2007 年情報論的学習理論ワークショップ
2007 Workshop on Information-Based Induction Sciences (IBIS2007)
Tokyo, Japan, November 5 - 7, 2007.

Semi-Supervised Local Fisher Discriminant Analysis for Dimensionality Reduction*

Masashi Sugiyama[†] Tsuyoshi Idé[‡] Shinichi Nakajima[§] Jun Sese[¶]

Abstract: When only a small number of labeled samples are available, supervised dimensionality reduction methods tend to perform poorly due to overfitting. In such cases, unlabeled samples could be useful in improving the performance. In this paper, we propose a semi-supervised dimensionality reduction method which preserves the global structure of unlabeled samples in addition to separating labeled samples in different classes from each other. The proposed method has an analytic form of the globally optimal solution which can be computed based on eigendecompositions. Therefore, the proposed method is computationally reliable and efficient. We show the effectiveness of the proposed method through extensive simulations with benchmark data sets.
Keywords: semi-supervised learning, dimensionality reduction, cluster assumption, local Fisher discriminant analysis, principal component analysis.

1 Introduction

The goal of dimensionality reduction is to obtain a low-dimensional representation of high-dimensional data samples while preserving most of 'intrinsic information' contained in the original data. Once dimensionality reduction is carried out appropriately, the compact representation of the data can be used for various succeeding tasks such as visualization and classification.

In supervised learning scenarios where data samples are accompanied with class labels, *Fisher discriminant analysis* (FDA) [3] is a popular dimensionality reduction method. FDA seeks an embedding transformation such that between-class scatter is maximized and within-class scatter is minimized. FDA works very well if samples in each class are Gaussian with the common covariance structure. However, it tends to give undesired results if samples in a class form several separate clusters or there exist outliers [3]. To overcome this drawback, *local FDA (LFDA)* has been proposed [5], which localizes the between-class and within-class scatter matrices. LFDA works well even when within-class multimodality or outliers exist. Furthermore, LFDA overcomes critical limitation of original FDA in dimensionality reduction—the dimension of the FDA embedding space should be less than the number of classes [3], while LFDA does not suffer from this restriction in general.

However, the performance of LFDA (and all other supervised dimensionality reduction methods) tend to be degraded when only a small number of labeled samples are available. Thus, the supervised methods overfit embedding spaces to the labeled samples. In such cases, it is effective to make use of *unlabeled* samples which are often available abundantly, i.e., *semi-supervised learning*. The book [2] showed through extensive simulations that *principal component analysis* (PCA), which is an unsupervised dimensionality reduction method for preserving the global data structure, works moderately well in semi-supervised learning scenarios.

Although PCA is reported to work well, it may not be the best choice in semi-supervised learning due to its unsupervised nature. In this paper, we propose a new semi-supervised dimensionality reduction method which smoothly bridges LFDA and PCA so that we can control our reliance on the global structure of unlabeled samples and information brought by (a small number of) labeled samples. We experimentally show that the proposed method, which we refer to as *semi-supervised LFDA* (SELF), compares favorably with other methods for various data sets. Note that SELF maintains the same computational advantage of LFDA and PCA, i.e., a global solution can be analytically computed based on eigendecompositions. Therefore, SELF is still computationally efficient and reliable.

2 Dimensionality Reduction

In this section, we formulate the linear dimensionality reduction problem and review existing methods, which

*A more detailed version of this paper is available from '`http://www.cs.titech.ac.jp/~tr/`' as TR07-0006.

[†]Tokyo Institute of Technology, 2-12-1 O-okayama, Meguro-ku, Tokyo 152-8552, Japan, `sugi@cs.titech.ac.jp`

[‡]IBM Research, 1623-14 Shimotsuruma, Yamato-shi, Kanagawa 242-8502, Japan, `goodidea@jp.ibm.com`

[§]Nikon Corporation, 201-9 Oaza-Miizugahara, Kumagaya-shi, Saitama 360-8559, Japan, `nakajima.s@nikon.co.jp`

[¶]Ochanomizu University, 2-1-1 Otsuka, Bunkyo, Tokyo 112-8610, Japan, `sesejun@is.ocha.ac.jp`

will be the basis for developing a new method in the following sections.

