



Face Recognition

University at Buffalo CSE666 Lecture Slides

Resources:

<http://www.face-rec.org/algorithms/>

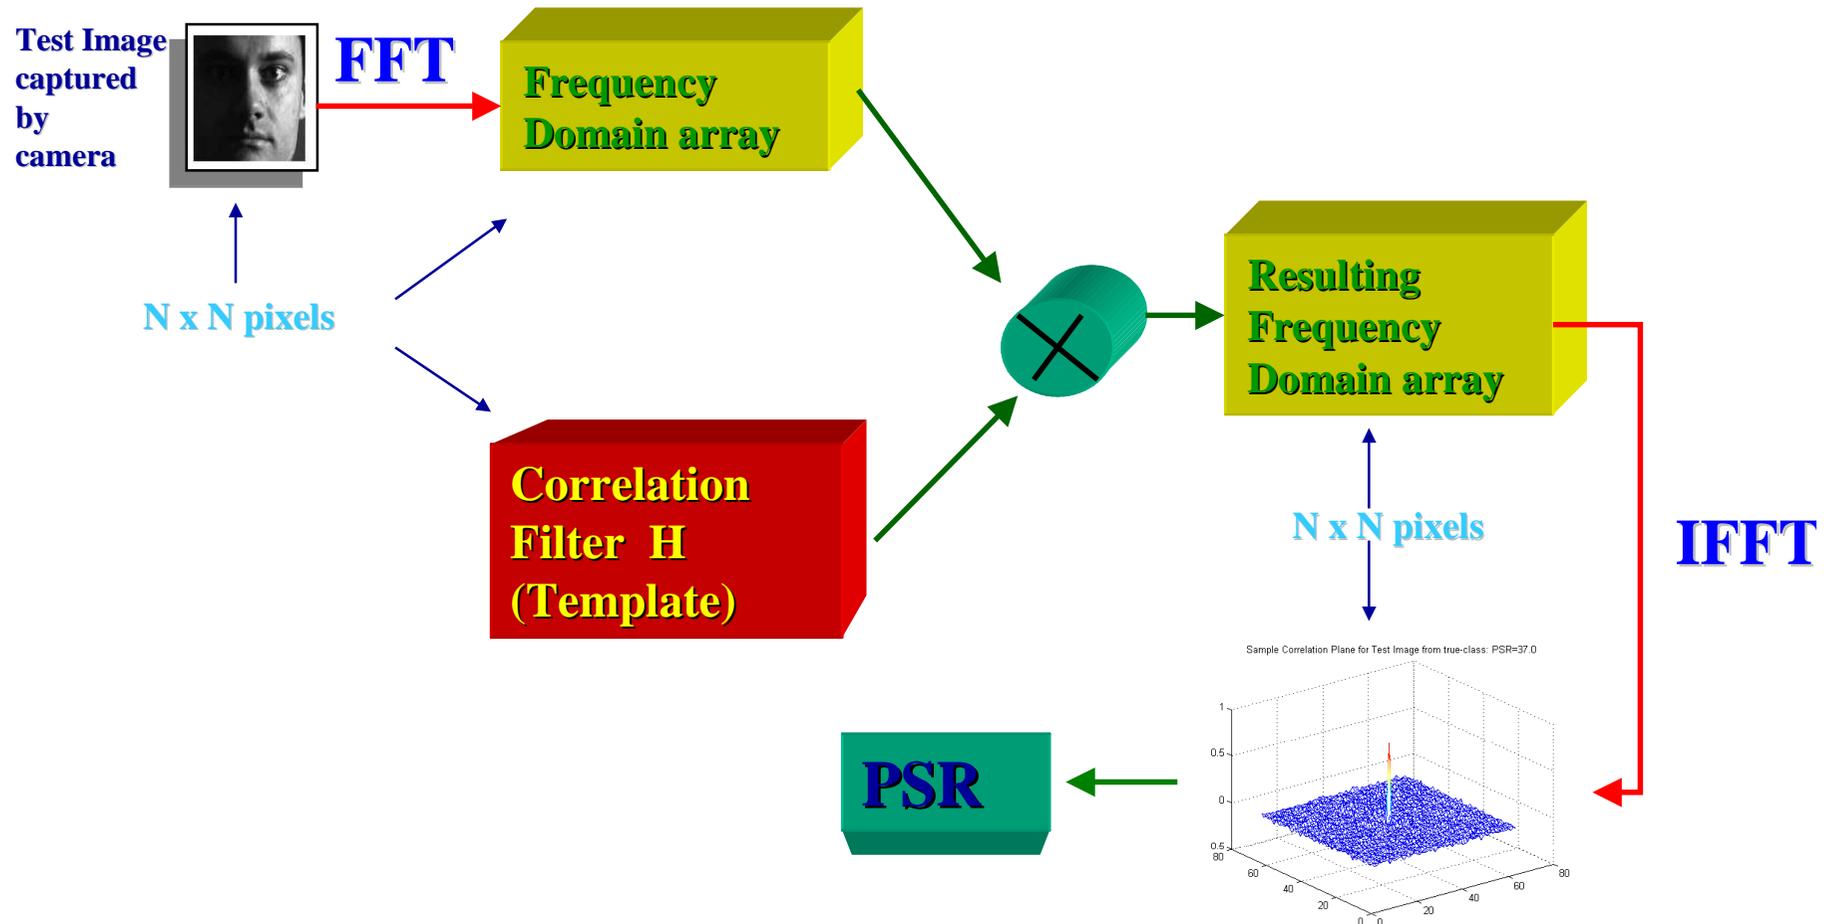


Overview of face recognition algorithms

- Correlation
 - Pixel based correspondence between two face images
- Structural
 - Based on identifying landmark points
- Linear models
 - PCA (Principal Component Analysis)
 - LDA (Linear Discriminant Analysis)
 - ICA (Independent Component Analysis)
 - Combinations of above
- Non-linear models
 - Kernel mapping (and using PCA, LDA, ICA)
 - Active shape/appearance models
 - Manifold mapping
- 2D Model vs. 3D Model
- Matching vs. Classification



Correlation Matching



*B.V.K. Vijaya Kumar, Marios Savvides, C. Xie, K. Venkataramani, J. Thornton and A. Mahalanobis, "Biometric Verification using Correlation Filters", Applied Optics, 2003

