Subspace Distribution Alignment for Unsupervised Domain Adaptation

Baochen Sun bsun@cs.uml.edu Kate Saenko saenko@cs.uml.edu



Figure 1: Existing approaches for domain adaptation via subspace mapping fail to properly align the source and target distributions. Our proposed approach improves domain alignment by taking the difference of the source and target distributions into account.

Abstract. Traditional machine learning algorithms often fail to generalize to new input distributions, causing reduced accuracy. Domain adaptation attempts to compensate for the performance degradation by transferring and adapting source knowledge to target domain. Existing unsupervised methods project domains into a lower-dimensional space and attempt to align the subspace bases, effectively learning a mapping from source to target points or vice versa. However, they fail to take into account the difference of the two distributions in the subspaces, resulting in misalignment even after adaptation. We present a unified view of existing subspace mapping based methods and develop a generalized approach that also aligns the distributions as well as the subspace bases.

Background. Domain adaptation, or covariate shift, is a fundamental problem in machine learning, and has attracted a lot of attention in the machine learning and computer vision community. Domain adaptation methods for visual data attempt to learn classifiers on a labeled source domain and transfer it to a target domain. There are two settings for visual domain adaptation: (1) unsupervised domain adaptation where there are no labeled examples available in the target domain; and (2) semi-supervised domain. Most existing algorithms operate in the semi-superised setting. However, in real world applications, unlabeled target data is often much more abundant and labeled examples are very limited, so the question of how to utilize the unlabeled target data is more important for practical visual domain adaptation. Thus, in this paper, we focus on the unsupervised scenario.

Most of the existing unsupervised approaches have pursued adaptation by separately projecting the source and target data into a lowerdimensional manifold, and finding a transformation that brings the subspaces closer together. This process is illustrated in Figure 1. Geodesic methods [2, 3] find a path along the subspace manifold, and either project source and target onto points along that path [3], or find a closed-form linear map that projects source points to target [2]. Alternatively, the subspaces can be aligned by computing the linear map that minimizes the Frobenius norm of the difference between them, a method known as Subspace Alignment [1].

Approach. The intuition behind our approach is that although the existing approaches might align the subspaces (the bases of the subspaces), it might not fully align the data distributions in the subspaces as illustrated in Figure 1. We use the first- and second-order statistics, namely the mean and the variance, to describe a distribution in this paper. Since the mean after data preprocessing (i.e. normalization) is zero and is not affected Computer Science Department University of Massachusetts Lowell Lowell, Massachusetts, US



Figure 2: Mean accuracy across all **12** experiment settings (domain shifts) of the k-NN Classifier on the **Office-Caltech10** dataset. Both our methods SDA-IS and SDA-TS outperform GFK and SA consistently. **Left:** k-NN Classifier with **k=1**; **Right:** k-NN Classifier with **k=3**.



Figure 3: Mean accuracy across all **6** experiment settings (domain shifts) of the k-NN Classifier on the **Office31** dataset. Both our methods SDA-IS and SDA-TS outperform GFK and SA consistently. **Left:** k-NN Classifier with **k=1**; **Right:** k-NN Classifier with **k=3**.

by subspace projection, there is no need to align it. However, neither normalization nor subspace basis alignment aligns the covariances.

We propose Subspace Distribution Alignment (SDA), which properly aligns the source and target data in the subspaces, namely, it aligns the distributions as well as the subspace bases. SDA constructs the mapping between source and target points to be

$$M_s = S_s T A S_t^T \tag{1}$$

where S_s and S_t are the *d* dimensional source and target subspaces respectively, *T* is the subspace basis alignment matrix, and *A* is the subspace distribution alignment matrix. Suppose W_s and W_t are the square roots of the covariance matrices of the source and target data in the subspaces respectively. Then *A* could simply be $W_s^{-1}W_t$ since W_s^{-1} transforms the source data into zero mean and unit covariance while W_t adds the covariance of target data. Thus, *A* transforms the source subspace distribution to the target one. We provide two variants of SDA: SDA-TS aligns the source and target data between two subspaces (source subspace and target subspace); SDA-IS aligns them between an infinite number of subspaces (all the subspaces on the geodesic flow). We show that Subspace Alignment (SA) [1] and Geodesic Flow Kernel (GFK) [2] are special cases of our generalized method.

Experiments. We follow the protocol of [1, 2, 3, 4] and compare our methods to their special cases as well as the no adaptation baseline (NA). Extensive experimental results on standard benchmarks (Figure 2 and 3) demonstrate the advantage of our approach over published methods.

- Basura Fernando, Amaury Habrard, Marc Sebban, and Tinne Tuytelaars. Unsupervised visual domain adaptation using subspace alignment. In *ICCV*, 2013.
- [2] B. Gong, Y. Shi, F. Sha, and K. Grauman. Geodesic flow kernel for unsupervised domain adaptation. In *CVPR*, 2012.
- [3] R. Gopalan, R. Li, and R. Chellappa. Domain adaptation for object recognition: An unsupervised approach. In *ICCV*, 2011.
- [4] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *ECCV*. 2010.