

An Extended Framework for Marginalized Domain Adaptation

Gabriela Csurka, Boris Chidlovski, Stéphane Clinchant and Sophia Michel
Xerox Research Center Europe (www.xrce.xerox.com)
6 chemin Maupertuis, 38240 Meylan, France
email:Firstname.Lastname@xrce.xerox.com

Abstract

We propose an extended framework for marginalized domain adaptation, aimed at addressing unsupervised, supervised and semi-supervised scenarios. We argue that the denoising principle should be extended to explicitly promote domain-invariant features as well as help the classification task. Therefore we propose to jointly learn the data auto-encoders and the target classifiers. First, in order to make the denoised features domain-invariant, we propose a domain regularization that may be either a domain prediction loss or a maximum mean discrepancy between the source and target data. The noise marginalization in this case is reduced to solving the linear matrix system $AX = B$ which has a closed-form solution. Second, in order to help the classification, we include a class regularization term. Adding this component reduces the learning problem to solving a Sylvester linear matrix equation $AX + BX = C$, for which an efficient iterative procedure exists as well. We did an extensive study to assess how these regularization terms improve the baseline performance in the three domain adaptation scenarios. We present experimental results on two image and one text benchmark datasets, conventionally used for validating domain adaptation methods. We report our findings and comparison with state-of-the-art methods.

1 Introduction

While huge volumes of unlabeled data are generated and made available in many domains, the cost of acquiring data labels remains high. Domain Adaptation (DA) problems arise each time we need to leverage labeled data in one or more related *source* domains, to learn a classifier for unseen or unlabeled data in a *target* domain. The domains are assumed to be related, but not identical and this *domain shift* occurs in multiple real-world applications, such as named entity recognition or opinion extraction across different text corpora, *etc.*

In this paper, we build on the domain adaptation work based on noise marginalization [9]. In deep learning, a denoising autoencoder (DA) learns a robust feature representation from training examples. In the case of domain adaptation, it takes unlabeled instances of both source and target data and learns a new feature representation by reconstructing the original features from their noised counterparts. A *marginalized denoising autoencoder* (MDA) marginalizes the noise at training time and thus does not require an optimization procedure using explicit data corruptions to learn the model parameters but computes the model in closed form. This makes MDAs scalable and computationally faster than the regular denoising autoencoders. The principle of noise marginalization has been successfully extended to learning with corrupted features [35], link prediction and multi-label learning [10], relational learning [11], collaborative filtering [31] and heterogeneous cross-domain learning [32, 44].

In this paper we extend the previous efforts and propose a larger framework for the marginalized domain adaptation. The *marginalized domain adaptation* refers to a denoising of source and target instances that explicitly makes their features *domain invariant* and eases the target prediction. We propose two extensions to the MDA. The first extension is a domain regularization, aimed at generating domain invariant features. Two families of such regularization are considered; one is based on the *domain prediction principle*, inspired by the adversarial learning of neural networks [24]; the second uses the maximum mean discrepancy (MMD) measure [30].

The second extension to the MDA is a class regularization; it allows to generate a classifier for target instances which can be learned jointly with the domain invariant representation.

Our framework works in *supervised, unsupervised and semi-supervised* settings, where the source data is completed with a few labeled target data, massive unlabeled target data or both, respectively. In all cases, the noise marginalization is maintained, thus ensuring the scalability and computational efficiency. We show how to jointly optimize the data denoising and the domain regularization, and how to marginalize the noise, which guarantees the closed-form solution and thus the computational efficiency. When the joint optimization is extended to the target prediction, the solution does not have a closed form, but is the solution of a Sylvester linear matrix equation $AX + XB = C$, for which efficient iterative methods can be used.

The remainder of the paper is organized as follows. In Section 2 we revise the prior art. Section 3 presents the components of the marginalized domain adaptation, including instance denoising, domain and class regularizations and target classifier learning. The joint loss minimization is detailed in Section 4. In Section 5 we describe two image and one text datasets we used, the experimental settings. We report the evaluation results which are grouped and analyzed by the three settings, namely, unsupervised, supervised and semi-supervised ones. Section 6 discusses the open questions and concludes the paper.

2 State of the art

Domain adaptation for text data has been studied for more than a decade, with applications in statistical machine translation, opinion mining, and document ranking [17, 45]. Most effective techniques include feature replication [16], pivot features [5] and finding topic models that are shared between source and target collections [12]. Domain adaptation has equally received a lot of attention in computer vision [1, 21, 22, 26, 27, 34, 40]. A considerable effort to systematize different domain adaptation and transfer learning techniques has been undertaken in [28, 37]. These studies distinguished three main categories of domain adaptation methods. The first category aims at correcting sampling bias [43]. The second category is in line with multi-task learning where a common predictor is learned for all domains, which makes it robust to domain shift [8]. The third family seeks to find a common representation for both source and target examples so that the classification task becomes easier [36]. Finally, an important research direction deals with the theory of domain adaptation, namely when adaptation can be effective and guaranteed with generalization bounds [3].

More recently, deep learning has been proposed as a generic solution to domain adaptation and transfer learning problems [13, 25, 33]. One successful method which aims to find common features between source and target collection relies on *denoising autoencoders*. In deep learning, a denoising autoencoder is a one-layer neural network trained to reconstruct input data from partial random corruption [42]. The denoisers can be stacked into multi-layered architectures where the weights are fine-tuned with costly back-propagation. Alternatively, outputs of intermediate layers can be used as input features to other learning algorithms. This learned feature representation was applied to domain adaptation [25], where stacked denoising autoencoders (SDA) achieved top performance in sentiment analysis tasks. The main drawback of SDAs is the long training time, and Chen et al. [9] proposed a variant of SDA where the random corruption is *marginalized out*. This crucial step yields a unique optimal solution which is computed in closed form and eliminates therefore the need for back-propagation. In addition, features learned with this approach lead to a classification accuracy comparable with SDAs, with a remarkable reduction of the training time [9].

More recently, deep learning architectures have demonstrated their ability to learn robust features and that good transfer performances could be obtained by just fine-tuning the neural network on the target task [13]. While such solutions perform relatively well on some tasks, the refinement may require a significant amount of new labeled data. More recent works proposed better strategies than fine-tuning, by designing deep architecture for the domain adaptation task. For example, Ganin *et al.* [24] has shown that adding a domain prediction task while learning the deep neural network leads to better domain-invariant feature representation. Long *et al.* [33] proposed to add a multi-layer adaptation regularizer, based on a multi-kernel maximum mean discrepancy (MMD). These

approaches obtained a significant performance gain which shows that transfer learning is not completely solved by fine-tuning and that transfer tasks should be addressed by appropriate deep learning representations.