*B.V.K. Vijaya Kumar, M. Savvides, K. Venkataramani, C. Xie, "Spatial frequency domain image processing for biometric recognition," IEEE Proc. of International Conference on Image Processing (ICIP), Vol. I, 53-56, 2002



Elastic Bunch Graph Matching

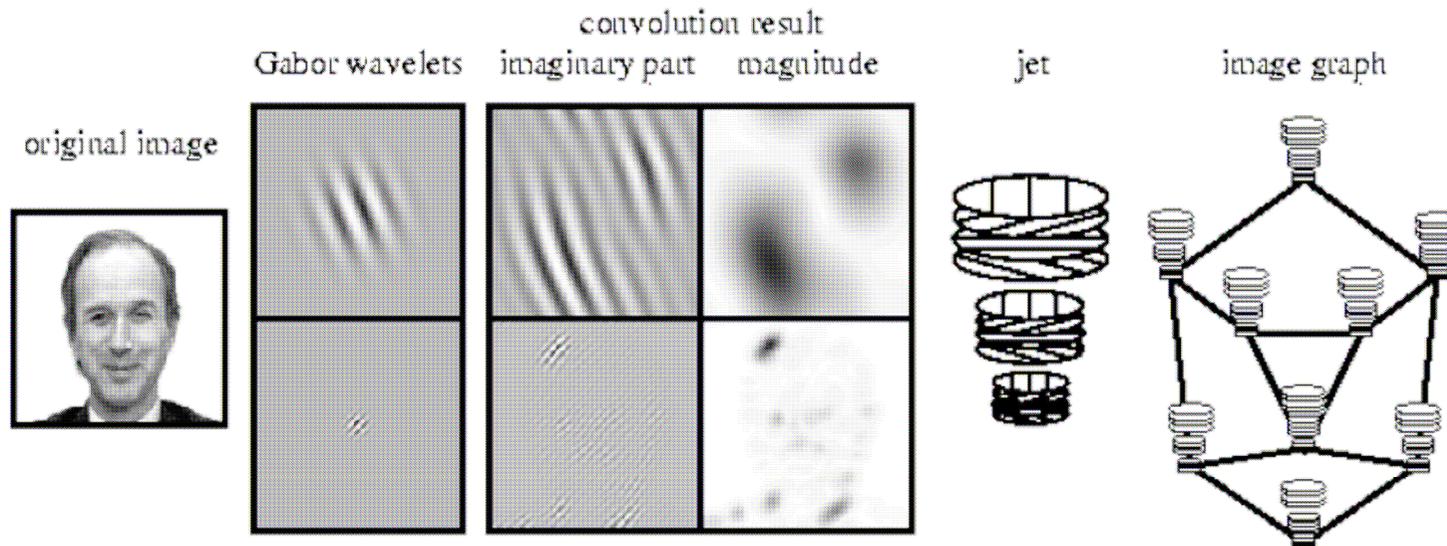


Figure 1: The graph representation of a face is based on the Gabor wavelet transform, a convolution with a set of wavelet kernels. These have the shape of plane waves restricted by a Gaussian envelope function. We compute 40 coefficients (5 frequencies \times 8 orientations). Coefficient phase varies approximately with wavelet frequency (see imaginary part), magnitude varies slowly. The set of 40 coefficients obtained for one image point is referred to as a jet (for clarity, only 3 frequencies and 4 orientations are represented in the figure). A sparse collection of such jets together with some information about their relative location constitutes an image graph, used to represent an object, such as a face.



Elastic Bunch Graph Matching

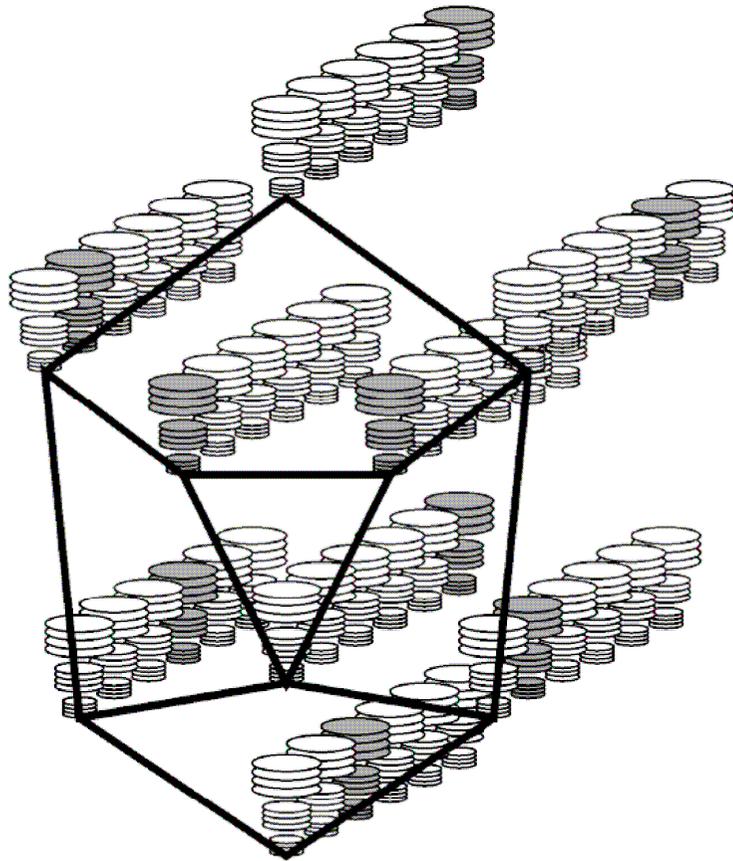


Fig. 3. The Face Bunch Graph (FBG) serves as a general representation of faces. Each stack of discs represents a jet. From a bunch of jets attached to a single node only the best fitting one is selected for a match, indicated by gray shading.

