

Classification Constrained Dimensionality Reduction

Raviv Raich, *Member, IEEE*, Jose A. Costa, *Member, IEEE*,

Steven B. Damelin, *Senior Member, IEEE*, and Alfred O. Hero III, *Fellow, IEEE*

Abstract

Dimensionality reduction is a topic of recent interest. In this paper, we present the classification constrained dimensionality reduction (CCDR) algorithm to account for label information. The algorithm can account for multiple classes as well as the semi-supervised setting. We present an out-of-sample expressions for both labeled and unlabeled data. For unlabeled data, we introduce a method of embedding a new point as preprocessing to a classifier. For labeled data, we introduce a method that improves the embedding during the training phase using the out-of-sample extension. We investigate classification performance using the CCDR algorithm on hyper-spectral satellite imagery data. We demonstrate the performance gain for both local and global classifiers and demonstrate a 10% improvement of the k -nearest neighbors algorithm performance. We present a connection between intrinsic dimension estimation and the optimal embedding dimension obtained using the CCDR algorithm.

Index Terms

Classification, Computational Complexity, Dimensionality Reduction, Embedding, High Dimensional Data, Kernel, K-Nearest Neighbor, Manifold Learning, Probability, Out-of-Sample Extension.

This work was partially funded by the DARPA Defense Sciences Office under Office of Naval Research contract #N00014-04-C-0437. Distribution Statement A. Approved for public release; distribution is unlimited. S. B. Damelin was supported in part by National Science Foundation grant no. NSF-DMS-0555839 and NSF-DMS-0439734 and by AFRL.

R. Raich is with the Oregon State University, Corvallis. A. O Hero III is with the University of Michigan, Ann Arbor. J. A. Costa is with the California Institute of Technology. S. B. Damelin is with the Georgia Southern University.

Classification Constrained Dimensionality Reduction

I. INTRODUCTION

In classification theory, the main goal is to find a mapping from an observation space \mathcal{X} consisting of a collection of points in some containing Euclidean space \mathbb{R}^d , $d \geq 1$ into a set consisting of several different integer valued hypotheses. In some problems, the observations from the set \mathcal{X} lie on a d -dimensional manifold \mathcal{M} and Whitney's theorem tells us that provided that this manifold is smooth enough, there exists an embedding of \mathcal{M} into \mathbb{R}^{2d+1} . This motivates the approach taken by kernel methods in classification theory, such as support vector machines [1] for example. Our interest is in finding an embedding of \mathcal{M} into a lower dimensional Euclidean space.

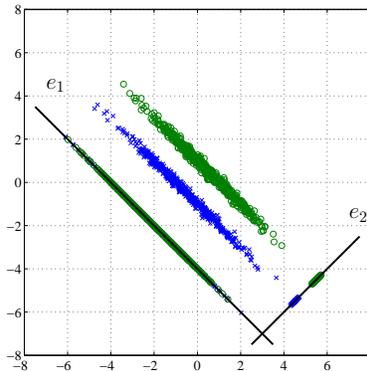


Fig. 1. PCA of a two-classes classification problem.

Dimensionality reduction of high dimensional data, was addressed in classical methods such as principal component analysis (PCA) [2] and multidimensional scaling (MDS) [3], [4]. In PCA, an eigendecomposition of the $d \times d$ empirical covariance matrix is performed and the data points are linearly projected along the $0 < m \leq d$ eigenvectors with the largest eigenvalues. A problem that may occur with PCA for classification is demonstrated in Fig. 1. When the information that is relevant for classification is present only in the eigenvectors associated with the small eigenvalues (e_2 in the figure), removal of such eigenvectors may result in severe degradation

in classification performance. In MDS, the goal is to find a lower dimensional embedding of the original data points that preserves the relative distances between all the data points. The two later methods suffer greatly when the manifold is nonlinear. For example, PCA will not be able to offer dimensionality reduction for classification of two classes lying each on one of two concentric circles.

In [5], a nonlinear extension to PCA is presented. The algorithm is based on the “kernel trick” [6]. Data points are nonlinearly mapped into a feature space, which in general has a higher (or even infinite) dimension as compared with the original space and then PCA is applied to the high dimensional data.

In the paper of Tenenbaum *et al* [7], Isomap, a global dimensionality reduction algorithm was introduced taking into account the fact that data points may lie on a lower dimensional manifold. Unlike MDS, geodesic distances (distances that are measured along the manifold) are preserved by Isomap. Isomap utilizes the classical MDS algorithm, but instead of using the matrix of Euclidean distances, it uses a modified version of it. Each point is connected only to points in its local neighborhood. A distance between a point and another point outside its local neighborhood is replaced with the sum of distances along the shortest path in graph. This procedure modifies the squared distances matrix replacing Euclidian with geodesic distances.

In [8], Belkin and Niyogi present a related Laplacian eigenmap dimensionality reduction algorithm. The algorithm performs a minimization on the weighted sum of squared-distances of the lower-dimensional data. Each weight multiplying the squared-distances of two low-dimensional data points is inversely related to distance between the corresponding two high-dimensional data points. Therefore, small distance between two high-dimensional data points results in small distance between two low-dimensional data points. To preserve the geodesic distances, the weight of the distance between two points that do not share a local neighborhood is set to zero.

We refer the interested reader to the references below and those cited therein for a list of some of the most commonly used additional algorithms within the class of *manifold learning* algorithms and their different advantages relevant to our work. Locally Linear Embedding (LLE) [9], Laplacian Eigenmaps [8], Hessian Eigenmaps (HLLE) [10], Local Space Tangent Analysis [11], Diffusion Maps [12] and Semidefinite Embedding (SDE) [13].

The algorithms mentioned above, consider the problem of learning a lower-dimensional em-

bedding of the data. In classification, such algorithms can be used to preprocess high-dimensional data before performing the classification. This could potentially allow for a lower computational complexity of the classifier. In some cases, dimensionality reduction results increase the computational complexity of the classifier. In fact, support vector machines suggest the opposing strategy: data points are projected onto a higher-dimensional space and classified by a low computational complexity classifier. To guarantee a low computational complexity of the classifier of the low-dimensional data, a classification constrained dimensionality reduction (CCDR) algorithm was introduced in [14]. The CCDR algorithm is an extension of Laplacian eigenmaps [8] and it incorporates class label information into the cost function, reducing the distance between points with similar label. Another algorithm that incorporates label information is the marginal fisher analysis (MFA) [15], in which a constraint on the margin between classes is used to enforce class separation.

In [14] the CCDR algorithm was only studied for two classes and its performance was illustrated for simulated data. In [16], a multi-class extension to the problem was presented. In this paper, we introduce two additional components that make the algorithm computationally viable. The first is an out-of-sample extension for classification of unlabeled test points. Similarly to the out-of-sample extension presented in [17], one can utilize the Nyström formula for classification problems in which label information is available. We study the algorithm performance as its various parameters, (e.g., dimension, label importance, and local neighborhood), are varied. We study the performance of CCDR as preprocessing prior to implementation of several classification algorithms such as k -nearest neighbors, linear classification, and neural networks. We demonstrate a 10% improvement over the k -nearest neighbors algorithm performance benchmark for this dataset. We address the issue of dimension estimation and its effect on classification performance.

The organization of this paper is as follows. Section III presents the multiple-class CCDR algorithm. Section ?? provides a study of the algorithm using the Landsat dataset and Section VI summaries our results.

II. DIMENSIONALITY REDUCTION

Let $\mathcal{X}_n = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$ be a set of n points constrained to lie on an m -dimensional submanifold $\mathcal{M} \subseteq \mathbb{R}^d$. In dimensionality reduction, our goal is to obtain a lower-dimensional

embedding $\mathcal{Y}_n = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n\}$ (where $\mathbf{y}_i \in \mathbb{R}^m$ with $m < d$) that preserves local geometry information such that processing of the lower dimensional embedding \mathcal{Y}_n yields comparable performance to processing of the original data points \mathcal{X}_n . Alternatively, we would like learn the mapping $f : \mathcal{M} \subseteq \mathbb{R}^d \rightarrow \mathbb{R}^m$ that maps every data point \mathbf{x}_i to $\mathbf{y}_i = f(\mathbf{x}_i)$ such that some geometric properties of the high-dimensional data are preserved in the lower dimensional embedding. The first question that comes to mind is how to select f , or more specifically how to restrict the function f so that we can still achieve our goal.

A. Linear dimensionality reduction

1) *PCA*: When principal component analysis (PCA) is used for dimensionality reduction, one considers a linear embedding of the form

$$\mathbf{y}_i = f(\mathbf{x}_i) = \mathbf{A}\mathbf{x}_i,$$

where \mathbf{A} is $m \times d$. This embedding captures the notion of proximity in the sense that close points in the high dimensional space map to close points in the lower dimensional embedding, i.e., $\|\mathbf{y}_i - \mathbf{y}_j\| = \|\mathbf{A}(\mathbf{x}_i - \mathbf{x}_j)\| \leq \|\mathbf{A}\|\|\mathbf{x}_i - \mathbf{x}_j\|$. Let

$$\bar{\mathbf{x}} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i$$

and

$$\mathcal{C}_x = \frac{1}{n} \sum_{i=1}^n (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})^T.$$

Similarly, let

$$\bar{\mathbf{y}} = \frac{1}{n} \sum_{i=1}^n \mathbf{y}_i$$

and

$$\mathcal{C}_y = \frac{1}{n} \sum_{i=1}^n (\mathbf{y}_i - \bar{\mathbf{y}})(\mathbf{y}_i - \bar{\mathbf{y}})^T.$$

Since $\mathbf{y}_i = \mathbf{A}\mathbf{x}_i$, we have $\bar{\mathbf{y}} = \mathbf{A}\bar{\mathbf{x}}$ and $\mathcal{C}_y = \mathbf{A}\mathcal{C}_x\mathbf{A}^T$. In PCA, the goal is to find the projection matrix \mathbf{A} that preserves most of the energy in the original data by solving

$$\max_{\mathbf{A}} \text{tr}\{\mathcal{C}_y(\mathbf{A})\} \quad \text{s.t.} \quad \mathbf{A}\mathbf{A}^T = \mathbf{I},$$

which is equivalent to

$$\max_{\mathbf{A}} \text{tr}\{\mathbf{A}\mathcal{C}_x\mathbf{A}^T\} \quad \text{s.t.} \quad \mathbf{A}\mathbf{A}^T = \mathbf{I}. \quad (1)$$

The solution to (1), is given by $\mathbf{A} = [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_m]^T$, where \mathbf{u}_i is the eigenvector of \mathcal{C}_x corresponding to its i th largest eigenvalue. When the data lies on an m -dimensional hyperplane, the matrix \mathcal{C}_x has only m positive eigenvalues and the rest are zero. Furthermore, every \mathbf{x}_i belongs to $\bar{\mathbf{x}} + \text{span}\{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_m\} \subseteq \mathbb{R}^d$. In this case, the mapping PCA finds $f(\mathbf{x}) = \mathbf{A}\mathbf{x}$ is one-to-one and satisfies $\|f(\mathbf{x}_i) - f(\mathbf{x}_j)\| = \|\mathbf{A}(\mathbf{x}_i - \mathbf{x}_j)\| = \|\mathbf{x}_i - \mathbf{x}_j\|$. Therefore, the lower embedding preserves all the geometry information in the original dataset \mathcal{X} . We would like to point out that PCA can be written as

$$\max_{\{\mathcal{Y}\}} \sum_{i=1}^n \|\mathbf{y}_i - \mathbf{y}_j\|^2 \quad \text{s.t.} \quad \mathbf{y}_i = \mathbf{A}\mathbf{x}_i \text{ and } \mathbf{A}\mathbf{A}^T = \mathbf{I},$$

2) *MDS*: Multidimensional Scaling (MDS) differs from PCA in the way the input is provided to it. While in PCA, the original data \mathcal{X} is provided, the classical MDS requires only the set of all Euclidean pairwise distances $\{\|\mathbf{x}_i - \mathbf{x}_j\|_2\}_{i=1, j>i}^{n-1}$. As MDS uses only pairwise distances, the solution it finds is given up to translation and unitary transformation. Let $\mathbf{x}'_i = \mathbf{x}_i - \mathbf{c}$, the Euclidean distance $\|\mathbf{x}'_i - \mathbf{x}'_j\|$ is the same as $\|\mathbf{x}_i - \mathbf{x}_j\|$. Let \mathbf{U} be an arbitrary unitary matrix \mathbf{U} satisfying $\mathbf{U}^T\mathbf{U} = \mathbf{I}$ and define $\mathbf{x}'_i = \mathbf{U}\mathbf{x}_i$. The distance $\|\mathbf{x}'_i - \mathbf{x}'_j\|$ is equal to $\|\mathbf{U}(\mathbf{x}_i - \mathbf{x}_j)\|$, which by the invariance of the Euclidean norm to a unitary transformation equals to $\|\mathbf{x}_i - \mathbf{x}_j\|$. Denote the pairwise squared-distance matrix by $[\mathbf{D}_2]_{ij} = \|\mathbf{x}_i - \mathbf{x}_j\|^2$. By the definition of Euclidean distance, the matrix \mathbf{D}_2 satisfies

$$\mathbf{D}_2 = \mathbf{1}\phi^T + \phi\mathbf{1}^T - 2\mathbf{X}^T\mathbf{X}, \quad (2)$$

where $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n]$ and $\phi = [\|\mathbf{x}_1\|^2, \|\mathbf{x}_2\|^2, \dots, \|\mathbf{x}_n\|^2]^T$. To verify (2), one can examine the ij -th term of \mathbf{D}_2 and compare with $\|\mathbf{x}_i - \mathbf{x}_j\|^2$. Denote the $n \times n$ matrix $\mathbf{H} = \mathbf{I} - \mathbf{1}\mathbf{1}^T/n$. Multiplying both sides of \mathbf{D}_2 with \mathbf{H} in addition to a factor of $-\frac{1}{2}$, yields

$$-\frac{1}{2}\mathbf{H}\mathbf{D}_2\mathbf{H} = (\mathbf{X}\mathbf{H})^T(\mathbf{X}\mathbf{H}),$$

which is key to MDS, i.e., Cholesky decomposition of $-\frac{1}{2}\mathbf{H}\mathbf{D}_2\mathbf{H}$ yields \mathbf{X} to within a translation and a unitary transformation. Consider the eigendecomposition $-\frac{1}{2}\mathbf{H}\mathbf{D}_2\mathbf{H} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T$. Therefore, a rank d \mathbf{X} can be obtained as $\mathbf{X} = \mathbf{\Lambda}_d^{\frac{1}{2}}\mathbf{U}_d^T$, where $\mathbf{\Lambda}_d = \text{diag}\{\lambda_1, \lambda_2, \dots, \lambda_d\}^{\frac{1}{2}}$ and $\mathbf{U}_d = [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_d]$. Note that $\mathbf{X}\mathbf{H}$ is a translated version of \mathbf{X} , in which every column \mathbf{x}_i is translated to $\mathbf{x}_i - \bar{\mathbf{x}}$.

