

# Zero-Shot Domain Adaptation via Kernel Regression on the Grassmannian

Yongxin Yang  
yongxin.yang@qmul.ac.uk  
Timothy Hospedales  
t.hospedales@qmul.ac.uk

RIM Group  
School of EECS  
Queen Mary, University of London

## Abstract

Most visual recognition methods implicitly assume the data distribution remains unchanged from training to testing. However, in practice *domain shift* often exists, where real-world factors such as lighting and sensor type change between train and test, and classifiers do not generalise from source to target domains. It is impractical to train separate models for all possible situations because collecting and labelling the data is expensive. Domain adaptation algorithms aim to ameliorate domain shift, allowing a model trained on a source to perform well on a different target domain. However, even for the setting of unsupervised domain adaptation, where the target domain is unlabelled, collecting data for every possible target domain is still costly. In this paper, we propose a new domain adaptation method that has no need to access either data or labels of the target domain when it can be described by a parametrised vector and there exists several related source domains within the same parametric space. It greatly reduces the burden of data collection and annotation, and our experiments show some promising results.

## 1 Introduction

Supervised learning usually assumes that the training and testing data are drawn from the same underlying distribution. This assumption is easily violated in many real-world problems. For example, our classifier might be trained with high-quality images captured by a HD camera in an ideal environment, while the trained model may be applied to images captured in poor lighting condition where the camera is not held still, so the image is blurred and dark. In this scenario, directly using the pre-trained model leads to poor performance.

The pervasiveness of this domain shift issue has motivated extensive research in domain adaptation [11, 11, 24, 26]. Domain adaptation (DA) aims to undo this distribution shift between the source domain (where the model is trained) and target domain (where the model is applied). DA has two main settings: supervised DA where the target domain has labels but its data volume is very small, and unsupervised DA where the target domain is completely unlabelled. In this paper, we will focus on unsupervised DA – and eventually a new problem setting of zero-shot domain adaptation – as they are more practically useful.

In most DA studies, ‘domain’ is often equivalent to ‘dataset’ for benchmarking convenience. In the classic Office dataset [26], images are split into three discrete domains/datasets based on their capture device. However, it is commonly overlooked by the DA community

that in many cases *domains are not determined by a single categorical variable, but rather a vector of continuous variables*. For example, consider visual surveillance cameras for person or event recognition: the camera angle  $\alpha$  affects the poses of people that it captures, and time of day  $\tau$  describes the illumination. In this case domains vary smoothly on a continuous manifold, and  $[\alpha, \tau]$  describes a specific domain. Given that both  $\alpha$  and  $\tau$  are continuous variables, we have, in fact, an infinite number of domains; and each parameter will have a certain degree of variation that prevents models trained on one domain from working well on another. Discretizing each factor allows conventional DA to be applied; but this inevitably leads to information loss, and to a large number of domains. We refer to the parameters such as  $\alpha$  as domain factors, and the vectors like  $[\alpha, \tau]$  as domain descriptors.

With this perspective it is clear that even unsupervised domain adaptation does not scale – since the number of datasets to collect grows exponentially in the number of domain factors. This motivates our search for a method that can learn from a few domains and then generalise well to an arbitrary parametrised domain *without collecting the data from it*. We refer to this scenario as zero-shot domain adaptation (ZSDA), distinct to supervised and unsupervised domain adaptation. ZSDA is possible if we can predict some pattern (e.g., the subspace that supports the data) of the target domain from the abstract description of the domain descriptor. We formulate this task as a problem of manifold-valued data regression.

ZSDA aims to enable the attractive use case where a model can consistently and instantaneously perform well on a varying test data distribution, by being automatically ‘calibrated’ on the fly by domain descriptor metadata. This would have important applications in a numerous areas including object [26], person [24] and audio [62] recognition. To this end, our contributions are three-fold: (i) We propose the novel problem of zero-shot domain adaptation, for domains parametrised by continuous vectors, (ii) We provide a solution to the proposed problem by designing a novel multivariate regression model for the Grassmannian, (iii) We show promising early results demonstrating the efficacy of our ZSDA.