3 Domain adaptation by feature denoising

We define a domain \mathcal{D} as the composition of a feature space $\mathcal{X} \subset \mathbb{R}^d$ and a label space \mathcal{Y} . A given task in the domain \mathcal{D} (classification, regression, ranking, etc.) is defined by a function $h : \mathcal{X} \rightarrow \mathcal{Y}$. In the domain adaptation setting, we assume working with a source domain \mathcal{D}^s represented by the feature matrix \mathbf{X}^s and the corresponding labels \mathbf{Y}^s , and a target domain \mathcal{D}^t with the features \mathbf{X}^t .

We distinguish among three scenarios of domain adaptation, depending on what is available in the target domain:

- *Unsupervised (US)* setting, where all available target instances are unlabeled. In this case, \mathbf{X}_l^t is empty and the labeled data, denoted by \mathbf{X}_l contain only the labeled source examples, $\mathbf{X}_l = \mathbf{X}^s$.
- *Supervised (SUP)* setting, where few labeled target instances \mathbf{X}_l^t are available at training time. In this case, we have $\mathbf{X}_l = [\mathbf{X}^s, \mathbf{X}_l^t]$.
- *Semi-supervised (SS)* setting, where massively unlabeled (\mathbf{X}_u^t) target data are available together with few labeled (\mathbf{X}_l^t) data at the training time.

In what follows we propose a framework to address all three scenarios in one uniform way. It aims at finding such a transformation of source and target data that minimizes the following loss function:

$$\mathcal{L} = \mathcal{L}_1 + \lambda \mathcal{L}_2 + \gamma \mathcal{L}_3, \quad (1)$$

where

- \mathcal{L}_1 is the data denoising loss on all data $\mathbf{X} = [\mathbf{X}^s, \mathbf{X}^t]$,
- \mathcal{L}_2 is the cross-domain classification loss on labeled data $\mathbf{X}_l = [\mathbf{X}^s, \mathbf{X}_l^t]$ with labels $\mathbf{Y}_l = [\mathbf{Y}^s, \mathbf{Y}_l^t]$,
- \mathcal{L}_3 is the domain regularization loss on source and target data \mathbf{X} .

Parameters λ and γ capture the trade-off between the three terms. All losses and parameters are described in the following subsections. Intuitively, minimizing the total loss (1) can help exploring the implicit dependencies between the data denoising, the domain regularization and the cross-domain classification.

In this paper we study the case when all three terms in (1) belong to the class of squared loss functions¹. More precisely:

- $\mathcal{L}_1 \equiv \mathcal{L}_1(\mathbf{X}, \mathbf{W})$ is the *instance denoising loss* under the dropout law; we minimize the square loss $\|\mathbf{X} - \tilde{\mathbf{X}}\mathbf{W}\|^2$ between the corrupted data $\tilde{\mathbf{X}}$ and the original data \mathbf{X} denoised with the linear transformation \mathbf{W} . This term is the core element of the marginalized denoising autoencoder (MDA) [9].
- $\mathcal{L}_2 \equiv \mathcal{L}_2(\mathbf{X}_l, \mathbf{Y}_l, \mathbf{W}, \mathbf{Z}_l)$ is the *class regularization loss*, aimed at learning a (multi-class) ridge classifier \mathbf{Z}_l from the available corrupted and denoised instances $\tilde{\mathbf{X}}_l\mathbf{W}$. The term is defined as $\|\mathbf{Y}_l - \tilde{\mathbf{X}}_l\mathbf{W}\mathbf{Z}_l\|^2$. It can be seen as a generalization of the Marginalized Corrupted Features (MCF) framework [35] with a square loss (the MCF corresponds to the case when $\mathbf{W} = \mathbf{I}_d$).

¹Other loss functions such as exponential, logistic, hinge losses are interesting to explore as well, but they are beyond of the scope of this paper.

- $\mathcal{L}_3 \equiv \mathcal{L}_3(\mathbf{X}^s, \mathbf{X}^t, \mathbf{W})$ is the *domain regularization loss* that expresses the discrepancy between the source and target domains. We explore two options for this term. One is based on the empirical maximum mean discrepancy (MMD), taking into account the class labels when available; the other uses a pre-trained domain classifier to regularize the total loss.

We follow the marginalized framework for optimizing the loss on corrupted data [9, 35], and minimize the loss expectation $\mathbb{E}[\mathcal{L}]$. To simplify the reading, we denote the expected loss values $\mathbb{E}[\mathcal{L}_i]$ also with \mathcal{L}_i .

By minimizing the marginalized expected loss (1), $\operatorname{argmin}_{\mathbf{W}, \mathbf{Z}_l} \mathcal{L}$, we obtain optimal solutions for the transformation matrix \mathbf{W} and classifier \mathbf{Z}_l . This can be achieved in two different ways, namely:

- \mathbf{W} and \mathbf{Z}_l are learned *sequentially*. In this case we first set $\lambda = 0$ and learn \mathbf{W} by minimizing $\mathcal{L}_1 + \gamma\mathcal{L}_3$. Then for the fixed \mathbf{W} we learn \mathbf{Z}_l from \mathcal{L}_2 . Except the supervised MMD in \mathcal{L}_3 , the learning of \mathbf{W} remains unsupervised, including the supervised and semi-supervised settings. The target labels in these cases are used at the second step, when the classifier \mathbf{Z}_l is learned.
- \mathbf{W} and \mathbf{Z}_l are learned *jointly*. In this case we iteratively optimize the joint loss with respect to \mathbf{W} and \mathbf{Z}_l . To initialize the iterative process, we set $\mathbf{W} = \mathbf{I}_d$ and minimize \mathcal{L} to compute \mathbf{Z}_l , then we fix \mathbf{Z}_l and optimize \mathcal{L} with respect to \mathbf{W} , and so on. The process is repeated until convergence. In practice we observed that the convergence is achieved after several iterations.

In the following subsections we describe in details and discuss each of the three loss terms, in Section 4 we address their different combinations.

3.1 Domain Instance Denoising

The first term we consider is the loss used by the *Marginalized Denoising Autoencoder* (MDA) [9]. Its basic idea is to reconstruct the input data from a partial random corruption [42] with a marginalization that yields optimal reconstruction weights in a closed form. The MDA loss can be written as

$$\mathcal{L}_1 \equiv \frac{1}{M} \sum_{m=1}^M \|\mathbf{X} - \tilde{\mathbf{X}}_m \mathbf{W}\|^2 + \omega \|\mathbf{W}\|^2, \quad (2)$$

where $\tilde{\mathbf{X}}_m \in \mathbb{R}^N \times \mathbb{R}^d$ is the m -th corrupted version of \mathbf{X} by random feature dropout with a probability p and $\omega \|\mathbf{W}\|^2$ is a regularization term. In order to avoid the explicit feature corruption and an iterative optimization, Chen *et al.* [9] showed that by considering the limit case $M \rightarrow \infty$, the weak law of large numbers allows to rewrite the loss \mathcal{L}_1 as its expectation and the optimal \mathbf{W} can be written as:

$$\mathbf{W} = (\mathbf{Q} + \omega \mathbf{I}_d)^{-1} \mathbf{P}, \quad (3)$$

where \mathbf{P} and \mathbf{Q} depend only on the covariance matrix $\mathbf{X}^\top \mathbf{X}$ and the noise level p .

