

Learning a Locality Preserving Subspace for Visual Recognition

Xiaofei He*, Shuicheng Yan#, Yuxiao Hu, and Hong-Jiang Zhang

Microsoft Research Asia, Beijing 100080, China

*Department of Computer Science, University of Chicago (xiaofei@cs.uchicago.edu)

#School of Mathematical Science, Peking University, China

Abstract

Previous works have demonstrated that the face recognition performance can be improved significantly in low dimensional linear subspaces. Conventionally, principal component analysis (PCA) and linear discriminant analysis (LDA) are considered effective in deriving such a face subspace. However, both of them effectively see only the Euclidean structure of face space. In this paper, we propose a new approach to mapping face images into a subspace obtained by Locality Preserving Projections (LPP) for face analysis. We call this Laplacianface approach. Different from PCA and LDA, LPP finds an embedding that preserves local information, and obtains a face space that best detects the essential manifold structure. In this way, the unwanted variations resulting from changes in lighting, facial expression, and pose may be eliminated or reduced. We compare the proposed Laplacianface approach with eigenface and fisherface methods on three test datasets. Experimental results show that the proposed Laplacianface approach provides a better representation and achieves lower error rates in face recognition.

1. Introduction

In recent years, computer vision research has witnessed a growing interest in subspace analysis techniques [1][6][14][16][20][21]. A face image can be represented as a point in the image space (given by the number of pixels in the image). Before we utilize any classification technique, it is beneficial to first perform dimensionality reduction to project an image into a low dimensional feature space or so-called face space, due to the consideration of learnability and computational efficiency. Specifically, learning from examples is computationally infeasible if it has to rely on high-dimensional representations. In particular, Principal Component Analysis (PCA) [16] and Linear Discriminant Analysis (LDA) [1] have been applied to face recognition with impressive results.

PCA is an eigenvector method designed to model linear variation in high-dimensional data. PCA performs dimensionality reduction by projecting the original n -dimensional data onto the k ($k \ll n$)-dimensional linear subspace spanned by the leading eigenvectors of the data's covariance matrix. Its goal is to find a set of mutually orthogonal basis functions that capture the directions of

maximum variance in the data and for which the coefficients are pairwise decorrelated. For linearly embedded manifolds, PCA is guaranteed to discover the dimensionality of the manifold and produces a compact representation.

LDA is a supervised learning algorithm. LDA searches for the projection axes on which the data points of different classes are far from each other and at the same time where the data points of a same class are close to each other. Unlike PCA which encodes information in an orthogonal linear space, LDA encodes discriminating information in a linear separable space whose bases are not necessarily orthogonal.

Recently, a number of research efforts have shown that the face images possibly reside on a nonlinear submanifold [9][10][11][15]. However, both PCA and LDA effectively see only the Euclidean structure. They fail to discover the underlying structure, if the face images lie on a nonlinear submanifold hidden in the image space. Some nonlinear techniques have been proposed to discover the nonlinear structure of the manifold, i.e. Isomap [15], LLE [9], and Laplacian eigenmaps [2]. These nonlinear methods do yield impressive results on some benchmark artificial data sets. However, they yield maps that are defined only on the training data points and how to evaluate the maps on new testing points remains unclear. Therefore, these nonlinear manifold learning techniques might not be suitable for some computer vision tasks, such as face recognition.

In the meantime, there has been some interest in the problem of developing low dimensional representations through kernel based techniques for face recognition [5][19]. These methods can discover the nonlinear structure of the face images. However, they are computationally expensive. Moreover, none of them explicitly considers the structure of the manifold on which the face images possibly reside.

In this paper, we propose a new approach to face representation and recognition, which explicitly considers the face manifold structure. To be specific, an adjacency graph is constructed to model the local structure of the face manifold. A *Locality Preserving Subspace* for face representation is learned by using *Locality Preserving Projections* (LPP). Each face image in the image space is mapped to a low-dimensional face subspace, which is

characterized by a set of feature images, called *Laplacian-faces*. The face subspace preserves local structure, and thus has more discriminating power than eigenfaces from the classification viewpoint. Moreover, the locality preserving property makes our algorithm insensitive to the unwanted variations due to changes in lighting, facial expression, and viewing points.