## 2 Related Work

### 2.1 Zero-Shot Learning

Zero-Shot Learning (ZSL) has received extensive attention in the computer vision community, such as character [18], object [8, 17], and action [19] recognition. Instead of building a map (classifier) directly from the image to label space, ZSL studies [18, 22] propose an intermediate representation. This representation may take the form of ‘attributes’, or more generally, a semantically meaningful descriptor. The motivation is that, assuming the mapping from image to attributes is sufficiently universal, it can be learned from a large amount of data [29, 32]. The mapping can then be used to build recognisers on-the-fly by giving the semantic descriptor for a new object. For example by assigning the attributes [‘black’, ‘white’, ‘stripes’] to the new object ‘zebra’.

Inspired by ZSL, we aim to achieve a similar on-the-fly capability for domain adaptation. In the same way that semantic representations such as attributes enable ZSL, we aim to exploit the often freely available domain-metadata to adapt a trained model to any target domain given only its descriptor. The closest work in this area is [62], which first mentioned the problem of zero-shot domain adaptation. However it has the drawback that the domain descriptor is limited to a vector of categorical factors only. We do not have this restriction.

## 2.2 Domain Adaptation

Domain Adaptation (DA) techniques reduce the divergence between source and target domains such that a model trained on the source performs well on the target. We focus on unsupervised domain adaptation here, as it is more practically valuable. However, it is challenging because we have zero knowledge of the conditional distribution of the target domain  $P(Y|X_{\mathcal{T}})$  thus a discriminative model trained on the source domain  $P(Y|X_{\mathcal{S}})$  can not be leveraged. The remaining option is to exploit the marginal distributions  $P(X_{\mathcal{T}})$  and  $P(X_{\mathcal{S}})$ .

There are two main approaches in this area, data and subspace centric. **Data:** This approach seeks a unified transformation  $\phi(\cdot)$  that projects two domains' data into a new space that reduces the discrepancy between the transformed target  $\phi(X_{\mathcal{T}})$  and source data  $\phi(X_{\mathcal{S}})$ . A typical pipeline is to perform PCA [10] or sparse coding [11] on the union of the domains with an additional objective that minimises maximum mean discrepancy (MMD) of the new representations of two domains in a reproducing kernel Hilbert space (RKHS)  $\mathcal{H}$ , i.e.,  $\|\mathbb{E}[\phi(X_{\mathcal{T}})] - \mathbb{E}[\phi(X_{\mathcal{S}})]\|_{\mathcal{H}}^2$ , where  $\mathbb{E}[\cdot]$  is the expectation operator. **Subspace:** This approach does not manipulate the data directly, instead it tries to make use of subspaces of two domains. Here, subspace refers to a  $D$ -by- $K$  matrix of the first  $K$  eigenvectors induced by PCA on the original  $D$ -dimensional data. We denote  $P_{\mathcal{S}}$  and  $P_{\mathcal{T}}$  as the subspaces of source and target domain learned by two separate PCAs. Subspace Alignment (SA) [6] learns a linear map  $M$  for  $P_{\mathcal{S}}$  that minimise the Bregman matrix divergence  $\|P_{\mathcal{S}}M - P_{\mathcal{T}}\|_F^2$ . [12] samples several intermediate subspaces  $P_1, P_2, \dots, P_N$  from  $P_{\mathcal{S}}$  to  $P_{\mathcal{T}}$ . That is achieved by thinking of  $P_{\mathcal{S}}$  and  $P_{\mathcal{T}}$  as two points on the Grassmann manifold (Grassmannian)  $\mathbb{G}(K, D)$  and finding a geodesic (shortest path on manifold) between them, then the points along the geodesic are meaningful subspaces. Then all subspaces are concatenated to form a richer linear operator  $[P_{\mathcal{S}}, P_1, P_2, \dots, P_N, P_{\mathcal{T}}]$  that projects two domains into a common space, where the source classifier generalises better to the target domain. A weakness of [12] is that the number of intermediate points is a hard-to-determine hyper-parameter. An elegant solution to this, [13] samples all the intermediate points. Although this produces infinitely long feature vectors, excluding conventional linear classifiers, their dot-product is still defined, and thus any kernelised classifier can be used. A recent study [14] considered the case where domains are associated with a *single* continuous variable (time) using a sequential PCA and subspace-based DA method. However, it does not extend to a vector domain descriptor, nor to ZSDA.