One main advantage of the MDA is that it requires no label and therefore can be applied in all three settings *US*, *SUP* and *SS*. Note in the supervised case $\mathbf{X}_l = [\mathbf{X}^s, \mathbf{X}^t]$ includes only few target examples to learn \mathbf{W} .

3.2 Learning with marginalized corrupted features

Inspired by the Marginalized Corrupted Features (MCF) approach [35], we propose to marginalize the following loss:

$$\mathcal{L}_2 \equiv \frac{1}{M} \sum_{m=1}^M \|\mathbf{Y}_l - \tilde{\mathbf{X}}_{lm} \mathbf{W} \mathbf{Z}_l\|^2 + \delta \|\mathbf{Z}_l\|^2, \quad (4)$$

where $\mathbf{Z}_l \in \mathbb{R}^d \times \mathbb{R}^C$ is a multi-class classifier (each column corresponds to one of the C classes), $\mathbf{Y}_l \in \mathbb{R}^N \times \mathbb{R}^C$ is a label matrix, where $y_{nc} = 1$ if \mathbf{x}_n belongs to class $c = 1, \dots, C$, and -1 otherwise, and $\delta \|\mathbf{Z}_l\|^2$ is a regularization term. When $\mathbf{W} = \mathbf{I}_d$, we obtain the MCF baseline where

Table 1: All models of the extended framework, with the corresponding notations, losses and solutions.

Method	Loss	Closed-form solution for \mathbf{W}
S1	\mathcal{L}_1	$(\mathbf{Q} + \omega \mathbf{I}_d)^{-1} \mathbf{P}$
S1M	$\mathcal{L}_1 + \gamma \mathcal{L}_m$	$(\mathbf{Q} + \omega \mathbf{I}_d + \gamma \mathbf{M})^{-1} \mathbf{P}$
S1C	$\mathcal{L}_1 + \gamma \mathcal{L}_c$	$(\mathbf{Q} + \omega \mathbf{I}_d + \gamma \mathbf{M}_c)^{-1} \mathbf{P}$
S1D	$\mathcal{L}_1 + \gamma \mathcal{L}_d$	$\mathbf{Q}^{-1} (\mathbf{P} + \gamma (1-p) \mathbf{X}^\top \mathbf{Y}_\mathcal{T} \mathbf{Z}_\mathcal{D}^\top) (\mathbf{I}_d + \gamma \mathbf{Z}_\mathcal{D} \mathbf{Z}_\mathcal{D}^\top)^{-1}$
		\mathbf{W} as solution of $\mathbf{A}\mathbf{W} + \mathbf{W}\mathbf{B} = \mathbf{C}$
J12	$\hat{\mathcal{L}}_1 + \lambda \hat{\mathcal{L}}_2$	$\mathbf{A} = \mathbf{A}_{12} = \mathbf{Q}_l^{-1} (\mathbf{Q} + \omega \mathbf{I}_d), \mathbf{B} = \mathbf{B}_{12} = \lambda \mathbf{Z}_l \mathbf{Z}_l^\top,$ $\mathbf{C} = \mathbf{C}_{12} = \mathbf{Q}_l^{-1} (\mathbf{P} + \lambda (1-p) \mathbf{X}_l^\top \mathbf{Y}_l \mathbf{Z}_l^\top)$
J12M	$\mathcal{L}_1 + \lambda \mathcal{L}_2 + \gamma \mathcal{L}_m$	$\mathbf{A} = \mathbf{Q}_l^{-1} (\mathbf{Q} + \omega \mathbf{I}_d + \gamma \mathbf{M}), \mathbf{B} = \mathbf{B}_{12}, \mathbf{C} = \mathbf{C}_{12}$
J12C	$\mathcal{L}_1 + \lambda \mathcal{L}_2 + \gamma \mathcal{L}_c$	$\mathbf{A} = \mathbf{Q}_l^{-1} (\mathbf{Q} + \omega \mathbf{I}_d + \gamma \mathbf{M}_c), \mathbf{B} = \mathbf{B}_{12}, \mathbf{C} = \mathbf{C}_{12}$
J12D	$\mathcal{L}_1 + \lambda \mathcal{L}_2 + \gamma \mathcal{L}_d$	$\mathbf{A} = \mathbf{Q}_l^{-1} \mathbf{Q}, \mathbf{B} = \mathbf{A}_{12} (\mathbf{I}_d + \gamma \mathbf{Z}_\mathcal{D} \mathbf{Z}_\mathcal{D}^\top)^{-1},$ $\mathbf{C} = (\mathbf{C}_{12} + \mathbf{Q}_l^{-1} \gamma (1-p) \mathbf{X}^\top \mathbf{Y}_\mathcal{T} \mathbf{Z}_\mathcal{D}^\top) (\mathbf{I}_d + \gamma \mathbf{Z}_\mathcal{D} \mathbf{Z}_\mathcal{D}^\top)^{-1}$

the classifier is learned directly with the corrupted features. Moreover, when $p = 0$, we obtain the ridge classifier learned with the original features.

Given \mathbf{W} , the multi-class classifier \mathbf{Z}_l can be computed in closed form using the expected loss of (4):

$$\mathbf{Z}_l^* = (1-p) (\mathbf{W}^\top \mathbf{Q}_l \mathbf{W} + \delta \mathbf{I}_d)^{-1} \mathbf{W}^\top \mathbf{X}_l^\top \mathbf{Y}_l. \quad (5)$$

The computation of \mathbf{Z}_l^* requires the labeled data \mathbf{X}_l , that contain the source (*US*) or possibly target data (*SUP, SS*).

3.3 Reducing the discrepancy between domains

The domain regularization term \mathcal{L}_3 in (1) is aimed at bringing the target domain closer to the source domain, by minimizing the discrepancy between the domains. In the following, we explore three options for the term \mathcal{L}_3 , namely (1) a classical empirical MMD using the linear kernel, (2) its supervised version where the discrepancy is minimized between the class means and (3) a domain classifier $\mathbf{Z}_\mathcal{D}$ trained on the uncorrupted data to distinguish between the source and target data.