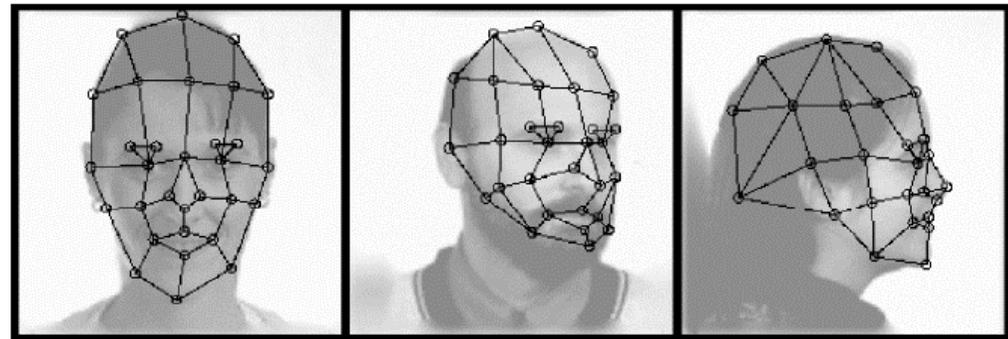


Fig. 2. Object-adapted grids for different poses. The nodes are positioned automatically by elastic bunch graph matching. (The grids used in Section 3 for the FERET database had about 14 additional nodes which are not shown here for simplicity.) One can see that, in general, the matching finds the fiducial points quite accurately. But mispositioning occurred, for example, for the face in the center. The chin was not found accurately; the leftmost node and the node below should be at the top and the bottom of the ear, respectively.

- Heuristic matching algorithm



Principal Component Analysis (PCA)

Previous discussion:

PCA as a feature extraction method

\mathbf{a}_i - eigenvectors of matrix $\mathbf{R}_x = E[\mathbf{x}\mathbf{x}^T]$

$\mathbf{A} = (\mathbf{a}_i)$ - Karhunen-Loeve transform

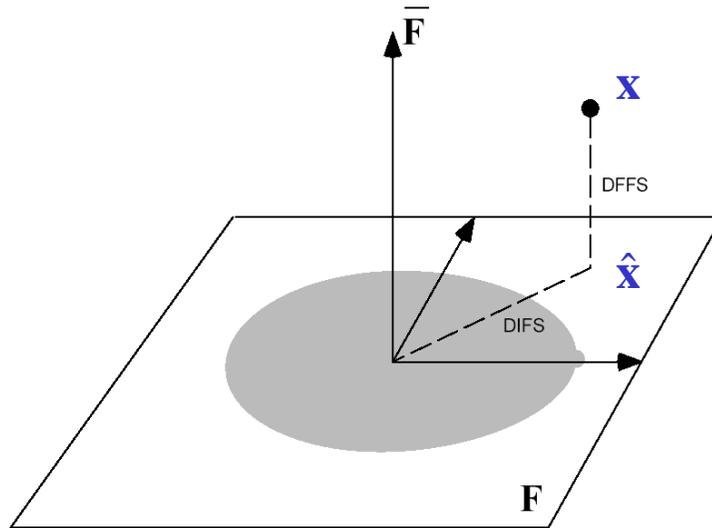
Features	Representation
$\mathbf{y} = \mathbf{A}\mathbf{x}$ - projection of original vector \mathbf{x} to KL basis	$\mathbf{x} = \sum_{i=1}^n y_i \mathbf{a}_i$ - representation of original vector \mathbf{x} in KL basis vectors
$\mathbf{y}_m = \mathbf{A}_m \mathbf{x}$ - m first PCA coefficients - 'best' features	$\hat{\mathbf{x}} = \sum_{i=1}^m y_i \mathbf{a}_i$ - the projection on subspace spanned by eigenvectors with largest eigenvalues

$$E[(\mathbf{x} - \hat{\mathbf{x}})^2] = \sum_{i=m+1}^n \lambda_i$$

- sum of the smallest $n - m$ eigenvalues



PCA as appearance model



$\mathbf{y}_m = \mathbf{A}_m \mathbf{x}$ - m first PCA coefficients -

'best' features - should provide a good approximation to all the samples (for which PCA was trained)

$$\hat{\mathbf{x}} = \mathbf{A}_m^T \mathbf{y}_m$$

• Distance from feature space (DFSS) should be small:

$$DFSS = |\mathbf{x} - \hat{\mathbf{x}}|$$

- PCA provides a model on typical faces
- DFSS can be used it to separate faces from non-faces

$$DIFS = |\hat{\mathbf{x}}_1 - \hat{\mathbf{x}}_2|$$

- DIFS (distance in feature space) can be used to match two faces



PCA for face recognition

- Collect a set of face images for training
- Train PCA
- During face recognition see if the distance between test and enrolled PCA feature vectors is less than threshold:

$$|\hat{\mathbf{x}}_1 - \hat{\mathbf{x}}_2| = |\mathbf{y}_1 - \mathbf{y}_2| < \theta$$

$\hat{\mathbf{x}}$ - reconstructed face (hyperplane in the original feature space)
 \mathbf{y} - projection into PCA feature space





PCA – subspace on training samples

If \mathbf{x} has dimension N^2 , then $\mathbf{R}_x = E[\mathbf{x}\mathbf{x}^T]$ has dimension $N^2 \times N^2$

Suppose we have M training samples $\mathbf{X} = [\mathbf{x}_1 \ \mathbf{x}_2 \ \dots \ \mathbf{x}_M]$

$$\mathbf{R}_x = \mathbf{X}\mathbf{X}^T \quad \text{matrix of dimension } N^2 \times N^2$$

Instead consider $\mathbf{L} = \mathbf{X}^T\mathbf{X}$ of dimension $M \times M$

If \mathbf{v} is an eigenvector of \mathbf{L} then $\mathbf{X}\mathbf{v}$ is an eigenvector of \mathbf{R}_x :

$$\mathbf{X}^T\mathbf{X}\mathbf{v} = \lambda\mathbf{v} \Rightarrow \mathbf{X}\mathbf{X}^T\mathbf{X}\mathbf{v} = \lambda\mathbf{X}\mathbf{v}$$

- The number of training samples is usually less than the dimension of image vectors $M < N^2$, so this procedure makes sense



Linear Discriminant Analysis

Covariance matrix for class i :

$$\mathbf{S}_i = E[(\mathbf{x} - \boldsymbol{\mu}_i)(\mathbf{x} - \boldsymbol{\mu}_i)^T]$$

Within-class scatter matrix: $\mathbf{S}_w = \sum_{i=1}^M P_i \mathbf{S}_i$

Between-class scatter matrix: $\mathbf{S}_b = \sum_{i=1}^M P_i (\boldsymbol{\mu}_i - \boldsymbol{\mu}_0)(\boldsymbol{\mu}_i - \boldsymbol{\mu}_0)^T$

where $\boldsymbol{\mu}_0 = \sum_{i=1}^M P_i \boldsymbol{\mu}_i$

Optimization criteria: maximize $\frac{\det(\mathbf{S}_b)}{\det(\mathbf{S}_w)} = \det(\mathbf{S}_w^{-1} \mathbf{S}_b)$