Reviewing these DA studies, one could easily conclude that target data are compulsory. It initially seems to be impossible to achieve DA without any data in the target domain. However, if we take a deeper look at the subspace approach, the key is the subspace rather than data. If each observed subspace (domain) is associated with an independent vector variable  $\mathbf{z}$  (domain descriptor), it is possible to predict a new subspace  $P^*$  given its corresponding  $\mathbf{z}^*$  via a regression model  $\mathbf{z} \rightarrow P$ . This is sometimes called manifold-valued data regression problem, where the output space is a manifold (e.g., the Grassmannian) and the input is in Euclidean space.

## 2.3 Manifold-valued Data Regression

There exists several studies addressing regression in the setting that the independent variable is a point in Euclidean space and the dependent variable is a point in non-flat manifold space such as Riemann and Grassmann manifolds. Based on their methodology, we can group these studies into three categories: (i) Parametric approaches like [8, 9, 9, 13, 15, 16] usually try to find a formulation for the geodesic and then provide a numerical solution for

its estimation. (ii) Semi-parametric approach, e.g., [27] uses a link function to map from Euclidean to Riemannian space. (ii) Non-parametric approaches such as [9, 9] adapt kernel regression to the manifold case by observing that they are all essentially about searching for a point for which the sum of its (reweighted) distances with all training points is minimised.

Note that for most parametric solutions, the independent variable is assumed to be a scalar (univariate regression). This is because: (i) In applications where these methods are popular, e.g., medical imaging, one usually wants to find a pattern against a single factor (e.g., age) and (ii) it is technically challenging to extend the method to multivariate case [16], because the prediction no longer corresponds to a single geodesic curve, which makes the gradient derivation problematic.

We therefore aim to find a solution based on a non-parametric method since the kernel function usually does not make assumptions on whether the input is a scalar or a vector.

## 3 Methodology

### 3.1 Kernel Regression on Grassmannian

Our goal is to build a regression model that takes an  $M$ -dimensional vector of independent variables as the input and predicts a point on Grassmannian (represented by a matrix with orthonormal columns). The output constraint means it can not be treated as conventional Euclidean regression. Therefore, we design a regression model for the Grassmannian.

#### 3.1.1 Kernel Regression Review

We first review kernel regression. Assume we are given a set of (data, label) pairs,

$$\{(\mathbf{z}_1, P_1), (\mathbf{z}_2, P_2), \dots, (\mathbf{z}_N, P_N)\} \quad (1)$$

where  $\mathbf{z} \in \mathbb{R}^M$  and  $P \in \mathbb{R}^1$ ; and a kernel function  $k(\mathbf{z}_1, \mathbf{z}_2)$  that measures the similarity between  $\mathbf{z}_1$  and  $\mathbf{z}_2$ . The kernel regression prediction of a test point  $\mathbf{z}$  is then estimated by,

$$P = \frac{\sum_{i=1}^N k(\mathbf{z}, \mathbf{z}_i) P_i}{\sum_{i=1}^N k(\mathbf{z}, \mathbf{z}_i)}. \quad (2)$$

#### 3.1.2 From Euclidean to Grassmannian regression

When  $P \in \mathcal{M}$  where  $\mathcal{M}$  is a non-flat manifold and  $P$  is no longer a scalar, Eq. 2 can be invalid. For example, suppose  $\mathcal{M}$  is a Grassmannian  $\mathbb{G}(K, D)$  so its members are now matrices  $P \in \mathbb{R}^{D \times K}$  with constrains that  $P^T P = I_K$ . Eq. 2 can be applied and matrices  $P_i$  added, but this is meaningless because adding two points on the Grassmann manifold does not necessarily give another point on Grassmann manifold.