3.3.1 Reducing the MMD between domains

The minimization of maximum mean discrepancy (MMD) [6] between the source and target domains is the state of art approach widely used in the literature. It is often integrated in feature transformation learning [2, 36] or used as a regularizer for the cross-domain classifier learning [19, 33, 41]. The MMD is defined as a distance in the reproducing kernel Hilbert space (RKHS). In practice, its empirical version is used, as it can be written as $Tr(\mathbf{KN})$, where

$$\mathbf{K} = \begin{bmatrix} \mathbf{K}^{s,s} & \mathbf{K}^{s,t} \\ \mathbf{K}^{t,s} & \mathbf{K}^{t,t} \end{bmatrix} \text{ and } \mathbf{N} = \begin{bmatrix} \frac{1}{N_s^2} \mathbf{1}^{s,s} & \frac{1}{N_s N_t} \mathbf{1}^{s,t} \\ \frac{1}{N_s N_t} \mathbf{1}^{s,t} & \frac{1}{N_t^2} \mathbf{1}^{t,t} \end{bmatrix},$$

where $\mathbf{K}^{a,b}$ is the kernel distance matrix between all elements of \mathbf{X}^a and \mathbf{X}^b , $\mathbf{1}^{a,b}$ is a constant matrix of size $N_a \times N_b$ with all elements being equal to 1, and N_s, N_t are the number of source and target examples.

We integrate this loss in the total one \mathcal{L} by considering the MMD between the source and target data after the denoising. To be able to marginalize out the loss and to keep our solution linearly solvable, we use the MMD with the linear kernel. Intuitively, this corresponds to minimizing the distance between the two centroids of the source and target data after denoising. The corresponding

loss can be expressed as follows

$$\mathcal{L}_m \equiv \frac{1}{M} \sum_{m=1}^M \text{Tr}(\mathbf{W}^\top \tilde{\mathbf{X}}_m^\top \mathbf{N} \tilde{\mathbf{X}}_m \mathbf{W}). \quad (6)$$

After marginalizing the expected loss, we obtain $\mathbb{E} \mathcal{L}_3 = \text{Tr}(\mathbf{W}^\top \mathbf{M} \mathbf{W})$, where $\mathbf{M} = \mathbb{E}[\tilde{\mathbf{X}}^\top \mathbf{N} \tilde{\mathbf{X}}]^2$.

3.3.2 Reducing the MMD between the domain class means

The MMD requires no labels and can be computed between all available source all target instances. If we have labeled source and labeled target examples we can go one step further and modify the MMD to measure the distance between the means (centroids) of corresponding classes in the source and target domains [34]. The corresponding loss is the following

$$\mathcal{L}_c \equiv \frac{1}{M} \sum_{m=1}^M \text{Tr}(\mathbf{W}^\top \tilde{\mathbf{X}}_m^\top \mathbf{C} \tilde{\mathbf{X}}_m \mathbf{W}), \quad (7)$$

where:

$$\mathbf{C}_{ij} = \begin{cases} \frac{1}{N_s^c N_s^c} & \text{if } \mathbf{x}_i, \mathbf{x}_j \in \mathcal{D}^s \quad \& \quad y_i = y_j = c \\ \frac{1}{N_t^c N_t^c} & \text{if } \mathbf{x}_i, \mathbf{x}_j \in \mathcal{D}^t \quad \& \quad y_i = y_j = c \\ \frac{-1}{N_s^c N_t^c} & \text{if } \mathbf{x}_i \in \mathcal{D}^s, \mathbf{x}_j \in \mathcal{D}^t \quad \& \quad y_i = y_j = c \\ \frac{-1}{N_t^c N_s^c} & \text{if } \mathbf{x}_i \in \mathcal{D}^t, \mathbf{x}_j \in \mathcal{D}^s \quad \& \quad y_i = y_j = c \\ 0 & \text{otherwise,} \end{cases}$$

N_s^c is the number of source instances from the class c and N_t^c is the number of target instances from the class c . Note that the "otherwise" item above includes all cases where $y_i \neq y_j$ and where either \mathbf{x}_i or \mathbf{x}_j is unlabeled.

Similarly to \mathcal{L}_m in (6), we can marginalize out the expected loss \mathcal{L}_c and obtain $\mathbb{E} \mathcal{L}_c = \text{Tr}(\mathbf{W}^\top \mathbf{M}_c \mathbf{W})$, where $\mathbf{M}_c = \mathbb{E}[\tilde{\mathbf{X}}^\top \mathbf{C} \tilde{\mathbf{X}}]$.

3.3.3 Learning a domain classifier

As the last option of the domain regularization \mathcal{L}_3 in (1), we explore a loss based on the domain classifier [14]. Inspired by [23] who proposed to regularize intermediate layers in a deep learning model with a domain prediction task, [14] combines the domain prediction regularization with the MDA. We develop a similar regularization term for our extended framework, and use it jointly with the feature denoising term \mathcal{L}_1 and the class regularization \mathcal{L}_2 .

The idea of this domain regularization is to denoise data in such a way that pushes source data towards the target and hence allows the cross-domain classifier to perform better on the target. This is done by first learning a domain classifier $\mathbf{Z}_D \in \mathbb{R}^N$ using a regularized ridge classifier learned on the uncorrupted data. The regularized loss is defined as $\|\mathbf{Y}_D - \mathbf{X} \mathbf{Z}_D\|^2 + \alpha \|\mathbf{Z}_D\|^2$, where $\mathbf{Y}_D \in \mathbb{R}^N$ are the domain labels (-1 for source and +1 for target). The closed form solution is the following

$$\mathbf{Z}_D = (\mathbf{X}^\top \mathbf{X} + \alpha \mathbf{I}_d)^{-1} (\mathbf{X}^\top \mathbf{Y}_D). \quad (8)$$

Then the loss we consider in our unified framework is:

$$\mathcal{L}_d \equiv \frac{1}{M} \sum_{m=1}^M \|\mathbf{Y}_T - \tilde{\mathbf{X}}_m \mathbf{W} \mathbf{Z}_D\|^2, \quad (9)$$

where $\mathbf{Y}_T = \mathbf{1}^N$ is a vector containing only ones (all denoised instances should be predicted as target).

²All the derivations are included in the full version.

4 Minimizing the total loss

In the previous section we described three terms of the loss function \mathcal{L} . Now we discuss two main cases of minimizing the total loss. First, we discuss the sequential case, where we first learn \mathbf{W} using only the data without labels (\mathcal{L}_1 or $\mathcal{L}_1 + \gamma\mathcal{L}_3$), and then we learn the classifier \mathbf{Z}_l or any other classifier. Second, we describe the joint case where \mathbf{W} and \mathbf{Z}_l are learned jointly, by iteratively minimizing the total loss $\mathcal{L}_1 + \lambda\mathcal{L}_2 + \gamma\mathcal{L}_3$. In both cases we consider three options for the domain regularization \mathcal{L}_3 and discuss three domain adaptation scenarios, *US*, *SUP* and *SS*.