To use MDS for dimensionality reduction, we can consider a two step process. First a squared-distance matrix \mathbf{D}_2 is obtained from the high-dimensional data \mathcal{X} . Then, MDS is applied to \mathbf{D}_2

to obtain a low-dimensional ($m < d$) embedding by $\mathbf{X}_m = \mathbf{\Lambda}_m^{\frac{1}{2}} \mathbf{U}_m^T = \mathbf{X} \mathbf{U}_m \mathbf{U}_m^T$. In the absence of noise, this procedure provides an affine transformation to the high-dimensional data and thus can be regarded as a linear method.

B. Nonlinear dimensionality reduction

Linear maps are limited as they cannot preserve the geometry of nonlinear manifolds.

1) *Kernel PCA*: Kernel PCA is one of the first methods in dimensionality reduction of data on nonlinear manifolds. The method combines the dimensionality reduction capabilities of PCA on linear manifolds with a nonlinear embedding of data points in a higher (or even infinite) dimensional space using “kernel trick” [6]. In PCA, one finds the eigenvectors satisfying: $\mathcal{C}_x \mathbf{v}_k = \lambda_k \mathbf{v}_k$. Since \mathbf{v}_k can be written as a linear combination of the \mathbf{x}_i ’s: $\mathbf{v}_k = \sum_i \alpha_{ki} (\mathbf{x}_i - \bar{\mathbf{x}})$, one can replace \mathbf{v}_k in the eigendecomposition, simplify, and obtain: $\mathbf{X}(\mathbf{K} \boldsymbol{\alpha}_k - \lambda_k \boldsymbol{\alpha}_k) = 0$, where $\mathbf{K}_{ij} = (\mathbf{x}_i - \bar{\mathbf{x}})^T (\mathbf{x}_j - \bar{\mathbf{x}})$. Consider the mapping $\phi : \mathcal{M} \rightarrow \mathcal{H}$ from the manifold to a Hilbert space. The “kernel trick” suggests replacing \mathbf{x}_i with $\phi(\mathbf{x}_i)$ and therefore rewriting the kernel as $\mathbf{K}_{ij} = \phi(\mathbf{x}_i)^T \phi(\mathbf{x}_j)$. Further generalization can be made by setting $\mathbf{K}_{ij} = K(\mathbf{x}_i, \mathbf{x}_j)$ where $K(\cdot, \cdot)$ is positive semidefinite. The resulting vectors are of the form $\mathbf{v}_k = \sum_i \alpha_{ki} \phi(\mathbf{x}_i)$ and thus implementing a nonlinear embedding into a nonlinear manifold.

2) *ISOMAP*: In [7], Tenenbaum *et al* find a nonlinear embedding that rather than preserving the Euclidean distance between points on a manifold, preserves the geodesic distance between points on the manifold. Similar to MDS where a lower dimensional embedding is found to preserve the Euclidean distances of high dimensional data, ISOMAP finds a lower dimensional embedding that preserves the geodesic distances between high-dimensional data points.

3) *Laplacian Eigenmaps*: Belkin and Niyogi’s Laplacian eigenmaps dimensionality reduction algorithm [8] takes a different approach. They consider a nonlinear mapping f that minimize the Laplacian

$$\arg \min_{\|f\|_{L^2(\mathcal{M})}=1} \int_{\mathcal{M}} \|\nabla f\|^2. \quad (3)$$

Since the manifold is not available but only data point on it are, the lower dimensional embedding is found by minimizing the graph Laplacian given by

$$\sum_{i=1}^n w_{ij} \|\mathbf{y}_i - \mathbf{y}_j\|^2, \quad (4)$$

where w_{ij} is the ij th element of the adjacency matrix which is constructed as follows: For $k \in \mathbb{N}$, a k -nearest neighbors graph is constructed with the points in \mathcal{X}_n as the graph vertices. Each point \mathbf{x}_i is connected to its k -nearest neighboring points. Note that it suffices that either \mathbf{x}_i is among \mathbf{x}_j 's k -nearest neighbors or \mathbf{x}_j is among \mathbf{x}_i 's k -nearest neighbors for \mathbf{x}_i and \mathbf{x}_j to be connected. For a fixed scale parameter $\epsilon > 0$, the weight associated with the two points \mathbf{x}_i and \mathbf{x}_j satisfies

$$w_{ij} = \begin{cases} \exp\{-\|\mathbf{x}_i - \mathbf{x}_j\|^2/\epsilon\} & \text{if } \mathbf{x}_i \text{ and } \mathbf{x}_j \text{ are connected} \\ 0 & \text{otherwise.} \end{cases}$$

III. CLASSIFICATION CONSTRAINED DIMENSIONALITY REDUCTION

A. Statistical framework

To put the problem in a classification context, we consider the following model. Let $\mathcal{X}_n = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$ be a set of n points sampled from an m -dimensional submanifold $\mathcal{M} \subseteq \mathbb{R}^d$. Each point $\mathbf{x}_i \in \mathcal{M}$ is associated with a class label $c_i \in \mathcal{A} = \{0, 1, 2, \dots, L\}$, where $c_i = 0$ corresponds to the case of unlabeled data. We assume that pairs $(\mathbf{x}_i, c_i) \in \mathcal{M} \times \mathcal{A}$ are i.i.d. drawn from a joint distribution

$$P(\mathbf{x}, c) = p_x(\mathbf{x}|c)P_c(c) = P_c(c|\mathbf{x})p_x(\mathbf{x}), \quad (5)$$

where $p_x(\mathbf{x}) > 0$ and $p_x(\mathbf{x}|c) > 0$ (for $\mathbf{x} \in \mathcal{M}$) are the marginal and the conditional probability density functions, respectively, satisfying $\int_{\mathcal{M}} p_x(\mathbf{x})d\mathbf{x} = 1$, $\int_{\mathcal{M}} p_x(\mathbf{x}|c)d\mathbf{x} = 1$ and $P_c(c) > 0$ and $P_c(c|\mathbf{x}) > 0$ are the a priori and a posteriori probability mass functions of the class label, respectively, satisfying $\sum_c P_c(c) = 1$ and $\sum_c P_c(c|\mathbf{x}) = 1$. While we consider unlabeled points of the form $(\mathbf{x}_i, 0)$ similar labeled points, we still make the following distinction. Consider the following mechanism for generating an unlabeled point. First, a class label $c \in \{1, 2, \dots, L\}$ is generated from the labeled a priori probability mass function $P'_c(c) = P(c|c \text{ is labeled}) = P_c(c)/\sum_{c'=1}^L P_c(c')$. Then \mathbf{x}_i is generated according to $p_x(\mathbf{x}|c)$. To treat c as an unobserved label, we marginalize $P(\mathbf{x}, c|c \text{ is labeled}) = p_x(\mathbf{x}|c)P'_c(c)$ over c :

$$p_x(\mathbf{x}|c = 0) = \sum_{q=1}^L p_x(\mathbf{x}|c = q)P'_c(q) = \frac{\sum_{q=1}^L p_x(\mathbf{x}|c = q)P_c(q)}{\sum_{c'=1}^L P_c(c')}. \quad (6)$$

This suggests that the conditional PDF of unlabeled points $f_x(\mathbf{x}|c = 0)$ is uniquely determined by the class priors and the conditionals for labeled point. We would like to point out that this is one

of few treatments that can be offered for unlabeled point. For example, in anomaly detection, one may want to associate the unlabeled point with contaminated data points, which can be represented as a density mixture of $p_x(\mathbf{x}|c=0)$ and $\gamma(\mathbf{x})$ (e.g., $\gamma(\mathbf{x})$ is uniform in \mathcal{X}).

In classification constraint dimensionality reduction, our goal is to obtain a lower-dimensional embedding $\mathcal{Y}_n = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n\}$ (where $\mathbf{y}_i \in \mathbb{R}^m$ with $m < d$) that preserves local geometry and that encourages clustering of points of the same class label. Alternatively, we would like to find a mapping $\mathbf{f}(\mathbf{x}, c) : \mathcal{M} \times \mathcal{A} \rightarrow \mathbb{R}^m$ for which $\mathbf{y}_i = \mathbf{f}(\mathbf{x}_i, c_i)$ that is smooth and that clusters points of the same label.

We introduce the class label indicator for data point \mathbf{x}_i as $c_{ki} = I(c_i = k)$, for $k = 1, 2, \dots, L$ and $i = 1, 2, \dots, n$. Note that when point \mathbf{x}_i is unlabeled $c_{ki} = 0$ for all k . Using the class indicator, we can write the number of point in class k as $n_k = \sum_{i=1}^n c_{ki}$. If all points are labeled, then $n = \sum_{k=1}^L n_k$.

B. Linear dimensionality reduction for classification

1) *LDA*: Restricting the discussion to linear maps, one can extend PCA to take into account label information using the multi-class extension to Fisher's linear discriminant analysis (LDA). Instead of maximizing the data covariance matrix, LDA maximizes the ratio of the between-class-covariance to within-class-covariance. In other words, we obtain a linear transformation $\mathbf{y}_i = f(\mathbf{x}_i, c_i) = \mathbf{A}\mathbf{x}_i$ with matrix \mathbf{A} that is the solution to the following maximization:

$$\max_{\mathbf{A}} \text{tr}\{\mathbf{A}\mathbf{C}_B\mathbf{A}^T\} \quad \text{s.t.} \quad \mathbf{A}\mathbf{C}_W\mathbf{A}^T = \mathbf{I}, \quad (7)$$

where

$$\mathbf{C}_B = \frac{1}{n} \sum_{k=1}^L n_k (\bar{\mathbf{x}}^{(k)} - \bar{\mathbf{x}})(\bar{\mathbf{x}}^{(k)} - \bar{\mathbf{x}})^T$$

is the between-class-covariance matrix, $\bar{\mathbf{x}}^{(k)} = \sum_i c_{ki} \mathbf{x}_i / n_k$ is the k th class center, $\bar{\mathbf{x}} = \sum_i \mathbf{x}_i / n$ is the center point of the dataset,

$$\mathbf{C}_W = \frac{1}{n} \sum_{k=1}^L n_k \mathbf{C}_W^{(k)},$$

is the within-class-covariance, and

$$\mathbf{C}_W^{(k)} = \frac{\sum_{i=1}^n c_{ki} (\mathbf{x}_i - \bar{\mathbf{x}}^{(k)})(\mathbf{x}_i - \bar{\mathbf{x}}^{(k)})^T}{n_k}$$

is within-class- k covariance matrix. In Fig. 1, LDA selects an embedding that projects the data onto \mathbf{e}_2 since the maximum distance between classes is achieved along with a minimum class variance when projecting the data onto \mathbf{e}_2 . We are interested in exploring a strategy that maximizes class separation in the lower dimensional embedding.

2) *Marginal Fisher Analysis*: Recent work [15], presents the marginal Fisher analysis (MFA), which is a method that minimizes the ratio between intraclass compactness and interclass separability. In its basic formulation MFA is a linear embedding, in which $\mathbf{y}_i = \mathbf{A}\mathbf{x}_i$. Another aspect of the method is that it considers two classes. The kernel trick is used to provide a nonlinear extension to MFA. To construct the cost function, two quantities are of interest: intraclass compactness and interclass separability. The intraclass compactness can be written as

$$\sum_{i,j} w_{ij} \|\mathbf{y}_i - \mathbf{y}_j\|^2, \quad (8)$$

where w_{ij} is given by

$$w_{ij} = \left(\sum_k c_{ki} c_{kj} \right) I(\mathbf{x}_i \in N_{k_1}^+(\mathbf{x}_j) \text{ or } \mathbf{x}_j \in N_{k_1}^+(\mathbf{x}_i)) \quad (9)$$

and $N_k^+(\mathbf{x})$ denote the k -nn neighborhood of \mathbf{x} within the same class as \mathbf{x} . Note that the term $\sum_k c_{ki} c_{kj}$ is one if \mathbf{x}_i and \mathbf{x}_j have the same label and zero otherwise. Similarly, the interclass separability can be written as

$$\sum_{i,j} w_{ij} \|\mathbf{y}_i - \mathbf{y}_j\|^2, \quad (10)$$

where w_{ij} is given by

$$w_{ij} = \left(1 - \sum_k c_{ki} c_{kj} \right) I(\mathbf{x}_i \in N_{k_2}^-(\mathbf{x}_j) \text{ or } \mathbf{x}_j \in N_{k_2}^-(\mathbf{x}_i)) \quad (11)$$

and $N_k^-(\mathbf{x})$ denote the k -nn neighborhood of \mathbf{x} outside the class of \mathbf{x} .