It is worthwhile to highlight several aspects of the proposed approach here:

1. While PCA aims to preserve the global structure of the image space, and LDA aims to preserve the discriminating information; LPP aims to preserve the local structure of the image space. In many real world classification problems, the local manifold structure is more important than the global Euclidean structure, especially when nearest-neighbor like classifiers are used for classification.
2. An efficient subspace learning algorithm for face recognition should be able to detect the nonlinear *manifold* structure of the face space. Our proposed Laplacianface method explicitly considers the manifold structure which is modeled by an adjacency graph.
3. LPP shares some similar properties with LLE, such as locality preserving character. However, their objective functions are totally different. LPP is obtained by finding the optimal linear approximations to the eigenfunctions of the Laplace Beltrami operator on the manifold [2][4]. LPP is linear, while LLE is nonlinear. Moreover, LPP is defined everywhere, while LLE is defined only on the training data points and it is unclear how to evaluate the map for new test points. In contrast, LPP may be simply applied to any new data point to locate it in the reduced representation space.

The rest of this paper is organized as follows: Section 2 describes the objective functions of PCA and LDA. The *Locality Preserving Projection* algorithm is described in section 3. In section 4, we present the manifold ways of face analysis. The experimental results are shown in Section 5. Finally, we give concluding remarks and future work in Section 6.

2. PCA and LDA

One approach to coping with the problem of excessive dimensionality of the image space is to reduce the dimensionality by combining features. Linear combinations are particularly attractive because they are simple to compute and analytically tractable. In effect, linear methods project the high-dimensional data onto a lower dimensional subspace.

Considering the problem of representing all of the vectors in a set of n d -dimensional samples $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n$, with zero mean, by a single vector $\mathbf{y} = \{y_1, y_2, \dots, y_n\}$ such that y_i represent \mathbf{x}_i . Specifically, we find a linear mapping from the d -dimensional space to a line. Without loss of generality, we denote the transformation vector by \mathbf{w} .

That is, $\mathbf{w}^T \mathbf{x}_i = y_i$. Actually, the magnitude of \mathbf{w} is of no real significance, because it merely scales y_i . In face recognition, each vector \mathbf{x}_i denotes a face image.

Different objective functions will yield different algorithms with different properties. PCA seeks a projection that best represents the data in a least-squares sense. The matrix $\mathbf{w}\mathbf{w}^T$ is a projection onto the principal component space spanned by $\{\mathbf{w}\}$ which minimizes the following objective function,

$$\min_{\mathbf{w}} \sum_{i=1}^n \left| \mathbf{x}_i - \mathbf{w}\mathbf{w}^T \mathbf{x}_i \right|^2$$

The output set of principal vectors $\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_k$ are an orthonormal set of vectors representing the eigenvectors of the sample covariance matrix associated with the $k < d$ largest eigenvalues.

While PCA seeks directions that are efficient for representation, LDA seeks directions that are efficient for discrimination. Suppose we have a set of n d -dimensional samples $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n$, belonging to l classes of faces. The objective function is as follows,

$$\max_{\mathbf{w}} \frac{\mathbf{w}^T S_B \mathbf{w}}{\mathbf{w}^T S_W \mathbf{w}}$$

$$S_B = \sum_{i=1}^l |C_i| (\mathbf{m}^{(i)} - \mathbf{m})(\mathbf{m}^{(i)} - \mathbf{m})^T$$

$$S_W = \sum_{i=1}^l |C_i| E \left[(\mathbf{x}^{(i)} - \mathbf{m}^{(i)})(\mathbf{x}^{(i)} - \mathbf{m}^{(i)})^T \right]$$

where \mathbf{m} is the total sample mean vector, $|C_i|$ is the number of samples in class C_i , $\mathbf{m}^{(i)}$ are the average vectors of C_i , and $\mathbf{x}^{(i)}$ are the sample vectors associated to C_i . We call S_W the *within-class scatter matrix* and S_B the *between-class scatter matrix*.