All mentioned combinations of the losses form different models; we denote them as follows. The sequential methods are prefixed by a character **S** followed by the indexes of the losses used. For example, when we learn \mathbf{W} with $\mathcal{L}_1 + \gamma\mathcal{L}_d$, the method is denoted **S1D**. When we learn \mathbf{W} and \mathbf{Z}_l jointly, we prefix the method by a character **J** followed by the loss indexes. For example, the method **J12C** means that we optimize $\mathcal{L}_1 + \lambda\mathcal{L}_2 + \gamma\mathcal{L}_c$. All the combinations are summarized in Table 1.

4.1 Sequential framework

In this case, we first obtain \mathbf{W} in an unsupervised manner and then learn a classifier \mathbf{Z}_l using the denoised features. To get \mathbf{W} , we set $\lambda = 0$ and minimize $\mathcal{L} = \mathcal{L}_1 + \gamma\mathcal{L}_3$. For each option of loss \mathcal{L}_3 , we get closed-form solutions for \mathbf{W} , denoted **S1**, **S1M**, and **S1D**. All the solutions are presented in Table 1. Any model can be deployed in three domain adaptation scenarios. In the *SUP* and *SS* cases, we additionally exploit the class labels using $\mathcal{L}_3 = \mathcal{L}_c$ (see **S1C** in Table 1).

Once \mathbf{W} is computed, we learn \mathbf{Z}_l using (5) or use any other classifier by feeding it with the denoised features $\mathbf{X}_l\mathbf{W}$. In the *US* case, the classifiers are learned with the denoised source features, while in the *SUP* and *SS* cases the classifier exploits additionally the labeled target data.

4.2 Joint framework

In this case, \mathbf{W} and \mathbf{Z}_l are learned *jointly*, by alternatively optimizing the total loss \mathcal{L} in variables \mathbf{W} and \mathbf{Z}_l . We start by initializing \mathbf{W} with \mathbf{I}_d and minimize the loss in \mathbf{Z}_l , then we fix \mathbf{Z}_l and compute \mathbf{W} , and so on. The process is repeated until convergence for a certain threshold.

The partial derivatives of \mathcal{L} with respect to \mathbf{Z}_l depend on \mathcal{L}_2 only, this makes solution (5) always valid. The partial derivatives with respect to \mathbf{W} can be written as a Sylvester linear matrix equation $\mathbf{AW} + \mathbf{WB} = \mathbf{C}$, that we solve using the Bartels-Stewart algorithm [39]. Depending on which loss is used as \mathcal{L}_3 , we obtain three versions of the Sylvester equation, denoted **J12**, **J12M** and **J12D** and detailed in Table 1.

Note that for **J12D** we do not use the regularizer term $\omega\|\mathbf{W}\|^2$ in the loss, in order to be able to reduce the partial derivatives to solving a Sylvester equation. Furthermore, in the *SUP* case, as $\mathbf{Q} = \mathbf{Q}_l$ and $\mathbf{P} = \mathbf{P}_l$, if we remove $\omega\|\mathbf{W}\|^2$ from **J12** we obtain a closed form solution³: $\mathbf{W} = \mathbf{Q}_l^{-1}(\mathbf{P}_l + \lambda(1-p)\mathbf{X}_l^T\mathbf{Y}_l\mathbf{Z}_l^T)(\mathbf{I}_d + \lambda\mathbf{Z}_l\mathbf{Z}_l^T)$.

5 Experimental Results

In the experimental section, we pursue a number of important goals. First, we want to assess all models proposed in the previous sections in three domain adaptation scenarios. Second, we evaluate the impact of the domain regularization \mathcal{L}_3 and the class regularization \mathcal{L}_2 on the denoising matrix \mathbf{W} and target classifier \mathbf{Z}_l . Finally, we report the performance of the sequential and joint learning cases, we analyze our results and compare them to the state-of-the art.

This section is organized as follows. In Section 5.1 we briefly describe three datasets used in the experiments, then Section 5.2 describes the experimental setting, including the methods and parameters used. In Section 5.3 we compare the sequential and joint models with different loss

³In the experiments, results with the term (by solving a Sylvester equation) and without (a closed form solution) are similar, but the latter case is much faster.

combinations in the *US*, *SUP* and *SS* settings, for all datasets. Finally, in Section 5.4 we compare our best performing models with the state-of-the art.

5.1 Datasets

All experiments are conducted on three domain adaptation datasets, well known in image processing and sentiment analysis communities.

OFF31 and OC10. Two most popular datasets used to compare visual domain adaptation methods are the Office31 dataset [38] (**OFF31**) and the Office+Caltech10 [27] (**OC10**). The former consists of three domains: Amazon (*A*), dsLR (*D*) and Webcam (*W*) with images of 31 products (classes). The latter contains 10 of the 31 classes for the same domains and includes an extra domain from the Caltech collection. For all images, we use the Decaf TF6 features [18] with the *full training* protocol [26] where all source data is used for training.

AMT. A standard dataset for textual domain adaptation is the Amazon dataset of text products reviews; it includes four domains: Books (*b*), DVD (*d*), Kitchen (*k*) and Electronics (*e*) preprocessed by Blitzer et al. [4]. Reviews are considered as positive if they have more than 3 stars, and negative otherwise. We adopt the experimental setting of [24] where documents are represented by a bag of uni-grams and bi-grams with the 5000 most frequent common words selected and a tf-idf weighting scheme.

5.2 Methods and settings

Most models proposed in the previous sections produce the denoising matrix \mathbf{W} and the target classifier \mathbf{Z}_l , which can be inferred sequentially or jointly. The sequential approach learns first the matrix \mathbf{W} and then a target classifier \mathbf{Z}_l for a fixed \mathbf{W} . The joint approach learns \mathbf{W} and \mathbf{Z}_l jointly⁴ by alternating the updates of \mathbf{W} and \mathbf{Z}_l using the same loss \mathcal{L} .

For each dataset, we consider all domains and take all possible source-target pairs as domain adaptation tasks. For example, for **OFF31** with three domains, we consider six following adaptation tasks: $D \rightarrow A$, $W \rightarrow A$, $A \rightarrow D$, $W \rightarrow D$, $A \rightarrow W$, and $D \rightarrow W$. Similarly, for **OC10** and **AMT** which include four domains each, we consider 12 different source-target pairs as adaptation tasks. To compare the different models, we report **averaged** accuracies over all adaptation tasks for a given dataset. In the supervised (*SUP*) and semi-supervised (*SS*) scenarios, we randomly select 3 target instances per class to form the target training set, and use the rest for the test.

In addition to the sequential and joint model learning, we include in our framework two standard classifiers, the nearest neighbor (NN) classifier and the Domain Specific Class Means (DSCM) classifier [15]⁵. They represent a valuable alternative to the ridge target \mathbf{Z}_l classifier. Fedded with the denoised instances, obtained with matrix \mathbf{W} , these classifiers help assess the value of our framework for domain adaptation tasks.

