

Supervised Probabilistic Principal Component Analysis

Shipeng Yu^{1,2}, Kai Yu², Volker Tresp², Hans-Peter Kriegel¹, Mingrui Wu³

¹Institute for Computer Science, University of Munich, Germany

²Siemens AG, Corporate Technology, Information and Communications, Munich, Germany

³MPI for Biological Cybernetics, Tübingen, Germany

spyu@dbs.ifi.lmu.de, kai.yu@siemens.com, volker.tresp@siemens.com,
kriegel@dbs.ifi.lmu.de, mingrui.wu@tuebingen.mpg.de

ABSTRACT

Principal component analysis (PCA) has been extensively applied in data mining, pattern recognition and information retrieval for unsupervised dimensionality reduction. When labels of data are available, e.g., in a classification or regression task, PCA is however not able to use this information. The problem is more interesting if only part of the input data are labeled, i.e., in a semi-supervised setting. In this paper we propose a supervised PCA model called SPPCA and a semi-supervised PCA model called S²PPCA, both of which are extensions of a probabilistic PCA model. The proposed models are able to incorporate the label information into the projection phase, and can naturally handle multiple outputs (i.e., in multi-task learning problems). We derive an efficient EM learning algorithm for both models, and also provide theoretical justifications of the model behaviors. SPPCA and S²PPCA are compared with other supervised projection methods on various learning tasks, and they show not only promising performance but also good scalability.

Categories and Subject Descriptors

H.3 [Information Storage and Retrieval]: Content Analysis and Indexing—*Indexing methods*

General Terms

Algorithms, Theory, Measurement, Performance

Keywords

Principal component analysis, Supervised projection

1. INTRODUCTION

Data mining problems often suffer from the high dimensionality of the data, for the reason of learnability or computational efficiency. Therefore dimensionality reduction, which is also called *feature transformation* or *document indexing*, is of great importance and has been extensively stud-

ied (see, e.g., [6, 2]). The most popular method is probably the principal component analysis (PCA), which performs a singular value decomposition (SVD) to the data matrix and obtains the sub-eigenspace with large singular values [8].

Traditional dimensionality reduction methods are *unsupervised*, i.e., they only focus on observations or input data. However, in discriminant analysis where the prediction value or output is available, it would be more helpful to incorporate this information into the mapping and derive a *supervised projection* for input data. Since this projection is designed for the specific prediction problem, it could be substantially different from unsupervised projection. A more interesting setting is *semi-supervised projection* where we have only part of input data labeled, along with a large number of unlabeled data. This is often true in real world problems because labeling is expensive, or unlabeled data are very easy to obtain. An ideal projection method should be able to take into account both the observed labeling information and the unlabeled inputs. There exist many supervised projection algorithms in the literature such as linear discriminative analysis (LDA), partial least square (PLS) and many others (see, e.g., [6] for an overview). However, these methods cannot incorporate the unlabeled data into the mapping, which will cause problems when we have only very few labeled points.

In this paper we propose a supervised PCA model called SPPCA, which extends the probabilistic PCA model [17] to incorporate label information into the projection. SPPCA takes into account not only the inter-covariance between inputs and outputs, but also the intra-covariance of both. It can naturally handle multiple outputs, and there is no limitation on the number of projected dimensions. More interestingly, the model can be further extended to model unlabeled data as well, which we call a semi-supervised PCA or S²PPCA. This model allows us to elegantly use all the information we have to define the mapping. We derive an efficient EM learning algorithm for both models, and provide some theoretical justifications for the model behaviors. Experimental results on various learning tasks show promising performance for both SPPCA and S²PPCA models.

This paper is organized as follows. After reviewing some previous work in Section 2, we formally introduce SPPCA and S²PPCA models in Section 3 and derive an EM learning algorithm in Section 4. We then present some theoretical justifications in Section 5 with further discussions. Finally Section 6 illustrates experimental results and Section 7 concludes the paper.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 2002 ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

2. PREVIOUS WORK

In this section we review some previous work on unsupervised and supervised projections. In what follows we consider a set of N objects (e.g., images), and each object n is described by an M -dimensional feature vector $\mathbf{x}_n \in \mathcal{X} \subset \mathbb{R}^M$. For dimensionality reduction we aim to derive a mapping $\Psi : \mathcal{X} \mapsto \mathcal{Z}$ which maps the input features into a K -dimensional space ($K < M$).

2.1 Unsupervised Projection

Probably the most popular unsupervised projection is principal component analysis (PCA). Let $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_N]^\top$ denote the input matrix after centralization, i.e., we subtract the sample mean from each input. In PCA we want to find the principal components which illustrate the directions with maximal variances of the data. Let $\mathbf{X} = \mathbf{V}\mathbf{D}\mathbf{U}^\top$ be the singular value decomposition (SVD) of \mathbf{X} , where \mathbf{V} and \mathbf{U} are $N \times N$ and $M \times M$ column orthogonal matrices, respectively, and \mathbf{D} is $N \times M$ diagonal matrix with singular values sorted in descending order along the diagonal. Then it is known that the first K columns of \mathbf{U} , which we denote \mathbf{U}_K , defines the mapping Ψ . The projections of \mathbf{X} onto the principal component space are given as $\mathbf{V}_K\mathbf{D}_K$, where \mathbf{V}_K contains the first K columns of \mathbf{V} , and \mathbf{D}_K is the top left $K \times K$ sub matrix of \mathbf{D} .

Unlike PCA which is maximizing the global covariance of the data, there also exist ‘‘local’’ projection algorithms such as locally linear embedding [15], which tries to preserve the local structure of the data after projecting into a low dimensional space.

2.2 Supervised Projection

When each data point n is associated with not only low level features \mathbf{x}_n but also some outputs $\mathbf{y}_n \in \mathcal{Y} \subset \mathbb{R}^L$ (i.e., classification or regression targets), unsupervised projection such as PCA may be not able to project the data into useful directions (see Figure 2). Therefore many *supervised projection* methods are introduced to make use of the output information. Linear discriminant analysis (LDA) focuses on multi-class classification and finds projection directions that separate the data best (see [3]). The number of projection dimensions is however limited by $L - 1$. Partial least squares (PLS) [18] originates from regression analysis and finds the directions of maximal covariance between inputs and outputs sequentially. It however ignores the intra covariance of either inputs or outputs, and its generalization performance on new dimensions of outputs is restricted (see discussions in [5]). Other related works include [7, 12, 1, 19, 10, ?] which consider this problem from different perspectives.

In the situations where we have few labeled points and a large amount of unlabeled data, all these supervised projection methods are however not able to use the unlabeled data. This is often the case in computer vision, information retrieval and bioinformatics, where labeling is expensive and unlabeled data are sufficient and cheap to obtain. We call this setting the *semi-supervised projection*, and we will propose a model in the next section which can deal with semi-supervised projection naturally.

3. THE SPPCA MODEL

We first review a probabilistic model for PCA, and then present our supervised models.

3.1 Probabilistic PCA (PPCA)

While PCA originates from the analysis of data variances, in statistics community there exists a probabilistic explanation for PCA, which is called *probabilistic PCA* or *PPCA* in the literature [17, 14]. PPCA is a latent variable model and defines a generative process for each object \mathbf{x} as (see Figure 1(a) for an illustration)

$$\mathbf{x} = \mathbf{W}_x \mathbf{z} + \boldsymbol{\mu}_x + \boldsymbol{\epsilon}_x,$$

where $\mathbf{z} \in \mathbb{R}^K$ are called the *latent variables*, and \mathbf{W}_x is a $M \times K$ matrix called *factor loadings*. In this probabilistic model, latent variables \mathbf{z} are conventionally assumed as a Gaussian distribution with zero mean and unit variance, i.e., $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, and $\boldsymbol{\epsilon}_x$ defines a noise process which also takes an isotropic Gaussian form as $\boldsymbol{\epsilon}_x \sim \mathcal{N}(\mathbf{0}, \sigma_x^2 \mathbf{I})$, with σ_x^2 the *noise level*. Additionally, we have parameters $\boldsymbol{\mu}_x \in \mathbb{R}^M$ which allow non-zero means for the data.

It is shown that PPCA has strong connections to PCA. In particular when $\sigma_x^2 \rightarrow 0$, the projections of data \mathbf{x} onto the K -dimensional principal subspace in PCA are the latent variables \mathbf{z} up to a rotation factor [17]. We summarize the related results in the following proposition without proof, since this is simply a corollary of Theorem 2 in Section 5.

PROPOSITION 1. *Let $\mathbf{S}_x = \frac{1}{N} \sum_{n=1}^N (\mathbf{x}_n - \boldsymbol{\mu}_x)(\mathbf{x}_n - \boldsymbol{\mu}_x)^\top$ be the sample covariance matrix for data $\{\mathbf{x}_n\}_{n=1}^N$, and $\lambda_1 \geq \dots \geq \lambda_M$ be its eigenvalues with eigenvectors $\mathbf{u}_1, \dots, \mathbf{u}_M$, then if the latent space in PPCA model is K -dimensional,*

- (i) *The maximal likelihood estimate of \mathbf{W}_x is given as*

$$\mathbf{W}_x = \mathbf{U}_K (\boldsymbol{\Lambda}_K - \sigma_x^2 \mathbf{I})^{\frac{1}{2}} \mathbf{R},$$

where $\boldsymbol{\Lambda}_K = \text{diag}(\lambda_1, \dots, \lambda_K)$, $\mathbf{U}_K = [\mathbf{u}_1, \dots, \mathbf{u}_K]$, and \mathbf{R} is an arbitrary $K \times K$ orthogonal matrix.