IV. DIMENSIONALITY REDUCTION FOR CLASSIFICATION ON NONLINEAR MANIFOLDS

Here, we review the CCDR algorithm [14] and its extension to multi-class classification.

To cluster lower dimensional embedded points of the same label we associate each class with a class center namely $\mathbf{z}_k \in \mathbb{R}^m$. We construct the following cost function:

$$J(\mathcal{Z}_L, \mathcal{Y}_n) = \sum_{ki} c_{ki} \|\mathbf{z}_k - \mathbf{y}_i\|^2 + \frac{\beta}{2} \sum_{ij} w_{ij} \|\mathbf{y}_i - \mathbf{y}_j\|^2, \quad (12)$$

where $\mathcal{Z}_L = \{z_1, \dots, z_L\}$ and $\beta \geq 0$ is a regularization parameter. We consider two terms on the RHS of (12). The first term corresponds to the concentration of points of the same label around their respective class center. The second term is as in (4) or as in Laplacian Eigenmaps [8] and controls the smoothness of the embedding over the manifold. Large values of β produce an embedding that ignores class labels and small values of β produce an embedding that ignores the manifold structure. Training data points will tend to collapse into the class centers, allowing many classifiers to produce perfect classification on the training data without being able to control the generalization error (i.e., classification error of the unlabeled data). Our goal is to find \mathcal{Z}_L and \mathcal{Y}_n that minimize the cost function in (12).

Let \mathbf{C} be the $L \times n$ class membership matrix with c_{ki} as its ki -th element, $\mathbf{Z} = [z_1, \dots, z_L, \mathbf{y}_1, \dots, \mathbf{y}_n]$, and $\mathbf{0}$ be the $L \times L$ all zeroes matrix and

$$\mathbf{G} = \begin{bmatrix} \mathbf{0} & \mathbf{C} \\ \mathbf{C}^T & \beta \mathbf{W} \end{bmatrix}.$$

Minimization over \mathbf{Z} of the cost function in (12) can be expressed as

$$\begin{aligned} \min_{\substack{\mathbf{Z} \mathbf{D} \mathbf{1} = \mathbf{0} \\ \mathbf{Z} \mathbf{D} \mathbf{Z}^T = \mathbf{I}}} \quad & \text{tr}(\mathbf{Z} \mathbf{L} \mathbf{Z}^T), \end{aligned} \quad (13)$$

where $\mathbf{D} = \text{diag}\{\mathbf{G} \mathbf{1}\}$ and $\mathbf{L} = \mathbf{D} - \mathbf{G}$. To prevent the lower-dimensional points and the class centers from collapsing into a single point at the origin, the regularization $\mathbf{Z} \mathbf{D} \mathbf{Z}^T = \mathbf{I}$ is introduced. The second constraint $\mathbf{Z} \mathbf{D} \mathbf{1} = \mathbf{0}$ is constructed to prevent a degenerate solution, e.g., $z_1 = \dots = z_L = \mathbf{y}_1 = \dots = \mathbf{y}_n$. This solution may occur since $\mathbf{1}$ is in the null-space of the Laplacian \mathbf{L} operator, i.e., $\mathbf{L} \mathbf{1} = \mathbf{0}$. The solution to (13) can be expressed in term of the following generalized eigendecomposition

$$\mathbf{L}^{(n)} \mathbf{u}_k^{(n)} = \lambda_k^{(n)} \mathbf{D}^{(n)} \mathbf{u}_k^{(n)}, \quad (14)$$

where $\lambda_k^{(n)}$ is the k th eigenvalue and $\mathbf{u}_k^{(n)}$ is its corresponding eigenvector. Note that we include $^{(n)}$ to emphasize the dependence on the n data points. Without loss of generality we assume $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_{n+L}$. Specifically, matrix \mathbf{Z} is given by $[\mathbf{u}_2, \mathbf{u}_3, \dots, \mathbf{u}_{m+1}]^T$, where the first L columns correspond to the coordinates of the class centers, i.e., $z_k = \mathbf{Z} \mathbf{e}_k$, and the following n columns determine the embedding of the n data points, i.e., $\mathbf{y}_t = \mathbf{Z} \mathbf{e}_{L+t}$. We use \mathbf{e}_i to denote the canonical vector such that $[\mathbf{e}_i]_s = 1$ for element $s = i$ and zero otherwise.

A. Classification and computational complexity

In classification, the goal is to find a classifier $a_x(\mathbf{x}) : \mathcal{M} \rightarrow \mathcal{A}$ based on the training data that minimizes the generalization error:

$$\hat{a} = \arg \min_{a \in \mathcal{F}} E[I(a(\mathbf{x}) \neq a)], \quad (15)$$

where the expectation is taken w.r.t. the pair (\mathbf{x}, a) . Since only samples from the joint distribution of \mathbf{x} and a are available, we replace the expectation with a sample average w.r.t. the training data $\frac{1}{n} \sum_{i=1}^n I(a(\mathbf{x}_i) \neq a_i)$. During the minimization, we search over a set of classifiers $a_x(\mathbf{x}) : \mathcal{M} \subseteq \mathbb{R}^d \rightarrow \mathcal{A}$, which is defined over a domain in \mathbb{R}^d . In our framework, we suggest replacing a classifier $a_x(\mathbf{x}) : \mathcal{M} \subseteq \mathbb{R}^d \rightarrow \mathcal{A}$ with dimensionality reduction via CCDR $f(\mathbf{x}) : \mathcal{M} \subseteq \mathbb{R}^d \rightarrow \mathbb{R}^m$ followed by a classifier on the lower-dimensional space $a_y(\mathbf{y}) : \mathbb{R}^m \rightarrow \mathcal{A}$, i.e., $a_x = a_y \circ f$. The first advantage is that the search space for the minimization in (15) defined over a d -dimensional space can be reduced to an m -dimensional space. This results in significant savings in computational complexity if the complexity associated with the process of obtaining f can be made low. In general, the classifier set \mathcal{F} has to be rich enough to attain a lower generalization error. The other advantage of our method lies in the fact that CCDR is designed to cluster points of the same label thus allowing for a linear classifier or other low complexity classifiers. Therefore, further reduction in the size of class \mathcal{F} can be achieved in addition to the reduction due to a lower-dimensional domain. To classify a new data point, one has to apply CCDR to a new data point. If it is done brute force, the point is added to the set of training points with no label a new matrix W' is formed and an eigendecomposition is carried out.

When performing CCDR, each of the $n(n-1)/2$ terms of the form $\{\|\mathbf{x}_i - \mathbf{x}_j\|^2\}$ requires one summation and d multiplications leading to computational complexity of the order $O(dn^2)$. Construction of a K -nearest neighbors graph requires $O(kn)$ comparisons per point and therefore a total of $O(kn^2)$. The total number of operations involved in constructing the graph is therefore $O((k+d)n^2)$. Next, an eigendecomposition is applied to W' , which is an $(L+n) \times (L+n)$ matrix. The associated computation complexity is $O(n^3)$. Therefore, the overall computational complexity of CCDR is $O(n^3)$. This holds for both training and classification as explained earlier. We are interested in reducing computational complexity in training the classifier and in classification. For that purpose, we consider an out-of-sample extension of CCDR.

V. OUT-OF-SAMPLE EXTENSION

We start by rearranging the generalized eigendecomposition of the Laplacian in (14) as

$$\mathbf{G}^{(n)} \mathbf{u}_l^{(n)} = (1 - \lambda_l^{(n)}) \mathbf{D}^{(n)} \mathbf{u}_l^{(n)}, \quad (16)$$

and recall that $\mathbf{u}_l^{(n)} = [\mathbf{z}_1(l), \mathbf{z}_1(l), \dots, \mathbf{z}_1(l), \mathbf{y}_1(l), \mathbf{y}_2(l), \dots, \mathbf{y}_n(l)]^T$. Since we consider an m -dimensional embedding, we are only interested in eigenvectors $\mathbf{u}_2, \dots, \mathbf{u}_{m+1}$. The $L+i$ equation (row) for $i = 1, 2, \dots, n$ in the eigendecomposition in (16) can be written as

$$\mathbf{y}_i^{(n)}(l) = \frac{1}{1 - \lambda_l^{(n)}} \frac{\sum_k c_{ki} \mathbf{z}_k^{(n)}(l) + \beta \sum_j K(\mathbf{x}_i, \mathbf{x}_j) \mathbf{y}_j^{(n)}(l)}{\sum_k c_{ki} + \beta \sum_j K(\mathbf{x}_i, \mathbf{x}_j)}. \quad (17)$$

Similarly, the k th equation (row) of (16) for $k = 1, 2, \dots, L$ is given by

$$\mathbf{z}_k^{(n)}(l) = \frac{\sum_i c_{ki} \mathbf{y}_i^{(n)}(l)}{(1 - \lambda_l^{(n)}) n_k}. \quad (18)$$

Our interest is in finding a mapping $\mathbf{f}(\mathbf{x}, c)$ that in addition to mapping every \mathbf{x}_i to \mathbf{y}_i , can perform an out-of-sample extension, i.e., is well-defined outside the set \mathcal{X} . We consider the following out-of-sample extension expression

$$\mathbf{f}_l^{(n)}(\mathbf{x}, c) = \frac{1}{1 - \lambda_l^{(n)}} \frac{I(c \neq 0) \mathbf{z}_c^{(n)}(l) + \beta \sum_j K(\mathbf{x}, \mathbf{x}_j) \mathbf{y}_j^{(n)}(l)}{I(c \neq 0) + \beta \sum_j K(\mathbf{x}, \mathbf{x}_j)}, \quad (19)$$

where $\mathbf{z}^{(n)}$ is the same as in (18). This formula can be explain as follows. First, the lower dimensional embedding $\mathbf{y}_1^{(n)}, \dots, \mathbf{y}_n^{(n)}$ and the class centers $\mathbf{z}_1^{(n)}, \dots, \mathbf{z}_L^{(n)}$ are obtained through an the eigendecomposition in (16). Then, the embedding outside the sample set \mathcal{X} is calculated via (19). By comparison of $\mathbf{f}_l^{(n)}(\mathbf{x}_i, c_i)$ evaluated through (19) with (17), we have $\mathbf{f}_l^{(n)}(\mathbf{x}_i, c_i) = \mathbf{y}_i^{(n)}(l)$. This suggests that the out-of-sample extension coincides with the solution, we already have for the mapping at the the data points \mathcal{X} . Moreover, using this result one can replace all $\mathbf{y}_i^{(n)}$ with $\mathbf{f}_l^{(n)}(\mathbf{x}_i, c_i)$ in (19) and obtain the following generalization of the eigendecomposition in (16):

$$\mathbf{f}_l^{(n)}(\mathbf{x}, c) = \frac{1}{1 - \lambda_l^{(n)}} \frac{I(c \neq 0) \mathbf{z}_c^{(n)}(l) + \beta \sum_j K(\mathbf{x}, \mathbf{x}_j) \mathbf{f}_l^{(n)}(\mathbf{x}_j, c_j)}{I(c \neq 0) + \beta \sum_j K(\mathbf{x}, \mathbf{x}_j)}, \quad (20)$$

and

$$\mathbf{z}_k^{(n)}(l) = \frac{\sum_i c_{ki} \mathbf{f}_l^{(n)}(\mathbf{x}_i, c_i)}{(1 - \lambda_l^{(n)}) n_k}. \quad (21)$$

In [18], it is proposed that if the out-of-sample solution to the eigendecomposition problem associated with kernel PCA converges, it is given by the solution to the asymptotic equivalent of the eigendecomposition. Using similar machinery, we can provide a similar result suggesting that if $\mathbf{f}_l^{(n)}(\mathbf{x}, c) \rightarrow \mathbf{f}_l^{(\infty)}(\mathbf{x}, c)$ as $n \rightarrow \infty$, then the asymptotic equivalent to (20) and (21) should provide the solution to the limit of $\mathbf{f}_l^{(n)}(\mathbf{x}, c)$. The asymptotic analogues to (17) and (18) are described in the following. The mapping for labeled data $f_l(\mathbf{x}, c) : \mathcal{M} \times \mathcal{A} \rightarrow \mathbb{R}$ for $c = 0, 1, 2, \dots, L$ equivalent to equation (17) is

$$f_l(\mathbf{x}, c) = \frac{1}{1 - \lambda_l} \frac{I(c \neq 0)z_c(l) + \beta' \sum_{c'=0}^L \int_{\mathcal{M}} K(\mathbf{x}, \mathbf{x}') f_l(\mathbf{x}', c') P(\mathbf{x}', c') d\mathbf{x}'}{I(c \neq 0) + \beta' \int_{\mathcal{M}} K(\mathbf{x}, \mathbf{x}') p(\mathbf{x}') d\mathbf{x}'} \quad (22)$$

where $z_c(l)$ for $c = 1, 2, \dots, L$ is equivalent to (18)

$$z_c(l) = \frac{\int_{\mathcal{M}} f_l(\mathbf{x}, c) p(\mathbf{x}|c) d\mathbf{x}}{1 - \lambda_l}, \quad (23)$$

and $\beta' = \beta n$. Since we are interested in an m -dimensional embedding, we consider only $l = 1, 2, \dots, m$, i.e., the eigenvectors that correspond to the m smallest eigenvalues. To guarantee that the relevant eigenvectors are unique (up to a multiplicative constant), we require $\lambda_1 < \lambda_2 < \dots < \lambda_{m+1} \leq \lambda_{m+2} \leq \dots \leq \lambda_n$.

