Machine Vision

Lecture 16

Face Recognition

Applications of face recognition

Table 1 Applications of Face Recognition Technology

Applications	Advantages	Disadvantages
1a. Credit Card, Driver's License, Passport, and Personal Identification	Controlled image Controlled segmentation Good quality images	No existing database Large potential database Rare search type
1b. Mug shots Matching	Mixed image quality More than one image available	
2. Bank/Store Security	High value Geographically localized search	Uncontrolled segmentation Low image quantity
3. Crowd Surveillance	High value Small file size Availability of video images	Uncontrolled segmentation Low image quality Real-time
4. Expert Identification	High value Enhancement possible	Low image quality Legal certainty required
5. Witness Face Reconstruction	Witness search limits	Unknown similarity
6. Electronic Mug Shots Book	Descriptor search limits	Viewer fatigue
7. Electronic Lineup	Descriptor search limits	Viewer fatigue
8. Reconstruction of Face from Remains	High value	Requires physiological input
9. Computerized Aging	High value	Requires example input

Main tasks in these applications

- Matching
  - 1, 2, 3
- Similarity detection
  - 4-7
- Transformation
  - 8, 9

Issues of FR by human

- Is face recognition a dedicated process?
- Is face perception the result of holistic or feature analysis
- Ranking of significance of facial features
- Caricatures
- Distinctiveness
- The role of spatial frequency analysis
- The role of the brain
- Face recognition by children
- Facial expression
- Role of race/gender
- Image quality

2D Face Recognition Approaches

- Neural networks
  - Back propagation techniques
  - Better for detection and localisation than identification
- Feature analysis
  - Localisation of features
  - Distance between features
  - Feature characteristics
- Graph matching
  - Construct a graph around the face
  - Possible need for feature localisation
  - Can include other data (colour, texture)
- Eigenface
  - Information Theory approach
  - Identify discriminating components
- Fisherface
  - Uses ‘within-class’ information to maximise class separation