- (ii) *The mean projections \mathbf{z}^* for new input \mathbf{x}^* is given as*

$$\mathbf{z}^* = \mathbf{R}^\top (\boldsymbol{\Lambda}_K - \sigma_x^2 \mathbf{I})^{\frac{1}{2}} \boldsymbol{\Lambda}_K^{-1} \mathbf{U}_K^\top (\mathbf{x}^* - \boldsymbol{\mu}_x).$$

As a probabilistic framework, PPCA provides additional benefits over PCA such as a fast EM learning procedure, a principled way of handling missing entries, and a possibility of considering mixture of PCA models. PPCA is also closely related to *factor analysis* models, but the modeling perspectives are different (see [17] for more discussions).

3.2 Supervised PPCA (SPPCA)

The key point of PPCA model is that all the M dimensions of \mathbf{x} are *conditionally independent* given the latent variables \mathbf{z} , due to the isotropic property of the noise process. This indicates that the principal components in PPCA are *the K latent variables which best explain the data covariance*.

When supervised information is available, each object \mathbf{x} is associated with an output value $y \in \mathcal{Y}$, e.g., $y \in \mathbb{R}$ for regression task and $y \in \{+1, -1\}$ for classification task. In general we believe there are *covariances* between input space \mathcal{X} and output space \mathcal{Y} (since otherwise the supervised learning task is not learnable), and it is reasonable to extend PPCA to model this covariance as well. Furthermore, when there are more than one learning tasks (i.e., in a multi-task learning setting [4]), the *covariances between different tasks* can also be modeled by latent variables.




Figure 1: Illustrations of the three models PPCA, SPPCA and S²PPCA. \mathbf{X} and \mathbf{Y} denote respectively the input and output matrices, where each row is one data point. f_x^1, \dots, f_x^M are the M input features, and f_y^1, \dots, f_y^L are the L outputs. On the top f_z^1, \dots, f_z^K are the K latent variables in each model. They are all in circles because they are variables in the probabilistic models. The arrows denote probabilistic dependency.

We now formally describe our proposed model family which we call the *supervised probabilistic principal component analysis* (SPPCA). Let the number of outputs be L , and each object \mathbf{x} be associated with an output vector $\mathbf{y} = [y_1, \dots, y_L]^\top \in \mathcal{Y} \subset \mathbb{R}^L$. In SPPCA the observed data (\mathbf{x}, \mathbf{y}) is generated from a latent variable model as

$$\begin{aligned} \mathbf{x} &= \mathbf{W}_x \mathbf{z} + \boldsymbol{\mu}_x + \boldsymbol{\epsilon}_x, \\ \mathbf{y} &= \mathbf{f}(\mathbf{z}, \boldsymbol{\Theta}) + \boldsymbol{\epsilon}_y, \end{aligned}$$

where $\mathbf{f}(\mathbf{z}, \boldsymbol{\Theta}) = [f_1(\mathbf{z}, \theta_1), \dots, f_L(\mathbf{z}, \theta_L)]^\top$ encode the values of L deterministic functions f_1, \dots, f_L with parameters $\boldsymbol{\Theta} = \{\theta_1, \dots, \theta_L\}$. Here $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ are the latent variables *shared by both inputs \mathbf{x} and outputs \mathbf{y}* , and the two noise models are independent to each other and both defined as isotropic Gaussians: $\boldsymbol{\epsilon}_x \sim \mathcal{N}(\mathbf{0}, \sigma_x^2 \mathbf{I})$, $\boldsymbol{\epsilon}_y \sim \mathcal{N}(\mathbf{0}, \sigma_y^2 \mathbf{I})$. We use two noise levels σ_x^2 and σ_y^2 for inputs and outputs, respectively, and it is also straightforward to define different noise levels for different outputs if desired. See Figure 1(b) for an illustration of the model.

In SPPCA model we keep the nice property of *conditional independence*, i.e., all the input and output dimensions are conditionally independent to each other given the latent variables. If we integrate out the latent variables \mathbf{z} , the likelihood of observation (\mathbf{x}, \mathbf{y}) is obtained as

$$P(\mathbf{x}, \mathbf{y}) = \int P(\mathbf{x}, \mathbf{y} | \mathbf{z}) P(\mathbf{z}) d\mathbf{z} = \int P(\mathbf{x} | \mathbf{z}) P(\mathbf{y} | \mathbf{z}) P(\mathbf{z}) d\mathbf{z},$$

where $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, and from the latent variable model,

$$\mathbf{x} | \mathbf{z} \sim \mathcal{N}(\mathbf{W}_x \mathbf{z} + \boldsymbol{\mu}_x, \sigma_x^2 \mathbf{I}), \quad \mathbf{y} | \mathbf{z} \sim \mathcal{N}(\mathbf{f}(\mathbf{z}, \boldsymbol{\Theta}), \sigma_y^2 \mathbf{I}). \quad (1)$$

After observing N pairs, the likelihood of all the observations $\mathcal{D} = \{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^N$, with i.i.d. assumption, is simply $P(\mathcal{D}) = \prod_{n=1}^N P(\mathbf{x}_n, \mathbf{y}_n)$.

In the following we consider the simplest model in this family, i.e., we assume each function f_ℓ , $\ell = 1, \dots, L$, is linear in \mathbf{z} :

$$f_\ell(\mathbf{z}, \theta_\ell) = \mathbf{w}_y^\ell \mathbf{z} + \mu_y^\ell,$$

where the parameters $\theta_\ell = \{\mathbf{w}_y^\ell, \mu_y^\ell\}$ include the linear coefficients and intercept. Then we can group all the f_ℓ 's and

write

$$\mathbf{f}(\mathbf{z}, \boldsymbol{\Theta}) = \mathbf{W}_y \mathbf{z} + \boldsymbol{\mu}_y,$$

a similar form as the generative model for \mathbf{x} where $\mathbf{W}_y = [\mathbf{w}_y^1, \dots, \mathbf{w}_y^L]^\top$ and $\boldsymbol{\mu}_y = [\mu_y^1, \dots, \mu_y^L]^\top$. The reason why we choose this form for \mathbf{f} is that the EM learning is simple (see the next section), and we have closed form solution (see Section 5). We will discuss other forms of \mathbf{f} in Section 5.3 which may need special approximation techniques.

Let us denote

$$\mathbf{W} = \begin{pmatrix} \mathbf{W}_x \\ \mathbf{W}_y \end{pmatrix}, \quad \boldsymbol{\mu} = \begin{pmatrix} \boldsymbol{\mu}_x \\ \boldsymbol{\mu}_y \end{pmatrix}, \quad \boldsymbol{\Phi} = \begin{pmatrix} \sigma_x^2 \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \sigma_y^2 \mathbf{I} \end{pmatrix},$$

then based on the model assumption, it is easily seen that (\mathbf{x}, \mathbf{y}) are jointly Gaussian distributed, with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Phi} + \mathbf{W}\mathbf{W}^\top$. All the parameters for the SPPCA model are $\boldsymbol{\Omega} = \{\mathbf{W}_x, \mathbf{W}_y, \boldsymbol{\mu}_x, \boldsymbol{\mu}_y, \sigma_x^2, \sigma_y^2\}$.

3.3 Semi-Supervised PPCA (S²PPCA)

In SPPCA model, we assume we observe both the inputs \mathbf{x} and outputs \mathbf{y} for every data point. In many real world problems, however, we may only observe the outputs for a small portion of data, and have many *unlabeled data* in which only inputs \mathbf{x} are known. This may be because some measures are unobservable, the labeling cost is too high, or simply we have too many unlabeled data available. Learning in this situation is in general called *semi-supervised learning*. For learning a projection, an ideal model would incorporate both the unlabeled inputs and the partially labeled outputs to define the mapping.

This can be easily done under the SPPCA framework. Let the number of labeled and unlabeled data points be N_1 and N_2 , respectively, with $N = N_1 + N_2$. The whole observation is now $\mathcal{D} = \mathcal{D}_1 \cup \mathcal{D}_2 = \{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^{N_1} \cup \{\mathbf{x}_{n'}\}_{n'=N_1+1}^N$. The likelihood, with the independence assumption of all the data points, is calculated as

$$P(\mathcal{D}) = P(\mathcal{D}_1)P(\mathcal{D}_2) = \prod_{n=1}^{N_1} P(\mathbf{x}_n, \mathbf{y}_n) \prod_{n'=N_1+1}^N P(\mathbf{x}_{n'}),$$




Figure 2: Projection directions for a toy data. They are fully labeled on the left, and only partially labeled on the right.

where $P(\mathbf{x}_n, \mathbf{y}_n)$ is calculated as in SPPCA model, and $P(\mathbf{x}_{n'}) = \int P(\mathbf{x}_{n'} | \mathbf{z}_{n'}) P(\mathbf{z}_{n'}) d\mathbf{z}_{n'}$. Due to its applicability to semi-supervised projection, we call it *semi-supervised PPCA* or S^2PPCA in this paper. Figure 1(c) illustrates this model.

Under the additional assumptions that all the f_ℓ 's are linear, it can be easily checked that all the likelihood terms in this product are Gaussians. This makes the model easy to learn. Other forms of \mathbf{f} will be discussed in Section 5.3.

When $N_2 = 0$, S^2PPCA degrades to SPPCA which is purely supervised. This means one can view SPPCA as a special case of S^2PPCA model with no unlabeled data. From the perspective of probabilistic modeling, S^2PPCA can also be viewed as an SPPCA model where all the \mathbf{y} 's for the N_2 unlabeled points are missing. Due to this close relationship, in the following we use SPPCA to denote both models unless clearly specified.

3.4 Projections in SPPCA Models

Analogous to the PPCA model, in SPPCA models the projection of data point \mathbf{x} is directly given in the latent variables \mathbf{z} . If we know all the parameters Ω , calculating this projection is simply an *inference* problem. To do this we can apply Bayes' rule and calculate the posterior distribution of \mathbf{z} . Therefore we can obtain not only the mean projection vector, but also the *uncertainty* of the projection.