Inspired by [9], we propose to think of kernel regression as the solution of the following optimisation problem:

$$\arg \min_{P \in \mathbb{R}^1} \sum_{i=1}^N w_i (P - P_i)^2 \quad (3)$$

where  $w_i = \frac{k(\mathbf{z}, \mathbf{z}_i)}{\sum_{i=1}^N k(\mathbf{z}, \mathbf{z}_i)}$ . More generally, we have

$$\arg \min_P \sum_{i=1}^N w_i d^2(P, P_i) \quad (4)$$

where  $d^2(\cdot, \cdot)$  is a metric (distance function).  $P$  is the Fréchet mean if the minimizer is unique (or Karcher mean when it is a local minimum). Note that the Fréchet mean is defined in general metric space, thus it provides a way to work with manifold-valued data as long as we can find a well defined distance function for the points on the manifold.

**Grassmann Manifold Background** We first review some basic concepts about the Grassmannian. Many distances on the Grassmannian are defined based on a key concept called ‘principal angle’, which can be calculated by SVD. E.g., for two points  $P_1$  and  $P_2$  on  $\mathbb{G}(K, D)$ ,

$$P_1^T P_2 = USV^T \quad (5)$$

where  $S = \text{diag}(\cos(\theta_1), \cos(\theta_2), \dots, \cos(\theta_K))$ . The angle  $\theta_k = \cos^{-1}(S_{k,k})$  is the  $k$ th principal angle. Table 1 lists some frequently used distance functions on the Grassmannian. The details on deriving these functions and their comparison are beyond the scope of this paper. Readers who are interested in this topic should see [10] and [13] for a good reference.

Table 1: Distances  $d^2(P_1, P_2)$  on  $\mathbb{G}(K, D)$  in terms of principal angles and orthonormal bases

	Principal angles	Orthonormal bases
Binet–Cauchy distance	$1 - \prod_{k=1}^K \cos^2 \theta_k$	$1 - (\det(P_1^T P_2))^2$
Chordal distance	$\sum_{k=1}^K \sin^2 \theta_k$	$\frac{1}{2} \ P_1 P_1^T - P_2 P_2^T\ _F^2$
Martin distance	$\log \prod_{k=1}^K (\cos^2 \theta_k)^{-1}$	$-2 \log(\det(P_1^T P_2))$
Procrustes distance	$4 \sum_{k=1}^K \sin^2 \frac{\theta_k}{2}$	$\ P_1 U - P_2 V\ _F^2$

**Manifold-valued data regression with vector input** For our manifold-valued data regression task, we choose Binet–Cauchy (BC) distance, because of its favourable sensitivity properties [10], and because it is amenable to deriving gradients. Substituting the BC distance into Eq. 4, we obtain the following objective function to optimise:

$$\arg \min_{P \in \mathbb{R}^{D \times K}} 1 - \sum_{i=1}^N w_i (\det(P^T P_i))^2, \quad (6)$$

which is subject to constraint  $P^T P = I_K$ . The gradient with respect to  $P$  is,

$$\nabla_P = \sum_{i=1}^N -w_i (\det(P^T P_i))^2 P_i (P^T P_i)^{-1}. \quad (7)$$

Vanilla gradient descent is not applicable because of the orthogonality constraints. It is a non-trivial optimisation problem as the constraints lead to non-convexity. A simple solution is to do gradient descent as usual and re-orthogonalise the matrix after each step, but it is

numerically very expensive. There are some studies on this topic, such as [60], [23], and [28]. We adopt the solution from [31]. It applies a Crank-Nicolson-like update scheme that preserves the constraints.

Given a feasible point  $P$  and the gradient  $G = \nabla_P$ , a skew-symmetric matrix  $A$  is defined as,

$$A := GP^T - PG^T. \quad (8)$$

The new trial point is determined by the Crank-Nicolson-like scheme,

$$P^{\text{Update}}(\eta) = P - \frac{\eta}{2}A(P + AP^{\text{Update}}(\eta)) \quad (9)$$

where  $\eta$  is the learning rate, and  $P^{\text{Update}}(\eta)$  is given by the closed form,

$$P^{\text{Update}}(\eta) = QP \quad \text{where} \quad Q = (I + \frac{\eta}{2}A)^{-1}(I - \frac{\eta}{2}A). \quad (10)$$

The details of deriving equations (8) – (10) can be found in [31].