The out-of-sample extension given by (19), can be useful in a couple of scenarios. The first, is in classification of new unlabeled samples. We assume that $\{\mathbf{y}_j\}_{j=1}^n$, $\{\mathbf{z}_k\}_{k=1}^L$, and $\{\lambda_l\}_{l=1}^m$ are already obtained based on labeled (or partially labeled) training data and we would like to embed a new unlabeled data point. We consider using (19) with $c = 0$, i.e., we can use $\mathbf{f}(\mathbf{x}, 0)$ to map a new sample \mathbf{x} to \mathbb{R}^m . The obvious immediate advantage is the savings in computational complexity as we avoid performing additional eigendecomposition that includes the new point.

The second scenario involves the out-of-sample extension for labeled data. The goal here is not to classify the data since the label is already available. Instead, we are interested in the training phase in the case of large n for which the eigendecomposition is infeasible. In this case, a large amount of labeled training data is available but due to the heavy computational complexity associated with the eigendecomposition in (14) (or by (16)), the data cannot be processed. In this case, we are interested in developing a resampling method, which integrates $\mathbf{f}_l^{(n)}(\mathbf{x}, c)$ obtained for different subsamples of the complete data set.

A. Classification Algorithms

We consider three widespread algorithms: k -nearest neighbors, linear classification, and neural networks. A standard implementation of k -nearest neighbors was used, see [1, p. 415]. The linear classifier we implemented is given by

$$\hat{c}(\mathbf{y}) = \arg \max_{c \in \{\mathcal{A}_1, \dots, \mathcal{A}_L\}} \mathbf{y}^T \boldsymbol{\alpha}^{(c)} + \alpha_0^{(c)}$$

$$[\boldsymbol{\alpha}^{(\mathcal{A}_k)}, \alpha_0^{(\mathcal{A}_k)}] = \arg \min_{[\boldsymbol{\alpha}, \alpha_0]} \sum_{i=1}^n (\mathbf{y}_i^T \boldsymbol{\alpha} + \alpha_0 - c_{ki})^2,$$

for $k = 1, \dots, L$. The neural network we implemented is a three-layer neural network with d elements in the input layer, $2d$ elements in the hidden layer, and 6 elements in the output layer (one for each class). Here d was selected using the common PCA procedure, as the smallest dimension that explains 99.9% of the energy of the data. A gradient method was used to train the network coefficients with 2000 iterations. The neural net is significantly more computationally burdensome than either linear or k -nearest neighbors classifications algorithms.

B. Data Description

In this section, we examine the performance of the classification algorithms on the benchmark label classification problem provided by the Landsat MSS satellite imagery database [19]. Each sample point consists of the intensity values of one pixel and its 8 neighboring pixels in 4 different spectral bands. The training data consists of 4435 36-dimensional points of which, 1072 are labeled as 1) red soil, 479 as 2) cotton crop, 961 as 3) grey soil, 415 as 4) damp grey soil, 470 are labeled as 5) soil with vegetation stubble, and 1038 are labeled as 6) very damp grey soil. The test data consists of 2000 36-dimensional points of which, 461 are labeled as 1) red soil, 224 as 2) cotton crop, 397 as 3) grey soil, 211 as 4) damp grey soil, 237 are labeled as 5) soil with vegetation stubble, and 470 are labeled as 6) very damp grey soil. In the following, each classifier is trained on the training data and its classification is evaluated based on the entire sample test data. In Table I, we present “best case” performance of neural networks, linear classifier, and k -nearest neighbors in three cases: no dimensionality reduction, dimensionality reduction via PCA, and dimensionality reduction via CCDR. The table presents the minimum probability of error achieved by varying the tuning parameters of the classifiers.

The benefit of using CCDR is obvious and we are prompted to further evaluate the performance gains attained using CCDR.

	Neural Net.	Lin.	k -nearest neigh.
No dim. reduc.	83 %	22.7 %	9.65 %
PCA	9.75 %	23 %	9.35 %
CCDR	8.95 %	8.95 %	8.1 %

TABLE I
CLASSIFICATION ERROR PROBABILITY

C. Regularization Parameter β

As mentioned earlier, the CCDR regularization parameter β controls the contribution of the label information versus the contribution of the geometry described by the sample. We apply CCDR to the 36-dimensional data to create a 14-dimensional embedding by varying β over a range of values. For justification of our choice of $d = 14$ dimensions see Section V-D. In the process of computing the weights w_{ij} for the algorithm, we use k -nearest neighbors with $k = 4$ to determine the local neighborhood. Fig. 2 shows the classification error probability (dashed lines) for the linear classifier vs. β after preprocessing the data using CCDR with $k = 4$ and dimension 14. We observe that for a large range of β the average classification error probability is greater than 0.09 but smaller than 0.095. This performance competes with the performance of k -nearest neighbors applied to the high-dimensional data, which is presented in [1] as the leading classifier for this benchmark problem. Another observation is that for small values of β (i.e., $\beta < 0.1$) the probability of error is constant. For such small value of β , classes in the lower-dimensional embedding are well-separated and are well-concentrated around the class centers. Therefore, the linear classifier yields perfect classification on the training set and fairly low constant probability of error on the test data is attained for low value of β . When β is increased, we notice an increase in the classification error probability. This is due to the fact that the training data become non separable by any linear classifier as β increases.

We perform a similar study of classification performance for k -nearest neighbors. In Fig. 2, classification probability error is plotted (dotted lines) vs. β . Here, we observed that an average

error probability of 0.086 can be achieved for $\beta \approx 0.5$. Therefore, k -nearest neighbors preceded by CCDD outperforms the straightforward k -nearest neighbors algorithm. We also observe that when β is decreased the probability of error is increased. This can be explained as due to the ability of k -nearest neighbors to utilize local information, i.e., local geometry. This information is discarded when β is decreased.

We conclude that CCDD can generate lower-dimensional data that is useful for global classifiers, such as the linear classifier, by using a small value of β , and also for local classifiers, such as k -nearest neighbors, by using a larger value β and thus preserving local geometry information.

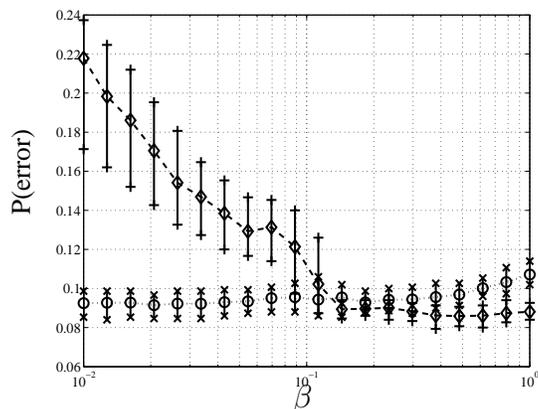


Fig. 2. Probability of incorrect classification vs. β for a linear classifier (dotted line \circ) and for the k -nearest neighbors algorithm (dashed line \diamond) preprocessed by CCDD. 80% confidence intervals are presented as \times for the linear classifier and as $+$ for the k -nearest neighbors algorithm.

D. Dimension Parameter

While the data points in \mathcal{X}_n may lie on a manifold of a particular dimension, the actual dimension required for classification may be smaller. Here, we examine classification performance as a function of the CCDD dimension. Using the entropic graph dimension estimation algorithm in [20], we obtain the following estimated dimension for each class:

class	1	2	3	4	5	6
dimension	13	7	13	10	6	13

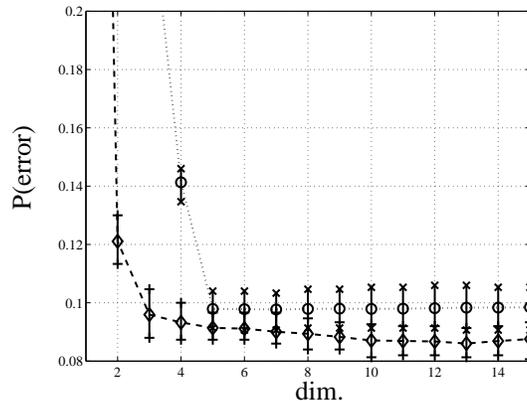


Fig. 3. Probability of incorrect classification vs. CCDR’s dimension for a linear classifier (dotted line \circ) and for the k -nearest neighbors algorithm (dashed line \diamond) preprocessed by CCDR. 80% confidence intervals are presented as \times for the linear classifier and as $+$ for the k -nearest neighbors algorithm.

Therefore, if an optimal nonlinear embedding of the data could be found, we suspect that a dimension greater than 13 may not yield significant improvement in classification performance. Since CCDR does not necessarily yield an optimal embedding, we choose CCDR embedding dimension as $d = 14$ in Section V-C.

In Fig. 3, we plot the classification error probability (dotted line) vs. CCDR dimension and its confidence interval for a linear classifier. We observed decrease in error probability as the dimension increases. When the CCDR dimension is greater than 5, the error probability seems fairly constant. This is an indication that CCDR dimension of 5 is sufficient for classification if one uses the linear classifier with $\beta = 0.5$, i.e., linear classifier cannot exploit geometry.

We also plot the classification error probability (dashed line) vs. CCDR dimension and its confidence interval for k -nearest neighbors classifier. Generally, we observe decrease in error probability as the dimension increases. When the CCDR dimension is greater than 5, the error probability seems fairly constant. When CCDR dimension is three, classifier error is below 0.1. On the other hand, minimum possibility of error obtained at CCDR dimension 12-14. This is remarkable agreement with the dimension estimate of 13 obtained using the entropic graph algorithm of [20].

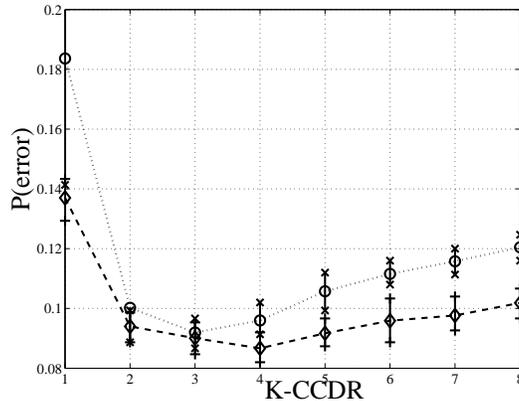


Fig. 4. Probability of incorrect classification vs. CCDR’s k -nearest neighbors parameter for a linear classifier (dotted line \circ) and for the k -nearest neighbors algorithm (dashed line \diamond) preprocessed by CCDR. 80% confidence intervals are presented as \times for the linear classifier and as $+$ for the k -nearest neighbors algorithm.

E. CCDR’s k -Nearest Neighbors Parameter

The last parameter we examine is the CCDR’s k -nearest neighbors parameter. In general, as k increases non-local distances are included in the lower-dimensional embedding. Hence, very large k prevents the flexibility necessary for dimensionality reduction on (globally) non-linear (but locally linear) manifolds.

In Fig. 4, the classification probability of error for the linear classifier (dotted line) is plotted vs. the CCDR’s k -nearest neighbors parameter. A minimum is obtained at $k = 3$ with probability of error of 0.092. The classification probability of error for k -nearest neighbors (dashed line) is plotted vs. the CCDR’s k -nearest neighbors parameter. A minimum is obtained at $k = 4$ with probability of error of 0.086.

VI. CONCLUSION

In this paper, we presented the CCDR algorithm for multiple classes. We examined the performance of various classification algorithms applied after CCDR for the Landsat MSS imagery dataset. We showed that for a linear classifier, decreasing β yields improved performance and for a k -nearest neighbors classifier, increasing β yields improved performance. We demonstrated that both classifiers have improved performance on the much smaller dimension of CCDR embedding space than when applied to the original high-dimensional data. We also explored the effect of k

in the k -nearest neighbors construction of CCDR weight matrix on classification performance. CCDR allows reduced complexity classification such as the linear classifier to perform better than more complex classifiers applied to the original data. We are currently pursuing an out-of-sample extension to the algorithm that does not require rerunning CCDR on test and training data to classify new test point.

REFERENCES

- [1] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning Data Mining, Inference, and Prediction*, ser. Springer Series in Statistics. New York: Springer Verlag, 2000.
- [2] A. K. Jain and R. C. Dubes, *Algorithms for clustering data*. New Jersey: Prentice Hall, 1998.
- [3] W. S. Torgerson, “Multidimensional scaling: I. theory and method,” *Psychometrika*, vol. 17, pp. 401–419, 1952.
- [4] T. F. Cox and M. A. A. Cox, *Multidimensional Scaling*, 2nd ed., ser. Monographs on Statistics and Applied Probability. London: Chapman & Hall/CRC, 2000, vol. 88.
- [5] B. Schölkopf, A. J. Smola, and K.-R. Müller, “Nonlinear component analysis as a kernel eigenvalue problem.” *Neural Computation*, vol. 10, no. 5, pp. 1299–1319, 1998.
- [6] A. Aizerman, E. M. Braverman, and L. I. Rozoner, “Theoretical foundations of the potential function method in pattern recognition learning,” *Automation and Remote Control*, vol. 25, pp. 821–837, 1964.
- [7] J. B. Tenenbaum, V. D. Silva, and J. C. Langford, “A global geometric framework for nonlinear dimensionality reduction.” *Science*, vol. 290, no. 5500, pp. 2319–2323, 2000.
- [8] M. Belkin and P. Niyogi, “Laplacian eigenmaps for dimensionality reduction and data representation,” *Neural Computation*, vol. 15, no. 6, pp. 1373–1396, June 2003.
- [9] S. T. Roweis and L. K. Saul, “Nonlinear dimensionality reduction by locally linear embedding,” *Science*, vol. 290, no. 5500, pp. 2323–2326, December 2000.
- [10] D. L. Donoho and C. Grimes, “Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 100, no. 10, pp. 5591–5596, May 2003.
- [11] Z. Zhang and H. Zha, “Principal manifolds and nonlinear dimensionality reduction via tangent space alignment,” *SIAM J. Sci. Comput.*, vol. 26, no. 1, pp. 313–338, 2005.
- [12] R. Coifman and S. Lafon, “Diffusion maps,” *Applied and Computational Harmonic Analysis: Special issue on Diffusion Maps and Wavelets*, vol. 21, pp. 5–30, July 2006.
- [13] K. Q. Weinberger and L. K. Saul, “Unsupervised learning of image manifolds by semidefinite programming,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR-04)*, vol. 2, Washington D.C., 2004, pp. 988–995.
- [14] J. A. Costa and A. O. H. III, “Classification constrained dimensionality reduction,” in *Proc. IEEE Intl. Conf. on Acoust., Speech, and Signal Processing*, vol. 5, March 2005, pp. 1077–1080.
- [15] S. Yan, D. Xu, B. Zhang, H.-J. Zhang, Q. Yang, and S. Lin, “Graph embedding and extensions: A general framework for dimensionality reduction,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 29, no. 1, pp. 40–51, Jan. 2007.