3.4.1 Projection for Fully Observed Data

When both inputs \mathbf{x} and outputs \mathbf{y} are observed, we can calculate the posterior distribution of \mathbf{z} given (\mathbf{x}, \mathbf{y}) as

$$P(\mathbf{z} | \mathbf{x}, \mathbf{y}) \propto P(\mathbf{x}, \mathbf{y} | \mathbf{z}) P(\mathbf{z}) = P(\mathbf{x} | \mathbf{z}) P(\mathbf{y} | \mathbf{z}) P(\mathbf{z}). \quad (2)$$

Since all the three terms on the right hand side are Gaussians, this distribution is also Gaussian $\mathcal{N}(\mu_{\mathbf{z}}, \Sigma_{\mathbf{z}})$ with

$$\mu_{\mathbf{z}} = \mathbf{A}^{-1} \left[\frac{1}{\sigma_x^2} \mathbf{W}_x^\top (\mathbf{x} - \mu_x) + \frac{1}{\sigma_y^2} \mathbf{W}_y^\top (\mathbf{y} - \mu_y) \right], \quad \Sigma_{\mathbf{z}} = \mathbf{A}^{-1},$$

where \mathbf{A} is a $K \times K$ matrix defined as

$$\mathbf{A} = \frac{1}{\sigma_x^2} \mathbf{W}_x^\top \mathbf{W}_x + \frac{1}{\sigma_y^2} \mathbf{W}_y^\top \mathbf{W}_y + \mathbf{I}. \quad (3)$$

This means that the projection is $\mu_{\mathbf{z}}$ with uncertainty $\Sigma_{\mathbf{z}}$.

3.4.2 Projection for Pure Input Data

For a test data \mathbf{x}^* that has no output information, what are the most likely latent variables \mathbf{z}^* ? This is our ultimate

goal in projection, and can also be done using Bayes' rule

$$P(\mathbf{z}^* | \mathbf{x}^*) \propto P(\mathbf{x}^* | \mathbf{z}^*) P(\mathbf{z}^*). \quad (4)$$

This turns out again to be a Gaussian $\mathcal{N}(\mu_{\mathbf{z} | \mathbf{x}}, \Sigma_{\mathbf{z} | \mathbf{x}})$, with

$$\mu_{\mathbf{z} | \mathbf{x}} = (\mathbf{W}_x^\top \mathbf{W}_x + \sigma_x^2 \mathbf{I})^{-1} \mathbf{W}_x^\top (\mathbf{x}^* - \mu_x),$$

$$\Sigma_{\mathbf{z} | \mathbf{x}} = \sigma_x^2 (\mathbf{W}_x^\top \mathbf{W}_x + \sigma_x^2 \mathbf{I})^{-1}.$$

This result looks similar as that in PPCA model, but the projection now is *supervised* because the learning of \mathbf{W}_x is influenced by those observed outputs. This is clarified in the next section and will be theoretically proven in Section 5.

4. LEARNING IN SPPCA MODEL

Learning in probabilistic models reduces to maximizing the data (log) likelihood with respect to all the model parameters. In the case of SPPCA model, the log likelihood of the whole observation \mathcal{D} is

$$\mathcal{L} = \sum_{n=1}^N \log \int P(\mathbf{x}_n | \mathbf{z}_n) P(\mathbf{y}_n | \mathbf{z}_n) P(\mathbf{z}_n) d\mathbf{z}_n,$$

which is however difficult to maximize directly. S^2PPCA has the same situation as well.

In this section we describe an expectation-maximization (EM) learning procedure for both SPPCA and S^2PPCA models. EM iterates the two steps *expectation* (E-step) and *maximization* (M-step) until convergence, and it is guaranteed to find a local minima of the data likelihood. In the E-step, we fix the model parameters (Ω for SPPCA models) and calculate the expected distributions of latent variables (all the \mathbf{z}_n 's for SPPCA models), and in the M-step we fix this distribution and maximize the *complete data likelihood* with respect to the model parameters. As will be discussed later, EM learning for SPPCA models is important because it can deal with very large data sets, and it has, in particular for SPPCA model with no unlabeled points, no local minima problem up to a rotation factor (see Section 5). For simplicity we only outline the update equations in the following and omit details (see [17] for a similar derivation).

4.1 EM Learning for SPPCA

In the E-step, for each data point n , we estimate the distribution of \mathbf{z}_n given observation $(\mathbf{x}_n, \mathbf{y}_n)$. This is done using (2), and we calculate the sufficient statistics as

$$\langle \mathbf{z}_n \rangle = \mathbf{A}^{-1} \left[\frac{1}{\sigma_x^2} \mathbf{W}_x^\top (\mathbf{x}_n - \mu_x) + \frac{1}{\sigma_y^2} \mathbf{W}_y^\top (\mathbf{y}_n - \mu_y) \right], \quad (5)$$

$$\langle \mathbf{z}_n \mathbf{z}_n^\top \rangle = \mathbf{A}^{-1} + \langle \mathbf{z}_n \rangle \langle \mathbf{z}_n \rangle^\top, \quad (6)$$

where $\langle \cdot \rangle$ denotes the expectation under the posterior distribution $P(\mathbf{z}_n | \mathbf{x}_n, \mathbf{y}_n)$ given in (2).

In the M-step, we maximize the complete log-likelihood

$$\tilde{\mathcal{L}} = \sum_{n=1}^N \int P(\mathbf{z}_n | \mathbf{x}_n, \mathbf{y}_n) \log \left(P(\mathbf{x}_n | \mathbf{z}_n) P(\mathbf{y}_n | \mathbf{z}_n) P(\mathbf{z}_n) \right) d\mathbf{z}_n$$

with respect to the model parameters, holding $P(\mathbf{z}_n | \mathbf{x}_n, \mathbf{y}_n)$ fixed from the E-step. This can be done by setting the partial derivatives with respect to each parameter to be zero. For means of \mathbf{x} and \mathbf{y} we have

$$\tilde{\mu}_x = \frac{1}{N} \sum_{n=1}^N \mathbf{x}_n, \quad \tilde{\mu}_y = \frac{1}{N} \sum_{n=1}^N \mathbf{y}_n, \quad (7)$$

Algorithm 1 Learning in SPPCA Model - Primal Form

Require: N data points $\{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^N$ with inputs $\mathbf{x}_n \in \mathbb{R}^M$ and outputs $\mathbf{y}_n \in \mathbb{R}^L$. A desired dimension $K < M$.

- 1: Calculate the sample means (7) and center the data by $\mathbf{x}_n \leftarrow \mathbf{x}_n - \boldsymbol{\mu}_x, \mathbf{y}_n \leftarrow \mathbf{y}_n - \boldsymbol{\mu}_y$.
- 2: Initialize model parameters $\boldsymbol{\Omega}$ randomly.
- 3: **repeat**
- 4: {E-step}
- 5: **for** $n = 1$ to N **do**
- 6: Calculate sufficient statistics (5) and (6);
- 7: **end for**
- 8: {M-step}
- 9: Update \mathbf{W}_x and \mathbf{W}_y via (8);
- 10: Update σ_x^2 and σ_y^2 via (9) and (10);
- 11: **until** the change of $\boldsymbol{\Omega}$ is smaller than a threshold.
- 12: **return** Parameters $\boldsymbol{\Omega}$ and projection vectors $\{\mathbf{z}_n\}_{n=1}^N$ which are obtained from E-step. For test data \mathbf{x}^* , the mean projection $\mathbf{z}^* = (\mathbf{W}_x^\top \mathbf{W}_x + \sigma_x^2 \mathbf{I})^{-1} \mathbf{W}_x^\top (\mathbf{x}^* - \boldsymbol{\mu}_x)$.

which are just sample means. Since they are always the same in all EM iterations, we can *center* the data by subtracting these means in the beginning and ignore these parameters in the learning process. So for simplicity we change the notations \mathbf{x}_n and \mathbf{y}_n to be the centered vectors in the following.

The mapping matrices \mathbf{W}_x and \mathbf{W}_y are updated as

$$\widetilde{\mathbf{W}}_x = \mathbf{X}^\top \mathbf{Z} \mathbf{C}^{-1}, \quad \widetilde{\mathbf{W}}_y = \mathbf{Y}^\top \mathbf{Z} \mathbf{C}^{-1}, \quad (8)$$

where for clarity we use matrix notations $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_N]^\top$, $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_N]^\top$ and $\mathbf{Z} = [\langle \mathbf{z}_1 \rangle, \dots, \langle \mathbf{z}_N \rangle]^\top$. Matrix \mathbf{C} is defined to be a summation of all second sufficient statistics of the data, i.e., $\mathbf{C} = \sum_{n=1}^N \langle \mathbf{z}_n \mathbf{z}_n^\top \rangle$. Finally the noise levels are updated as

$$\tilde{\sigma}_x^2 = \frac{1}{MN} \left[\sum_{n=1}^N \|\mathbf{x}_n\|^2 + \text{tr}(\widetilde{\mathbf{W}}_x^\top \widetilde{\mathbf{W}}_x \mathbf{C}) - 2 \text{tr}(\mathbf{X} \widetilde{\mathbf{W}}_x \mathbf{Z}^\top) \right] \quad (9)$$

$$\tilde{\sigma}_y^2 = \frac{1}{LN} \left[\sum_{n=1}^N \|\mathbf{y}_n\|^2 + \text{tr}(\widetilde{\mathbf{W}}_y^\top \widetilde{\mathbf{W}}_y \mathbf{C}) - 2 \text{tr}(\mathbf{Y} \widetilde{\mathbf{W}}_y \mathbf{Z}^\top) \right] \quad (10)$$

where $\|\cdot\|$ denotes vector 2-norm, and $\text{tr}(\cdot)$ denotes matrix trace. The whole algorithm is summarized in Algorithm 1 for clarity.