3. Learning a Locality Preserving Subspace

Both PCA and LDA aim to preserve the global structure. However, in many real world applications, the local structure is more important, especially when nearest-neighbor search needs to be performed. In this section, we describe how to learn a Locality Preserving Subspace by using Locality Preserving Projections (LPP) [4]. LPP is a linear approximation of the nonlinear Laplacian Eigenmaps [2]. It seeks to preserve the intrinsic geometry of the data and the local structures. The objective function of LPP is as follows:

$$\min_{\mathbf{y}} \sum_{ij} (y_i - y_j)^2 S_{ij}$$

The objective function with our choice of symmetric weights S_{ij} ($S_{ij} = S_{ji}$) incurs a heavy penalty if neighboring points \mathbf{x}_i and \mathbf{x}_j are mapped far apart. Therefore, minimizing it is an attempt to ensure that if \mathbf{x}_i and \mathbf{x}_j are “close” then y_i and y_j are close as well. S_{ij} can be thought of as a similarity measure between objects. Let \mathbf{w} denote the

transformation vector. By simple algebra formulation, we can reduce the above objective function as follows:

$$\begin{aligned}
& \frac{1}{2} \sum_{ij} (y_i - y_j)^2 S_{ij} \\
&= \sum_{ij} (\mathbf{w}^T \mathbf{x}_i - \mathbf{w}^T \mathbf{x}_j)^2 S_{ij} \\
&= \sum_i \mathbf{w}^T \mathbf{x}_i D_{ii} \mathbf{w}^T \mathbf{x}_i - \sum_{ij} \mathbf{w}^T \mathbf{x}_i S_{ij} \mathbf{w}^T \mathbf{x}_j \\
&= \mathbf{w}^T X(D - S)X^T \mathbf{w} \\
&= \mathbf{w}^T XLX^T \mathbf{w}
\end{aligned}$$

where $X = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n]$, and D is a diagonal matrix; its entries are column (or row, since S is symmetric) sums of S , $D_{ii} = \sum_j S_{ji}$. $L = D - S$ is the Laplacian matrix [3]. Matrix D provides a natural measure on the data points. The bigger the value D_{ii} (corresponding to y_i) is, the more “important” is y_i . Therefore, we impose a constraint as follows:

$$\begin{aligned}
& \mathbf{y}^T D \mathbf{y} = 1 \\
\Rightarrow & \mathbf{w}^T XDX^T \mathbf{w} = 1
\end{aligned}$$

Finally, the minimization problem reduces to finding:

$$\arg \min_{\mathbf{w}^T XDX^T \mathbf{w} = 1} \mathbf{w}^T XLX^T \mathbf{w}$$

The transformation vector \mathbf{w} that minimizes the objective function is given by the minimum eigenvalue solution to the generalized eigenvalue problem:

$$XLX^T \mathbf{w} = \lambda XDX^T \mathbf{w}$$

Note that the two matrices XLX^T and XDX^T are both symmetric and positive semi-definite.

The derivation reflects the intrinsic geometric structure of the manifold. The theoretical justification for LPP can be traced back to [4].

4. Manifold Ways of Face Analysis

In the above two sections, we have described three different linear subspace learning algorithm. The key difference between PCA, LDA and LPP is that, PCA and LDA aim to discover Euclidean structure, while LPP aims to discover manifold structure. In this Section, we discuss the manifold ways of face analysis.

4.1. Manifold Learning via Dimensionality Reduction

In many cases, face images may be visualized as points drawn on a low-dimensional manifold hidden in a high-dimensional Euclidean space. Specially, we can consider that a sheet of rubber is crumpled into a (high dimensional)

ball. The objective of a dimensionality-reducing mapping is to unfold the sheet and to make its low-dimensional structure explicit. If the sheet is not torn in the process, the mapping is topology-preserving. Moreover, if the rubber is not stretched or compressed, the mapping preserves the metric structure of the original space. In this paper, our objective is to discover the face manifold by a locality-preserving mapping for face representation and recognition.

4.2. Learning Laplacianfaces for Representation

In section 3, we have described LPP, a method for learning a locality preserving subspace. It is obtained by finding the optimal linear approximations to the eigenfunctions of the Laplace Betrami operator on the manifold [4]. Base on LPP, we describe our Laplacianface method for face representation and recognition.