## 3.2 Zero-Shot Domain Adaptation

Using the methodology developed in Sec. 3.1, our ultimate goal of zero-shot domain adaptation becomes possible. Assuming that we have  $N$  observed domains given by

$$\{(X_1, [\mathbf{y}_1], \mathbf{z}_1), (X_2, [\mathbf{y}_2], \mathbf{z}_2), \dots, (X_N, [\mathbf{y}_N], \mathbf{z}_N)\}, \quad (11)$$

where  $X_i$  is the feature matrix, from which we can learn a subspace  $P_i$  by PCA.  $\mathbf{z}_i$  and  $\mathbf{y}_i$  are the domain descriptor and label vector respectively for domain  $i$ .  $[\mathbf{y}_i]$  indicates that we do not assume all observed domains have been labelled.

For an unseen domain with descriptor  $\mathbf{z}^*$ , we can predict its subspace  $P^*$  based on the proposed method (Eq. 6) and the training data

$$\{(\mathbf{z}_1, P_1), (\mathbf{z}_2, P_2), \dots, (\mathbf{z}_N, P_N)\}. \quad (12)$$

Once  $P_*$  is obtained, any subspace-based DA method (e.g., [6, 11, 11]) can be applied to align the unseen (target) domain to any labelled source domain where a classifier was trained.

## 4 Experiments

**Dataset:** To test our algorithm, we need a dataset which exhibits a range of continuously parametrised domains (Sec. 3.1). However, most existing DA and more general vision datasets are grouped into discrete domains/datasets. We therefore alter an existing dataset for our purposes. We use the Office Dataset [26], which collects the images of office supplies from three sources: *Amazon*, *webcam* and *DSLR*. It is a classic dataset to evaluate domain adaptation algorithms. The typical experimental design for this dataset is to evaluate recognition performance when a model is trained on one domain (e.g., Amazon) and tested on another (e.g., webcam). To test our algorithm, we create a new dataset based on Office.

**Settings:** We use all the Amazon images, which contain 31 categories and with an average of 90 images each. Then, we simulate a range of continuously parametrised domains by degrading each image by two means: lowering the resolution and reducing the brightness. Specifically, we apply a Gaussian filter for simulating low resolution and divide every pixel



Figure 1: Example images in category Backpack. (Left) Original. (Right) Nine simulated domains

value by a factor for simulating poor lighting. The size of Gaussian filter and the darkening factor provide two factors of the domain descriptor. In this experiment, we generate nine distinct domains by three levels of degradation for each parameter as shown in Table 2. An example of each domain along with the original image is shown in Fig. 1. For each domain, we split the training and testing set equally, and the assignment of train versus test set is consistent for all domains. This guarantees that when an image appears in training set, its other versions of degradation will not appear in the testing set.

Table 2: Nine domains generated by the degradation of Office/Amazon.

Domain Index	1	2	3	4	5	6	7	8	9
Gaussian Filter Size	5	5	5	10	10	10	15	15	15
Brightness Factor	1.5	2	3	1.5	2	3	1.5	2	3

**Features:** We use the state-of-the-art Convolutional Neural Network (CNN) model VGG-full [14] as the feature extractor. The image is first preprocessed: rescaled into  $224 \times 224$  and mean subtracted. Then it is fed into the CNN, where the value in the penultimate layer (4096 neurons) is used as the feature vector for further experiments.

#### 4.1 Demonstrating the Domain Shift Challenge

In the first experiment, we demonstrate how domain shift effects recognition performance, and to what degree domain adaptation methods can alleviate it. A linear SVM classifier is trained on the original domain’s training data and we evaluate the performance of the classifier measured by accuracy ( $\frac{\text{Num. of Correctly Classified Examples}}{\text{Num. of All Examples}}$ ) in the nine degraded domains’ test data. Then we apply a popular domain adaptation algorithm – Geodesic Flow Kernel (GFK) [15] that aims to manipulate the target subspace so that domain shift is reduced.