- [16] R. Raich, J. A. Costa, and A. O. Hero III, "On dimensionality reduction for classification and its application," in *Proc. IEEE Intl. Conf. Acoust., Speech, Signal Processing*, vol. 5, Toulouse, France, May 2006, pp. 917–920.
- [17] Y. Bengio, J.-F. Paiement, and P. Vincent, "Out-of-sample extensions for LLE, isomap, MDS, eigenmaps, and spectral clustering," Département d'Informatique et Recherche Opérationnelle Université de Montréal Montréal, Québec, Canada, H3C 3J7, Tech. Rep., 2003, technical Report 1238, Département d'Informatique et Recherche Opérationnelle.
- [18] Y. Bengio, O. Delalleau, N. Le Roux, J. F. Paiement, P. Vincent, and M. Ouimet, "Learning eigenfunctions links spectral embedding and kernel PCA," *Neural Computation*, vol. 16, no. 10, pp. 2197–2219, 2004.
- [19] "Satellite image data," available at <http://www.liacc.up.pt/ML/statlog/datasets/satimage/satimage.doc.html>.
- [20] J. A. Costa and A. O. Hero, "Geodesic entropic graphs for dimension and entropy estimation in manifold learning," *IEEE Trans. Signal Processing*, vol. 52, no. 8, pp. 2210–2221, Aug. 2004.

Classification Constrained Dimensionality Reduction

Raviv Raich, *Member, IEEE*, Jose A. Costa, *Member, IEEE*,

Steven B. Damelin, *Senior Member, IEEE*, and Alfred O. Hero III, *Fellow, IEEE*

Abstract

Dimensionality reduction is a topic of recent interest. In this paper, we present the classification constrained dimensionality reduction (CCDR) algorithm to account for label information. The algorithm can account for multiple classes as well as the semi-supervised setting. We present an out-of-sample expressions for both labeled and unlabeled data. For unlabeled data, we introduce a method of embedding a new point as preprocessing to a classifier. For labeled data, we introduce a method that improves the embedding during the training phase using the out-of-sample extension. We investigate classification performance using the CCDR algorithm on hyper-spectral satellite imagery data. We demonstrate the performance gain for both local and global classifiers and demonstrate a 10% improvement of the k -nearest neighbors algorithm performance. We present a connection between intrinsic dimension estimation and the optimal embedding dimension obtained using the CCDR algorithm.

Index Terms

Classification, Computational Complexity, Dimensionality Reduction, Embedding, High Dimensional Data, Kernel, K-Nearest Neighbor, Manifold Learning, Probability, Out-of-Sample Extension.

This work was partially funded by the DARPA Defense Sciences Office under Office of Naval Research contract #N00014-04-C-0437. Distribution Statement A. Approved for public release; distribution is unlimited. S. B. Damelin was supported in part by National Science Foundation grant no. NSF-DMS-0555839 and NSF-DMS-0439734 and by AFRL.

R. Raich is with Oregon State University, Corvallis, OR. A. O Hero III is with the University of Michigan, Ann Arbor. J. A. Costa is with the California Institute of Technology. S. B. Damelin is with the Georgia Southern University.

Classification Constrained Dimensionality Reduction

I. INTRODUCTION

In classification theory, the main goal is to find a mapping from an observation space \mathcal{X} consisting of a collection of points in some containing Euclidean space \mathbb{R}^d , $d \geq 1$ into a set consisting of several different integer valued hypotheses. In some problems, the observations from the set \mathcal{X} lie on a d -dimensional manifold \mathcal{M} and Whitney's theorem tells us that provided that this manifold is smooth enough, there exists an embedding of \mathcal{M} into \mathbb{R}^{2d+1} . This motivates the approach taken by kernel methods in classification theory, such as support vector machines [?] for example. Our interest is in finding an embedding of \mathcal{M} into a lower dimensional Euclidean space.

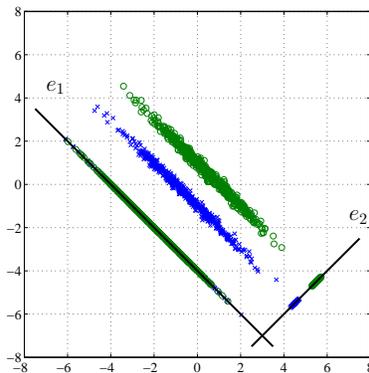


Fig. 1. PCA of a two-classes classification problem.

Dimensionality reduction of high dimensional data, was addressed in classical methods such as principal component analysis (PCA) [?] and multidimensional scaling (MDS) [?], [?]. In PCA, an eigendecomposition of the $d \times d$ empirical covariance matrix is performed and the data points are linearly projected along the $0 < m \leq d$ eigenvectors with the largest eigenvalues. A problem that may occur with PCA for classification is demonstrated in Fig. 1. When the information that is relevant for classification is present only in the eigenvectors associated with the small eigenvalues (e_2 in the figure), removal of such eigenvectors may result in severe degradation

in classification performance. In MDS, the goal is to find a lower dimensional embedding of the original data points that preserves the relative distances between all the data points. The two later methods suffer greatly when the manifold is nonlinear. For example, PCA will not be able to offer dimensionality reduction for classification of two classes lying each on one of two concentric circles.

In [?], a nonlinear extension to PCA is presented. The algorithm is based on the “kernel trick” [?]. Data points are nonlinearly mapped into a feature space, which in general has a higher (or even infinite) dimension as compared with the original space and then PCA is applied to the high dimensional data.

In the paper of Tenenbaum *et al* [?], Isomap, a global dimensionality reduction algorithm was introduced taking into account the fact that data points may lie on a lower dimensional manifold. Unlike MDS, geodesic distances (distances that are measured along the manifold) are preserved by Isomap. Isomap utilizes the classical MDS algorithm, but instead of using the matrix of Euclidean distances, it uses a modified version of it. Each point is connected only to points in its local neighborhood. A distance between a point and another point outside its local neighborhood is replaced with the sum of distances along the shortest path in graph. This procedure modifies the squared distances matrix replacing Euclidian with geodesic distances.

In [?], Belkin and Niyogi present a related Laplacian eigenmap dimensionality reduction algorithm. The algorithm performs a minimization on the weighted sum of squared-distances of the lower-dimensional data. Each weight multiplying the squared-distances of two low-dimensional data points is inversely related to distance between the corresponding two high-dimensional data points. Therefore, small distance between two high-dimensional data points results in small distance between two low-dimensional data points. To preserve the geodesic distances, the weight of the distance between two points that do not share a local neighborhood is set to zero.

We refer the interested reader to the references below and those cited therein for a list of some of the most commonly used additional algorithms within the class of *manifold learning* algorithms and their different advantages relevant to our work. Locally Linear Embedding (LLE) [?], Laplacian Eigenmaps [?], Hessian Eigenmaps (HLLE) [?], Local Space Tangent Analysis [?], Diffusion Maps [?] and Semidefinite Embedding (SDE) [?].

The algorithms mentioned above, consider the problem of learning a lower-dimensional em-

bedding of the data. In classification, such algorithms can be used to preprocess high-dimensional data before performing the classification. This could potentially allow for a lower computational complexity of the classifier. In some cases, dimensionality reduction results increase the computational complexity of the classifier. In fact, support vector machines suggest the opposing strategy: data points are projected onto a higher-dimensional space and classified by a low computational complexity classifier. To guarantee a low computational complexity of the classifier of the low-dimensional data, a classification constrained dimensionality reduction (CCDR) algorithm was introduced in [?]. The CCDR algorithm is an extension of Laplacian eigenmaps [?] and it incorporates class label information into the cost function, reducing the distance between points with similar label. Another algorithm that incorporates label information is the marginal fisher analysis (MFA) [?], in which a constraint on the margin between classes is used to enforce class separation.

In [?] the CCDR algorithm was only studied for two classes and its performance was illustrated for simulated data. In [?], a multi-class extension to the problem was presented. In this paper, we introduce two additional components that make the algorithm computationally viable. The first is an out-of-sample extension for classification of unlabeled test points. Similarly to the out-of-sample extension presented in [?], one can utilize the Nyström formula for classification problems in which label information is available. We study the algorithm performance as its various parameters, (e.g., dimension, label importance, and local neighborhood), are varied. We study the performance of CCDR as preprocessing prior to implementation of several classification algorithms such as k -nearest neighbors, linear classification, and neural networks. We demonstrate a 10% improvement over the k -nearest neighbors algorithm performance benchmark for this dataset. We address the issue of dimension estimation and its effect on classification performance.

The organization of this paper is as follows. Section III presents the multiple-class CCDR algorithm. Section ?? provides a study of the algorithm using the Landsat dataset and Section VI summaries our results.

II. DIMENSIONALITY REDUCTION

Let $\mathcal{X}_n = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$ be a set of n points constrained to lie on an m -dimensional submanifold $\mathcal{M} \subseteq \mathbb{R}^d$. In dimensionality reduction, our goal is to obtain a lower-dimensional

embedding $\mathcal{Y}_n = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n\}$ (where $\mathbf{y}_i \in \mathbb{R}^m$ with $m < d$) that preserves local geometry information such that processing of the lower dimensional embedding \mathcal{Y}_n yields comparable performance to processing of the original data points \mathcal{X}_n . Alternatively, we would like learn the mapping $f : \mathcal{M} \subseteq \mathbb{R}^d \rightarrow \mathbb{R}^m$ that maps every data point \mathbf{x}_i to $\mathbf{y}_i = f(\mathbf{x}_i)$ such that some geometric properties of the high-dimensional data are preserved in the lower dimensional embedding. The first question that comes to mind is how to select f , or more specifically how to restrict the function f so that we can still achieve our goal.

A. Linear dimensionality reduction

1) *PCA*: When principal component analysis (PCA) is used for dimensionality reduction, one considers a linear embedding of the form

$$\mathbf{y}_i = f(\mathbf{x}_i) = \mathbf{A}\mathbf{x}_i,$$

where \mathbf{A} is $m \times d$. This embedding captures the notion of proximity in the sense that close points in the high dimensional space map to close points in the lower dimensional embedding, i.e., $\|\mathbf{y}_i - \mathbf{y}_j\| = \|\mathbf{A}(\mathbf{x}_i - \mathbf{x}_j)\| \leq \|\mathbf{A}\|\|\mathbf{x}_i - \mathbf{x}_j\|$. Let

$$\bar{\mathbf{x}} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i$$

and

$$\mathcal{C}_x = \frac{1}{n} \sum_{i=1}^n (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})^T.$$

Similarly, let

$$\bar{\mathbf{y}} = \frac{1}{n} \sum_{i=1}^n \mathbf{y}_i$$

and

$$\mathcal{C}_y = \frac{1}{n} \sum_{i=1}^n (\mathbf{y}_i - \bar{\mathbf{y}})(\mathbf{y}_i - \bar{\mathbf{y}})^T.$$

Since $\mathbf{y}_i = \mathbf{A}\mathbf{x}_i$, we have $\bar{\mathbf{y}} = \mathbf{A}\bar{\mathbf{x}}$ and $\mathcal{C}_y = \mathbf{A}\mathcal{C}_x\mathbf{A}^T$. In PCA, the goal is to find the projection matrix \mathbf{A} that preserves most of the energy in the original data by solving

$$\max_{\mathbf{A}} \text{tr}\{\mathcal{C}_y(\mathbf{A})\} \quad \text{s.t.} \quad \mathbf{A}\mathbf{A}^T = \mathbf{I},$$

which is equivalent to

$$\max_{\mathbf{A}} \text{tr}\{\mathbf{A}\mathcal{C}_x\mathbf{A}^T\} \quad \text{s.t.} \quad \mathbf{A}\mathbf{A}^T = \mathbf{I}. \quad (1)$$

The solution to (1), is given by $\mathbf{A} = [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_m]^T$, where \mathbf{u}_i is the eigenvector of \mathcal{C}_x corresponding to its i th largest eigenvalue. When the data lies on an m -dimensional hyperplane, the matrix \mathcal{C}_x has only m positive eigenvalues and the rest are zero. Furthermore, every \mathbf{x}_i belongs to $\bar{\mathbf{x}} + \text{span}\{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_m\} \subseteq \mathbb{R}^d$. In this case, the mapping PCA finds $f(\mathbf{x}) = \mathbf{A}\mathbf{x}$ is one-to-one and satisfies $\|f(\mathbf{x}_i) - f(\mathbf{x}_j)\| = \|\mathbf{A}(\mathbf{x}_i - \mathbf{x}_j)\| = \|\mathbf{x}_i - \mathbf{x}_j\|$. Therefore, the lower embedding preserves all the geometry information in the original dataset \mathcal{X} . We would like to point out that PCA can be written as

$$\max_{\{\mathcal{Y}\}} \sum_{i=1}^n \|\mathbf{y}_i - \mathbf{y}_j\|^2 \quad \text{s.t.} \quad \mathbf{y}_i = \mathbf{A}\mathbf{x}_i \text{ and } \mathbf{A}\mathbf{A}^T = \mathbf{I},$$