In the face analysis and recognition problems one is confronted with the difficulty that the matrix XDX^T is sometimes singular. This stems from the fact that, sometimes the number of images in the training set (m) is much smaller than the number of pixels in each image (n). In such case, the rank of XDX^T is at most m , while XDX^T is an $n \times n$ matrix, which implies that XDX^T is singular. To overcome the complication of a singular XDX^T , we first project the image set to a PCA subspace so that the resulting matrix XDX^T is nonsingular. Another consideration of using PCA as preprocessing is for noise reduction. This method, we call Laplacianface, can learn an optimal subspace for face representation and recognition.

The algorithmic procedure of Laplacianface is formally stated below:

1. **PCA projection:** We project the image set $\{\mathbf{x}_i\}$ into the PCA subspace by throwing away the smallest principal components. In our experiments, we kept 98% information in the sense of reconstruction error. For the sake of simplicity, we still use \mathbf{x} to denote the images in the PCA subspace in the following steps. We denote the transformation matrix of PCA by W_{PCA} .
2. **Constructing the nearest-neighbor graph:** Let G denote a graph with n nodes. The i^{th} node corresponds to the face image \mathbf{x}_i . We put an edge between nodes i and j if \mathbf{x}_i and \mathbf{x}_j are “close”, i.e. \mathbf{x}_i is among k nearest neighbors of \mathbf{x}_j or \mathbf{x}_j is among k nearest neighbors of \mathbf{x}_i . Note that, one might take a more utilitarian perspective and construct a nearest-neighbor graph based on the class labels. That is, we put an edge between two nodes if and only if they have the same class label. The constructed nearest-neighbor graph is an approximation of the local manifold structure.
3. **Choosing the weights:** If node i and j are connected, put




Figure 1. Two-dimensional linear embedding of face images by Locality Preserving Projection. As can be seen, the face images are divided into two parts, the faces with open mouth and the faces with closed mouth. Moreover, it can be clearly seen that the pose and expression of human faces change continuously and smoothly, from top to bottom, from left to right. The bottom images correspond to points along the right path (linked by solid line), illustrating one particular mode of variability in pose.

$$S_{ij} = e^{-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|^2}{t}}$$

where t is a suitable constant. Otherwise, put $S_{ij} = 0$. The weight matrix S of graph G models the face manifold structure by preserving local structure. The justification for this choice of weights can be traced back to [2].

4. **Eigenmap:** Compute the eigenvectors and eigenvalues for the generalized eigenvector problem:

$$XLX^T \mathbf{w} = \lambda XD X^T \mathbf{w} \quad (1)$$

where D is a diagonal matrix whose entries are column (or row, since S is symmetric) sums of S , $D_{ii} = \sum_j S_{ji}$. $L = D - S$ is the Laplacian matrix. The i^{th} column of the matrix X is \mathbf{x}_i .

Let $\mathbf{w}_0, \mathbf{w}_1, \dots, \mathbf{w}_{k-1}$ be the solutions of equation (1), ordered according to their eigenvalues, $\lambda_0 < \lambda_1 < \dots < \lambda_{k-1}$. Thus, the embedding is as follows:

$$\mathbf{x} \rightarrow \mathbf{y} = W^T \mathbf{x}$$

$$W = W_{PCA} W_{LPP}$$

$$W_{LPP} = [\mathbf{w}_0, \mathbf{w}_1, \dots, \mathbf{w}_{k-1}]$$

where \mathbf{y} is a k -dimensional vector. W is the transformation matrix. This linear mapping best preserves the manifold's estimated intrinsic geometry in a linear sense. The column vectors of W are the so called *Laplacian faces* which span the face subspace.

4.3 Face Manifold Analysis

Now consider a simple example of image variability, a set of face images are generated while the human face rotates slowly. Intuitively, the set of face images correspond to a continuous curve in image space, since there is only one degree of freedom, i.e. the angle of rotation. Thus, we can say that the set of face images are intrinsically one-dimensional. Actually, much recent work [9][10][11][15] has shown that the face images do reside on a low-dimensional submanifold embedded in high-dimensional image space. Therefore, an effective subspace learning algorithm should be able to detect the nonlinear manifold structure. The conventional algorithms, such as PCA and LDA, model the face images in Euclidean space. They effectively see only the Euclidean struc-




Figure 3. The first 10 Eigenfaces (first row), Fisherfaces (second row) and Laplacianfaces (third row) calculated from the face images in the YALE database.