2.1 Formulation and Notation

Let $\boldsymbol{x}_i \in \mathbb{R}^d$ $(i = 1, 2, \ldots, n)$ be d-dimensional samples, and let $\boldsymbol{z} \in \mathbb{R}^r$ $(1 \leq r \leq d)$ be a low-dimensional representation of a high-dimensional sample $\boldsymbol{x} \in \mathbb{R}^d$, where r is the dimensionality of the reduced space. Through the paper, we focus on linear dimensionality reduction, i.e., using a $d \times r$ transformation matrix $\boldsymbol{T}$, an embedded representation $\boldsymbol{z}$ of a sample $\boldsymbol{x}$ is obtained as

$$\boldsymbol{z} = \boldsymbol{T}^\top \boldsymbol{x}, \tag{1}$$

where $^\top$ denotes the transpose of a matrix or a vector.

2.2 Principal Component Analysis

A fundamental unsupervised dimensionality reduction method is *principal component analysis* (PCA).

Let $\boldsymbol{S}^{(t)}$ be the *total scatter matrix*:

$$\boldsymbol{S}^{(t)} \equiv \sum_{i=1}^{n} (\boldsymbol{x}_i - \boldsymbol{\mu})(\boldsymbol{x}_i - \boldsymbol{\mu})^\top, \tag{2}$$

where $\boldsymbol{\mu} \equiv \frac{1}{n}\sum_{i=1}^{n} \boldsymbol{x}_i$. The PCA transformation matrix $\boldsymbol{T}_{PCA}$ is defined as

$$\boldsymbol{T}_{PCA} \equiv \operatorname*{argmax}_{\boldsymbol{T} \in \mathbb{R}^{d \times r}} \left[\mathrm{tr}\left(\boldsymbol{T}^\top \boldsymbol{S}^{(t)} \boldsymbol{T} (\boldsymbol{T}^\top \boldsymbol{T})^{-1} \right) \right]. \tag{3}$$

That is, PCA seeks a transformation matrix $\boldsymbol{T}$ such that scatter in the embedding space is maximized.

2.3 Local Fisher Discriminant Analysis

Local Fisher discriminant analysis (LFDA) is a supervised dimensionality reduction method [5] which overcomes vulnerability of original FDA against within-class multimodality or outliers [3]. When discussing supervised learning problems, we assume that class labels $y_i \in \{1, 2, \ldots, c\}$ associated with the samples $\boldsymbol{x}_i$ are available, where c is the number of classes. Let n'_m be the number of samples in class $m \in \{1, 2, \ldots, c\}$.

Let $\boldsymbol{A}$ be the *affinity matrix*, i.e., the n'-dimensional square matrix with $A_{i,j}$ being the affinity between $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$. We assume that $A_{i,j} \in [0, 1]$; $A_{i,j}$ is large if $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$ are 'close' and $A_{i,j}$ is small if $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$ are 'far apart'. There are several different manners of defining $\boldsymbol{A}$, e.g., based on nearest neighbors or the heat kernel. Through the paper, we use the *local scaling heuristic* [7] as the definition of the affinity matrix $\boldsymbol{A}$, i.e., $A_{i,j} = \exp\left(-\|\boldsymbol{x}_i - \boldsymbol{x}_j\|^2 / \sigma_i \sigma_j\right)$. σ_i is the local scaling around $\boldsymbol{x}_i$ defined by $\sigma_i = \|\boldsymbol{x}_i - \boldsymbol{x}_i^{(k)}\|$, where $\boldsymbol{x}_i^{(k)}$ is the k-th nearest neighbor of $\boldsymbol{x}_i$. A heuristic choice of $k = 7$ has shown to be useful through extensive simulations [7, 5].

Let $\boldsymbol{S}^{(lb)}$ and $\boldsymbol{S}^{(lw)}$ be the *local* between-class scatter matrix and the *local* within-class scatter matrix:

$$\boldsymbol{S}^{(lb)} \equiv \frac{1}{2} \sum_{i,j=1}^{n'} W_{i,j}^{(lb)} (\boldsymbol{x}_i - \boldsymbol{x}_j)(\boldsymbol{x}_i - \boldsymbol{x}_j)^\top, \tag{4}$$

$$\boldsymbol{S}^{(lw)} \equiv \frac{1}{2} \sum_{i,j=1}^{n'} W_{i,j}^{(lw)} (\boldsymbol{x}_i - \boldsymbol{x}_j)(\boldsymbol{x}_i - \boldsymbol{x}_j)^\top, \tag{5}$$

$$W_{i,j}^{(lb)} \equiv \begin{cases} A_{i,j}(1/n' - 1/n'_{y_i}) & \text{if } y_i = y_j, \\ 1/n' & \text{if } y_i \neq y_j, \end{cases} \tag{6}$$

$$W_{i,j}^{(lw)} \equiv \begin{cases} A_{i,j}/n'_{y_i} & \text{if } y_i = y_j, \\ 0 & \text{if } y_i \neq y_j. \end{cases} \tag{7}$$