2) *MDS*: Multidimensional Scaling (MDS) differs from PCA in the way the input is provided to it. While in PCA, the original data \mathcal{X} is provided, the classical MDS requires only the set of all Euclidean pairwise distances $\{\|\mathbf{x}_i - \mathbf{x}_j\|_2\}_{i=1, j>i}^{n-1}$. As MDS uses only pairwise distances, the solution it finds is given up to translation and unitary transformation. Let $\mathbf{x}'_i = \mathbf{x}_i - \mathbf{c}$, the Euclidean distance $\|\mathbf{x}'_i - \mathbf{x}'_j\|$ is the same as $\|\mathbf{x}_i - \mathbf{x}_j\|$. Let \mathbf{U} be an arbitrary unitary matrix \mathbf{U} satisfying $\mathbf{U}^T\mathbf{U} = \mathbf{I}$ and define $\mathbf{x}'_i = \mathbf{U}\mathbf{x}_i$. The distance $\|\mathbf{x}'_i - \mathbf{x}'_j\|$ is equal to $\|\mathbf{U}(\mathbf{x}_i - \mathbf{x}_j)\|$, which by the invariance of the Euclidean norm to a unitary transformation equals to $\|\mathbf{x}_i - \mathbf{x}_j\|$. Denote the pairwise squared-distance matrix by $[\mathbf{D}_2]_{ij} = \|\mathbf{x}_i - \mathbf{x}_j\|^2$. By the definition of Euclidean distance, the matrix \mathbf{D}_2 satisfies

$$\mathbf{D}_2 = \mathbf{1}\phi^T + \phi\mathbf{1}^T - 2\mathbf{X}^T\mathbf{X}, \quad (2)$$

where $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n]$ and $\phi = [\|\mathbf{x}_1\|^2, \|\mathbf{x}_2\|^2, \dots, \|\mathbf{x}_n\|^2]^T$. To verify (2), one can examine the ij -th term of \mathbf{D}_2 and compare with $\|\mathbf{x}_i - \mathbf{x}_j\|^2$. Denote the $n \times n$ matrix $\mathbf{H} = \mathbf{I} - \mathbf{1}\mathbf{1}^T/n$. Multiplying both sides of \mathbf{D}_2 with \mathbf{H} in addition to a factor of $-\frac{1}{2}$, yields

$$-\frac{1}{2}\mathbf{H}\mathbf{D}_2\mathbf{H} = (\mathbf{X}\mathbf{H})^T(\mathbf{X}\mathbf{H}),$$

which is key to MDS, i.e., Cholesky decomposition of $-\frac{1}{2}\mathbf{H}\mathbf{D}_2\mathbf{H}$ yields \mathbf{X} to within a translation and a unitary transformation. Consider the eigendecomposition $-\frac{1}{2}\mathbf{H}\mathbf{D}_2\mathbf{H} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T$. Therefore, a rank d \mathbf{X} can be obtained as $\mathbf{X} = \mathbf{\Lambda}_d^{\frac{1}{2}}\mathbf{U}_d^T$, where $\mathbf{\Lambda}_d = \text{diag}\{\lambda_1, \lambda_2, \dots, \lambda_d\}^{\frac{1}{2}}$ and $\mathbf{U}_d = [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_d]$. Note that $\mathbf{X}\mathbf{H}$ is a translated version of \mathbf{X} , in which every column \mathbf{x}_i is translated to $\mathbf{x}_i - \bar{\mathbf{x}}$.

To use MDS for dimensionality reduction, we can consider a two step process. First a squared-distance matrix \mathbf{D}_2 is obtained from the high-dimensional data \mathcal{X} . Then, MDS is applied to \mathbf{D}_2

to obtain a low-dimensional ($m < d$) embedding by $\mathbf{X}_m = \mathbf{\Lambda}_m^{\frac{1}{2}} \mathbf{U}_m^T = \mathbf{X} \mathbf{U}_m \mathbf{U}_m^T$. In the absence of noise, this procedure provides an affine transformation to the high-dimensional data and thus can be regarded as a linear method.

B. Nonlinear dimensionality reduction

Linear maps are limited as they cannot preserve the geometry of nonlinear manifolds.

1) *Kernel PCA*: Kernel PCA is one of the first methods in dimensionality reduction of data on nonlinear manifolds. The method combines the dimensionality reduction capabilities of PCA on linear manifolds with a nonlinear embedding of data points in a higher (or even infinite) dimensional space using “kernel trick” [?]. In PCA, one finds the eigenvectors satisfying: $\mathcal{C}_x \mathbf{v}_k = \lambda_k \mathbf{v}_k$. Since \mathbf{v}_k can be written as a linear combination of the \mathbf{x}_i ’s: $\mathbf{v}_k = \sum_i \alpha_{ki} (\mathbf{x}_i - \bar{\mathbf{x}})$, one can replace \mathbf{v}_k in the eigendecomposition, simplify, and obtain: $\mathbf{X}(\mathbf{K} \boldsymbol{\alpha}_k - \lambda_k \boldsymbol{\alpha}_k) = 0$, where $\mathbf{K}_{ij} = (\mathbf{x}_i - \bar{\mathbf{x}})^T (\mathbf{x}_j - \bar{\mathbf{x}})$. Consider the mapping $\phi : \mathcal{M} \rightarrow \mathcal{H}$ from the manifold to a Hilbert space. The “kernel trick” suggests replacing \mathbf{x}_i with $\phi(\mathbf{x}_i)$ and therefore rewriting the kernel as $\mathbf{K}_{ij} = \phi(\mathbf{x}_i)^T \phi(\mathbf{x}_j)$. Further generalization can be made by setting $\mathbf{K}_{ij} = K(\mathbf{x}_i, \mathbf{x}_j)$ where $K(\cdot, \cdot)$ is positive semidefinite. The resulting vectors are of the form $\mathbf{v}_k = \sum_i \alpha_{ki} \phi(\mathbf{x}_i)$ and thus implementing a nonlinear embedding into a nonlinear manifold.

2) *ISOMAP*: In [?], Tenenbaum *et al* find a nonlinear embedding that rather than preserving the Euclidean distance between points on a manifold, preserves the geodesic distance between points on the manifold. Similar to MDS where a lower dimensional embedding is found to preserve the Euclidean distances of high dimensional data, ISOMAP finds a lower dimensional embedding that preserves the geodesic distances between high-dimensional data points.

3) *Laplacian Eigenmaps*: Belkin and Niyogi’s Laplacian eigenmaps dimensionality reduction algorithm [?] takes a different approach. They consider a nonlinear mapping f that minimize the Laplacian

$$\arg \min_{\|f\|_{L^2(\mathcal{M})}=1} \int_{\mathcal{M}} \|\nabla f\|^2. \quad (3)$$

Since the manifold is not available but only data point on it are, the lower dimensional embedding is found by minimizing the graph Laplacian given by

$$\sum_{i=1}^n w_{ij} \|\mathbf{y}_i - \mathbf{y}_j\|^2, \quad (4)$$

where w_{ij} is the ij th element of the adjacency matrix which is constructed as follows: For $k \in \mathbb{N}$, a k -nearest neighbors graph is constructed with the points in \mathcal{X}_n as the graph vertices. Each point \mathbf{x}_i is connected to its k -nearest neighboring points. Note that it suffices that either \mathbf{x}_i is among \mathbf{x}_j 's k -nearest neighbors or \mathbf{x}_j is among \mathbf{x}_i 's k -nearest neighbors for \mathbf{x}_i and \mathbf{x}_j to be connected. For a fixed scale parameter $\epsilon > 0$, the weight associated with the two points \mathbf{x}_i and \mathbf{x}_j satisfies

$$w_{ij} = \begin{cases} \exp\{-\|\mathbf{x}_i - \mathbf{x}_j\|^2/\epsilon\} & \text{if } \mathbf{x}_i \text{ and } \mathbf{x}_j \text{ are connected} \\ 0 & \text{otherwise.} \end{cases}$$

III. CLASSIFICATION CONSTRAINED DIMENSIONALITY REDUCTION

A. Statistical framework

To put the problem in a classification context, we consider the following model. Let $\mathcal{X}_n = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$ be a set of n points sampled from an m -dimensional submanifold $\mathcal{M} \subseteq \mathbb{R}^d$. Each point $\mathbf{x}_i \in \mathcal{M}$ is associate with a class label $c_i \in \mathcal{A} = \{0, 1, 2, \dots, L\}$, where $c_i = 0$ corresponds to the case of unlabeled data. We assume that pairs $(\mathbf{x}_i, c_i) \in \mathcal{M} \times \mathcal{A}$ are i.i.d. drawn from a joint distribution

$$P(\mathbf{x}, c) = p_x(\mathbf{x}|c)P_c(c) = P_c(c|\mathbf{x})p_x(\mathbf{x}), \quad (5)$$

where $p_x(\mathbf{x}) > 0$ and $p_x(\mathbf{x}|c) > 0$ (for $\mathbf{x} \in \mathcal{M}$) are the marginal and the conditional probability density functions, respectively, satisfying $\int_{\mathcal{M}} p_x(\mathbf{x})d\mathbf{x} = 1$, $\int_{\mathcal{M}} p_x(\mathbf{x}|c)d\mathbf{x} = 1$ and $P_c(c) > 0$ and $P_c(c|\mathbf{x}) > 0$ are the a priori and a posteriori probability mass functions of the class label, respectively, satisfying $\sum_c P_c(c) = 1$ and $\sum_c P_c(c|\mathbf{x}) = 1$. While we consider unlabeled points of the form $(\mathbf{x}_i, 0)$ similar labeled points, we still make the following distinction. Consider the following mechanism for generating an unlabeled point. First, a class label $c \in \{1, 2, \dots, L\}$ is generated from the labeled a priori probability mass function $P'_c(c) = P(c|c \text{ is labeled}) = P_c(c)/\sum_{c'=1}^L P_c(c')$. Then \mathbf{x}_i is generated according to $p_x(\mathbf{x}|c)$. To treat c as an unobserved label, we marginalize $P(\mathbf{x}, c|c \text{ is labeled}) = p_x(\mathbf{x}|c)P'_c(c)$ over c :

$$p_x(\mathbf{x}|c = 0) = \sum_{q=1}^L p_x(\mathbf{x}|c = q)P'_c(q) = \frac{\sum_{q=1}^L p_x(\mathbf{x}|c = q)P_c(q)}{\sum_{c'=1}^L P_c(c')}. \quad (6)$$

This suggests that the conditional PDF of unlabeled points $f_x(\mathbf{x}|c = 0)$ is uniquely determined by the class priors and the conditionals for labeled point. We would like to point out that this is one

of few treatments that can be offered for unlabeled point. For example, in anomaly detection, one may want to associate the unlabeled point with contaminated data points, which can be represented as a density mixture of $p_x(\mathbf{x}|c=0)$ and $\gamma(\mathbf{x})$ (e.g., $\gamma(\mathbf{x})$ is uniform in \mathcal{X}).

In classification constraint dimensionality reduction, our goal is to obtain a lower-dimensional embedding $\mathcal{Y}_n = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n\}$ (where $\mathbf{y}_i \in \mathbb{R}^m$ with $m < d$) that preserves local geometry and that encourages clustering of points of the same class label. Alternatively, we would like to find a mapping $\mathbf{f}(\mathbf{x}, c) : \mathcal{M} \times \mathcal{A} \rightarrow \mathbb{R}^m$ for which $\mathbf{y}_i = \mathbf{f}(\mathbf{x}_i, c_i)$ that is smooth and that clusters points of the same label.

We introduce the class label indicator for data point \mathbf{x}_i as $c_{ki} = I(c_i = k)$, for $k = 1, 2, \dots, L$ and $i = 1, 2, \dots, n$. Note that when point \mathbf{x}_i is unlabeled $c_{ki} = 0$ for all k . Using the class indicator, we can write the number of point in class k as $n_k = \sum_{i=1}^n c_{ki}$. If all points are labeled, then $n = \sum_{k=1}^L n_k$.

B. Linear dimensionality reduction for classification

1) *LDA*: Restricting the discussion to linear maps, one can extend PCA to take into account label information using the multi-class extension to Fisher's linear discriminant analysis (LDA). Instead of maximizing the data covariance matrix, LDA maximizes the ratio of the between-class-covariance to within-class-covariance. In other words, we obtain a linear transformation $\mathbf{y}_i = f(\mathbf{x}_i, c_i) = \mathbf{A}\mathbf{x}_i$ with matrix \mathbf{A} that is the solution to the following maximization:

$$\max_{\mathbf{A}} \text{tr}\{\mathbf{A}\mathbf{C}_B\mathbf{A}^T\} \quad \text{s.t.} \quad \mathbf{A}\mathbf{C}_W\mathbf{A}^T = \mathbf{I}, \quad (7)$$

where

$$\mathbf{C}_B = \frac{1}{n} \sum_{k=1}^L n_k (\bar{\mathbf{x}}^{(k)} - \bar{\mathbf{x}})(\bar{\mathbf{x}}^{(k)} - \bar{\mathbf{x}})^T$$

is the between-class-covariance matrix, $\bar{\mathbf{x}}^{(k)} = \sum_i c_{ki} \mathbf{x}_i / n_k$ is the k th class center, $\bar{\mathbf{x}} = \sum_i \mathbf{x}_i / n$ is the center point of the dataset,

$$\mathbf{C}_W = \frac{1}{n} \sum_{k=1}^L n_k \mathbf{C}_W^{(k)},$$

is the within-class-covariance, and

$$\mathbf{C}_W^{(k)} = \frac{\sum_{i=1}^n c_{ki} (\mathbf{x}_i - \bar{\mathbf{x}}^{(k)})(\mathbf{x}_i - \bar{\mathbf{x}}^{(k)})^T}{n_k}$$

is within-class- k covariance matrix. In Fig. 1, LDA selects an embedding that projects the data onto \mathbf{e}_2 since the maximum distance between classes is achieved along with a minimum class variance when projecting the data onto \mathbf{e}_2 . We are interested in exploring a strategy that maximizes class separation in the lower dimensional embedding.