As we can see in Fig. 2, the CNN feature is very discriminative. The accuracy in the original domain’s test data is 81.68% (blue horizontal line in Fig. 2). However, the performance




Figure 2: Recognition rate in Office/Amazon. The domain shift problem is significant despite the use of state of the art CNN features.

Table 3: Recognition accuracy on Domain 5 after zero-shot domain adaptation.

Source ID	1	2	3	4	6	7	8	9	Avg.
No DA	66.12	67.33	66.69	71.73	69.74	69.82	68.68	62.93	67.88
ZSDA→GFK	67.90	69.03	69.03	74.43	71.80	72.66	71.09	66.41	70.29

‘Source ID’ indicates *which one* of the source domains is used to train the classifier. ‘No DA’ means applying the trained model directly to the target domain. ‘ZSDA→GFK’ is the proposed method: the subspace of the target domain is predicted by *all* source domains, and the domain adaptation is conducted by GFK from each source domain in turn to the target domain.

inevitably drops when the quality of images gets lower. For the last domain, domain\_9, the accuracy has declined to 24.29%. Nevertheless, unsupervised domain adaptation by GFK does reduce the performance drop. It improves the recognition rate by 7.42% on average.

From these results we conclude that: (i) state of the art features do not eliminate the domain-shift problem (contrary to some claims [10]), and (ii) domain adaptation methods still play an important role in the era of deep learning.

## 4.2 Zero-Shot Domain Adaptation

In the second experiment, we validate the proposed algorithm’s ability to undo domain shift given only a target domain descriptor.

**Target domain performance for different sources of labels:** The fifth domain, domain\_5, is chosen to be the target domain, and the evaluation is run on its testing data. Unlike traditional domain adaptation approaches [10], *none of the data (either feature or label) in target domain is given*. All the known information about the target domain is the descriptor of domain\_5, i.e.,  $\mathbf{z}_5 = [10, 2]$ . Besides this, we assume all the training data from the other domains as well as their domain descriptors are given, from which we can learn 8 subspaces via PCA (the reduced dimension is  $K = 512$ ). Then the subspace of domain\_5 is predicted by the proposed algorithm so that the subspace based domain adaptation [10] can be applied. Note that we rescale each factor in  $\mathbf{z}$  to  $[0, 1]$  to reduce the effect of different ranges of factors, and RBF kernel with  $\sigma = 0.1$  is chosen to be  $k(\cdot, \cdot)$ .

The results are summarised in Table 3. These are generated by predicting the subspace  $P_5$  of the unseen domain (domain\_5) based *all* the remaining eight source domains and testing the classifier (with or without DA) trained from the labeled data of *each* source domain in turn (i.e., single source domain to target). The source ID indicates which source domain is



Table 4: Average accuracy on every target domain after zero-shot domain adaptation.

Target ID	1	2	3	4	5	6	7	8	9
No DA	69.10	69.15	68.31	68.18	67.88	63.23	59.14	56.86	46.41
ZSDA→GFK	72.06	72.07	71.24	71.07	70.29	65.94	61.37	59.54	50.76

used to train the classifier and optionally apply DA. Note the distinction between source of unlabelled data for subspace learning, and source of labeled data for classifier learning.

In every case, the proposed ZSDA method generates a useful subspace which allows domain adaptation to improve the performance. It is as expected that using ‘closer’ domains (measured by distance between  $\mathbf{z}$ ’s) such as domain\_2, \_4, \_6, and \_8 to learn the classifier should lead to better performance compared to source domains that are ‘far-away’ from the target. This is indeed the case: if we focus on the nearby domains\_{2,4,6,8}, then the average accuracy in Table 3 is 71.59%. Given that the accuracy obtained within domain\_5 only (i.e. training and testing on domain\_5) is 72.94%, and the average accuracy of those sources without DA is 69.37%, this is a very encouraging result: *we have reduced the cross-domain performance drop by over half without relying on any target domain data.*