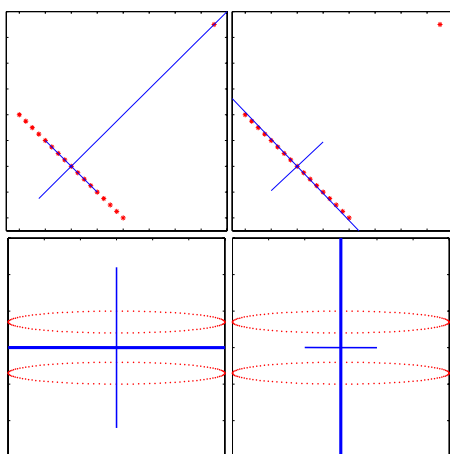


Figure 2. The left plots show the results of PCA. The right plots show the results of LPP. The first basis is shown as a longer line segment, and the second basis is shown as a shorter line segment. Clearly, LPP has more discriminating power than PCA, and is less sensitive to outliers.

ture; thus, they fail to detect the intrinsic low dimensionality.

With neighborhood preserving character, the LPP algorithm is capable of capturing the intrinsic manifold structure to a large extent. Figure 1 shows an example that the face images with various pose and expression of a person are mapped into a two-dimensional subspace. The face image data set used here is the same as that used in [9]. The representative face images are shown in the different parts of the space. As can be seen, the face images are divided into two parts. The left part includes the face images with open mouth, and the right part includes the face images with closed mouth. This is because that, by trying to preserve local structure in the embedding, LPP implicitly emphasizes the natural clusters in the data. Specifically, it makes the neighboring points in the image space nearer in the face space, and faraway points in the image space farther in the face space. Some theoretical analysis can be found in [2][4][12]. Moreover, we can see from the figure that the pose and expression of the faces

change continuously and smoothly. The bottom images correspond to points along the right path (linked by solid line), illustrating one particular mode of variability in pose. This observation tells us that LPP is capable of capturing the intrinsic face manifold structure.

5. Experimental Results

In this section, several experiments are carried out to show the effectiveness of our proposed Laplacianface method for face representation and recognition. We begin with two simple synthetic examples to compare LPP and PCA.

5.1 Simple Synthetic Examples

Two simple synthetic examples are given in Fig. 2. Both of the two data sets correspond to an essentially one-dimensional manifold. Projection of the data points onto the first basis would then correspond to a one-dimensional linear manifold representation. The second basis, shown as a shorter line segment in the figure, would be discarded in this low-dimensional example. As can be seen, PCA captures the direction of maximum variance in the data. LPP finds direction which preserves local structure and the discriminating power. Moreover, PCA is sensitive to outliers while LPP is not.

5.2 Face Representation Using Laplacianfaces

As we described previously, a face image can be represented as a point in image space. A typical image of size $m \times n$ describes a point in $m \times n$ -dimensional image space. However, due to the unwanted variations resulting from changes in lighting, facial expression, and pose, the image space might not be an optimal space for visual representation and recognition.

In section 2, we have discussed how to learn a locality preserving face subspace which is insensitive to outlier and noise. The images of faces in the training set are used to learn such a face subspace. The subspace is spanned by the Laplacianfaces as described in section 4.2. We can display the Laplacianfaces as a sort of feature images. Using the Yale face database as the training set, we present the first 10 Laplacianfaces in Figure 3, together with



Figure 4. The sample cropped face images of one individual from PIE database. The original face images are taken under varying pose, illumination, and expression.

eigenfaces and fisherfaces. Thus, a face image can be mapped into the locality preserving subspace spanned by the Laplacianfaces.

5.3 Face Recognition

Once the Laplacianfaces are created, face recognition [1][16][17] becomes a pattern classification task. In this section, we investigate the performance of our proposed Laplacianface method for face recognition. The system performance is compared with the eigenface method [16] and the fisherface method [1], two of the most popular methods in face recognition.