Linear Discriminant Analysis

Solution: projection is determined by the eigenvectors of $\mathbf{S}_w^{-1}\mathbf{S}_b$

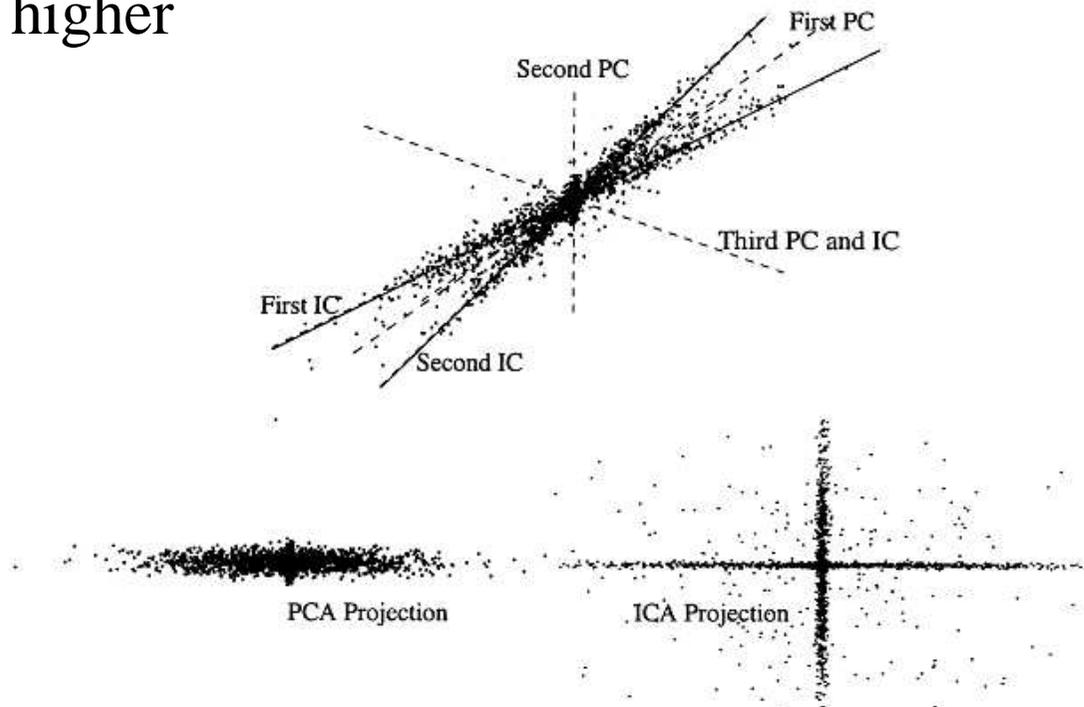
- Usually has better performance than PCA
- Requires samples of the same class (same person face) to train \mathbf{S}_w
- Need to make sure that \mathbf{S}_w is non-singular

Solution: apply PCA first to reduce the number of dimensions, then perform LDA.