Neural Network Based Face Detection

Henry A. Rowley, Shumeet Baluja, Takeo Kanade

– **CMU, Pittsburgh**

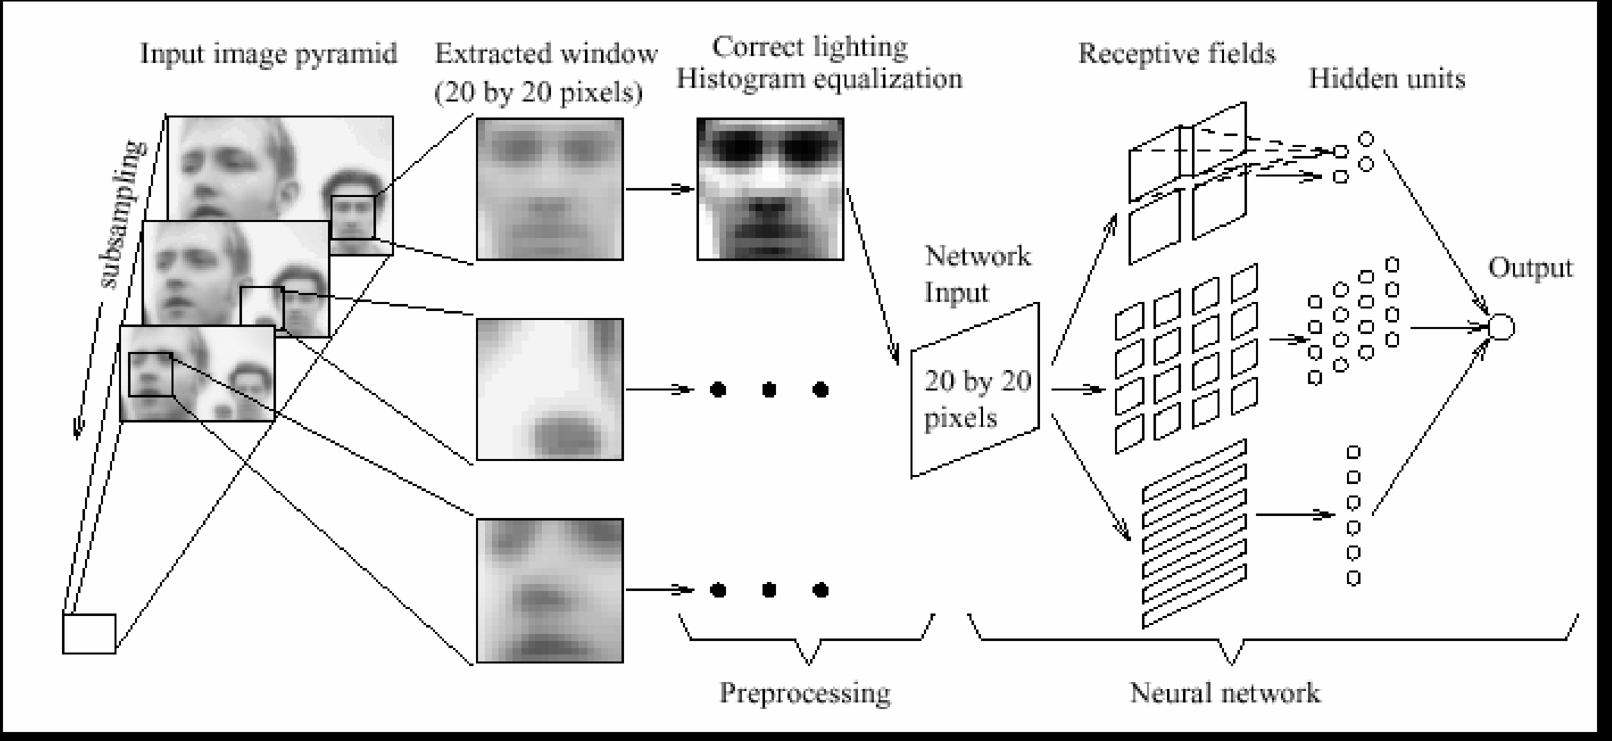

- Large training set of faces and small set of non-faces
- Training set of non-faces automatically built up:
  - Set of images with no faces
  - Every 'face' detected is added to the non-face training set.

Extraction of Facial Features for Recognition Using Neural Networks

– Nathan Intrator, Daniel Reisfeld, Yehezkel Yeshurun

– Tel-Aviv University

- Assigns a symmetry magnitude to each pixel, to create a symmetry map(right)
- Applying geometric constrains, locates regions of interest.
- Several neural networks are trained using various back-propagation methods.
- The ensemble network results are used to classify features.

Face Recognition through Geometric Features

R. Brunelli, Istituto per la Ricerca Scientifica e Technologica
T. Poggio, MIT

- Uses vertical and horizontal integral projections of edge maps.
- The nose is found by searching for peaks in the vertical projection.
- 22 Geometrical features used.
- Recognition performed by nearest neighbour.
- Only useful for small databases, or preliminary step.

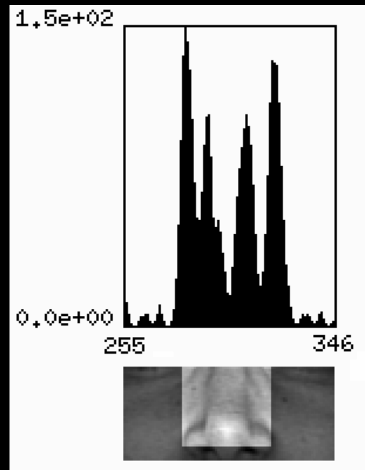

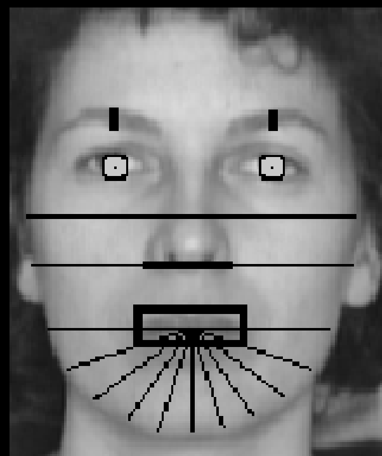

Face Recognition by Elastic Bunch Graph Matching

– L. Wiskott, N. Kruger, C. Malsburg Ruhr-University, Germany
– J. Fellous, University of Southern California, USA

•Uses a Gabor wavlet transform on images of faces.