The LFDA transformation matrix $\boldsymbol{T}_{LFDA}$ is defined as

$$\boldsymbol{T}_{LFDA} \equiv \operatorname*{argmax}_{\boldsymbol{T} \in \mathbb{R}^{d \times r}} \left[\mathrm{tr}\left(\boldsymbol{T}^\top \boldsymbol{S}^{(lb)} \boldsymbol{T} (\boldsymbol{T}^\top \boldsymbol{S}^{(lw)} \boldsymbol{T})^{-1} \right) \right]. \tag{8}$$

In Eqs.(4) and (5), $A_{i,j}(1/n' - 1/n'_{y_i})$ is negative while $A_{i,j}/n'_{y_i}$ and $1/n'$ are non-negative. Thus, LFDA seeks a transformation matrix $\boldsymbol{T}$ such that nearby data pairs in the same class are made close and the data pairs in different classes are made apart; far apart data pairs in the same class are not imposed to be close. Samples in different classes are separated from each other irrespective of their affinity.

3 Semi-Supervised LFDA

In this section, we propose a new dimensionality reduction method for semi-supervised learning scenarios. From here on, we consider the case where, among all samples $\{\boldsymbol{x}_i\}_{i=1}^{n}$, only $\{\boldsymbol{x}_i\}_{i=1}^{n'}$ $(1 \leq n' \leq n)$ are labeled and the rest are unlabeled.

3.1 Basic Idea

When only a small number of labeled samples are available, supervised dimensionality reduction methods tend to find embedding spaces which are overfitted to the labeled samples. In such situations, using unlabeled samples is often effective—indeed, the book [2] showed through extensive simulations that PCA works well on the whole; our experimental results in Section 4 also show that PCA is sometimes better than LFDA. This means that preserving the global structure of all samples in an unsupervised manner can be better than strongly relying on class information provided by a small number of labeled samples.

Figure 1 depicts 2-dimensional 2-class examples; circle/triangle symbols denote samples in different classes and filled/unfilled symbols denote labeled/unlabeled samples; solid and dashed lines denote 1-dimensional embedding spaces found by LFDA and PCA, respectively (onto which data samples will be

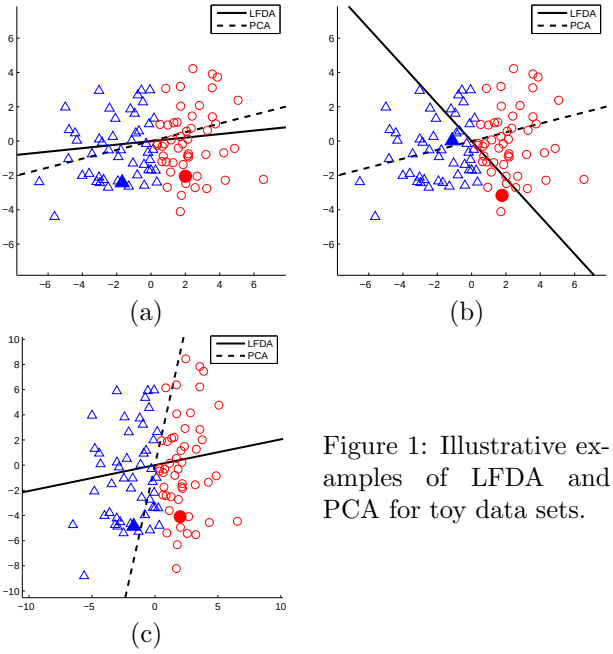

Figure 1: Illustrative examples of LFDA and PCA for toy data sets.

projected). For the data set in Figure 1(a), both LFDA and PCA can find good embedding spaces, which separate unlabeled samples in different classes from each other. However, for the data set in Figure 1(b), LFDA finds an embedding space that is overfitted to the labeled samples. On the other hand, in the case of Figure 1(c), PCA does not work well due to its unsupervised nature.