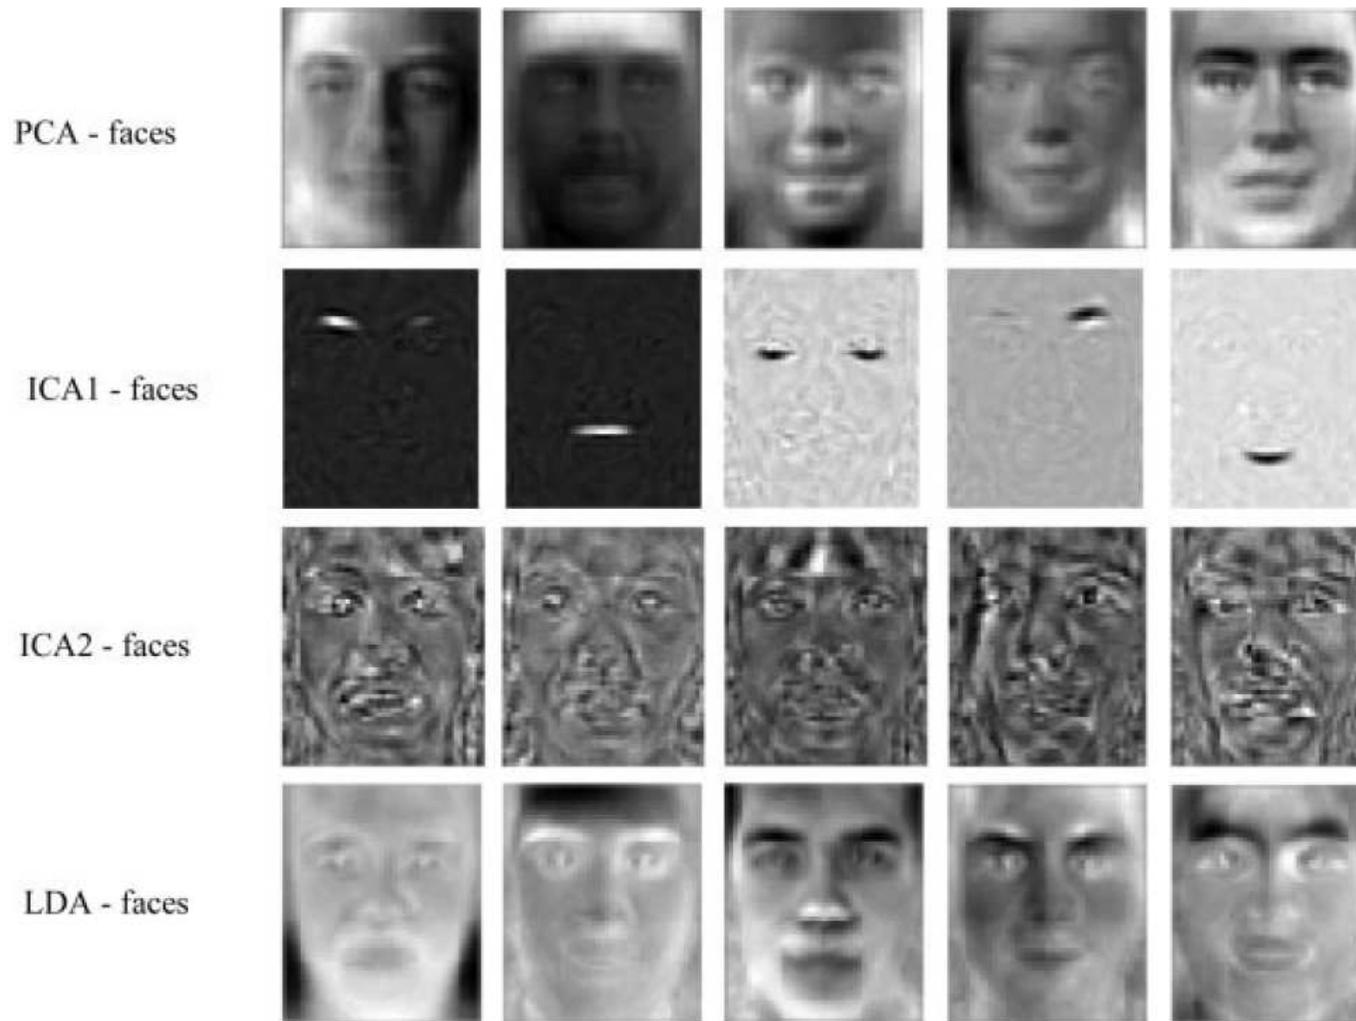
Independent Component Analysis

- Whereas PCA makes uncorrelated features, there might be still dependent
- ICA tries to reduce higher order dependence
- Search for proper projection is more difficult
- Use approximation approaches





Comparisons of subspace methods





Comparisons of subspace methods

Table I. Performance across four projection methods and four metrics.

Projection	Results at Rank 1 (%)			
	Metric			
	L1	L2	MAH	COS
				<i>Fb</i>
PCA	82.26	82.18	64.94	81.00
ICA1	81.00	81.51	64.94	80.92
ICA2	64.94	74.31	64.94	83.85
LDA	78.08	82.76	70.88	81.51
				<i>Fc</i>
PCA	55.67	25.26	32.99	18.56
ICA1	18.04	17.53	32.99	12.89
ICA2	15.98	44.85	32.99	64.95
LDA	26.80	26.80	41.24	20.62
				<i>Dup1</i>
PCA	36.29	33.52	25.62	33.52
ICA1	32.55	31.86	25.62	32.27
ICA2	28.81	31.99	25.62	42.66
LDA	34.76	32.96	27.70	33.38
				<i>Dup2</i>
PCA	17.09	10.68	14.53	11.11
ICA1	8.97	7.69	14.53	8.97
ICA2	16.24	19.66	14.53	28.21
LDA	16.24	10.26	16.67	10.68



Kernel Methods

- Kernel mapping:
 - map original image feature vectors into higher dimensional space using some kernel functions:

$$\Phi : R^N \rightarrow R^f, \quad f \gg N$$

- Covariance matrix in kernel space has elements

$$\Phi(\mathbf{x}_i) \cdot \Phi(\mathbf{x}_j)$$

- Kernel trick: there exists a function

$$k(\mathbf{x}_i, \mathbf{x}_j) = \Phi(\mathbf{x}_i) \cdot \Phi(\mathbf{x}_j) = K_{ij}$$

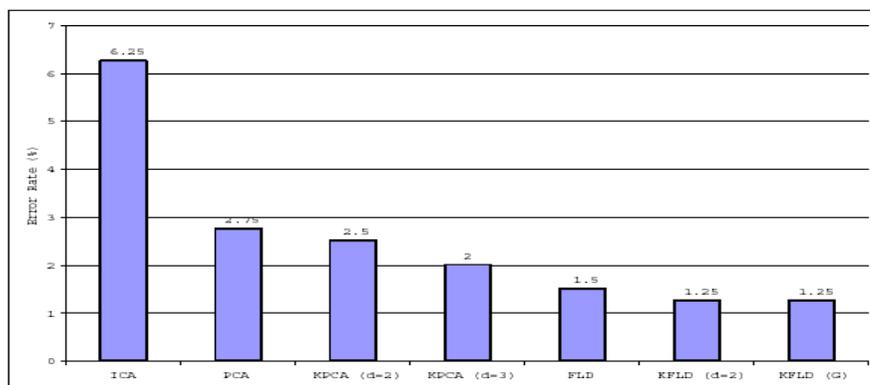
- No need to explicitly calculate $\Phi(\mathbf{x}_i)$ - use some chosen kernel function:

$$k(\mathbf{x}_i, \mathbf{x}_j) = e^{-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|^2}{2\sigma^2}}$$

- The number of eigenvectors is still limited by the number of training samples as in regular PCA

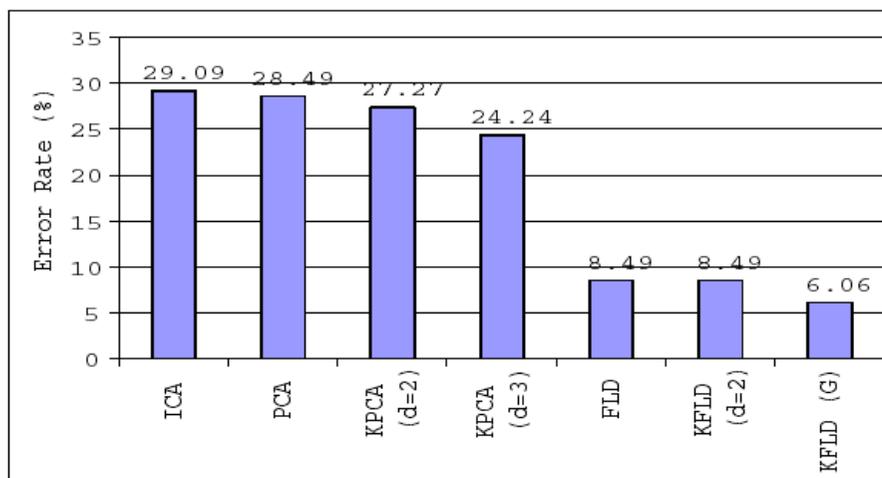


Kernel Methods



Method	Reduced Space	Error Rate (%)
ICA	40	6.25 (25/400)
Eigenface	30	2.75 (11/400)
Fisherface	14	1.50 (6/400)
Kernel Eigenface, d=2	50	2.50 (10/400)
Kernel Eigenface, d=3	50	2.00 (8/400)
Kernel Fisherface (P)	14	1.25 (5/400)
Kernel Fisherface (G)	14	1.25 (5/400)

Figure 2: Experimental results on AT&T database.