2) *Marginal Fisher Analysis*: Recent work [?], presents the marginal Fisher analysis (MFA), which is a method that minimizes the ratio between intraclass compactness and interclass separability. In its basic formulation MFA is a linear embedding, in which $\mathbf{y}_i = \mathbf{A}\mathbf{x}_i$. Another aspect of the method is that it considers two classes. The kernel trick is used to provide a nonlinear extension to MFA. To construct the cost function, two quantities are of interest: intraclass compactness and interclass separability. The intraclass compactness can be written as

$$\sum_{i,j} w_{ij} \|\mathbf{y}_i - \mathbf{y}_j\|^2, \quad (8)$$

where w_{ij} is given by

$$w_{ij} = \left(\sum_k c_{ki} c_{kj} \right) I(\mathbf{x}_i \in N_{k_1}^+(\mathbf{x}_j) \text{ or } \mathbf{x}_j \in N_{k_1}^+(\mathbf{x}_i)) \quad (9)$$

and $N_k^+(\mathbf{x})$ denote the k -nn neighborhood of \mathbf{x} within the same class as \mathbf{x} . Note that the term $\sum_k c_{ki} c_{kj}$ is one if \mathbf{x}_i and \mathbf{x}_j have the same label and zero otherwise. Similarly, the interclass separability can be written as

$$\sum_{i,j} w_{ij} \|\mathbf{y}_i - \mathbf{y}_j\|^2, \quad (10)$$

where w_{ij} is given by

$$w_{ij} = \left(1 - \sum_k c_{ki} c_{kj} \right) I(\mathbf{x}_i \in N_{k_2}^-(\mathbf{x}_j) \text{ or } \mathbf{x}_j \in N_{k_2}^-(\mathbf{x}_i)) \quad (11)$$

and $N_k^-(\mathbf{x})$ denote the k -nn neighborhood of \mathbf{x} outside the class of \mathbf{x} .

IV. DIMENSIONALITY REDUCTION FOR CLASSIFICATION ON NONLINEAR MANIFOLDS

Here, we review the CCDD algorithm [?] and its extension to multi-class classification.

To cluster lower dimensional embedded points of the same label we associate each class with a class center namely $\mathbf{z}_k \in \mathbb{R}^m$. We construct the following cost function:

$$J(\mathcal{Z}_L, \mathcal{Y}_n) = \sum_{ki} c_{ki} \|\mathbf{z}_k - \mathbf{y}_i\|^2 + \frac{\beta}{2} \sum_{ij} w_{ij} \|\mathbf{y}_i - \mathbf{y}_j\|^2, \quad (12)$$

where $\mathcal{Z}_L = \{z_1, \dots, z_L\}$ and $\beta \geq 0$ is a regularization parameter. We consider two terms on the RHS of (12). The first term corresponds to the concentration of points of the same label around their respective class center. The second term is as in (4) or as in Laplacian Eigenmaps [?] and controls the smoothness of the embedding over the manifold. Large values of β produce an embedding that ignores class labels and small values of β produce an embedding that ignores the manifold structure. Training data points will tend to collapse into the class centers, allowing many classifiers to produce perfect classification on the training data without being able to control the generalization error (i.e., classification error of the unlabeled data). Our goal is to find \mathcal{Z}_L and \mathcal{Y}_n that minimize the cost function in (12).

Let \mathbf{C} be the $L \times n$ class membership matrix with c_{ki} as its ki -th element, $\mathbf{Z} = [z_1, \dots, z_L, \mathbf{y}_1, \dots, \mathbf{y}_n]$, and $\mathbf{0}$ be the $L \times L$ all zeroes matrix and

$$\mathbf{G} = \begin{bmatrix} \mathbf{0} & \mathbf{C} \\ \mathbf{C}^T & \beta \mathbf{W} \end{bmatrix}.$$

Minimization over \mathbf{Z} of the cost function in (12) can be expressed as

$$\begin{aligned} \min_{\mathbf{Z}} \quad & \text{tr}(\mathbf{Z}\mathbf{L}\mathbf{Z}^T), \\ \text{subject to} \quad & \mathbf{Z}\mathbf{D}\mathbf{1} = \mathbf{0} \\ & \mathbf{Z}\mathbf{D}\mathbf{Z}^T = \mathbf{I} \end{aligned} \quad (13)$$

where $\mathbf{D} = \text{diag}\{\mathbf{G}\mathbf{1}\}$ and $\mathbf{L} = \mathbf{D} - \mathbf{G}$. To prevent the lower-dimensional points and the class centers from collapsing into a single point at the origin, the regularization $\mathbf{Z}\mathbf{D}\mathbf{Z}^T = \mathbf{I}$ is introduced. The second constraint $\mathbf{Z}\mathbf{D}\mathbf{1} = \mathbf{0}$ is constructed to prevent a degenerate solution, e.g., $z_1 = \dots = z_L = \mathbf{y}_1 = \dots = \mathbf{y}_n$. This solution may occur since $\mathbf{1}$ is in the null-space of the Laplacian \mathbf{L} operator, i.e., $\mathbf{L}\mathbf{1} = \mathbf{0}$. The solution to (13) can be expressed in term of the following generalized eigendecomposition

$$\mathbf{L}^{(n)} \mathbf{u}_k^{(n)} = \lambda_k^{(n)} \mathbf{D}^{(n)} \mathbf{u}_k^{(n)}, \quad (14)$$

where $\lambda_k^{(n)}$ is the k th eigenvalue and $\mathbf{u}_k^{(n)}$ is its corresponding eigenvector. Note that we include $^{(n)}$ to emphasize the dependence on the n data points. Without loss of generality we assume $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_{n+L}$. Specifically, matrix \mathbf{Z} is given by $[\mathbf{u}_2, \mathbf{u}_3, \dots, \mathbf{u}_{n+1}]^T$, where the first L columns correspond to the coordinates of the class centers, i.e., $z_k = \mathbf{Z}\mathbf{e}_k$, and the following n columns determine the embedding of the n data points, i.e., $\mathbf{y}_t = \mathbf{Z}\mathbf{e}_{L+t}$. We use \mathbf{e}_i to denote the canonical vector such that $[\mathbf{e}_i]_s = 1$ for element $s = i$ and zero otherwise.

A. Classification and computational complexity

In classification, the goal is to find a classifier $a_x(\mathbf{x}) : \mathcal{M} \rightarrow \mathcal{A}$ based on the training data that minimizes the generalization error:

$$\hat{a} = \arg \min_{a \in \mathcal{F}} E[I(a(\mathbf{x}) \neq a)], \quad (15)$$

where the expectation is taken w.r.t. the pair (\mathbf{x}, a) . Since only samples from the joint distribution of \mathbf{x} and a are available, we replace the expectation with a sample average w.r.t. the training data $\frac{1}{n} \sum_{i=1}^n I(a(\mathbf{x}_i) \neq a_i)$. During the minimization, we search over a set of classifiers $a_x(\mathbf{x}) : \mathcal{M} \subseteq \mathbb{R}^d \rightarrow \mathcal{A}$, which is defined over a domain in \mathbb{R}^d . In our framework, we suggest replacing a classifier $a_x(\mathbf{x}) : \mathcal{M} \subseteq \mathbb{R}^d \rightarrow \mathcal{A}$ with dimensionality reduction via CCDR $f(\mathbf{x}) : \mathcal{M} \subseteq \mathbb{R}^d \rightarrow \mathbb{R}^m$ followed by a classifier on the lower-dimensional space $a_y(\mathbf{y}) : \mathbb{R}^m \rightarrow \mathcal{A}$, i.e., $a_x = a_y \circ f$. The first advantage is that the search space for the minimization in (15) defined over a d -dimensional space can be reduced to an m -dimensional space. This results in significant savings in computational complexity if the complexity associated with the process of obtaining f can be made low. In general, the classifier set \mathcal{F} has to be rich enough to attain a lower generalization error. The other advantage of our method lies in the fact that CCDR is designed to cluster points of the same label thus allowing for a linear classifier or other low complexity classifiers. Therefore, further reduction in the size of class \mathcal{F} can be achieved in addition to the reduction due to a lower-dimensional domain. To classify a new data point, one has to apply CCDR to a new data point. If it is done brute force, the point is added to the set of training points with no label a new matrix W' is formed and an eigendecomposition is carried out.

When performing CCDR, each of the $n(n-1)/2$ terms of the form $\{\|\mathbf{x}_i - \mathbf{x}_j\|^2\}$ requires one summation and d multiplications leading to computational complexity of the order $O(dn^2)$. Construction of a K -nearest neighbors graph requires $O(kn)$ comparisons per point and therefore a total of $O(kn^2)$. The total number of operations involved in constructing the graph is therefore $O((k+d)n^2)$. Next, an eigendecomposition is applied to W' , which is an $(L+n) \times (L+n)$ matrix. The associated computation complexity is $O(n^3)$. Therefore, the overall computational complexity of CCDR is $O(n^3)$. This holds for both training and classification as explained earlier. We are interested in reducing computational complexity in training the classifier and in classification. For that purpose, we consider an out-of-sample extension of CCDR.

V. OUT-OF-SAMPLE EXTENSION

We start by rearranging the generalized eigendecomposition of the Laplacian in (14) as

$$\mathbf{G}^{(n)} \mathbf{u}_l^{(n)} = (1 - \lambda_l^{(n)}) \mathbf{D}^{(n)} \mathbf{u}_l^{(n)}, \quad (16)$$

and recall that $\mathbf{u}_l^{(n)} = [\mathbf{z}_1(l), \mathbf{z}_1(l), \dots, \mathbf{z}_1(l), \mathbf{y}_1(l), \mathbf{y}_2(l), \dots, \mathbf{y}_n(l)]^T$. Since we consider an m -dimensional embedding, we are only interested in eigenvectors $\mathbf{u}_2, \dots, \mathbf{u}_{m+1}$. The $L+i$ equation (row) for $i = 1, 2, \dots, n$ in the eigendecomposition in (16) can be written as

$$\mathbf{y}_i^{(n)}(l) = \frac{1}{1 - \lambda_l^{(n)}} \frac{\sum_k c_{ki} \mathbf{z}_k^{(n)}(l) + \beta \sum_j K(\mathbf{x}_i, \mathbf{x}_j) \mathbf{y}_j^{(n)}(l)}{\sum_k c_{ki} + \beta \sum_j K(\mathbf{x}_i, \mathbf{x}_j)}. \quad (17)$$

Similarly, the k th equation (row) of (16) for $k = 1, 2, \dots, L$ is given by

$$\mathbf{z}_k^{(n)}(l) = \frac{\sum_i c_{ki} \mathbf{y}_i^{(n)}(l)}{(1 - \lambda_l^{(n)}) n_k}. \quad (18)$$

Our interest is in finding a mapping $\mathbf{f}(\mathbf{x}, c)$ that in addition to mapping every \mathbf{x}_i to \mathbf{y}_i , can perform an out-of-sample extension, i.e., is well-defined outside the set \mathcal{X} . We consider the following out-of-sample extension expression

$$\mathbf{f}_l^{(n)}(\mathbf{x}, c) = \frac{1}{1 - \lambda_l^{(n)}} \frac{I(c \neq 0) \mathbf{z}_c^{(n)}(l) + \beta \sum_j K(\mathbf{x}, \mathbf{x}_j) \mathbf{y}_j^{(n)}(l)}{I(c \neq 0) + \beta \sum_j K(\mathbf{x}, \mathbf{x}_j)}, \quad (19)$$

where $\mathbf{z}^{(n)}$ is the same as in (18). This formula can be explain as follows. First, the lower dimensional embedding $\mathbf{y}_1^{(n)}, \dots, \mathbf{y}_n^{(n)}$ and the class centers $\mathbf{z}_1^{(n)}, \dots, \mathbf{z}_L^{(n)}$ are obtained through an the eigendecomposition in (16). Then, the embedding outside the sample set \mathcal{X} is calculated via (19). By comparison of $\mathbf{f}_l^{(n)}(\mathbf{x}_i, c_i)$ evaluated through (19) with (17), we have $\mathbf{f}_l^{(n)}(\mathbf{x}_i, c_i) = \mathbf{y}_i^{(n)}(l)$. This suggests that the out-of-sample extension coincides with the solution, we already have for the mapping at the the data points \mathcal{X} . Moreover, using this result one can replace all $\mathbf{y}_i^{(n)}$ with $\mathbf{f}_l^{(n)}(\mathbf{x}_i, c_i)$ in (19) and obtain the following generalization of the eigendecomposition in (16):

$$\mathbf{f}_l^{(n)}(\mathbf{x}, c) = \frac{1}{1 - \lambda_l^{(n)}} \frac{I(c \neq 0) \mathbf{z}_c^{(n)}(l) + \beta \sum_j K(\mathbf{x}, \mathbf{x}_j) \mathbf{f}_l^{(n)}(\mathbf{x}_j, c_j)}{I(c \neq 0) + \beta \sum_j K(\mathbf{x}, \mathbf{x}_j)}, \quad (20)$$

and

$$\mathbf{z}_k^{(n)}(l) = \frac{\sum_i c_{ki} \mathbf{f}_l^{(n)}(\mathbf{x}_i, c_i)}{(1 - \lambda_l^{(n)}) n_k}. \quad (21)$$

In [?], it is propose that if the out-of-sample solution to the eigendecomposition problem associated with kernel PCA converge, it is given by the solution to the asymptotic equivalent of the eigendecomposition. Using similar machinery, we can provide a similar result suggesting that if $\mathbf{f}_l^{(n)}(\mathbf{x}, c) \rightarrow \mathbf{f}_l^{(\infty)}(\mathbf{x}, c)$ as $n \rightarrow \infty$, then the asymptotic equivalents to (20) and (21) should provide the solution to the limit of $\mathbf{f}_l^{(n)}(\mathbf{x}, c)$. The asymptotic analogues to (17) and (18) are described in the following. The mapping for labeled data $f_l(\mathbf{x}, c) : \mathcal{M} \times \mathcal{A} \rightarrow \mathbb{R}$ for $c = 0, 1, 2, \dots, L$ equivalent to equation (17) is

$$f_l(\mathbf{x}, c) = \frac{1}{1 - \lambda_l} \frac{I(c \neq 0)z_c(l) + \beta' \sum_{c'=0}^L \int_{\mathcal{M}} K(\mathbf{x}, \mathbf{x}') f_l(\mathbf{x}', c') P(\mathbf{x}', c') d\mathbf{x}'}{I(c \neq 0) + \beta' \int_{\mathcal{M}} K(\mathbf{x}, \mathbf{x}') p(\mathbf{x}') d\mathbf{x}'} \quad (22)$$

where $z_c(l)$ for $c = 1, 2, \dots, L$ is equivalent to (18)

$$z_c(l) = \frac{\int_{\mathcal{M}} f_l(\mathbf{x}, c) p(\mathbf{x}|c) d\mathbf{x}}{1 - \lambda_l}, \quad (23)$$

and $\beta' = \beta n$. Since we are interested in an m -dimensional embedding, we consider only $l = 1, 2, \dots, m$, i.e., the eigenvectors that correspond to the m smallest eigenvalues. To guarantee that the relevant eigenvectors are unique (up to a multiplicative constant), we require $\lambda_1 < \lambda_2 < \dots < \lambda_{m+1} \leq \lambda_{m+2} \leq \dots \lambda_n$.