In this study, three face databases were tested. The first one is the Yale database [18], the second one is the PIE (pose, illumination, and expression) database from CMU [13], and the third one is the MSRA database collected by our own. In all the experiments, preprocessing to locate the faces was applied. Original images were normalized (in scale and orientation) such that the two eyes were aligned at the same position. Then, the facial areas were cropped into the final images for matching. The size of each cropped image in all the experiments is 32×32 pixels, with 256 grey levels per pixel. Thus, each image can be represented by a 1024-dimensional vector in image space. No further preprocessing is done. Figure 5 shows an example of the original face image and the cropped image. Different pattern classifiers have been applied for face recognition, including nearest-neighbor [16], Bayesian [7], and support vector machine [8], etc. In this paper, we apply nearest-neighbor classifier for its simplicity.

The recognition process has three steps. First, we calculate the Laplacianfaces from the training set of face images; then, the new face image to be identified is projected into the face subspace spanned by the Laplacianfaces; finally, the new face image is identified by a nearest-neighbor classifier.

For Yale and PIE database, a random subset of a fixed size is taken with labels to form the training set. The rest of the database is considered to be the testing set.

5.3.1 Yale Database

The Yale face database [18] is constructed at the Yale Center for Computational Vision and Control. It contains 165 grayscale images of 15 individuals. The images demonstrate variations in lighting condition (left-light, center-light, right-light), facial expression (normal, happy, sad,




Figure 5. The original face image and the cropped image.

sleepy, surprised, and wink), and with/without glasses. For each individual, 6 faces are used for training, and the rest 5 are used for testing.

The face subspace is constructed by our Laplacianfaces method to best preserve the local structure while reducing the dimensionality of the image space. For each face image, it can be projected into the face subspace by the transformation matrix W , i.e. Laplacianfaces.

The recognition results are shown in Table 1. It is found that the Laplacianface approach significantly outperforms both eigenface and fisherface approaches. The error rate is 11.3%, 20.0% and 25.3% for Laplacianface, fisherface, and eigenface methods, respectively. The corresponding face subspaces are called optimal face subspaces for each method. There is no significant improvement if more dimensions are used. Figure 7 shows a plot of error rate vs. dimensionality reduction. Note that, the upper bound of the dimensionality of fisherfaces is $c-1$ where c is the number of individuals.

5.3.2 PIE Database

The CMU PIE face database contains 68 subjects with 41,368 face images as a whole. The face images were captured by 13 synchronized cameras and 21 flashes, under varying pose, illumination and expression. We use 170 near frontal face images for each individual in our experiment, 85 for training and the other 85 for testing. Figure 4 shows some of the faces with pose, illumination and expression variations in the PIE database.

Table 2 shows the recognition results. As can be seen, Laplacianface method performed better than eigenface and fisherface methods. Figure 8 shows a plot of error rate vs. dimensionality reduction.

5.3.3. MSRA Database

This database was collected at the Microsoft Research Asia. It contains 12 individuals, captured in two different sessions with different backgrounds and illuminations. 64 to 80 face images are collected for each individual in each session. All the faces are frontal. Figure 6 shows the sample cropped face images from this database. In this test, one session is used for training and the other is used for testing. Table 3 shows the recognition results. Laplacianface approach has lower error rate (8.2%) than those of eigenface (35.4%) and fisherface (26.5%). Figure 9 shows a plot of error rate vs. dimensionality reduction.



Figure 6. The sample cropped face images of 8 individuals from MSRA database. The face images in the first row are taken in the first session, which are used for training. The face images in the second row are taken in the second session, which are used for testing. The two images in the same column are corresponding to the same individual.

5.4. Discussions

Three experiments have been systematically performed. These experiments reveal a number of interesting points:

1. All these three approaches performed better in the optimal face subspace than in the original image space.
2. In all the three experiments, the Laplacianface approach consistently performed better than the eigenface and fisherface approaches. Especially, it significantly outperformed the fisherface and eigenface approaches on Yale database and MSRA database.
3. Though Laplacianface does not explicitly consider the classification problem, it still outperforms fisherfaces, which is based on discriminant analysis. This is because that, Laplacianface approach encodes more discriminating information in the low-dimensional face subspace by preserving local structure which is more important than the global structure for classification, especially when nearest neighbor like classifiers are used. In fact, if there is a reason to believe that Euclidean distances ($\|x_i - x_j\|$) are meaningful only if they are small (local), then the LPP algorithm finds a projection that respects such a belief. Another reason is that, as we show in Fig. 1, the face images probably reside on a nonlinear manifold. Therefore, an efficient and effective subspace representation of face images should be capable of charactering the nonlinear manifold structure, while the Laplacianfaces are exactly derived by finding the optimal linear approximations to the eigenfunctions of the Laplace Beltrami operator on the face manifold [2][4]. By discovering the face manifold structure, our Laplacianface approach can identify the person with various pose, illumination and expression.
4. The Laplacianface approach appears to be the best at simultaneously handling variation in lighting, pose and expression.