4.2 EM Learning for S²PPCA

The log likelihood of the observations in S²PPCA model is a sum of two parts: $\mathcal{L}_1 = \log P(\mathcal{D}_1)$ which contains all the labeled points, and $\mathcal{L}_2 = \log P(\mathcal{D}_2)$ which includes all unlabeled points. Therefore in E-step we need to deal with them differently. For a labeled points $(\mathbf{x}_n, \mathbf{y}_n) \in \mathcal{D}_1$, the latent variables \mathbf{z}_n are estimated as (5) and (6), the same as in SPPCA model. For an unlabeled point $\mathbf{x}_{n'} \in \mathcal{D}_2$, the distribution of $\mathbf{z}_{n'}$ is only conditioned on input $\mathbf{x}_{n'}$, which can be calculated via (4), with sufficient statistics (the data are assumed centered already):

$$\langle \mathbf{z}_{n'} \rangle = (\mathbf{W}_x^\top \mathbf{W}_x + \sigma_x^2 \mathbf{I})^{-1} \mathbf{W}_x^\top \mathbf{x}_{n'}, \quad (11)$$

$$\langle \mathbf{z}_{n'} \mathbf{z}_{n'}^\top \rangle = (\mathbf{W}_x^\top \mathbf{W}_x + \sigma_x^2 \mathbf{I})^{-1} + \langle \mathbf{z}_{n'} \rangle \langle \mathbf{z}_{n'} \rangle^\top, \quad (12)$$

where here $\langle \cdot \rangle$ denotes the expectation under the posterior distribution $P(\mathbf{z}_{n'} | \mathbf{x}_{n'})$ given in (4).

The M-step is similarly obtained by setting the partial derivatives of the complete log likelihood with respect to

Algorithm 2 Learning in S²PPCA Model - Primal Form

Require: N_1 labeled data points $\{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^{N_1}$ and N_2 unlabeled points $\{\mathbf{x}_{n'}\}_{n'=N_1+1}^N$, with inputs $\mathbf{x} \in \mathbb{R}^M$ and observed outputs $\mathbf{y} \in \mathbb{R}^L$. A desired dimension $K < M$.

- 1: Calculate the sample means (7) and center the data by $\mathbf{x}_n \leftarrow \mathbf{x}_n - \boldsymbol{\mu}_x, \mathbf{y}_n \leftarrow \mathbf{y}_n - \boldsymbol{\mu}_y, \mathbf{x}_{n'} \leftarrow \mathbf{x}_{n'} - \boldsymbol{\mu}_x$.
- 2: Initialize model parameters $\boldsymbol{\Omega}$ randomly.
- 3: **repeat**
- 4: {E-step}
- 5: **for** $n = 1$ to N_1 **do**
- 6: Calculate (5) and (6) for labeled data n ;
- 7: **end for**
- 8: **for** $n' = N_1 + 1$ to N **do**
- 9: Calculate (11) and (12) for unlabeled data n' ;
- 10: **end for**
- 11: {M-step}
- 12: Update \mathbf{W}_x and \mathbf{W}_y via (13) and (14);
- 13: Update σ_x^2 and σ_y^2 via (15) and (16);
- 14: **until** the change of $\boldsymbol{\Omega}$ is smaller than a threshold.
- 15: **return** Parameters $\boldsymbol{\Omega}$ and projection vectors $\{\mathbf{z}_n\}_{n=1}^N$ which are obtained from E-step. For test data \mathbf{x}^* , the mean projection $\mathbf{z}^* = (\mathbf{W}_x^\top \mathbf{W}_x + \sigma_x^2 \mathbf{I})^{-1} \mathbf{W}_x^\top (\mathbf{x}^* - \boldsymbol{\mu}_x)$.

each parameter to zero. For the two mapping matrices, we have the updates

$$\widetilde{\mathbf{W}}_x = (\mathbf{X}_1^\top \mathbf{Z}_1 + \mathbf{X}_2^\top \mathbf{Z}_2)(\mathbf{C}_1 + \mathbf{C}_2)^{-1}, \quad (13)$$

$$\widetilde{\mathbf{W}}_y = \mathbf{Y}^\top \mathbf{Z}_1 \mathbf{C}_1^{-1}, \quad (14)$$

where $\mathbf{X}_1, \mathbf{Z}_1, \mathbf{C}_1$ are defined for labeled data, i.e., $\mathbf{X}_1 = [\mathbf{x}_1, \dots, \mathbf{x}_{N_1}]^\top$, $\mathbf{Z}_1 = [\langle \mathbf{z}_1 \rangle, \dots, \langle \mathbf{z}_{N_1} \rangle]^\top$, $\mathbf{C}_1 = \sum_{n=1}^{N_1} \langle \mathbf{z}_n \mathbf{z}_n^\top \rangle$, and $\mathbf{X}_2, \mathbf{Z}_2, \mathbf{C}_2$ are similarly defined for unlabeled data. It is seen that the update for \mathbf{W}_x depends on both labeled data and unlabeled data, while \mathbf{W}_y only depends on the labeled data. Updates for the noise levels are similar to those in SPPCA model, except that for σ_x^2 we need to consider both labeled data and unlabeled data:

$$\tilde{\sigma}_x^2 = \frac{1}{MN} \left[\sum_{n=1}^N \|\mathbf{x}_n\|^2 + \text{tr}(\widetilde{\mathbf{W}}_x^\top \widetilde{\mathbf{W}}_x (\mathbf{C}_1 + \mathbf{C}_2)) - 2 \text{tr}(\widetilde{\mathbf{W}}_x (\mathbf{Z}_1^\top \mathbf{X}_1 + \mathbf{Z}_2^\top \mathbf{X}_2)) \right] \quad (15)$$

$$\tilde{\sigma}_y^2 = \frac{1}{LN_1} \left[\sum_{n=1}^{N_1} \|\mathbf{y}_n\|^2 + \text{tr}(\widetilde{\mathbf{W}}_y^\top \widetilde{\mathbf{W}}_y \mathbf{C}_1) - 2 \text{tr}(\mathbf{Y} \widetilde{\mathbf{W}}_y \mathbf{Z}_1^\top) \right] \quad (16)$$

The whole algorithm is summarized in Algorithm 2. When $N_2 = 0$, i.e., we have no unlabeled data, the learning algorithm reduces to SPPCA learning.

4.3 EM Learning in Dual Form

It is known that when the number of data points is less than the number of features, i.e., $N < M$, it is more efficient to consider the *dual* solution for PCA in which we perform SVD to the *Gram matrix* $\mathbf{K} = \mathbf{X} \mathbf{X}^\top$. The canonical PCA is sometimes called the *primal* solution. For SPPCA we have similar dual solution, and it can be directly derived from the EM learning in previous subsections. To avoid the tedious mathematics in the main text, we put the derivation details into Appendix and summarize the algorithm in Algorithm 3.

Algorithm 3 Learning in S²PPCA Model - Dual Form

Require: N_1 labeled data points $\{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^{N_1}$ and N_2 unlabeled points $\{\mathbf{x}_{n'}\}_{n'=N_1+1}^N$, with inputs $\mathbf{x} \in \mathbb{R}^M$ and observed outputs $\mathbf{y} \in \mathbb{R}^L$. A desired dimension $K < M$.

- 1: Calculate Gram matrix with $\mathbf{K}_{ij} = \mathbf{x}_i^\top \mathbf{x}_j$ and center it using (27). Let $\mathbf{y}_n \leftarrow \mathbf{y}_n - \boldsymbol{\mu}_y$.
- 2: Initialize \mathbf{Z} , \mathbf{C} and model parameters $\boldsymbol{\Omega}$ randomly.
- 3: **repeat** {The EM-step}
- 4: Calculate \mathbf{Z}_1 and \mathbf{C}_1 using (21) and (22);
- 5: Calculate \mathbf{Z}_2 and \mathbf{C}_2 using (23) and (24);
- 6: Update σ_x^2 and σ_y^2 via (25) and (26);
- 7: **until** the change of $\boldsymbol{\Omega}$ is smaller than a threshold.
- 8: **return** Parameters $\boldsymbol{\Omega}$ and projection vectors $\{\mathbf{z}_n\}_{n=1}^N$.
For test data \mathbf{x}^* , we first calculate the inner product vector $\mathbf{k}(\mathbf{X}, \mathbf{x}^*) = [\mathbf{x}_1^\top \mathbf{x}^*, \dots, \mathbf{x}_N^\top \mathbf{x}^*]^\top$ and center it using (28). Then the mean projection $\mathbf{z}^* = \mathbf{C}(\mathbf{Z}^\top \mathbf{K} \mathbf{Z} + \sigma_x^2 \mathbf{C}^2)^{-1} \mathbf{Z}^\top \mathbf{k}(\mathbf{X}, \mathbf{x}^*)$.

Since SPPCA can be viewed as a special case of S²PPCA, here we only give the algorithm for S²PPCA model.

One important observation in the dual solution is that all the calculation involving input data \mathbf{X} can be done via inner product, e.g., in the Gram matrix \mathbf{K} we have $\mathbf{K}_{ij} = \mathbf{x}_i^\top \mathbf{x}_j$. This motivates us to consider *non-linear PCA* where we first map the data into a new feature space (via, e.g., basis functions), and then perform PCA in that space with a proper definition of inner product. This is the idea behind kernel PCA [16], and we put detailed discussion into Appendix.

4.4 Computational Issues

In the primal form (i.e., Algorithm 1 and 2), the time complexity for both algorithms is $\mathcal{O}(m(M+L)NK)$, with m the number of iterations.¹ It is linear in the number of data points N and the input dimension M . The space complexity is $\mathcal{O}((M+L)N)$, which is also linear both in N and M . The projection for a test data point is just a linear operation and costs $\mathcal{O}(MK)$ time.