Method	Reduced Space	Error Rate (%)
ICA	30	29.09 (48/165)
Eigenface	30	28.48 (47/165)
Fisherface	14	8.48 (14/165)
Kernel Eigenface, d=2	80	27.27 (45/165)
Kernel Eigenface, d=3	60	24.24 (40/165)
Kernel Fisherface (P)	14	6.67 (11/165)
Kernel Fisherface (G)	14	6.06 (10/165)

Figure 3: Experimental results on Yale database.



Local Linear Embedding

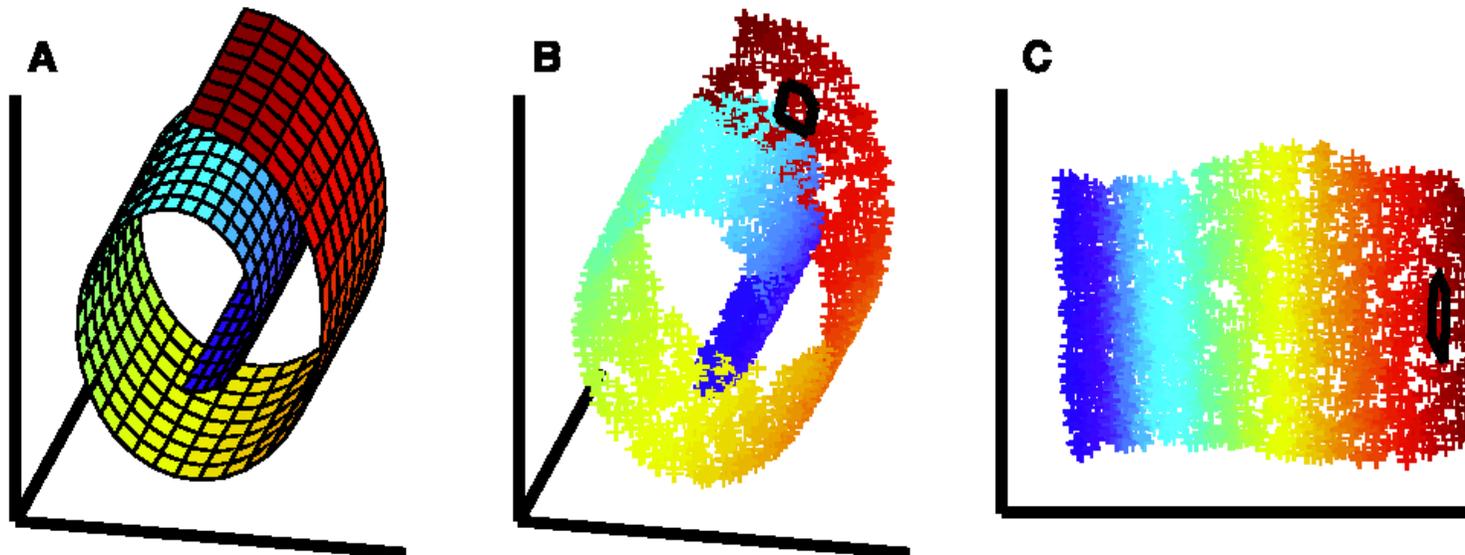


Fig. 1. The problem of nonlinear dimensionality reduction, as illustrated (10) for three-dimensional data (B) sampled from a two-dimensional manifold (A). An unsupervised learning algorithm must discover the global internal coordinates of the manifold without signals that explicitly indicate how the data should be embedded in two dimensions. The color coding illustrates the neighborhood-preserving mapping discovered by LLE; black outlines in (B) and (C) show the neighborhood of a single point. Unlike LLE, projections of the data by principal component analysis (PCA) (28) or classical MDS (2) map faraway data points to nearby points in the plane, failing to identify the underlying structure of the manifold. Note that mixture models for local dimensionality reduction (29), which cluster the data and perform PCA within each cluster, do not address the problem considered here: namely, how to map high-dimensional data into a single global coordinate system of lower dimensionality.



Local Linear Embedding

- Step 1: find optimal local linear representation of point \vec{X}_i as a combination of its neighbors \vec{X}_j

$$\mathcal{E}(W) = \sum_i \left| \vec{X}_i - \sum_j W_{ij} \vec{X}_j \right| \rightarrow \min$$

$$\sum_j W_{ij} = 1$$

$$W_{ij} = 0 \quad \text{if } j \text{ is not a neighbor of } i$$

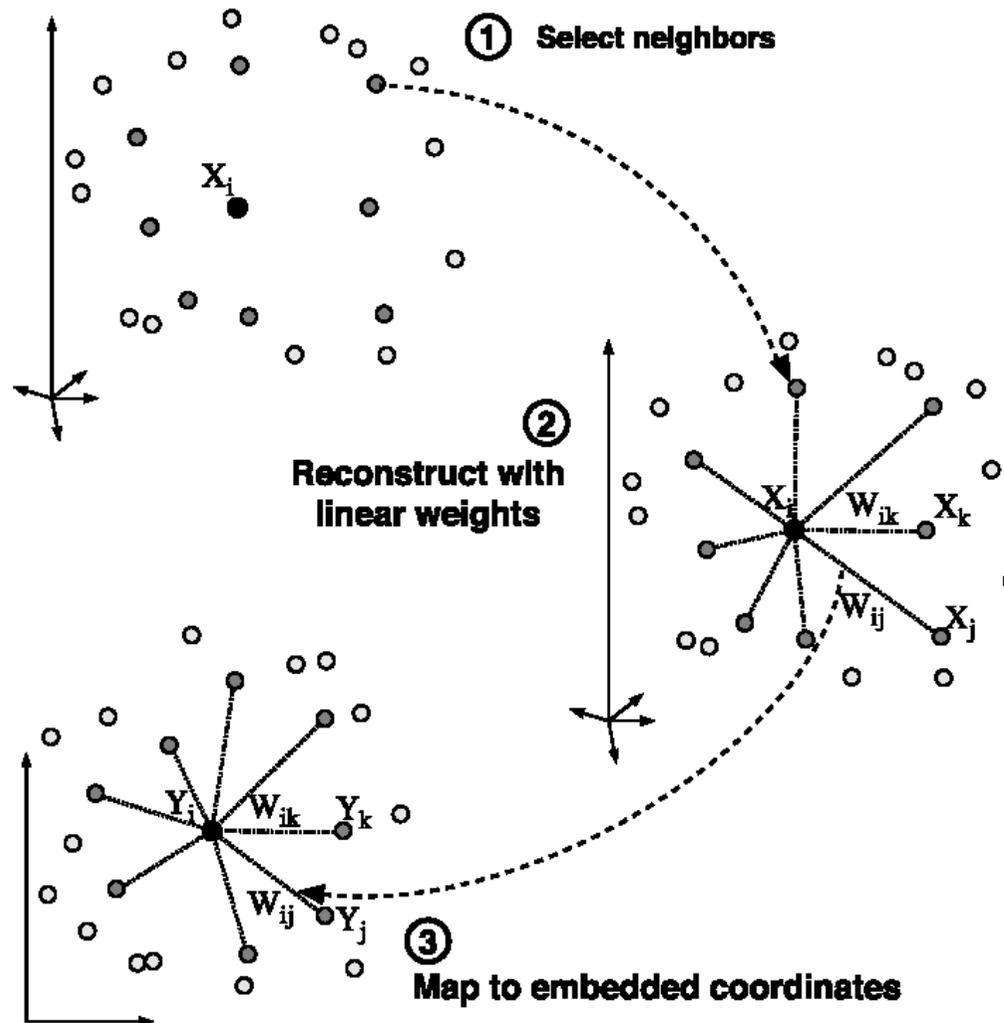
- Step 2: find lower dimensional representation \vec{Y}_i of \vec{X}_i