6. Conclusion and Future Work

The manifold ways of face representation and recognition is introduced in this paper in order to detect the underlying nonlinear manifold structure in the manner of subspace learning. To the best of our knowledge, this is the first devoted work on face representation and recognition which explicitly considers manifold structure in a linear manner. The manifold structure is approximated by the nearest-neighbor graph computed from the data points. Using the notion of the Laplacian of the graph, we then compute a transformation matrix which maps the face images into the face subspace. We call this Laplacianfaces approach. The Laplacianfaces are obtained by finding the optimal linear approximations to the eigenfunctions of the Laplace Beltrami operator of the face manifold [2][4]. This linear transformation optimally preserves local manifold structure. Experimental results on the Yale database, CMU PIE database, and MSRA database show the effectiveness of our method.

One of the central problems in face manifold learning is to estimate the intrinsic dimensionality of the nonlinear manifold, or, degrees of freedom. Moreover, by using kernel methods, the linear projective maps can be easily extended to nonlinear maps, i.e. kernel Laplacianfaces which might be able to detect the nonlinear face manifold structure. We are currently exploring these problems in theory and practice.

Table 1. Performance comparison on the Yale database

Approach	Dims	Error Rate
Eigenfaces	33	25.3%
Fisherfaces	14	20.0%
Laplacianfaces	28	11.3%

Table 2. Performance comparison on the PIE database

Approach	Dims	Error Rate
Eigenfaces	150	20.6%
Fisherfaces	67	5.7%
Laplacianfaces	110	4.6%

Table 3. Performance comparison on MSRA database

Approach	Dims	Error Rate
Eigenfaces	142	35.4%
Fisherfaces	11	26.5%
Laplacianfaces	66	8.2%