The above result implies that LFDA and PCA can compensate for the weakness of each other, i.e., LFDA can utilize label information, while PCA can avoid overfitting. Our simulation results with benchmark data sets in Section 4 also show that LFDA and PCA work in a complementary manner. Motivated by these facts, we propose *bridging* LFDA and PCA so that we can smoothly control our reliance on the global structure of unlabeled samples and class information brought by labeled samples. We refer to the proposed method as *semi-supervised LFDA* (SELF).

The embedding transformations of LFDA and PCA can be analytically computed based on eigendecompositions. So we combine the eigenvalue problems of LFDA and PCA and solve them together. This allows us to maintain the computational efficiency and reliability of LFDA and PCA.

3.2 Definition

More specifically, we propose solving the following generalized eigenvalue problem:

$$\boldsymbol{S}^{(rlb)}\boldsymbol{\varphi} = \lambda \boldsymbol{S}^{(rlw)}\boldsymbol{\varphi}, \tag{9}$$

where $\boldsymbol{S}^{(rlb)}$ and $\boldsymbol{S}^{(rlw)}$ are *regularized* local between-class scatter matrix and *regularized* local within-class scatter matrix defined by

$$\boldsymbol{S}^{(rlb)} \equiv (1-\beta)\boldsymbol{S}^{(lb)} + \beta \boldsymbol{S}^{(t)}, \tag{10}$$

$$\boldsymbol{S}^{(rlw)} \equiv (1-\beta)\boldsymbol{S}^{(lw)} + \beta \boldsymbol{I}_d. \tag{11}$$

β ($\in [0,1]$) is a trade-off parameter—SELF is reduced to LFDA when $\beta = 0$, and SELF is reduced to PCA when $\beta = 1$. In general, SELF inherits characteristics of both LFDA and PCA.

The optimization problem of SELF is expressed as

$$\boldsymbol{T}_{SELF} \equiv \underset{\boldsymbol{T}\in\mathbb{R}^{d\times r}}{\operatorname{argmax}} \left[\operatorname{tr}\left(\boldsymbol{T}^\top \boldsymbol{S}^{(rlb)} \boldsymbol{T} (\boldsymbol{T}^\top \boldsymbol{S}^{(rlw)} \boldsymbol{T})^{-1} \right) \right]. \tag{12}$$

A solution $\boldsymbol{T}_{SELF}$ is analytically computed as $(\boldsymbol{\varphi}_1 | \boldsymbol{\varphi}_2 | \cdots | \boldsymbol{\varphi}_r)$, where $\{\boldsymbol{\varphi}_k\}_{k=1}^d$ are the generalized eigenvectors of Eq.(9) associated with the generalized eigenvalues $\{\lambda_k\}_{k=1}^d$. We assume that $\{\lambda_k\}_{k=1}^d$ are sorted as $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_d$ and $\{\boldsymbol{\varphi}_k\}_{k=1}^d$ are normalized as $\boldsymbol{\varphi}_k^\top \boldsymbol{S}^{(rlw)} \boldsymbol{\varphi}_k = 1$.

It can be confirmed that Eq.(12) is invariant under linear transformations [3, 5], i.e., for any r-dimensional invertible matrix $\boldsymbol{U}$, $\boldsymbol{T}_{OPT}\boldsymbol{U}$ is also a global solution. This implies that the *range* of the embedding space can be uniquely determined by Eq.(12), but the *metric* in the embedding space is arbitrary. A practically useful heuristic (e.g, [5]) is to set $\boldsymbol{U} = \operatorname{diag}(\sqrt{\lambda_1}, \sqrt{\lambda_2}, \ldots, \sqrt{\lambda_r})$, where $\operatorname{diag}(\cdot)$ denotes a diagonal matrix. Then the solution $\boldsymbol{T}_{SELF}$ becomes

$$\boldsymbol{T}_{SELF} = (\sqrt{\lambda_1}\boldsymbol{\varphi}_1 | \sqrt{\lambda_2}\boldsymbol{\varphi}_2 | \cdots | \sqrt{\lambda_r}\boldsymbol{\varphi}_r), \tag{13}$$

We can prove that all the generalized eigenvalues are non-negative, which guarantees that the solution (13) is always valid.

