.8
WDGR	87.0	93.7	89.5	93.7	100	89.4	91.7	97.9	90.2	93.5	91.6	94.7	92.7
SMAUI	99.9	100.0	96.7	96.8	100	94.4	84.8	93.4	91.8	98.8	80.4	91.4	94.0

Table 4: Performance comparison for UDA tasks of Office-Caltech dataset.

tion purpose. Figure 2 represents the testing data splits of the source and the target domains and samples of the internal distribution for the $\mathcal{S} \rightarrow \mathcal{M}$ digit recognition task. In this

figure, each point represents one data point and each color represents one of the ten digit classes. Comparing Figures 2a and 2b, we can see that the high-confidence internal GMM




Figure 2: UMAP visualization for the $\mathcal{S} \rightarrow \mathcal{M}$ task: (a) the source domain testing split, (b) the internal distribution samples, (c) the target domain testing split prior to adaptation, and (d) post adaptation. (Best viewed in color).

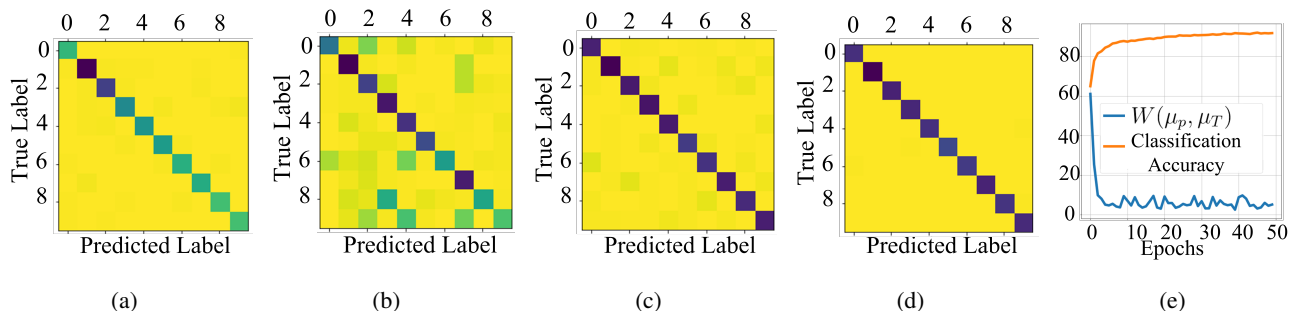


Figure 3: Confusion matrices for the $\mathcal{S} \rightarrow \mathcal{M}$ task: (a) the source domain (b) the target domain prior to adaptation, (c) the target domain after adaptation, (d) the target domain with a model trained using the target labeled data, (e) the cross-domain distribution & the test error vs epochs.

Task	JAN	DJT	GtA	SimNet	CDAN	DANCE	MCD	SMAUI
Syn. \rightarrow Re.	61.6	66.9	69.5	69.6	70.0	70.2	71.9	76.9 \pm 0.7

Table 5: Classification accuracy for the VisDA UDA task.

distribution samples approximate the source domain distribution reasonably well. Figure 2c denotes that the target domain samples are separable prior to adaptation to some extent due to domain similarity, but we can observe more regions with overlapped classes, i.e., less separability, that lead to performance degradation. This is due to distributional gap between the two domains. Figure 2d denotes that SMAUI algorithm has successfully aligned the source and the target distributions using the intermediate internal distribution. Figure 2 empirically confirms the results that we deduced from Theorem 1 because we observe that the upperbound minimization has led to improved generalization on the target domain.

For a class-level analysis of effect of model adaptation, Figures 3a–3d visualize the confusion matrices for the classifier with explanations in the caption. We can see in Figure 3b that domain shift causes confusion between digit classes that are in visually similar classes, e.g., digits “3” and “8” or digits “4”, “7”, and “9”. As seen in Figure 3c, the confusion is reduced for all classes using SMAUI algorithm. Comparing Figure 3c with Figures 3a and 3d, we see that the initial confusions between the classes in the source domain translate into the target domain, despite the fact that the target domain (\mathcal{M}) is an easier problem (Figures 3d). This observation is

predictable from Theorem 1 because the model performance on the target domain is upperbounded by the source domain error, even if it is a simpler problem. We again conclude that empirical results support our theorem.

We have also validated Theorem 1 in terms of effect of SMAUI on performance in the target domain in Figure 3e. We have plotted the test error versus optimization epochs during model adaptation process. We observe that as more training epochs are performed and the distributions are aligned progressively, i.e., domain discrepancy is minimized, the testing accuracy on the target domain accuracy constantly increases. This observation accords with theoretical prediction by Eq. (4) because as SMAUI gradually aligns the target domain distribution with the source distribution using the internal distribution, the upperbound in Eq. (4) becomes tighter.