$$\Phi(Y) = \sum_i \left| \vec{Y}_i - \sum_j W_{ij} \vec{Y}_j \right| \rightarrow \min$$



Local Linear Embedding

Fig. 2. Steps of locally linear embedding: (1) Assign neighbors to each data point \bar{X}_i (for example by using the K nearest neighbors). (2) Compute the weights W_{ij} that best linearly reconstruct \bar{X}_i from its neighbors, solving the constrained least-squares problem in Eq. 1. (3) Compute the low-dimensional embedding vectors \bar{Y}_i best reconstructed by W_{ij} , minimizing Eq. 2 by finding the smallest eigenmodes of the sparse symmetric matrix in Eq. 3. Although the weights W_{ij} and vectors Y_i are computed by methods in linear algebra, the constraint that points are only reconstructed from neighbors can result in highly nonlinear embeddings.





Local Linear Embedding

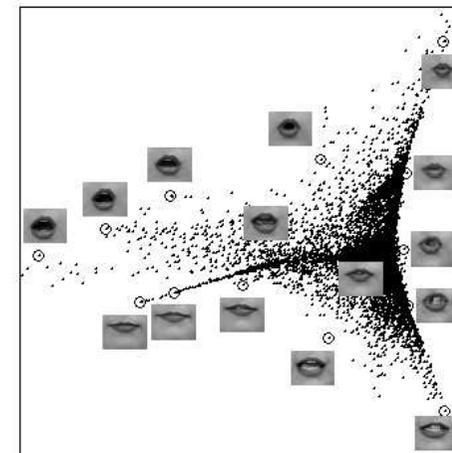
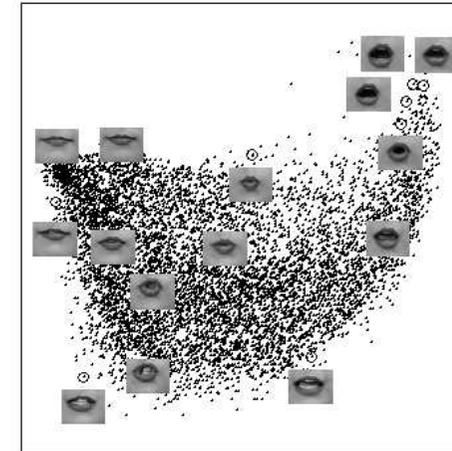
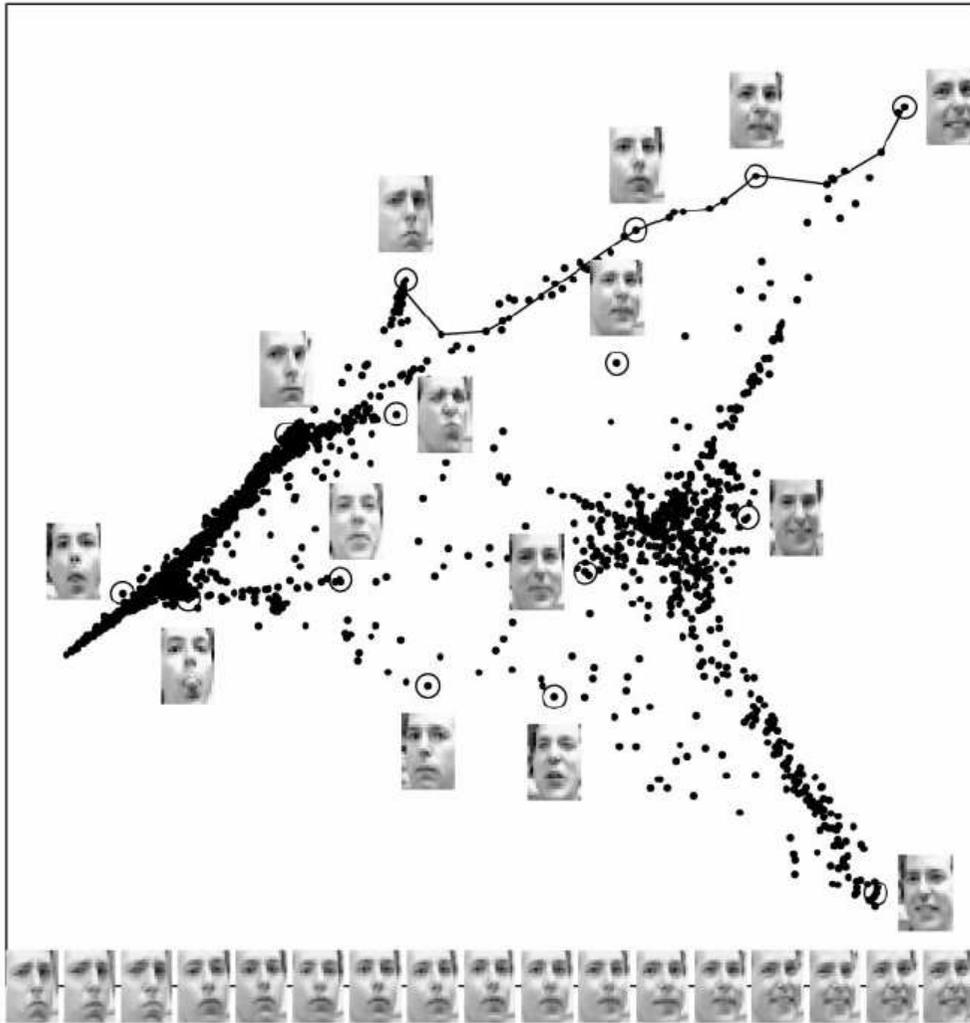