3.3 Properties

First, we give an interpretation of $\boldsymbol{S}^{(rlb)}$. The matrix $\boldsymbol{S}^{(rlb)}$ can be expressed as

$$\boldsymbol{S}^{(rlb)} \equiv \frac{1}{2}\sum_{i,j=1}^{n} W_{i,j}^{(rlb)} (\boldsymbol{x}_i - \boldsymbol{x}_j)(\boldsymbol{x}_i - \boldsymbol{x}_j)^\top, \tag{14}$$

$$W_{i,j}^{(rlb)} \equiv \begin{cases} (1-\beta)A_{i,j}(1/n' - 1/n'_{y_i}) + \beta/n & \text{if } y_i = y_j, \\ (1-\beta)/n' + \beta/n & \text{if } y_i \neq y_j, \\ \beta/n & \text{otherwise.} \end{cases} \tag{15}$$

The first case in Eq.(15) is negative if

$$\beta < \frac{A_{i,j} n (n' - n'_{y_i})}{A_{i,j} n (n' - n'_{y_i}) + n' n'_{y_i}}. \tag{16}$$

This implies that SELF tries to make sample pairs in the same class close if β is small, while it separates them from each other if β is large. Thus the local data structure in the same class tends to be preserved when β is small, but it is no longer preserved when β is large.

The second case in Eq.(15) is always positive for any $\beta \in [0, 1]$, implying that SELF always tries to make sample pairs in different classes apart for any β. This would be natural in semi-supervised learning scenarios. The third case in Eq.(15) is always non-negative, implying that unlabeled samples are separated from each other for preserving the global data structure.

Finally, we give an interpretation of $\boldsymbol{S}^{(rlw)}$. When $\beta = 0$, $\boldsymbol{S}^{(rlw)}$ $(= \boldsymbol{S}^{(lw)})$ could be ill-conditioned—this is crucial particularly when the dimension d of the original data space is larger than the number n' of labeled samples. In such situations, $\beta \boldsymbol{I}_d$ included in $\boldsymbol{S}^{(rlw)}$ works as a *regularizer* and SELF can avoid overfitting to the labeled samples. Therefore, SELF is regarded as a regularized variant of LFDA and would be more stable and reliable than original LFDA particularly when the number of labeled samples is small. Note that unlike Eq.(14), $\boldsymbol{S}^{(rlw)}$ does not have a pairwise expression since $\boldsymbol{I}_d$ can not be expressed in a pairwise form.

3.4 Numerical Examples

For illustrating how SELF behaves, let us use the *Olivetti face* data set[1]. The data set consists of 400 gray-scale face images (40 people, 10 images per person); each image consists of 4096 $(= 64 \times 64)$ pixels and each pixel takes an integer value between 0 and 255 as the intensity level. In this simulation, we use the image samples of only 10 subjects (i.e., totally 100 images) for making the visualization results clear. We note that the result does not change essentially (but visually denser) when all 400 images are used.

Among 10 people used for the experiments, 3 subjects are with glasses and other 7 are without glasses (see the left-most pictures of Figure 2). Our task is to embed the face images into a two-dimensional space so that the subjects *with* and *without* glasses are separated from each other. We treat 1 image per person as labeled (i.e., totally 3 faces with glasses and 7 faces without glasses) and the rest are treated as unlabeled. Since each class contains several different subjects, this data set is thought to possess within-class multimodality.

The embedded results are shown in Figure 2, where circle/triangle symbols are faces with/without glasses and filled/unfilled symbols are labeled/unlabeled samples. The figure shows that FDA and LFDA perfectly separate the labeled samples in different classes from each other. However, unlabeled samples tend to be mixed due to an overfitting phenomenon. PCA tends to mix the labeled samples in different classes due to the unsupervised nature. Consequently, unlabeled samples in different classes are also mixed. On the other hand, SELF with $\beta = 0.5$ clearly separates the labeled samples in different classes from each other, and at the same time, it also nicely separates the unlabeled samples in different classes from each other. We note that, in this visualization simulation, the result of SELF is not sensitive to the choice of the trade-off parameter β; the results are almost unchanged for $0.01 \leq \beta \leq 0.99$.

[1] 'http://www.cs.toronto.edu/~roweis/data.html'

3.5 Distance Metric Learning

The performance of distance-based learning methods such as nearest neighbor classifiers depend heavily on the definition of distances between samples. The idea of *distance metric learning* is to optimize a metric $\boldsymbol{M}$ used for computing distances between samples:

$$\mathrm{dist}(\boldsymbol{x}_i, \boldsymbol{x}_j; \boldsymbol{M}) = (\boldsymbol{x}_i - \boldsymbol{x}_j)^\top \boldsymbol{M} (\boldsymbol{x}_i - \boldsymbol{x}_j). \qquad (17)$$

The metric matrix $\boldsymbol{M}$ is symmetric and positive semi-definite by definition. For this reason, metric learning is typically formulated as a *semi-definite programming* (SDP) problem, which is a convex optimization problem and the unique global solution can be obtained [1, 6].