In the dual form, the time complexity is $\mathcal{O}(mN^2K)$ plus $\mathcal{O}(N^2M)$ which is the one-time calculation of Gram matrix, and the space complexity is $\mathcal{O}(N^2)$. Both of them are now quadratic in the number of data points N . The time for projecting a test data point is now $\mathcal{O}(NM)$. As the case for PCA, in situations where $M > N$, i.e., we have more features than the number of data points, the dual form is more efficient than the primal form.

5. THEORETICAL JUSTIFICATION

In this section we provide some theoretical analysis for SPPCA model and show how the supervised information influences the projection. For clarity we put all the proofs in Appendix.

5.1 Primal Form Solution

Recall that matrix $\boldsymbol{\Phi}$ is a $(M+L) \times (M+L)$ diagonal matrix with all the noise levels in diagonal, i.e., $\boldsymbol{\Phi} = \text{diag}(\sigma_x^2, \dots, \sigma_x^2, \sigma_y^2, \dots, \sigma_y^2)$. We obtain the following theorem for mapping matrix \mathbf{W}_x and \mathbf{W}_y in SPPCA model.

¹Note that we only need to calculate the diagonal entries for matrix trace in the updates for noise levels.

This makes it easier to compare SPPCA with related models such as PCA.

THEOREM 2. Let \mathbf{S} denote the normalized sample covariance matrix for centered observations $\{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^N$,

$$\mathbf{S} = \frac{1}{N} \sum_{n=1}^N \boldsymbol{\Phi}^{-\frac{1}{2}} \begin{pmatrix} \mathbf{x}_n \\ \mathbf{y}_n \end{pmatrix} \begin{pmatrix} \mathbf{x}_n \\ \mathbf{y}_n \end{pmatrix}^\top \boldsymbol{\Phi}^{-\frac{1}{2}} = \begin{pmatrix} \frac{1}{\sigma_x^2} \mathbf{S}_x & \frac{1}{\sigma_x \sigma_y} \mathbf{S}_{xy} \\ \frac{1}{\sigma_x \sigma_y} \mathbf{S}_{yx} & \frac{1}{\sigma_y^2} \mathbf{S}_y \end{pmatrix},$$

and $\lambda_1 \geq \dots \geq \lambda_{(M+L)}$ be its eigenvalues with eigenvectors $\mathbf{u}_1, \dots, \mathbf{u}_{(M+L)}$, then if the latent space is K -dimensional, the following results hold:

(i) In SPPCA model \mathbf{W}_x and \mathbf{W}_y are calculated as

$$\mathbf{W}_x = \sigma_x \mathbf{U}_x (\boldsymbol{\Lambda}_K - \mathbf{I})^{\frac{1}{2}} \mathbf{R}, \quad (17)$$

$$\mathbf{W}_y = \sigma_y \mathbf{U}_y (\boldsymbol{\Lambda}_K - \mathbf{I})^{\frac{1}{2}} \mathbf{R},$$

where $\boldsymbol{\Lambda}_K = \text{diag}(\lambda_1, \dots, \lambda_K)$, \mathbf{U}_x (\mathbf{U}_y) contains the first M (last L) rows of $[\mathbf{u}_1, \dots, \mathbf{u}_K]$, and \mathbf{R} is an arbitrary $K \times K$ orthogonal rotation matrix.

(ii) Projection \mathbf{z}^* for centered new input \mathbf{x}^* is given as

$$\mathbf{z}^* = \frac{1}{\sigma_x} \mathbf{R}^\top (\boldsymbol{\Lambda}_K - \mathbf{I})^{-\frac{1}{2}} \left[\mathbf{U}_x^\top \mathbf{U}_x + (\boldsymbol{\Lambda}_K - \mathbf{I})^{-1} \right]^{-1} \mathbf{U}_x^\top \mathbf{x}^*.$$

In the special case that $L = 0$, the model is unsupervised and $\mathbf{S} = \frac{1}{\sigma_x^2} \mathbf{S}_x$ holds. Then (17) degrades to $\sigma_x \mathbf{U}_x (\boldsymbol{\Lambda}_K - \mathbf{I})^{\frac{1}{2}} \mathbf{R}$, which recovers the PPCA solution. \mathbf{U}_x is seen to be column orthogonal in this case, and the mapping \mathbf{z}^* of \mathbf{x}^* is standard PCA mapping when $\sigma_x^2 \rightarrow 0$ and $\mathbf{R} = \mathbf{I}$. This proves Proposition 1 which is a corollary of this theorem.

When $L > 0$, SPPCA solutions explain not only the sample covariance of inputs \mathbf{S}_x , but also the *intra-covariance* of outputs \mathbf{S}_y (if $L > 1$) and the *inter-correlations* between inputs and outputs, \mathbf{S}_{xy} and \mathbf{S}_{yx} . Therefore, one column of \mathbf{W}_x is the direction that best explains the whole system from the perspective of inputs, and thus are *biased* by the outputs. Unlike the case of PCA, the learned \mathbf{W}_x in SPPCA needs not to be column orthogonal. This means we are only learning an *affine* mapping for \mathbf{x} . If necessary, it is straightforward to find the orthogonal basis by performing SVD to matrix \mathbf{W}_x .

In both cases the learned \mathbf{W}_x has an arbitrary rotation factor \mathbf{R} . This is due to the spherical noise model for \mathbf{x} , and the mapping is invariant under a rotation of latent space, as can be seen from the equation for \mathbf{z}^* . Therefore the SPPCA model can only find the *latent principal subspace*, which has been mentioned in [17] for PPCA. Thus the EM algorithm in Section 4 can find different mappings with different initializations, but they define the same subspace and do not change the structure of projected data. If necessary, this ambiguity can be removed by eigen-decomposing

$$\mathbf{W}_x^\top \mathbf{W}_x + \mathbf{W}_y^\top \mathbf{W}_y = \mathbf{R}^\top (\boldsymbol{\Lambda}_K - \mathbf{I}) \mathbf{R}$$

and uncovering the rotation factor \mathbf{R} . A final comment is that it is of theoretical importance but may be not applicable to applications, since we have to solve an eigenvalue problem for a square matrix of size $M+L$.

5.2 Dual Form Solution

In the dual form, we do not obtain the mapping matrix \mathbf{W}_x , but the projected vectors directly. The following theorem gives the solution in the dual form.

THEOREM 3. Let $\hat{\mathbf{K}} = \frac{1}{\sigma_x^2} \mathbf{K} + \frac{1}{\sigma_y^2} \mathbf{Y} \mathbf{Y}^\top$, and $\lambda_1 \geq \dots \geq \lambda_N$ be its eigenvalues with eigenvectors $\mathbf{v}_1, \dots, \mathbf{v}_N$, then if the latent space in SPPCA is K -dimensional,

- (i) The projections of training data, which are encoded in rows of matrix \mathbf{Z} , are calculated as

$$\mathbf{Z} = \sqrt{N} \mathbf{V}_K (\mathbf{I} - N \mathbf{\Lambda}_K^{-1})^{\frac{1}{2}} \mathbf{R} \quad (18)$$

where $\mathbf{\Lambda}_K = \text{diag}(\lambda_1, \dots, \lambda_K)$, $\mathbf{V}_K = [\mathbf{v}_1, \dots, \mathbf{v}_K]$, and \mathbf{R} is an arbitrary $K \times K$ orthogonal matrix.

- (ii) Projections \mathbf{z}^* for new input \mathbf{x}^* is given as

$$\mathbf{z}^* = \sqrt{N} \mathbf{R}^\top \mathbf{D}^{-\frac{1}{2}} \left(\mathbf{V}_K^\top \mathbf{K} \mathbf{V}_K + \mathbf{D} \right)^{-1} \mathbf{V}_K^\top \mathbf{k}(\mathbf{X}, \mathbf{x}^*),$$

where $\mathbf{D} = \mathbf{I} - N \mathbf{\Lambda}_K^{-1}$.

It is seen from this theorem that when there is no output in SPPCA, i.e., $\hat{\mathbf{K}} = \frac{1}{\sigma_x^2} \mathbf{K}$, SPPCA reduces to the dual form of PCA as desired. This theorem directly applies for non-linear mappings if the inner product is defined in a reproducing kernel Hilbert space (RKHS) [16], and leads to the kernel PCA solution when $L = 0$.

Theorem 3 presents a nice explanation for SPPCA model: We just use the output information to modify the Gram matrix of input data, and control the trade-off via the ratio of the noise levels. The complexity of the model remains the same (i.e., quadratic in N) no matter how many output dimensions we have.

5.3 Discussion

Previous two subsections give some theoretical results for SPPCA model. There exists however no such a closed-form solution for S^2 PPCA. One can only empirically analyze the behavior of this model.

In the EM learning algorithm we are learning the maximal likelihood (ML) estimate for the two mapping matrices \mathbf{W}_x and \mathbf{W}_y . In probabilistic modeling we can assign a prior to them to reduce overfitting. For instance, we can assign an isotropic Gaussian prior for each column of \mathbf{W}_x , and if we consider the maximal-a-posterior (MAP) estimate this prior corresponds to a smooth term in the update equations. For simplicity we do not consider this prior here.

In this paper we mainly discuss the simplest form for function \mathbf{f} , which is linear in \mathbf{z} . One can define other forms for specific tasks, but then we lose the nice closed-form solutions described in Theorem 2 and 3, and in the E-step of the EM learning the posterior distribution of \mathbf{z} is no longer a Gaussian (see (2)). To solve this problem we can apply the EM-EP learning algorithm [9] to approximate each likelihood $P(y_\ell | \mathbf{z}) = P(y_\ell | f_\ell(\mathbf{z}, \Theta), \sigma_y^2)$ as a Gaussian for \mathbf{z} sequentially. Then the learning has higher time complexity due to the expectation-propagation (EP) [13] step. Empirically comparing this algorithm with the basic ones would be part of the future work.