To ensure the fair comparison of all methods, we run all experiments with a unique parameter setting. All selected parameter values are explained below:

- we set $\lambda = \gamma = 1$ as term weights, this corresponds to the equal weighting in the global loss (1);
- we set $\omega = 10^{-2}$ as \mathbf{W} norm regularization in (3) (as in [9]);
- we set the dropping noise level for \mathbf{P} and \mathbf{Q} in (3): $p = 0.5$ for image datasets and $p = 0.9$ for AMT, as text representations are initially very sparse and a higher noise level is required;
- we set $\alpha = \delta = 1$ for the classifier regularization terms in (4) and (8);

⁴In this case while we do not have the guarantee a global minimum, we observed in general to quick convergence of the loss (only a few iterations).

⁵In DSCM, a target test example is assigned to a class by using a soft-max distance to the domain specific class means. The main reason for choosing these classifiers is that NN is related to retrieval (equivalent to precision at 1) and NCM with clustering, so the impact of \mathbf{W} on these two extra tasks is indirectly assessed.

Table 2: Unsupervised domain adaptation. Bold indicates the best result per dataset, underlined are improvements over **S1**.

	OC10			OFF31			AMT
	nn	dscm	Z_l	nn	dscm	Z_l	Z_l
BL	84.5	78.7	82.6	65.2	61.6	62.8	76.7
S1	85.2	79.9	84	65.9	62.1	66.5	81.6
S1D	84.9	<u>80.9</u>	85.6	65.8	<u>62.6</u>	67.1	82.2
J12	84.4	<u>82.3</u>	82.3	67.7	65	65	76.0
J12D	82.8	<u>83.2</u>	83.4	64.4	<u>65.2</u>	65.5	76.0

- we consider a single layer MDA only, to enable a fair comparison of different loss combinations and learning methods⁶.

To reveal all strong and weak points of our framework, we compare all models and classifiers with two baselines. The first baseline is denoted **BL** and provided by the classifier learned on the original features, without denoising. The classifier Z_l is learned using (5) with $p = 0$ and hence $Q = S$. The second baseline refers to the original MDA method [9] and corresponds to **S1** method in our framework. It uses the loss \mathcal{L}_1 to build W in an unsupervised manner and learns a classifier on the denoised features.

In the following subsections, we compare the methods of our framework to the baselines, for all domain adaptation settings and all datasets.

5.3 Comparing domain adaptation methods

5.3.1 Unsupervised Setting

In this case, labeled data are available from the source domain only, $X_l = X^s$. We compare the different models described in Section 4 (see the summary in Table 1). For each method and each domain adaptation task, we learn the Z_l classifier and the NN and DSCM classifiers applied to the denoised features. The accuracy results are averaged per dataset and reported in Table 2.

The first observation is that all **BL** baselines get improved by the MDA (**S1**). Second, on the text data (AMT), the Z_l classifier performs the best⁷, Third, on the image collections the picture is more complex, with the NN showing the globally highest accuracy. If we compare the baseline **S1** and extended methods, we can conclude the following. In *sequential framework*, the domain regularization **S1D** often improves over **S1** for the linear classifier Z_l and the DSCM, but not of the NN. In the *joint framework*, the class regularization \mathcal{L}_2 degrades the linear classifier Z_l results but improves the DSCM classifier.

To conclude, in the unsupervised domain adaptation, the best strategy may depend on the data type, regularization and classification method. If we compare methods by averaging their results over the rows in Table 2, **S1D** with the average 75.7% appears as the best strategy over all classifiers and datasets, followed by **J12**. Both outperform the baseline **S1** with the average 74.5%. This suggests that the best strategy is to learn the denoising W with the domain regularization and then to learn any classifier from denoised features.

5.3.2 Supervised Setting

In this case we have $X = X_l = [X^s, X_l^t]$, $Q_l = Q$ and unlabeled target data is unavailable. In Table 3 we report the evaluation results for the different models using Z_l , NN and DSCM classifiers. It is easy to note that all models in the supervised case behave quite different from the unsupervised one.

⁶Similarly to the stacked MDA, our framework can be extended to a stacked version, where the denoised features XW of one layer become the input of a next layer and nonlinear functions such as tangent hyperbolic are applied between the two.

⁷NN and DSCM results are not included in the table, due to space limitation.

Table 3: Supervised domain adaptation. Bold indicates best result per dataset, underlined are improvements over **S1**.

	OC10			OFF31			AMT
	nn	dscm	Z_l	nn	dscm	Z_l	Z_l
BL	90.8	91.6	88.1	77.6	76.6	70.3	77.3
S1	90.6	92.1	89.7	77.2	76.2	72.1	78.4
S1D	89.8	91.6	89	76.7	75.9	71.9	74.3
S1C	<u>90.7</u>	92	<u>90.1</u>	76.8	<u>77</u>	<u>72.7</u>	78.4
J12	<u>90.7</u>	92	<u>90.1</u>	<u>79.1</u>	<u>80.7</u>	<u>78.0</u>	71.1
J12D	<u>90.7</u>	90.1	<u>89.9</u>	<u>79.1</u>	<u>80.6</u>	<u>78.0</u>	71.1
J12C	92.1	92.1	<u>91.1</u>	80.8	<u>80.7</u>	<u>78.1</u>	73.9

Table 4: Semi-supervised domain adaptation on OC10, OFF31 and AMT. Bold indicates the best result per dataset, underline indicates the improvement over **L1**.

	OC10			OFF31			AMT
	nn	dscm	Z_l	nn	dscm	Z_l	Z_l
BL	90.8	91.6	88.1	77.6	76.6	70.3	77.3
S1	91.1	91.9	90.4	78.4	76.9	74.6	82.4
S1D	<u>91.2</u>	91.9	<u>90.7</u>	78.1	<u>77.3</u>	74.5	82.7
S1C	<u>91.3</u>	<u>92.1</u>	90.4	<u>78.5</u>	<u>77.6</u>	74.5	82.4
J12	92.3	91.8	89.1	80.8	<u>80.5</u>	74.7	76.6
J12D	87.9	89.8	89.1	80.8	<u>80.4</u>	<u>74.7</u>	76.6
J12C	<u>92.1</u>	91.6	89.1	80.8	<u>80.2</u>	74.9	76.6

Sequential framework. The domain regularization \mathcal{L}_d seems to harm the baseline performance. One possible explanation is that only few target examples are available, comparing to the much larger source set. In contrast, adding the regularization \mathcal{L}_c that exploits the class labels improves in many cases the results.