**Performance for each target domain:** The previous analysis fixed domain\_5 as the target. We next extend this analysis and consider each domain in turn as the target. In each case we use the proposed algorithm to infer the subspace of the target domain given the eight sources, and we evaluate the recognition accuracies with and without DA. For conciseness, we just report the average accuracy over all possible label sources for each target (cf. the last column in Table 3). Table 4 summarises the result and it demonstrates the effectiveness and robustness of the proposed method. It shows performance improvements on all choices of target domains, and on average, it boosts the accuracy by a factor of 4.75%.

## 5 Conclusion

We proposed the problem of continuously-parametrised zero-shot domain adaptation and developed a solution based on manifold-valued data regression. This allows us to predict the subspace to use at test-time and thus align a source classifier to a test-domain in advance of seeing any data. Preliminary results demonstrate the value of our approach. This approach is highly promising for its potential impact on a variety of areas where it would be useful to be able to ‘calibrate’ a recognition model on the fly based on metadata.

There are numerous areas for future work. A more thorough evaluation that also covers a wider range of applications is clearly needed. Our current kernel regression-based method is weak in domain extrapolation in contrast to interpolation. This is important for some kinds of subspace prediction, e.g., predicting future subspaces when one domain factor is time. Thus a generalisation for extrapolation is of interest. While the proposed method can and must use a set of available source domains to learn the subspace regressor, it can only exploit a single source domain’s labels. A useful extension would therefore be to exploit multiple source domains worth of labels. Besides, it is interesting to see if the predicted subspace can act as a regulariser so that it still helps when the target data are available but limited. Finally, a key assumption of this paper (in common with most other domain adaptation work) is that the domain descriptor is always observed and accurate. To relax this assumption and enable the model to deal with missing or noisy descriptors is also an interesting direction.

## References

- [1] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman. Return of the devil in the details: Delving deep into convolutional nets. In *British Machine Vision Conference (BMVC)*, 2014.
- [2] B. C. Davis and S. Lazebnik. Analysis of human attractiveness using manifold kernel regression. In *International Conference on Image Processing (ICIP)*, 2008.
- [3] Bradley C. Davis, P. Thomas Fletcher, Elizabeth Bullitt, and Sarang Joshi. Population shape regression from random design data. In *International Conference on Computer Vision (ICCV)*, 2007.
- [4] 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. In *International Conference on Machine Learning (ICML)*, 2015.
- [5] F. Silva Leite E. Batzies, L. Machado. The geometric mean and the geodesic fitting problem on the grassmann manifold. 2013. URL <http://www.mat.uc.pt/preprints/ps/p1322.pdf>.
- [6] Basura Fernando, Amaury Habrard, Marc Sebban, and Tinne Tuytelaars. Unsupervised visual domain adaptation using subspace alignment. In *International Conference on Computer Vision (ICCV)*, 2013.
- [7] P. Thomas Fletcher. Geodesic regression and the theory of least squares on riemannian manifolds. *International Journal of Computer Vision (IJCV)*, 105:171–185, 2013.
- [8] Y. Fu, T. Hospedales, T. Xiang, Z. Fu, and S. Gong. Transductive multi-view embedding for zero-shot recognition and annotation. In *European Conference on Computer Vision (ECCV)*, 2014.
- [9] K.A. Gallivan, A. Srivastava, Xiuwen Liu, and P. Van Dooren. Efficient algorithms for inferences on grassmann manifolds. In *IEEE Workshop on Statistical Signal Processing*, 2003.
- [10] Boqing Gong, Yuan Shi, Fei Sha, and K. Grauman. Geodesic flow kernel for unsupervised domain adaptation. In *Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [11] R. Gopalan, Ruonan Li, and R. Chellappa. Domain adaptation for object recognition: An unsupervised approach. In *International Conference on Computer Vision (ICCV)*, 2011.
- [12] Jihun Ham and Daniel D. Lee. Grassmann discriminant analysis: a unifying view on subspace-based learning. In *International Conference on Machine Learning (ICML)*, 2008.
- [13] Jacob Hinkle, Prasanna Muralidharan, P. Thomas Fletcher, and Sarang C. Joshi. Polynomial regression on riemannian manifolds. In *European Conference on Computer Vision (ECCV)*, 2012.