The out-of-sample extension given by (19), can be useful in a couple of scenario. The first, is in classification of new unlabeled samples. We assume that $\{\mathbf{y}_j\}_{j=1}^n$, $\{\mathbf{z}_k\}_{k=1}^L$, and $\{\lambda_l\}_{l=1}^m$ are already obtained based on labeled (or partially labeled) training data and we would like to embed a new unlabeled data point. We consider using (19) with $c = 0$, i.e., we can use $\mathbf{f}(\mathbf{x}, 0)$ to map a new sample \mathbf{x} to \mathbb{R}^m . The obvious immediate advantage is the savings in computational complexity as we avoid performing addition eigendecomposition that includes the new point.

The second scenario involves the out-of-sample extension for labeled data. The goal here is not to classify the data since the label is already available. Instead, we are interested in the training phase in the case of large n for which the eigendecomposition is infeasible. In this case, a large amount of labeled training data is available but due to the heavy computational complexity associated with the eigendecomposition in (14) (or by (16)), the data cannot be processed. In this case, we are interested in developing a resampling method, which integrates $\mathbf{f}_l^{(n)}(\mathbf{x}, c)$ obtained for different subsamples of the complete data set.

A. Classification Algorithms

We consider three widespread algorithms: k -nearest neighbors, linear classification, and neural networks. A standard implementation of k -nearest neighbors was used, see [?, p. 415]. The linear classifier we implemented is given by

$$\hat{c}(\mathbf{y}) = \arg \max_{c \in \{\mathcal{A}_1, \dots, \mathcal{A}_L\}} \mathbf{y}^T \boldsymbol{\alpha}^{(c)} + \alpha_0^{(c)}$$

$$[\boldsymbol{\alpha}^{(\mathcal{A}_k)}, \alpha_0^{(\mathcal{A}_k)}] = \arg \min_{[\boldsymbol{\alpha}, \alpha_0]} \sum_{i=1}^n (\mathbf{y}_i^T \boldsymbol{\alpha} + \alpha_0 - c_{ki})^2,$$

for $k = 1, \dots, L$. The neural network we implemented is a three-layer neural network with d elements in the input layer, $2d$ elements in the hidden layer, and 6 elements in the output layer (one for each class). Here d was selected using the common PCA procedure, as the smallest dimension that explains 99.9% of the energy of the data. A gradient method was used to train the network coefficients with 2000 iterations. The neural net is significantly more computationally burdensome than either linear or k -nearest neighbors classifications algorithms.

B. Data Description

In this section, we examine the performance of the classification algorithms on the benchmark label classification problem provided by the Landsat MSS satellite imagery database [?]. Each sample point consists of the intensity values of one pixel and its 8 neighboring pixels in 4 different spectral bands. The training data consists of 4435 36-dimensional points of which, 1072 are labeled as 1) red soil, 479 as 2) cotton crop, 961 as 3) grey soil, 415 as 4) damp grey soil, 470 are labeled as 5) soil with vegetation stubble, and 1038 are labeled as 6) very damp grey soil. The test data consists of 2000 36-dimensional points of which, 461 are labeled as 1) red soil, 224 as 2) cotton crop, 397 as 3) grey soil, 211 as 4) damp grey soil, 237 are labeled as 5) soil with vegetation stubble, and 470 are labeled as 6) very damp grey soil. In the following, each classifier is trained on the training data and its classification is evaluated based on the entire sample test data. In Table I, we present “best case” performance of neural networks, linear classifier, and k -nearest neighbors in three cases: no dimensionality reduction, dimensionality reduction via PCA, and dimensionality reduction via CCDR. The table presents the minimum probability of error achieved by varying the tuning parameters of the classifiers.

The benefit of using CCDR is obvious and we are prompted to further evaluate the performance gains attained using CCDR.

	Neural Net.	Lin.	k -nearest neigh.
No dim. reduc.	83 %	22.7 %	9.65 %
PCA	9.75 %	23 %	9.35 %
CCDR	8.95 %	8.95 %	8.1 %

TABLE I
CLASSIFICATION ERROR PROBABILITY

C. Regularization Parameter β

As mentioned earlier, the CCDR regularization parameter β controls the contribution of the label information versus the contribution of the geometry described by the sample. We apply CCDR to the 36-dimensional data to create a 14-dimensional embedding by varying β over a range of values. For justification of our choice of $d = 14$ dimensions see Section V-D. In the process of computing the weights w_{ij} for the algorithm, we use k -nearest neighbors with $k = 4$ to determine the local neighborhood. Fig. 2 shows the classification error probability (dashed lines) for the linear classifier vs. β after preprocessing the data using CCDR with $k = 4$ and dimension 14. We observe that for a large range of β the average classification error probability is greater than 0.09 but smaller than 0.095. This performance competes with the performance of k -nearest neighbors applied to the high-dimensional data, which is presented in [?] as the leading classifier for this benchmark problem. Another observation is that for small values of β (i.e., $\beta < 0.1$) the probability of error is constant. For such small value of β , classes in the lower-dimensional embedding are well-separated and are well-concentrated around the class centers. Therefore, the linear classifier yields perfect classification on the training set and fairly low constant probability of error on the test data is attained for low value of β . When β is increased, we notice an increase in the classification error probability. This is due to the fact that the training data become non separable by any linear classifier as β increases.

We perform a similar study of classification performance for k -nearest neighbors. In Fig. 2, classification probability error is plotted (dotted lines) vs. β . Here, we observed that an average

error probability of 0.086 can be achieved for $\beta \approx 0.5$. Therefore, k -nearest neighbors preceded by CCDD outperforms the straightforward k -nearest neighbors algorithm. We also observe that when β is decreased the probability of error is increased. This can be explained as due to the ability of k -nearest neighbors to utilize local information, i.e., local geometry. This information is discarded when β is decreased.

We conclude that CCDD can generate lower-dimensional data that is useful for global classifiers, such as the linear classifier, by using a small value of β , and also for local classifiers, such as k -nearest neighbors, by using a larger value β and thus preserving local geometry information.

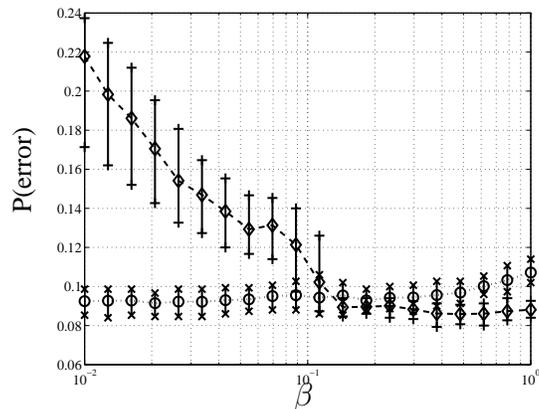


Fig. 2. Probability of incorrect classification vs. β for a linear classifier (dotted line \circ) and for the k -nearest neighbors algorithm (dashed line \diamond) preprocessed by CCDD. 80% confidence intervals are presented as \times for the linear classifier and as $+$ for the k -nearest neighbors algorithm.

D. Dimension Parameter

While the data points in \mathcal{X}_n may lie on a manifold of a particular dimension, the actual dimension required for classification may be smaller. Here, we examine classification performance as a function of the CCDD dimension. Using the entropic graph dimension estimation algorithm in [?], we obtain the following estimated dimension for each class:

class	1	2	3	4	5	6
dimension	13	7	13	10	6	13

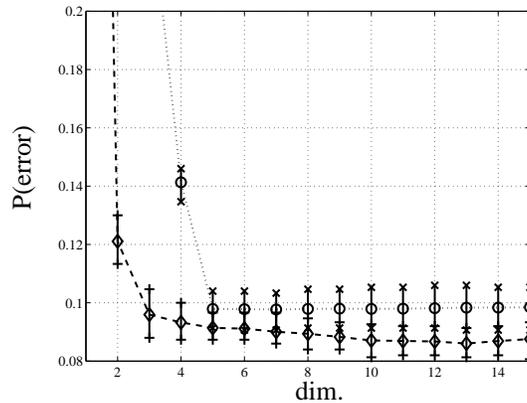


Fig. 3. Probability of incorrect classification vs. CCDR's dimension for a linear classifier (dotted line \circ) and for the k -nearest neighbors algorithm (dashed line \diamond) preprocessed by CCDR. 80% confidence intervals are presented as \times for the linear classifier and as $+$ for the k -nearest neighbors algorithm.

Therefore, if an optimal nonlinear embedding of the data could be found, we suspect that a dimension greater than 13 may not yield significant improvement in classification performance. Since CCDR does not necessarily yield an optimal embedding, we choose CCDR embedding dimension as $d = 14$ in Section V-C.

In Fig. 3, we plot the classification error probability (dotted line) vs. CCDR dimension and its confidence interval for a linear classifier. We observed decrease in error probability as the dimension increases. When the CCDR dimension is greater than 5, the error probability seems fairly constant. This is an indication that CCDR dimension of 5 is sufficient for classification if one uses the linear classifier with $\beta = 0.5$, i.e., linear classifier cannot exploit geometry.

We also plot the classification error probability (dashed line) vs. CCDR dimension and its confidence interval for k -nearest neighbors classifier. Generally, we observe decrease in error probability as the dimension increases. When the CCDR dimension is greater than 5, the error probability seems fairly constant. When CCDR dimension is three, classifier error is below 0.1. On the other hand, minimum possibility of error obtained at CCDR dimension 12-14. This is remarkable agreement with the dimension estimate of 13 obtained using the entropic graph algorithm of [?].

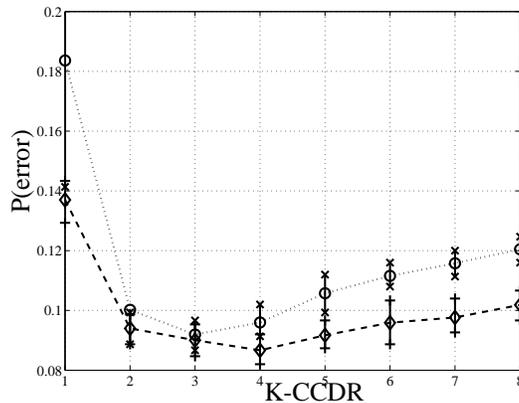


Fig. 4. Probability of incorrect classification vs. CCDDR’s k -nearest neighbors parameter for a linear classifier (dotted line \circ) and for the k -nearest neighbors algorithm (dashed line \diamond) preprocessed by CCDDR. 80% confidence intervals are presented as \times for the linear classifier and as $+$ for the k -nearest neighbors algorithm.

E. CCDDR’s k -Nearest Neighbors Parameter

The last parameter we examine is the CCDDR’s k -nearest neighbors parameter. In general, as k increases non-local distances are included in the lower-dimensional embedding. Hence, very large k prevents the flexibility necessary for dimensionality reduction on (globally) non-linear (but locally linear) manifolds.

In Fig. 4, the classification probability of error for the linear classifier (dotted line) is plotted vs. the CCDDR’s k -nearest neighbors parameter. A minimum is obtained at $k = 3$ with probability of error of 0.092. The classification probability of error for k -nearest neighbors (dashed line) is plotted vs. the CCDDR’s k -nearest neighbors parameter. A minimum is obtained at $k = 4$ with probability of error of 0.086.

VI. CONCLUSION

In this paper, we presented the CCDDR algorithm for multiple classes. We examined the performance of various classification algorithms applied after CCDDR for the Landsat MSS imagery dataset. We showed that for a linear classifier, decreasing β yields improved performance and for a k -nearest neighbors classifier, increasing β yields improved performance. We demonstrated that both classifiers have improved performance on the much smaller dimension of CCDDR embedding space than when applied to the original high-dimensional data. We also explored the effect of k

in the k -nearest neighbors construction of CCDR weight matrix on classification performance. CCDR allows reduced complexity classification such as the linear classifier to perform better than more complex classifiers applied to the original data. We are currently pursuing an out-of-sample extension to the algorithm that does not require rerunning CCDR on test and training data to classify new test point.