Joint framework. Adding \mathcal{L}_2 improves the results in general, except when using Z_l on AMT dataset. While **J12** and **J1D** perform rather similarly, **J12C** outperforms them globally on all datasets and classifiers. Note however, that the jointly learned Z_l performs less well than using the jointly learned \mathbf{W} to denoise the data on which a standard classifier is learned. This is the case when the framework implicitly combines the benefit of the joint learning with any classifier; such a strategy seems to work the best.

5.3.3 Semi-supervised Setting

In this case, a large set of unlabeled target data is available together with a small set of labeled target data. We explore if the proposed methods are able to take advantage of the two. Table 4 shows the results of different models on the three datasets, and them to the baselines.

Sequential framework. Compared to the supervised case, having more target examples makes the domain regularization \mathcal{L}_3 to either have no effect or slightly improve the results.

Joint framework. Adding class regularization \mathcal{L}_2 either improves or does not change the results, except learning Z_l for AMT, where a significant drop is observed. Adding \mathcal{L}_d often decreases the performance, while adding \mathcal{L}_c is less harmful. In general, in the semi-supervised case **J12** is the best strategy for the image datasets. For the text dataset, like in the unsupervised case, using **S1D** with \mathbf{W} , followed by learning a classifier on the denoised features seems to be a better option.

5.4 Comparison with the state of the art

We complete the experimental section by comparing our results to the state of art results, in the unsupervised and semi-supervised cases⁸.

⁸We exclude the supervised scenario as rarely addressed in the literature.

OC10. Using the FC6 features, our **S1M** accuracy (86.5%) in Table 2 outperforms⁹ the domain adaptive SVM [7] (70.3%) and domain adaptation via auxiliary classifiers [20] (84%), and slightly underperforms the more complex JDA [34] (87.5%) and TTM [21] (87.5%) methods.

OFF31. Table 5 compares our **S1D+Z_l** and **S12+NN** results with the feature transformation methods, using the same deep FC6 features, and with recent deep learning methods. It reports the results for six domain adaptation tasks on **OFF31** and their average. We can see that our methods behave similarly or better than feature transformation methods, but below the deep adaptation ones. Designed to be fast, our methods solve a few linear systems at training time and a simple matrix multiplication at test time, while the deep architectures have thousands of parameters to be tuned with the back propagation and GPU computations at the training time.

AMT. Our **L1D+Z_l** results (82.2%, see Table 2) is similar to the state-of-the art results obtained with the Domain-Adversarial Neural Networks (DANN) [24], despite the fact that DANN uses a 5 layer stacked MDA where the 5 outputs are concatenated with input to generate 30,000 dimensional features, on which the network is trained.

Table 5: Unsupervised adaptation on the **OFF31**: (top) feature transformation methods with FC6 features; (middle) our results with using FC6 features; (bottom) the deep domain adaptation methods. Best results per task are in bold. Underline indicates best results on FC6 features.

	A→W	D→W	W→D	D→A	W→A	A→D	Mean
GFK+SVM [40]	37.8	81.0	86.9	34.8	31.4	44.8	49.1
SA+SVM [40]	35	74.5	81.5	32.3	30.1	41.3	49.1
TCA+SVM [40]	36.8	82.3	84.1	32.9	28.9	40.6	50.9
CORAL+SVM [40]	48.4	<u>96.5</u>	99.2	44.4	41.9	53.7	64.0
L12 + NN	50.5	94.3	98.0	46.4	42.8	53.8	64.3
L1D + RDG	<u>53.9</u>	91.2	95.1	<u>47.3</u>	49.0	<u>55.5</u>	<u>65.3</u>
DDC [41]	59.4	92.5	91.7	52.1	52.2	64.4	68.7
RGrad [24]	67.3	94.0	93.7	-	-	-	-
DAN [33]	68.5	96.0	99	54	53.1	67.0	72.9

Concerning the semi-supervised scenario, it is much less used and most papers report results with SURF BOV features and the sampling protocol [38, 27]. We therefore tested our methods on **OC10** with **L12C+DSCM** and BOV features averaged over the 20 random samples; and we get an accuracy of 55.8% that is above most state of art results, including GFK [27] (48.6%), SA [22] (53.6%), MMDT [29] (52.5%).

6 Conclusion

We proposed an extended framework for domain adaptation, where the state-of-the-art marginalized denoising autoencoder is extended with domain and class regularization terms, aimed at addressing unsupervised, supervised and semi-supervised scenarios. The domain regularization drives the denoising of both source and target data toward domain invariant features. Two families of domain regularization, based on domain prediction and the maximum mean discrepancy, are proposed. The class regularization learns a cross-domain classifier jointly with the common representation learning. In all cases, the models can be reduced to solving a linear matrix equation or its Sylvester version, for which efficient algorithms exist.

We presented the results of an extensive set of experiments on two image and one text benchmark datasets, where the proposed framework is tested in different settings. We showed that adding the new regularization terms allow to outperform the baselines and help design best performing strategies for each adaptation scenarios and data types. Compared to the state of art we showed that despite of their speed and relatively low cost, our models yield comparable or better results than existing

⁹We report results from [21].

feature transformation methods but below highly expensive non-linear methods with additional data processing such as the landmark selection or those using deep architectures requiring costly operations both at training and at test time. Furthermore, similarly to the stacked MDA framework, we can easily stack several layers together with only forward learning, where the denoised features of the previous layer become the input of a new layer and nonlinear functions can be applied between the layers.

References

- [1] Rahaf Aljundi, R'emi Emonet, Damien Muselet, and Marc Sebban. Landmarks-based kernelized subspace alignment for unsupervised domain adaptation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [2] Mahsa Baktashmotlagh, Mehrtash Harandi, Brian Lovell, and Mathieu Salzmann. Unsupervised domain adaptation by domain invariant projection. In *IEEE International Conference on Computer Vision (ICCV)*, 2013.
- [3] Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for domain adaptation. In *Annual Conference on Neural Information Processing Systems (NIPS)*, 2006.
- [4] John Blitzer, Sham Kakade, and Dean P. Foster. Domain adaptation with coupled subspaces. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2011.
- [5] John Blitzer, Ryan McDonald, and Fernando Pereira. Domain adaptation with structural correspondence learning. In *International Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2006.
- [6] Karsten M. Borgwardt, Arthur Gretton, Malte J. Rasch, Hans-Peter Kriegel, Bernhard Schölkopf, and Alex J. Smola. Integrating structured biological data by kernel maximum mean discrepancy. *Bioinformatics*, 22:49–57, 2006.
- [7] Lorenzo Bruzzone and Mattia Marconcini. Domain adaptation problems: A dasvm classification technique and a circular validation strategy. *Transactions of Pattern Recognition and Machine Analyses (PAMI)*, 32:770–787, 2010.
- [8] Minmin Chen, Kilian Q. Weinberger, and John Blitzer. Co-training for domain adaptation. In *Annual Conference on Neural Information Processing Systems (NIPS)*, 2011.
- [9] Minmin Chen, Zhixiang Xu, Kilian Q. Weinberger, and Fei Sha. Marginalized denoising autoencoders for domain adaptation. In *International Conference on Machine Learning (ICML)*, 2012.
- [10] Zheng Chen, Minmin Chen, Kilian Q. Weinberger, and Weixiong Zhang. Marginalized denoising for link prediction and multi-label learning. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2015.
- [11] Zheng Chen and Weixiong Zhang. A marginalized denoising method for link prediction in relational data. In *SIAM International Conference on Data Mining (SDM)*, 2014.
- [12] Zhiyuan Chen and Bing Liu. Topic modeling using topics from many domains, lifelong learning and big data. In *International Conference on Machine Learning (ICML)*, 2014.
- [13] Sumit Chopra, Suhrid Balakrishnan, and Raghuraman Gopalan. DLID: Deep learning for domain adaptation by interpolating between domains. In *ICML Workshop on Challenges in Representation Learning (WREPL)*, 2013.