- [14] Judy Hoffman, Trevor Darrell, and Kate Saenko. Continuous manifold based adaptation for evolving visual domains. In *Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [15] Yi Hong, Roland Kwitt, Nikhil Singh, Brad Davis, Nuno Vasconcelos, and Marc Niethammer. Geodesic regression on the grassmannian. In *European Conference on Computer Vision (ECCV)*, 2014.
- [16] Hyunwoo J. Kim, Barbara B. Bendlin, Nagesh Adluru, Maxwell D. Collins, Moo K. Chung, Sterling C. Johnson, Richard J. Davidson, and Vikas Singh. Multivariate general linear models (MGLM) on riemannian manifolds with applications to statistical analysis of diffusion weighted images. In *Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [17] Christoph H. Lampert, Hannes Nickisch, and Stefan Harmeling. Learning to detect unseen object classes by between-class attribute transfer. In *Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [18] Hugo Larochelle, Dumitru Erhan, and Yoshua Bengio. Zero-data learning of new tasks. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2008.
- [19] Jingen Liu, Benjamin Kuipers, and Silvio Savarese. Recognizing human actions by attributes. In *Computer Vision and Pattern Recognition (CVPR)*, 2011.
- [20] Mingsheng Long, Guiguang Ding, Jianmin Wang, Jianguang Sun, Yuchen Guo, and P.S. Yu. Transfer sparse coding for robust image representation. In *Computer Vision and Pattern Recognition (CVPR)*, 2013.
- [21] Mingsheng Long, Jianmin Wang, Guiguang Ding, Jianguang Sun, and Philip S Yu. Transfer feature learning with joint distribution adaptation. In *International Conference on Computer Vision (ICCV)*, 2013.
- [22] Mark Palatucci, Dean Pomerleau, Geoffrey Hinton, and Tom Mitchell. Zero-shot learning with semantic output codes. In *Neural Information Processing Systems (NIPS)*, 2009.
- [23] Mark D. Plumbley. Lie group methods for optimization with orthogonality constraints. In *Independent Component Analysis and Blind Signal Separation*, 2004.
- [24] Qiang Qiu, Vishal M. Patel, Pavan Turaga, and Rama Chellappa. Domain adaptive dictionary learning. In *European Conference on Computer Vision (ECCV)*, 2012.
- [25] Quentin Rentmeesters. A gradient method for geodesic data fitting on some symmetric riemannian manifolds. In *Conference on Decision and Control and European Control Conference (CDC-ECC)*, 2011.
- [26] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *European Conference on Computer Vision (ECCV)*, 2010.
- [27] Xiaoyan Shi, Martin Styner, Jeffrey A. Lieberman, Joseph G. Ibrahim, Weili Lin, and Hongtu Zhu. Intrinsic regression models for manifold-valued data. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2009.

- [28] S. T. Smith. Optimization Techniques on Riemannian Manifolds. *ArXiv e-prints*, 2014.
- [29] Richard Socher, Milind Ganjoo, Christopher D Manning, and Andrew Y Ng. Zero-shot learning through cross-modal transfer. In *Neural Information Processing Systems (NIPS)*, 2013.
- [30] V. M. Tkachuk. Supersymmetric Method for Constructing Quasi-Exactly Solvable Potentials. *eprint arXiv:quant-ph/9806030*, 1998.
- [31] Zaiwen Wen and Wotao Yin. A feasible method for optimization with orthogonality constraints. Technical report, Rice University, 2010.
- [32] Yongxin Yang and Timothy M. Hospedales. A unified perspective on multi-domain and multi-task learning. In *International Conference on Learning Representations (ICLR)*, 2015.
- [33] K. Ye and L.-H. Lim. Distance between subspaces of different dimensions. *ArXiv e-prints*, 2014.