If the rank of the d-dimensional matrix $\boldsymbol{M}$ is constrained to r, distance metric learning methods automatically carry out dimensionality reduction in an implicit manner. More specifically, symmetricity and positive semi-definiteness of the metric matrix $\boldsymbol{M}$ implies that $\boldsymbol{M}$ can be decomposed as $\boldsymbol{M} = \boldsymbol{T}\boldsymbol{T}^\top$. Thus $\boldsymbol{T}^\top \boldsymbol{x}_i$ is an explicit expression of a sample $\boldsymbol{x}_i$ after dimensionality reduction. However, simultaneously reducing the dimensionality of samples and learning the distance metric is usually hard since the rank constraint is non-convex [1]. Thus it may not be possible to obtain the global optimal solution.

On the other hand, our approach to dimensionality reduction allows us to obtain the global solution in terms of the range of the embedding space. This means that we can obtain the unique solution of the metric matrix by combining SELF with a convex metric learning method (e.g., [6]), i.e., a two-stage procedure of first reducing the dimensionality (i.e., determining the range of the embedding space) by SELF and then learning the metric of the embedding space without the rank constraint. We expect that this procedure is practically useful.

4 Simulations

In this section, we experimentally evaluate the performance of relevant dimensionality reduction methods using standard classification benchmark data sets.

The book [2] conducted systematic experiments for comparing semi-supervised learning methods. The results showed that each method performs very well for a particular type of data sets. However, at the same time, it tends to be poor for other kinds of data sets. Thus, the performance of semi-supervised learning methods is highly dependent on the type of data

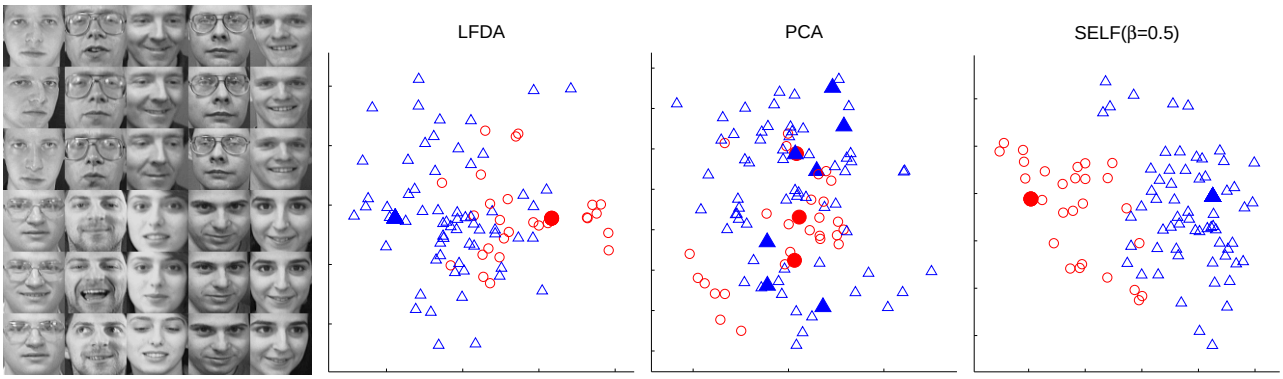

Figure 2: Embedded face samples (glasses vs. non-glasses). Circle/triangle symbols are faces with/without glasses and filled/unfilled symbols are labeled/unlabeled samples.

sets and there seems to be no single best method. On the other hand, 1-*nearest neighbor classifier* is shown to be stable for various data sets, although it may not be the best possible method in semi-supervised classification. For avoiding the bias caused by the choice of the learning methods, we decided to use the 1-nearest neighbor classifier in our experiments.

The misclassification rate is sometimes monotone increasing as the dimensionality is reduced. In such cases, if the best dimensionality is chosen, e.g., by cross-validation, the largest dimension is mostly chosen (i.e., no dimensionality reduction). Then we may not be able to compare the performance of dimensionality reduction methods in a meaningful way. Prefixing the reduced dimensionality r to some number is a possible option for avoiding the above problem, but the evaluation results can significantly depend on the choice of the dimensionality. Based on this argument, we decided to use the *average* misclassification rate over reduced dimensions (or equivalently the area under the classification error curve) as our error metric, which we believe to be reasonable in the current experiments.