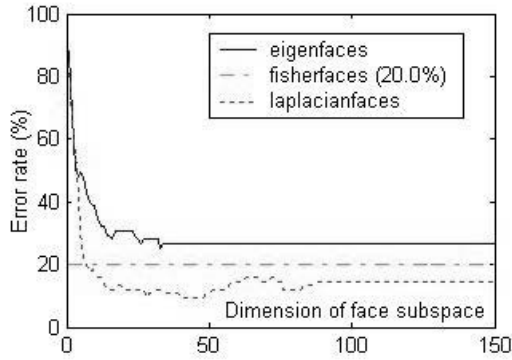


Figure 7. Error rate vs. dimensionality reduction on Yale database

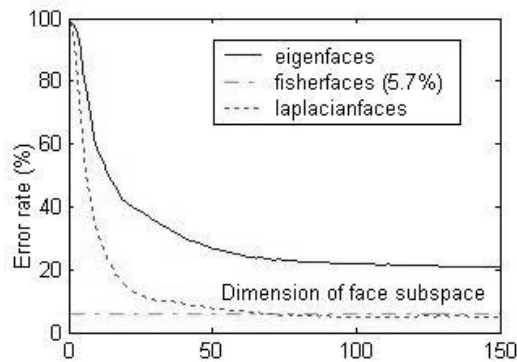


Figure 8. Error rate vs. dimensionality reduction on Pie database

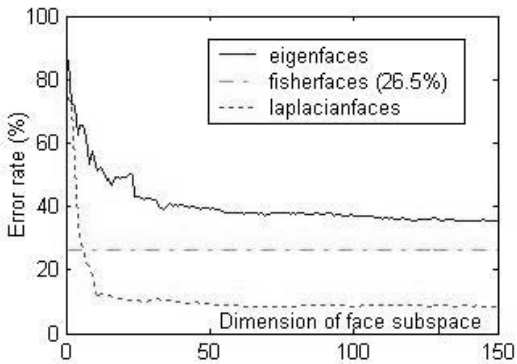


Figure 9. Error rate vs. dimensionality reduction on Our Own database

References

- [1] P. N. Belhumeur, J. P. Hespanha and D. J. Kriegman, "Eigenfaces vs. Fisherfaces: Recognition Using Class Specific Linear Projection", *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 19, No. 7, 1997, pp. 711-720.
- [2] M. Belkin and P. Niyogi, "Laplacian Eigenmaps and Spectral Techniques for Embedding and Clustering", *Advances in Neural Information Processing System 15*, Vancouver, British Columbia, Canada, 2001.
- [3] Fan R. K. Chung, Spectral Graph Theory, *Regional Conferences Series in Mathematics*, number 92, 1997.
- [4] Xiaofei He and Partha Niyogi, "Locality Preserving Projection", Technical Report TR-2002-09, Department of Computer Science, the University of Chicago.
- [5] Q. Liu, R. Huang, H. Lu, and S. Ma, "Face Recognition Using Kernel Based Fisher Discriminant Analysis", in *Proc. of the fifth Int. Conf. on Automatic Face and Gesture Recognition*, Washington D. C., May 2002.
- [6] A. M. Martinez and A. C. Kak, "PCA versus LDA", *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 23(2):228-233, 2001.
- [7] B. Moghaddam and A. Pentland, "Probabilistic Visual Learning for Object Representation", *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol 19, pp. 696-710, 1997.
- [8] P. J. Phillips, "Support Vector Machines Applied to Face Recognition". In *Advances in Neural Information Processing Systems 11*, pp. 803-809, 1998.
- [9] Sam T. Roweis, and Lawrence K. Saul, "Nonlinear Dimensionality Reduction by Locally Linear Embedding", *Science*, vol 290, 22 December 2000.
- [10] Sam Roweis, Lawrence Saul and Geoff Hinton, "Global Coordination of Local Linear Models", *Advances in Neural Information Processing System 14*, 2001.
- [11] H. Sebastian Seung and Daniel D. Lee, "The Manifold Ways of Perception", *Science*, vol 290, 22 December 2000.
- [12] J. Shi and J. Malik, "Normalized Cuts and Image Segmentation", *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 22 (2000), 888-905.
- [13] T. Sim, S. Baker, and M. Bsat, "The CMU Pose, Illumination, and Expression (PIE) Database", in *Proceedings of the IEEE International Conference on Automatic Face and Gesture Recognition*, May, 2002.
- [14] L. Sirovich and M. Kirby, "Low-Dimensional Procedure for the Characterization of Human Faces", *Journal of the Optical Society of America A*, vol. 4, pp. 519-524, 1987.
- [15] Joshua B. Tenenbaum, Vin de Silva, and John C. Langford, "A Global Geometric Framework for Nonlinear Dimensionality Reduction", *Science*, vol 290, 22 December 2000.
- [16] M. Turk and Pentland, "Face Recognition Using Eigenfaces", in *Proc. IEEE International Conference on Computer Vision and Pattern Recognition*, Maui, Hawaii, 1991.
- [17] L. Wiskott, J.M. Fellous, N. Kruger, and C.v.d. Malsburg, "Face Recognition by Elastic Bunch Graph Matching", *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 19:775-779, 1997.
- [18] Yale Univ. Face Database, <http://cvc.yale.edu/projects/yalefaces/yalefaces.html>, 2002.
- [19] Ming-Hsuan Yang, "Kernel Eigenfaces vs. Kernel Fisherfaces: Face Recognition Using Kernel Methods", in *Proc. of the fifth Int. Conf. on Automatic Face and Gesture Recognition*, Washington D. C., May 2002.
- [20] J. Yang, Y. Yu, and W. Kunz, "An Efficient LDA Algorithm for Face Recognition", *The sixth International Conference on Control, Automation, Robotics and Vision*, Singapore, 2000.
- [21] W. Zhao, R. Chellappa, and P. J. Phillips, "Subspace Linear Discriminant Analysis for Face Recognition", *Technical Report CAR-TR-914*, Center for Automation Research, University of Maryland, 1999.