6. EXPERIMENTS

In this section we empirically investigate the performance of SPPCA models. The supervised tasks here are multi-class classification and multi-label classification. Our basic setting is that we train a supervised projection model using the input features and label information, and then test the classification performance for test data using the projected

Table 1: Statistics of the multi-class data sets

	CATEGORY	# DATA	# DIM	# CLASS
YALE	FACE	165	1024	15
ORL	FACE	400	1024	40
PIE	FACE	11554	1024	68
YALEB	FACE	2414	1024	38
11_TUMORS	GENE	174	12533	11
14_TUMORS	GENE	308	15009	26
LUNG_CANCER	GENE	203	12600	5
20NEWSGROUP	TEXT	19928	25284	20
TDT2	TEXT	8692	35452	20

Table 2: Statistics of the multi-label data sets

	CATEGORY	# DATA	# DIM	# CLASS
YEAST	GENE	2417	103	14
RCV1	TEXT	23149	15500	103

features. Since the test data are assumed known in the training phase, for S^2 PPCA we will be able to use these unlabeled data to train the mapping.

6.1 Data Sets

We test the proposed model on 9 multi-class and 2 multi-label classification problems. These problems include face recognition, gene classification and text categorization. Some statistics of these data sets are shown in Table 1 and 2.

For face recognition we use four data sets Yale, ORL, PIE and YaleB (the extended Yale Face Database B).² The Yale data set contains 165 grayscale images in GIF format of 15 individuals. There are 11 images per subject, one per different facial expression or configuration such as center-light, left-light, happy or surprised. The ORL database contains 10 different images of each of 40 distinct subjects. For some subjects, the images were taken at different times with varying lighting and facial details. The PIE databases we use contains 170 images for each of 68 people. These images are the five near frontal poses under different illuminations and expressions. For YaleB we have 38 individuals and around 64 near frontal images under different illuminations per individual. All the faces are manually aligned, cropped and resized to 32×32 pixels. We then normalize each image to have Euclidean distance 1.

We consider three gene expression datasets 11_Tumors, 14_Tumors and Lung_Cancer for gene classification.³ 11_Tumors describes 11 various human tumor types, and 14_Tumors describes 14 tumor types with 12 normal tissue types. For Lung_Cancer we need to classify 4 lung cancer types and normal tissues. The characteristic of these data is that the number of data points is small, but the input dimensionality is very high.

The two textual datasets we use are taken from 20Newsgroup and TDT2. 20Newsgroup contains 20,000 news articles posted in 20 news groups. We remove the words that occur less than 5 times, and obtain 19,928 documents with 25,284 words. The TDT2 corpus we use consists of the documents collected during the first half of 1998 and taken from 6 sources, including 2 newswires (APW, NYT), 2 radio programs (VOA, PRI) and 2 television programs (CNN, ABC).

²These data sets can be downloaded from <http://www.ews.uiuc.edu/dengcai2/Data/data.html>.

³They are available at <http://www.gems-system.org>.

Table 3: Results for Multi-class Classification Tasks. Bold face indicates lowest error rate. Symbols \star indicate that the best method is significantly better than the competitors (p-value 0.01 in Wilcoxon rank sum test).

TASK	FULL	PCA	LDA	PLS	SPPCA	S ² PPCA
YALE	0.5656 ± 0.0394	0.6690 ± 0.0333	0.6133 ± 0.0471	0.6440 ± 0.0383	0.7007 ± 0.0402	0.7121 ± 0.0393
ORL	0.3308 ± 0.0347	0.5593 ± 0.0263	0.5302 ± 0.0444	0.5505 ± 0.0294	0.5459 ± 0.0305	0.5287 ± 0.0286
PIE	0.6988 ± 0.0085	0.9325 ± 0.0032	0.7066 ± 0.0177	0.8781 ± 0.0058	0.8780 ± 0.0116	0.8452 ± 0.0037
YALEB	0.6360 ± 0.0160	0.9895 ± 0.0023	\star0.5328 ± 0.0251	0.9546 ± 0.0066	0.9701 ± 0.0088	0.9800 ± 0.0034
11_TUMORS	0.3161 ± 0.0566	0.5409 ± 0.0490	0.4505 ± 0.0755	N/A	0.5226 ± 0.0636	0.5130 ± 0.0491
14_TUMORS	0.6084 ± 0.0360	0.7363 ± 0.0286	0.7161 ± 0.0481	N/A	0.7312 ± 0.0371	0.7138 ± 0.0296
LUNG_CANCER	0.3680 ± 0.1148	0.3768 ± 0.0939	0.3225 ± 0.1658	N/A	0.4287 ± 0.1338	0.3896 ± 0.0923
20NEWSGROUP	0.6135 ± 0.0155	0.9070 ± 0.0177	0.9140 ± 0.0208	N/A	0.9030 ± 0.0162	0.9126 ± 0.0116
TDT2	0.1875 ± 0.0233	0.6664 ± 0.0657	0.7834 ± 0.0782	N/A	0.6236 ± 0.0739	\star0.3686 ± 0.0349

(a) Projection dimension $K = 5$.

TASK	FULL	PCA	LDA	PLS	SPPCA	S ² PPCA
YALE	0.5656 ± 0.0394	0.5993 ± 0.0312	0.5279 ± 0.0460	0.5698 ± 0.0386	0.6101 ± 0.0447	0.5916 ± 0.0433
ORL	0.3308 ± 0.0347	0.4049 ± 0.0293	0.3625 ± 0.0468	0.4048 ± 0.0349	0.3832 ± 0.0409	0.3509 ± 0.0287
PIE	0.6988 ± 0.0085	0.8573 ± 0.0051	0.5496 ± 0.0185	0.8062 ± 0.0068	0.7105 ± 0.0161	0.6942 ± 0.0047
YALEB	0.6360 ± 0.0160	0.9308 ± 0.0046	\star0.3846 ± 0.0282	0.8762 ± 0.0108	0.7976 ± 0.0242	0.7986 ± 0.0117
11_TUMORS	0.3161 ± 0.0566	0.3682 ± 0.0655	0.3926 ± 0.0667	N/A	0.3801 ± 0.0624	\star0.3297 ± 0.0664
14_TUMORS	0.6084 ± 0.0360	0.6868 ± 0.0288	0.6212 ± 0.0430	N/A	0.6322 ± 0.0363	0.6120 ± 0.0331
LUNG_CANCER	0.3680 ± 0.1148	0.3493 ± 0.0996	0.3225 ± 0.1658	N/A	0.6235 ± 0.1520	0.3517 ± 0.1063
20NEWSGROUP	0.6135 ± 0.0155	0.9039 ± 0.0172	0.8943 ± 0.0292	N/A	0.8931 ± 0.0242	0.8548 ± 0.0138
TDT2	0.1875 ± 0.0233	0.5531 ± 0.0742	0.6878 ± 0.1068	N/A	0.5346 ± 0.0885	\star0.2794 ± 0.0327

(b) Projection dimension $K = 10$.

TASK	FULL	PCA	LDA	PLS	SPPCA	S ² PPCA
YALE	0.5656 ± 0.0394	0.5437 ± 0.0414	0.5793 ± 0.0438	0.5216 ± 0.0435	0.5093 ± 0.0391	0.5001 ± 0.0589
ORL	0.3308 ± 0.0347	0.3323 ± 0.0310	0.2944 ± 0.0398	0.3366 ± 0.0331	0.3271 ± 0.0372	0.2755 ± 0.0286
PIE	0.6988 ± 0.0085	0.7999 ± 0.0060	0.4352 ± 0.0186	0.7454 ± 0.0092	0.5912 ± 0.0146	0.5361 ± 0.0090
YALEB	0.6360 ± 0.0160	0.8304 ± 0.0096	\star0.3004 ± 0.0227	0.7695 ± 0.0148	0.5619 ± 0.0276	0.5652 ± 0.0172
11_TUMORS	0.3161 ± 0.0566	0.3267 ± 0.0635	0.3926 ± 0.0667	N/A	0.4470 ± 0.0691	0.3012 ± 0.0582
14_TUMORS	0.6084 ± 0.0360	0.6379 ± 0.0360	0.5822 ± 0.0388	N/A	0.5669 ± 0.0347	0.5674 ± 0.0372
LUNG_CANCER	0.3680 ± 0.1148	0.3584 ± 0.0953	0.3225 ± 0.1658	N/A	0.6487 ± 0.1540	0.4092 ± 0.1107
20NEWSGROUP	0.6135 ± 0.0155	0.9160 ± 0.0220	0.8001 ± 0.0425	N/A	0.6254 ± 0.0420	0.6568 ± 0.0146
TDT2	0.1875 ± 0.0233	0.4582 ± 0.1441	0.1524 ± 0.0622	N/A	0.1566 ± 0.0509	0.1520 ± 0.0210

(c) Projection dimension $K = 20$.

It consists of 11,021 documents which are classified into 96 semantic categories. In our experiments, we keep the largest 20 categories and remove those documents that are assigned to more than one categories. This leaves us 8,692 documents with totally 35,452 words. For both of these datasets we use TF-IDF features and normalize each document to have Euclidean distance 1.

For multi-label classification we use Yeast and RCV1. The Yeast dataset is formed by micro-array expression data and phylogenetic profiles with 2,417 genes in total and 103 input dimensions. There are 14 groups and each gene can belong to multiple groups. The other data is a subset of the RCV1-v2 text data set, provided by Reuters and corrected by Lewis et al. [11]. We use the training set provided by Lewis, which contains 103 labels, 23,149 documents with 15,500 words after we remove words that occur less than 5 times. We also extract TF-IDF features and normalize each document to have length 1.