- [14] Stéphane Clinchant, Gabriela Csurka, and Boris Chidlovskii. A domain adaptation regularization for denoising autoencoders. In *Annual Meeting of the Association for Computational Linguistics(ACL)*, 2016.
- [15] Gabriela Csurka, Boris Chidlovskii, and Florent Perronnin. Domain adaptation with a domain specific class means classifier. In *ECCV Workshop on Transferring and Adapting Source Knowledge in Computer Vision (TASK-CV)*, 2014.
- [16] Hal Daumé III. Frustratingly easy domain adaptation. *CoRR*, arXiv:0907.1815, 2009.
- [17] Hal Daumé III and Daniel Marcu. Domain adaptation for statistical classifiers. *Journal of Artificial Intelligence Research*, 26(1):101–126, 2006.
- [18] Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. *CoRR*, arXiv:1310.1531, 2013.
- [19] Lixin Duan, Ivor W. Tsang, and Dong Xu. Domain transfer multiple kernel learning. *Transactions of Pattern Recognition and Machine Analyses (PAMI)*, 34(3):465–479, 2012.
- [20] Lixin Duan, Ivor W. Tsang, Dong Xu, and Tat-Seng Chua. Domain adaptation from multiple sources via auxiliary classifiers. In *International Conference on Machine Learning (ICML)*, 2009.
- [21] Nazli FarajiDavar, Teofilo deCampos, and Josef Kittler. Transductive transfer machines. In *Asian Conference on Computer Vision (ACCV)*, 2014.
- [22] Basura Fernando, Amaury Habrard, Marc Sebban, and Tinne Tuytelaars. Unsupervised visual domain adaptation using subspace alignment. In *IEEE International Conference on Computer Vision (ICCV)*, 2013.
- [23] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International Conference on Machine Learning (ICML)*, 2015.
- [24] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor S. Lempitsky. Domain-adversarial training of neural networks. *CoRR*, arXiv:1505.07818, 2015.
- [25] Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Domain adaptation for large-scale sentiment classification: A deep learning approach. In *International Conference on Machine Learning (ICML)*, 2011.
- [26] Boqing Gong, Kristen Grauman, and Fei Sha. Connecting the dots with landmarks: Discriminatively learning domain invariant features for unsupervised domain adaptation. In *International Conference on Machine Learning (ICML)*, 2013.
- [27] Boqing Gong, Yuan Shi, Fei Sha, and Kristen Grauman. Geodesic flow kernel for unsupervised domain adaptation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [28] Raghuraman Gopalan, Ruonan Li, Vishal M. Patel, and Rama Chellappa. Domain adaptation for visual recognition. *Foundations and Trends in Computer Graphics and Vision*, 8(4), 2015.
- [29] Judy Hoffman, Erik Rodner, Jeff Donahue, Trevor Darrell, and Kate Saenko. Efficient learning of domain-invariant image representations. In *International Conference on Learning representations (ICLR)*, 2013.
- [30] Jiayuan Huang, Alex Smola, Arthur Gretton, Karsten Borgwardt, and Bernhard Schölkopf. Correcting sample selection bias by unlabeled data. In *Annual Conference on Neural Information Processing Systems (NIPS)*, 2007.

- [31] Sheng Li, Jaya Kawale, and Yun Fu. Deep collaborative filtering via marginalized denoising auto-encode. In *ACM International Conference on Information and Knowledge Management (CIKM)*, 2015.
- [32] Yingming Li, Ming Yang, Zenglin Xu, and Zhang Zhongfei. Learning with marginalized corrupted features and labels together. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2016.
- [33] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael I. Jordan. Learning transferable features with deep adaptation networks. In *International Conference on Machine Learning (ICML)*, 2015.
- [34] Mingsheng Long, Jianmin Wang, Guiguang Ding, Jianguang Sun, and Philip S. Yu. Transfer feature learning with joint distribution adaptation. In *IEEE International Conference on Computer Vision (ICCV)*, 2013.
- [35] Laurens van der Maaten, Minmin Chen, Stephen Tyree, and Kilian Weinberger. Learning with marginalized corrupted features. In *International Conference on Machine Learning (ICML)*, 2013.
- [36] Sinno J. Pan, James T. Tsang, Ivor W. and Kwok, and Qiang Yang. Domain adaptation via transfer component analysis. *Transactions on Neural Networks*, 22(2):199 – 210, 2011.
- [37] Sinno J. Pan and Qiang Yang. A survey on transfer learning. *Transactions on Knowledge and Data Engineering*, 22(10):1345–1359, 2010.
- [38] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *European Conference on Computer Vision (ECCV)*, 2010.
- [39] Danny C. Sorensen and Yunkai Zhou. Direct methods for matrix sylvester and lyapunov equations. *Journal of Applied Mathematics*, 6:277–303, 2003.
- [40] Baochen Sun, Jiashi Feng, and Kate Saenko. Return of frustratingly easy domain adaptation. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2016.
- [41] Eric Tzeng, Judy Hoffman, Ning Zhang, Kate Saenko, and Trevor Darrell. Deep domain confusion: Maximizing for domain invariance. *CoRR*, arXiv:1412.3474, 2014.
- [42] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *International Conference on Machine Learning (ICML)*, 2008.
- [43] Zhijie Xu and Shiliang Sun. Multi-source transfer learning with multi-view adaboost. In *Annual Conference on Neural Information Processing Systems (NIPS)*, 2010.
- [44] Joey Tianyi Zhou, Sinno Jialin Pan, Ivor W. Tsang, and Yan Yan. Hybrid heterogeneous transfer learning through deep learning. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2014.
- [45] Mianwei Zhou and Kevin C. Chang. Unifying learning to rank and domain adaptation: Enabling cross-task document scoring. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining (SIGKDD)*, 2014.