Conclusions

We addressed the problem of model adaptation in a sequential task learning based on minimizing the distributional domain discrepancy in a shared embedding space using an intermediate internal distributional estimate. We estimate the source domain internal distribution using a GMM distribution. The internal distribution encodes what has been learned from the source domain. Hence, we can adapt the source-trained classifier using this distribution to improve generalization on the target domain. A future direction is equipping the algorithm with class-conditional alignment mechanisms.

References

- Bhushan Damodaran, B.; Kellenberger, B.; Flamary, R.; Tuia, D.; and Courty, N. 2018. Deepjdot: Deep joint distribution optimal transport for unsupervised domain adaptation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 447–463.
- Bonnotte, N. 2013. *Unidimensional and evolution methods for optimal transportation*. Ph.D. thesis, Paris 11.
- Bousmalis, K.; Trigeorgis, G.; Silberman, N.; Krishnan, D.; and Erhan, D. 2016. Domain separation networks. In *Advances in Neural Information Processing Systems*, 343–351.
- Chen, C.; Chen, Z.; Jiang, B.; and Jin, X. 2019. Joint domain alignment and discriminative feature learning for unsupervised deep domain adaptation. In *Proc. of the AAAI Conf. on Artificial Intelligence*, volume 33, 3296–3303.
- Chen, Z.; and Liu, B. 2018. Lifelong machine learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 12(3): 1–207.
- 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.
- Daumé III, H. 2007. Frustratingly Easy Domain Adaptation. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, 256–263.
- Dong, J.; Cong, Y.; Sun, G.; Fang, Z.; and Ding, Z. 2021. Where and how to transfer: knowledge aggregation-induced transferability perception for unsupervised domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Dong, J.; Cong, Y.; Sun, G.; Zhong, B.; and Xu, X. 2020. What can be transferred: Unsupervised domain adaptation for endoscopic lesions segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4023–4032.
- 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.
- Ganin, Y.; and Lempitsky, V. 2015. Unsupervised domain adaptation by backpropagation. In *Proceedings of ICML*, 1180–1189.
- Ganin, Y.; Ustinova, E.; Ajakan, H.; Germain, P.; Larochelle, H.; Laviolette, F.; Marchand, M.; and Lempitsky, V. 2016. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1): 2096–2030.
- Ghifary, M.; Kleijn, W. B.; Zhang, M.; Balduzzi, D.; and Li, W. 2016. Deep reconstruction-classification networks for unsupervised domain adaptation. In *European Conference on Computer Vision*, 597–613. Springer.
- Glorot, X.; Bordes, A.; and Bengio, Y. 2011. Domain adaptation for large-scale sentiment classification: a deep learning approach. In *Proceedings of the 28th International Conference on Machine Learning*, 513–520.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- Jain, V.; and Learned-Miller, E. 2011. Online domain adaptation of a pre-trained cascade of classifiers. In *Proceedings of CVPR*, 577–584.
- Kang, G.; Jiang, L.; Yang, Y.; and Hauptmann, A. G. 2019. Contrastive adaptation network for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4893–4902.
- Kirkpatrick, J.; Pascanu, R.; Rabinowitz, N.; Veness, J.; Desjardins, G.; Rusu, A. A.; Milan, K.; Quan, J.; Ramalho, T.; Grabska-Barwinska, A.; et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13): 3521–3526.
- Kolouri, S.; Rohde, G. K.; and Hoffman, H. 2018. Sliced Wasserstein Distance for Learning Gaussian Mixture Models. In *IEEE Conference on Computer Vision and Pattern Recognition*, 3427–.
- Kundu, J. N.; Venkat, N.; Babu, R. V.; et al. 2020. Universal Source-Free Domain Adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4544–4553.
- 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. 2020a. Enhanced transport distance for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition*, 13936–13944.
- Li, R.; Jiao, Q.; Cao, W.; Wong, H.-S.; and Wu, S. 2020b. Model Adaptation: Unsupervised Domain Adaptation without Source Data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9641–9650.
- Liang, J.; Hu, D.; and Feng, J. 2020. Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In *International Conference on Machine Learning*, 6028–6039. PMLR.
- Liu, Y.; Zhang, W.; and Wang, J. 2021. Source-free domain adaptation for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1215–1224.
- 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, M.; Cao, Z.; Wang, J.; and Jordan, M. I. 2018. Conditional adversarial domain adaptation. In *Advances in Neural Information Processing Systems*, 1640–1650.
- Long, M.; Zhu, H.; Wang, J.; and Jordan, M. I. 2017. Deep transfer learning with joint adaptation networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 2208–2217. JMLR. org.