We employ the benchmark data sets taken from the book [2], which consist of 9 semi-supervised data sets. We refer to them as the *SSL* data sets. We did not test the *SSL8* and *SSL9* data sets since they are too huge. Note that the *SSL6* data set contains 6 classes, while the other data sets have 2 classes. Table 1 describes the mean and standard deviation of the misclassification rate over repetitions. Since we had a numerical problem when computing LFDA, we slightly regularized it and consider SELF with $\beta = 0.001$ as LFDA. The fulfillment of the *cluster assumption* [2] is described as 'CA', which is the correct classification rate by the 1-nearest-neighbor classifier when both training and test labels are used for classifying all the training and test samples. Note that CA is computed *before* dimensionality reduction is applied, so it represents the fulfillment of the cluster assumption of the original data samples. The larger the value of CA is, the more reliable the cluster assumption would be (although the values are coarse).

When the number of labeled samples is 100 (see the upper half of the table), LFDA and PCA tend to work well in a complementary way—LFDA works well if CA is small while PCA works well if CA is large. SELF with $\beta = 0.5$ tends to make up the deficit of each method; moreover it can outperform both LFDA and PCA for some cases. We also test 'SELF(CV)', where β in SELF is chosen from $\{0, 0.25, 0.5, 0.75, 1\}$ by 10-fold cross validation. The results shown in the table show that SELF(CV) further improves the performance over SELF with $\beta = 0.5$. *Locality preserving projection* (LPP) [4], an unsupervised dimensionality reduction method which tends to preserve cluster structures, does not work so well on the whole. The combination of LFDA and LPP (indicated by SELF'(CV) in the table) also does not perform as good as SELF(CV). We also tested the combination of LFDA, PCA, and LPP, but this did not further improve the performance over SELF so we omit the detail.

When the number of labeled samples is only 10 (see the lower half of Table 1), the difference of the performance among the methods shrinks but SELF(CV) is still slightly better than the other methods.

5 Conclusions

Our approach to dimensionality reduction in this paper is called the *filter* approach, i.e., the dimensionality reduction procedure is independent of subsequent classification algorithms. Our experimental results showed that the proposed method, *SELF*, works well when it is combined with the 1-nearest-neighbor classifier. An important future direction is to develop a *wrapper* method of semi-supervised dimensionality reduction, which explicitly takes properties of subsequent classification algorithms into account. We expect that a wrapper approach is promising in semi-supervised learning since the performance of elaborate semi-supervised learning methods is highly dependent on the reliability of the assumption behind unlabeled samples such as the cluster or manifold structure [2].

Table 1: Misclassification rate for the SSL data sets. The numbers in the bracket are the standard deviation over repetitions. For each data set, the best method and comparable ones based on the *t-test* at the significance level 5% are described in bold face. SELF(CV) denotes SELF with β chosen by cross validation. SELF' denotes the combination of LFDA and LPP in a similar manner.