6.2 Experimental Setting

For the multi-class classification tasks, we randomly pick up a small number of labeled data points for training (2 for those datasets with less than 500 data points, and 5 for the others), and test the classification error rate on the unlabeled data. We will in general compare the following six algorithms if applicable:

- PCA: Unsupervised projection. Note that in our experiments we use both the labeled and unlabeled data to perform the mapping.

- LDA: Linear discriminant analysis.
- PLS: Partial least squares.
- SPPCA: Supervised probabilistic PCA.
- S²PPCA: Semi-supervised probabilistic PCA. We allow S²PPCA to use the test data to train the mapping.
- FULL: All the features are used without projection.

For all the projection methods, we project the data onto 5, 10 and 20 dimensional space, and train a nearest neighbor classifier for the test points using new features with Euclidean distance. For FULL we directly train the nearest neighbor classifier using original features. For PLS, SPPCA and S²PPCA, we translate the one column output to the “One of C” setting, i.e., each class has one column with binary labels.

For multi-label classification, we pick up 5 positive examples from each label to obtain the training data. For all projection methods we project to 5, 10 and 20 dimensions, and then train a linear SVM classifier for each label. The comparison metrics are F1-Macro, F1-Micro and AUC (Area Under ROC Curve) score. The candidate algorithms are almost the same as multi-class setting, except LDA which is not applicable to this task. The C in SVM is fixed as 100, and from our experience it is not sensible for all these algorithms.

In all these comparisons, the iteration number for SPPCA and S²PPCA is set to 1000. Both the noise levels σ_x^2 and σ_y^2 are set to 10^{-5} initially. It turns out that PLS gets memory

Table 4: Results for Multi-label Classification Tasks. Bold face indicates best performance.

K	MODEL	YEAST			RCV1		
		F1-MACRO	F1-MICRO	AUC	F1-MACRO	F1-MICRO	AUC
5	FULL	0.3813 ± 0.0102	0.5161 ± 0.0154	0.5571 ± 0.0094	0.2796 ± 0.0055	0.5053 ± 0.0095	0.6030 ± 0.0063
	PCA	0.2318 ± 0.0354	0.5600 ± 0.0220	0.5279 ± 0.0108	0.0467 ± 0.0063	0.3540 ± 0.0106	0.5208 ± 0.0072
	PLS	0.3432 ± 0.0231	0.5795 ± 0.0233	0.5556 ± 0.0094	N/A	N/A	N/A
	SPPCA	0.3823 ± 0.0120	0.5332 ± 0.0188	0.5641 ± 0.0087	0.1155 ± 0.0071	0.4433 ± 0.0089	0.5568 ± 0.0079
	S ² PPCA	0.3927 ± 0.0134	0.5890 ± 0.0126	0.5842 ± 0.0104	0.1312 ± 0.0118	0.4620 ± 0.0236	0.5762 ± 0.0162
10	FULL	0.3813 ± 0.0102	0.5161 ± 0.0154	0.5571 ± 0.0094	0.2796 ± 0.0055	0.5053 ± 0.0095	0.6030 ± 0.0063
	PCA	0.3113 ± 0.0304	0.5916 ± 0.0146	0.5493 ± 0.0101	0.0843 ± 0.0124	0.4003 ± 0.0162	0.5347 ± 0.0072
	PLS	0.3756 ± 0.0154	0.5517 ± 0.0177	0.5610 ± 0.0095	N/A	N/A	N/A
	SPPCA	0.3924 ± 0.0117	0.5459 ± 0.0180	0.5685 ± 0.0084	0.1797 ± 0.0112	0.4474 ± 0.0124	0.5872 ± 0.0087
	S ² PPCA	0.3985 ± 0.0103	0.5914 ± 0.0106	0.5896 ± 0.0107	0.1956 ± 0.0110	0.4735 ± 0.0198	0.6012 ± 0.0098
20	FULL	0.3813 ± 0.0102	0.5161 ± 0.0154	0.5571 ± 0.0094	0.2796 ± 0.0055	0.5053 ± 0.0095	0.6030 ± 0.0063
	PCA	0.3723 ± 0.0171	0.5537 ± 0.0204	0.5614 ± 0.0097	0.1320 ± 0.0086	0.4419 ± 0.0095	0.5504 ± 0.0061
	PLS	0.3799 ± 0.0123	0.5208 ± 0.0158	0.5585 ± 0.0102	N/A	N/A	N/A
	SPPCA	0.3859 ± 0.0133	0.5517 ± 0.0151	0.5640 ± 0.0097	0.2297 ± 0.0119	0.4690 ± 0.0126	0.6044 ± 0.0054
	S ² PPCA	0.3976 ± 0.0142	0.6012 ± 0.0190	0.5921 ± 0.0119	0.2536 ± 0.0117	0.4921 ± 0.0102	0.6090 ± 0.0076

problems when applied to large dimensions. We repeat each experiment 50 times independently,⁴ and the results are illustrated in .

6.3 Analysis of Results

The first observation is that in most cases the supervised PCA model is better than unsupervised PCA model. This means by using the output information, we are able to derive a more meaningful projection for the supervised tasks. When the dimensionality is larger (e.g., 20), SPPCA and S²PPCA obtain the best results for most of the tasks.

When we compare SPPCA model with other supervised projection methods, SPPCA is consistently better than PLS, but in some tasks worse than LDA (e.g., for YaleB). The reasons may be that SPPCA models are still based on the PCA assumptions for input features, so the mapping is strongly biased by PCA. When PCA projection directions are almost useless for classification, like for YaleB dataset, the SPPCA projections are also not informative enough. In this case discriminative methods like LDA can often do a good job. In other situations where PCA does help, SPPCA can in general be better than pure discriminative methods.

When we compare SPPCA and S²PPCA, in most cases S²PPCA gets better results. For some tasks the difference is very big (e.g., for TDT2). This indicates that by incorporating the unlabeled data we can learn a better mapping. But S²PPCA is in general slower than SPPCA because it has to consider all the test data in the training phase. In this case 1000 iterations may be not enough to get the algorithm converge. This may also be part of the reason why S²PPCA is inferior to other methods for some tasks like PIE.

Most of the supervised projection methods can get a better performance than FULL even if they only project the data into a very low dimensional space. This is important because we can not only speed up the system, but also improve the performance. PCA in our experiments uses the input features of both labeled and unlabeled data, thus it sometimes can get better results than FULL method (e.g., for Yale and Lung_Cancer).

The results for multi-label classification show that S²PPCA is consistently better than other methods. S²PPCA also shows very good scalability in our experiments, since for 20Newsgroup and RCV1 it need to handle 20,000 docu-

⁴For the four face recognition tasks we use the available split versions.

ments with more than 15,000 features. Most of the other algorithms fail on these large datasets.

7. CONCLUSION

We proposed a supervised PCA and a semi-supervised PCA in this paper, and derived an efficient EM algorithm for model learning. Empirical results show that the proposed model obtain good performance and scale well for large datasets.

8. REFERENCES

- [1] E. Bair, T. Hastie, D. Paul, and R. Tibshirani. Prediction by supervised principal components. Technical report, 2004.
- [2] S. C. Deerwester, S. T. Dumais, T. K. Landauer, G. W. Furnas, and R. A. Harshman. Indexing by latent semantic analysis. *Journal of the American Society of Information Science*, 41(6):391–407, 1990.
- [3] R.O. Duda, P.E. Hart, and D.G. Stork. *Pattern Classification*. Wiley, 2000.
- [4] T. Evgeniou and M. Pontil. Regularized multi-task learning. In *Proceedings SIGKDD*, 2004.
- [5] K. Fukumizu, F. R. Bach, and M. I. Jordan. Dimensionality reduction for supervised learning with reproducing kernel Hilbert spaces. *Journal of Machine Learning Research*, 5(Jan):73–99, 2004.
- [6] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning*. Springer Verlag, 2001.
- [7] H. Hotelling. Relations between two sets of variables. *Biometrika*, 28:321–377, 1936.
- [8] I. T. Jolliffe. *Principal Component Analysis*. Springer Verlag, 2002.
- [9] H. Kim and Z. Ghahramani. The EM-EP algorithm for Gaussian process classification. In *Proceedings of the Workshop on Probabilistic Graphical Models for Classification at ECML*, 2003.
- [10] H. Kim, P. Howland, and H. Park. Dimension reduction in text classification with support vector machines. *Journal of Machine Learning Research*, 6:37–53, 2005.
- [11] D. D. Lewis, Y. Yang, T. Rose, and F. Li. RCV1: A new benchmark collection for text categorization

research. *Journal of Machine Learning Research*, 5:361–397, 2004.

- [12] R. G. Miller. *Beyond Anova: Basics of Applied Statistics*. John Wiley, 1986.
- [13] T. P. Minka. *A family of algorithms for approximate Bayesian inference*. PhD thesis, Massachusetts Institute of Technology, 2001.
- [14] S. Roweis and Z. Ghahramani. A unifying review of linear Gaussian models. *Neural Computation*, 11(2):305–345, 1999.
- [15] S. Roweis and L. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, pages 2323–2326, 2000.
- [16] B. Schölkopf and A. J. Smola. *Learning with Kernels*. MIT Press, 2002.
- [17] M. E. Tipping and C. M. Bishop. Probabilistic principal component analysis. *Journal of the Royal Statistical Society*, B(61):611–622, 1999.
- [18] H. Wold. Soft modeling by latent variables; the nonlinear iterative partial least squares approach. *Perspectives in Probability and Statistics, Papers in Honour of M.S. Bartlett*, 1975.
- [19] J. Ye, R. Janardan, Q. Li, and H. Park. Feature extraction via generalized uncorrelated linear discriminant analysis. In *The Twenty-First International Conference on Machine Learning*, 2004.