- Ma, X.; Zhang, T.; and Xu, C. 2019. Deep multi-modality adversarial networks for unsupervised domain adaptation. *IEEE Transactions on Multimedia*, 21(9): 2419–2431.
- McInnes, L.; Healy, J.; Saul, N.; and Großberger, L. 2018. UMAP: Uniform Manifold Approximation and Projection. *Journal of Open Source Soft.*, 3(29): 861.
- Mirza, M. J.; Micorek, J.; Possegger, H.; and Bischof, H. 2022. The Norm Must Go On: Dynamic Unsupervised Domain Adaptation by Normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14765–14775.
- Oliveira, G. H.; Minku, L. L.; and Oliveira, A. L. 2019. GMM-VRD: A Gaussian Mixture model for dealing with virtual and real concept drifts. In *2019 International Joint Conference on Neural Networks (IJCNN)*, 1–8. IEEE.
- Pei, Z.; Cao, Z.; Long, M.; and Wang, J. 2018. Multi-adversarial domain adaptation. In *Proceedings Thirty-Second AAAI Conference on Artificial Intelligence*, 3934–3941.
- Pföhl, B.; and Gepperth, A. 2021. Overcoming catastrophic forgetting with gaussian mixture replay. In *2021 International Joint Conference on Neural Networks (IJCNN)*, 1–9. IEEE.
- Rostami, M. 2021. Lifelong domain adaptation via consolidated internal distribution. In *Advances in Neural Information Processing Systems*, volume 34, 11172–11183.
- Rostami, M.; Isele, D.; and Eaton, E. 2020. Using task descriptions in lifelong machine learning for improved performance and zero-shot transfer. *Journal of Artificial Intelligence Research*, 67: 673–704.
- 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.; Kim, K.; and Eaton, E. 2018. Multi-Agent Distributed Lifelong Learning for Collective Knowledge Acquisition. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, 712–720.
- Rostami, M.; Kolouri, S.; Pilly, P.; and McClelland, J. 2020. Generative continual concept learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 5545–5552.
- Saito, K.; Kim, D.; Sclaroff, S.; and Saenko, K. 2020. Universal domain adaptation through self supervision. *Advances in neural information processing systems*, 33: 16282–16292.
- Sankaranarayanan, S.; Balaji, Y.; Castillo, C. D.; and Chellappa, R. 2018. Generate to adapt: Aligning domains using generative adversarial networks. In *Proceedings of CVPR*, 8503–8512.
- Shen, J.; Qu, Y.; Zhang, W.; and Yu, Y. 2018. Wasserstein distance guided representation learning for domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Stan, S.; and Rostami, M. 2021. Unsupervised model adaptation for continual semantic segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 2593–2601.
- Stan, S.; and Rostami, M. 2022a. Domain Adaptation for the Segmentation of Confidential Medical Images. In *33rd British Machine Vision Conference 2022, BMVC 2022, London, UK, November 21–24, 2022*. BMVA Press.
- Stan, S.; and Rostami, M. 2022b. Secure Domain Adaptation with Multiple Sources. *Transactions on Machine Learning Research*.
- Sun, B.; Feng, J.; and Saenko, K. 2016. Return of frustratingly easy domain adaptation. In *Proceedings of the AAAI Conf. on Artificial Intelligence*, volume 30.
- Sun, B.; and Saenko, K. 2016. Deep coral: Correlation alignment for deep domain adaptation. In *European conference on computer vision*, 443–450. Springer.
- Tzeng, E.; Hoffman, J.; Saenko, K.; and Darrell, T. 2017. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7167–7176.
- Vorburger, P.; and Bernstein, A. 2006. Entropy-based concept shift detection. In *Sixth International Conference on Data Mining (ICDM'06)*, 1113–1118. IEEE.
- Wang, R.; Wu, Z.; Weng, Z.; Chen, J.; Qi, G.-J.; and Jiang, Y.-G. 2022. Cross-domain contrastive learning for unsupervised domain adaptation. *IEEE Transactions on Multimedia*.
- Wang, S.; and Zhang, L. 2020. Self-adaptive reweighted adversarial domain adaptation. *arXiv preprint arXiv:2006.00223*.
- Wu, D. 2016. Online and offline domain adaptation for reducing BCI calibration effort. *IEEE Transactions on Human-Machine Systems*, 47(4): 550–563.
- Wu, J.; Hua, X.-S.; and Zhang, B. 2005. Tracking concept drifting with Gaussian mixture model. In *Visual Communications and Image Processing 2005*, volume 5960, 1562–1570. SPIE.
- Zhang, Y.; Tang, H.; Jia, K.; and Tan, M. 2019. Domain-symmetric networks for adversarial domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5031–5040.