Data	Dim	Lab	Unlab	Rep	CA	LFDA	SELF ($\beta = 0.5$)	PCA	SELF (CV)	LPP	SELF' (CV)
SSL1	241	100	1400	12	0.98	14.9(1.8)	**6.0(1.3)**	**6.2(1.1)**	**6.0(1.4)**	27.4(1.4)	28.4(2.6)
SSL2	241	100	1400	12	0.97	15.7(0.9)	**9.6(1.1)**	11.2(0.8)	**10.3(2.4)**	24.1(2.2)	21.9(1.9)
SSL3	241	100	1400	12	1.00	21.1(3.9)	**14.3(1.8)**	15.5(1.0)	**14.1(1.4)**	18.0(2.4)	18.5(2.4)
SSL4	117	100	300	12	0.58	**33.4(3.5)**	36.6(2.4)	48.7(2.4)	**33.4(3.7)**	46.7(1.7)	**36.0(4.7)**
SSL5	241	100	1400	12	0.64	**27.5(2.3)**	**27.2(2.3)**	31.0(1.9)	**27.3(2.9)**	37.0(1.3)	35.3(1.9)
SSL6	241	100	1400	12	0.98	38.1(1.5)	35.4(2.4)	**27.3(2.7)**	**27.0(2.7)**	35.2(1.7)	36.9(3.2)
SSL7	241	100	1400	12	0.68	29.4(2.4)	**29.1(2.4)**	29.3(1.6)	**27.7(1.4)**	32.0(0.9)	32.8(1.5)
# Bests						2	5	2	7	0	1
SSL1	241	10	1490	12	0.98	**22.9(5.1)**	26.3(6.1)	**19.2(4.2)**	**22.3(5.4)**	45.9(2.3)	48.5(2.4)
SSL2	241	10	1490	12	0.97	**22.3(3.0)**	**21.3(2.9)**	25.8(4.2)	**21.5(2.5)**	31.2(7.5)	**21.4(0.8)**
SSL3	241	10	1490	12	1.00	**42.7(2.9)**	**42.9(3.0)**	**42.7(4.2)**	43.6(3.2)	**40.4(4.1)**	**41.0(5.2)**
SSL4	117	10	390	12	0.58	**47.3(2.9)**	**47.7(2.7)**	49.9(2.2)	**48.3(3.3)**	**49.5(2.5)**	**48.5(1.9)**
SSL5	241	10	1490	12	0.64	45.4(4.4)	45.4(4.4)	**36.3(5.5)**	**40.2(6.9)**	41.2(3.3)	44.5(3.6)
SSL6	241	10	1490	12	0.98	**67.7(4.6)**	**67.0(4.0)**	**67.7(4.1)**	**67.6(4.6)**	71.4(4.0)	73.7(2.9)
SSL7	241	10	1490	12	0.68	43.6(5.2)	43.6(5.2)	**38.9(5.7)**	**40.1(7.1)**	**40.3(4.2)**	**42.7(5.3)**
# Bests						5	4	5	6	3	4

Although we focused on linear dimensionality reduction, we can easily obtain a non-linear variant of SELF by employing the standard *kernel trick*. However, the kernelized SELF shares the common difficulty in kernel methods, i.e., how to choose the kernel functions. This needs to be investigated in the context of semi-supervised dimensionality reduction. In the future work, we will also explore semi-supervised dimensionality reduction of structured data using the kernel SELF.

Although we focused on linear dimensionality reduction, it is straightforward to show that a non-linear variant of SELF can be obtained by employing the standard *kernel trick*. However, the kernelized variant shares the problem that is common to all kernel methods, i.e., how to choose the kernel functions. This needs to be investigated in the context of semi-supervised dimensionality reduction. The kernelized variant can also be used for dimensionality reduction of *non-vectorial structured data* such as strings, trees, and graphs. In the future work, we will explore semi-supervised dimensionality reduction of such structured data.

A remaining important issue to be discussed—which is common to all semi-supervised learning techniques—is how to optimize tuning parameters. We may simply employ cross-validation for this purpose, but it has two potential problems. The first problem is that the number of labeled samples is typically small in semi-supervised learning scenarios and thus cross-validation is not reliable [2]. Fortunately, our experiments showed that SELF is not so sensitive to the trade-off parameter β in small sample cases, but there is still room for further improvement. The second problem is that labeled samples and unlabeled samples can have different (input) distributions. Such a situation is referred to as *covariate shift* in statistics and ordinary cross-validation is known to be significantly biased; *importance-weighted* cross-validation is unbiased under covariate shift. In the future work, we will investigate how the covariate shift adaptation techniques could be employed in the context of semi-supervised dimensionality reduction.

Acknowledgments: MS acknowledges financial support from MEXT (17700142 and 18300057) and Tateishi Science and Technology Foundation.

References

[1] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[2] O. Chapelle, B. Schölkopf, and A. Zien, editors. *Semi-Supervised Learning*. MIT Press, 2006.

[3] K. Fukunaga. *Introduction to Statistical Pattern Recognition*. Academic Press, 1990.

[4] X. He and P. Niyogi. Locality preserving projections. *NIPS16*, 2004.

[5] M. Sugiyama. Dimensionality reduction of multimodal labeled data by local fisher discriminant analysis. *JMLR*, vol.8, pp.1027–1061, 2007.

[6] K. Weinberger, J. Blitzer, and L. Saul. Distance metric learning for large margin nearest neighbor classification. *NIPS18*, 2006.

[7] L. Zelnik-Manor and P. Perona. Self-tuning spectral clustering. *NIPS17*, 2005.