APPENDIX

A. DERIVATION OF DUAL FORM

We only focus on S^2 PPCA model here. Let $\mathbf{C} = \mathbf{C}_1 + \mathbf{C}_2$, and

$$\mathbf{X} = \begin{pmatrix} \mathbf{X}_1 \\ \mathbf{X}_2 \end{pmatrix}, \quad \mathbf{Z} = \begin{pmatrix} \mathbf{Z}_1 \\ \mathbf{Z}_2 \end{pmatrix}, \quad \mathbf{K} = \mathbf{X}\mathbf{X}^\top = \begin{pmatrix} \mathbf{X}_1\mathbf{X}_1^\top & \\ & \mathbf{X}_2\mathbf{X}_2^\top \end{pmatrix} = \begin{pmatrix} \mathbf{K}_1 & \\ & \mathbf{K}_2 \end{pmatrix}$$

then (13) can be written as $\tilde{\mathbf{W}}_x = \mathbf{X}^\top \mathbf{Z} \mathbf{C}^{-1}$. This leads to

$$\tilde{\mathbf{W}}_x^\top \tilde{\mathbf{W}}_x = \mathbf{C}^{-1} \mathbf{Z}^\top \mathbf{K} \mathbf{Z} \mathbf{C}^{-1}, \quad (19)$$

$$\tilde{\mathbf{W}}_x^\top \mathbf{x} = \mathbf{C}^{-1} \mathbf{Z}^\top \mathbf{k}(\mathbf{X}, \mathbf{x}), \quad (20)$$

which are the building blocks for the non-linear extension. The matrix $\mathbf{K} = \mathbf{X}\mathbf{X}^\top$ has each entry the inner product of data points of corresponding row and column, $\mathbf{K}_{ij} = \mathbf{x}_i^\top \mathbf{x}_j$, and $\mathbf{k}(\mathbf{X}, \mathbf{x}) = \mathbf{X}\mathbf{x}$ is a N -dimensional column vector with the i -th entry $\mathbf{x}_i^\top \mathbf{x}$.

In the E-step, we first rewrite \mathbf{A} as

$$\tilde{\mathbf{A}} = \frac{1}{\sigma_x^2} \mathbf{C}^{-1} \mathbf{Z}^\top \mathbf{K} \mathbf{Z} \mathbf{C}^{-1} + \frac{1}{\sigma_y^2} \mathbf{C}_1^{-1} \mathbf{Z}_1^\top \mathbf{Y} \mathbf{Y}^\top \mathbf{Z}_1 \mathbf{C}_1^{-1} + \mathbf{I}.$$

Applying (19) and (20) in sufficient statistics (5), we get

$$\tilde{\mathbf{Z}}_1 = \left[\frac{1}{\sigma_x^2} \mathbf{K}_1 \mathbf{Z} \mathbf{C}^{-1} + \frac{1}{\sigma_y^2} \mathbf{Y} \mathbf{Y}^\top \mathbf{Z}_1 \mathbf{C}_1^{-1} \right] \tilde{\mathbf{A}}^{-1}, \quad (21)$$

by collecting $\langle \mathbf{z}_n \rangle$ in columns and transposing it. Sufficient statistics (6) can be written in terms of \mathbf{C}_1 :

$$\tilde{\mathbf{C}}_1 = \sum_{n=1}^{N_1} \langle \mathbf{z}_n \mathbf{z}_n^\top \rangle = N_1 \tilde{\mathbf{A}}^{-1} + \tilde{\mathbf{Z}}_1^\top \tilde{\mathbf{Z}}_1. \quad (22)$$

Similarly, we obtain the following two updates for unlabeled data:

$$\tilde{\mathbf{Z}}_2 = \frac{1}{\sigma_x^2} \mathbf{K}_2 \mathbf{Z} \mathbf{C}^{-1} \left[\frac{1}{\sigma_x^2} \mathbf{C}^{-1} \mathbf{Z}^\top \mathbf{K} \mathbf{Z} \mathbf{C}^{-1} + \mathbf{I} \right]^{-1}, \quad (23)$$

$$\tilde{\mathbf{C}}_2 = N_2 \left[\frac{1}{\sigma_x^2} \mathbf{C}^{-1} \mathbf{Z}^\top \mathbf{K} \mathbf{Z} \mathbf{C}^{-1} + \mathbf{I} \right]^{-1} + \tilde{\mathbf{Z}}_2^\top \tilde{\mathbf{Z}}_2. \quad (24)$$

In M-step, we only need to update variances σ_x^2 and σ_y^2 as

$$\tilde{\sigma}_x^2 = \frac{1}{MN} \left[\text{tr}(\mathbf{K}) - \text{tr}(\tilde{\mathbf{Z}}^\top \mathbf{K} \tilde{\mathbf{Z}} \tilde{\mathbf{C}}^{-1}) \right], \quad (25)$$

$$\tilde{\sigma}_y^2 = \frac{1}{LN} \left[\text{tr}(\mathbf{Y} \mathbf{Y}^\top) - \text{tr}(\tilde{\mathbf{Z}}^\top \mathbf{Y} \mathbf{Y}^\top \tilde{\mathbf{Z}} \tilde{\mathbf{C}}^{-1}) \right], \quad (26)$$

which can be easily verified from (9) and (10).

Therefore, it can be seen that all interesting terms in the EM algorithm take input data into account *only* via the inner product. This nice property allows us to extend the SPPCA model to non-linear mappings by defining a *kernel function* $\kappa(\cdot, \cdot)$ for each pair of input data instead of the normal inner product [16].

Since we are now working on the *centered* data in the feature space, we can achieve this by modifying the kernel matrix $\tilde{\mathbf{K}}$ as

$$\tilde{\mathbf{K}} = \mathbf{K} - \frac{1}{N} \mathbf{1} \mathbf{1}^\top \mathbf{K} - \frac{1}{N} \mathbf{K} \mathbf{1} \mathbf{1}^\top + \frac{1}{N^2} \mathbf{1} \mathbf{1}^\top \mathbf{K} \mathbf{1} \mathbf{1}^\top, \quad (27)$$

where $\mathbf{1}$ denotes the all one column vector of dimension N . For the kernel vector $\mathbf{k}(\mathbf{X}, \mathbf{x}^*)$ given test data \mathbf{x}^* , it can also be centered by

$$\tilde{\mathbf{k}} = \mathbf{k} - \frac{1}{N} \mathbf{1} \mathbf{1}^\top \mathbf{k} - \frac{1}{N} \mathbf{K} \mathbf{1} + \frac{1}{N^2} \mathbf{1} \mathbf{1}^\top \mathbf{K} \mathbf{1}. \quad (28)$$

B. PROOF OF THEOREM 2

We give a sketch here. The mapping matrices are obtained by finding a fixed point in the EM algorithm in Section 4. For this proof we use notation $\mathbf{W} := \left(\frac{1}{\sigma_x} \mathbf{W}_x^\top, \frac{1}{\sigma_y} \mathbf{W}_y^\top \right)^\top$, and (3) can be rewritten as $\mathbf{A} = \mathbf{W}^\top \mathbf{W} + \mathbf{I}$. Plugging this and (5), (6) into (8) yields an update equation only related to \mathbf{W} : $\tilde{\mathbf{W}} = \mathbf{S} \mathbf{W} [\mathbf{W}^\top \mathbf{S} \mathbf{W} + \mathbf{W}^\top \mathbf{W} + \mathbf{I}]^{-1} (\mathbf{W}^\top \mathbf{W} + \mathbf{I})$. At the fixed point, this simplifies to $\mathbf{S} \mathbf{W} = \mathbf{W} \mathbf{W}^\top \mathbf{W} + \mathbf{W}$. Let $\mathbf{W} = \mathbf{U} \mathbf{D} \mathbf{V}^\top$ be the SVD of \mathbf{W} . Then each column \mathbf{u} of \mathbf{U} satisfies $\mathbf{S} \mathbf{u} = (1 + d^2) \mathbf{u}$, with d the corresponding singular value. Therefore, solving an eigenvalue problem for \mathbf{S} gives the mapping matrix \mathbf{W} , and plugging them into (4) gives the mapping \mathbf{t}^* for \mathbf{x}^* . \square

C. PROOF OF THEOREM 3

We give a sketch here. We define $\mathbf{B} = \mathbf{Z} \mathbf{C}^{-1}$ and rewrite (21) and (22) using only \mathbf{B} . This leads to $N \mathbf{B} (\mathbf{I} + \mathbf{B}^\top \mathbf{K} \mathbf{B}) = \mathbf{K} \mathbf{B}$. To solve \mathbf{B} we denote the SVD of $\mathbf{B}^\top \mathbf{K} \mathbf{B}$ as $\mathbf{Q} \mathbf{D} \mathbf{Q}^\top$, then we can obtain $\mathbf{B}^\top \mathbf{B} = \frac{1}{N} \mathbf{Q} \mathbf{D} (\mathbf{I} + \mathbf{D})^{-1} \mathbf{Q}^\top$ after some mathematics. Then define $\mathbf{U} := \mathbf{B} \mathbf{Q} \sqrt{N} \mathbf{D}^{-1/2} (\mathbf{I} + \mathbf{D})^{1/2}$, we have $\mathbf{U}^\top \mathbf{U} = \mathbf{I}$ and $\mathbf{U}^\top \mathbf{K} \mathbf{U} = N (\mathbf{I} + \mathbf{D})$. This clearly defines a SVD for \mathbf{K} , so by definition we have $\mathbf{\Lambda} = N (\mathbf{I} + \mathbf{D})$, and then we can solve for \mathbf{B} from \mathbf{U} and $\mathbf{\Lambda}$, which yields $\mathbf{B} = \frac{1}{\sqrt{N}} \mathbf{U} (\mathbf{I} - N \mathbf{\Lambda}^{-1})^{1/2} \mathbf{Q}^\top$. This recovers (18) with $\mathbf{R} = \mathbf{Q}^\top$ and the update equation for new test data can be easily obtained. \square