Fig. 3. Images of faces (11) mapped into the embedding space described by the first two coordinates of LLE. Representative faces are shown next to circled points in different parts of the space. The bottom images correspond to points along the top-right path (linked by solid line), illustrating one particular mode of variability in pose and expression.

Figure 4: Images of lips mapped into the embedding space described by the first two coordinates of PCA (top) and LLE (bottom). Representative lips are shown next to circled points in different parts of each space. The differences between the two embeddings indicate the presence of nonlinear structure in the data.



3D Model Based Recognition

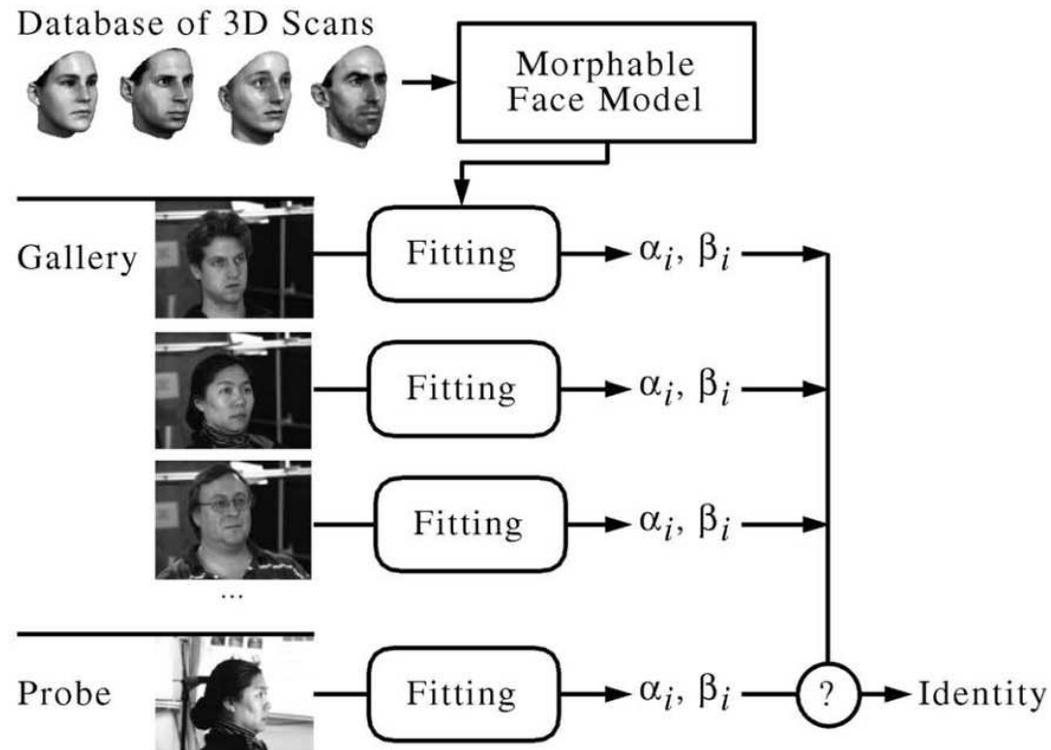


Fig. 1. Derived from a database of laser scans, the 3D morphable face model is used to encode gallery and probe images. For identification, the model coefficients α_i, β_i of the probe image are compared with the stored coefficients of all gallery images.



3D Model Based Recognition

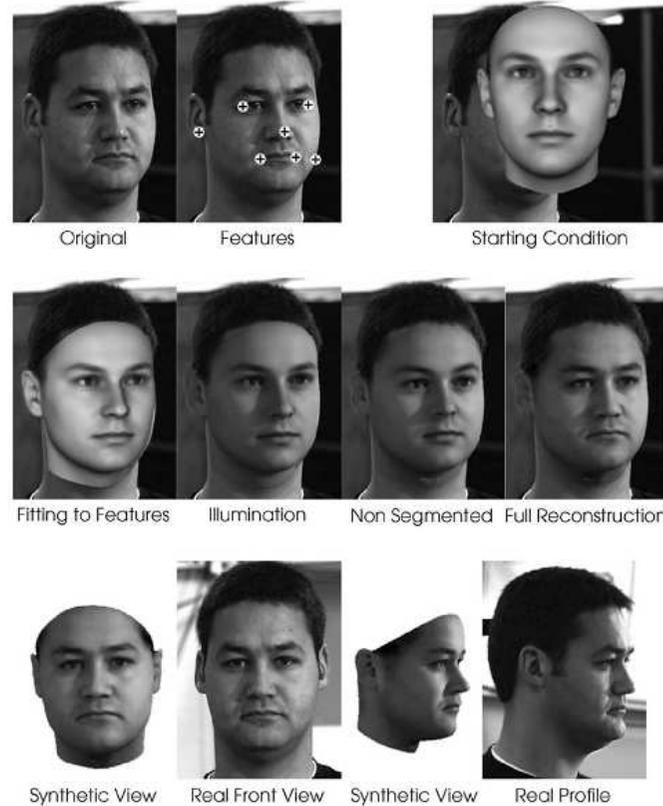


Fig. 6. Face reconstruction from a single image (top, left) and a set of feature points (top, center): Starting from standard pose and illumination (top, right), the algorithm computes a rigid transformation and a slight deformation to fit the features. Subsequently, illumination is estimated. Shape, texture, transformation, and illumination are then optimized for the entire face and refined for each segment (second row). From the reconstructed face, novel views can be generated (bottom row).



3D Model Based Recognition



(a)



(b)



(c)

Fig. 8. Reconstructions of 3D shape and texture from FERET images (top row). In the second row, results are rendered into the original images with pose and illumination recovered by the algorithm. The third row shows novel views.