•A face graph is a sparse collection of jets:

A set of (40) Gabor kernel coefficients for a single point in an image.

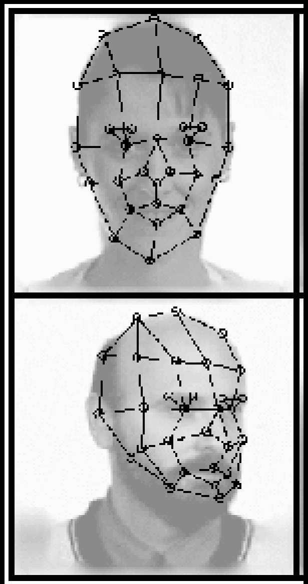

•A face bunch graph is a combination of various face graphs (A set of jets at each node – called a bunch).

•A graph is created for a specific face by selecting the best matching jets from each bunch.

•Recognition is performed by comparing graph similarity.

The Eigenface Method

- *Eigenfaces for Recognition*
  - Matthew Turk, Alex Pentland
    - MIT
- *Face Recognition Using Eigenfaces*
  - Matthew Turk, Alex Pentland
    - MIT

- Use PCA to determine the most discriminating features between images of faces.
- Create an image subspace (face space) which best discriminates between faces.
- Like faces occupy near points in face space.
- Compare two faces by projecting the images into face space and measuring the distance between them.

Image space

- Similarly the following 1x2 pixel images are converted into the vectors shown.

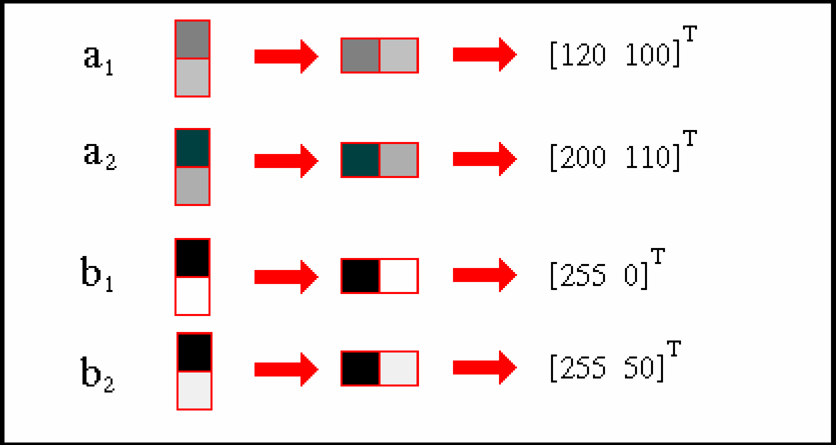

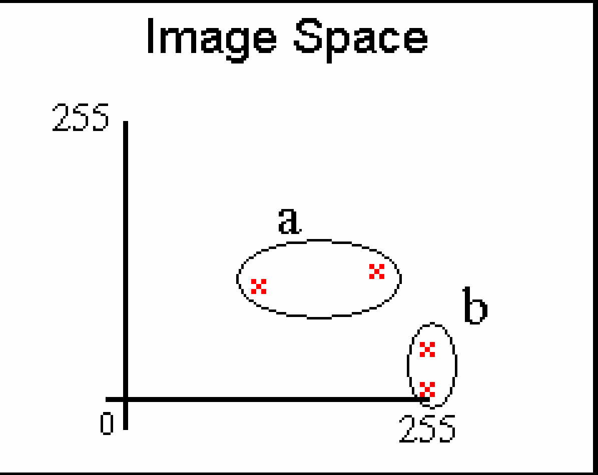

- Each image occupies a different point in image space.
- Similar images are near each other in image space.
- Different images are far from each other in image space.

Applying the same principal to faces

- A 256x256 pixel image of a face occupies a single point in 65,536-dimensional image space.
- Images of faces occupy a small region of this large image space.
- Similarly, different faces should occupy different areas of this smaller region.
- We can identify a face by finding the nearest 'known' face in image space.

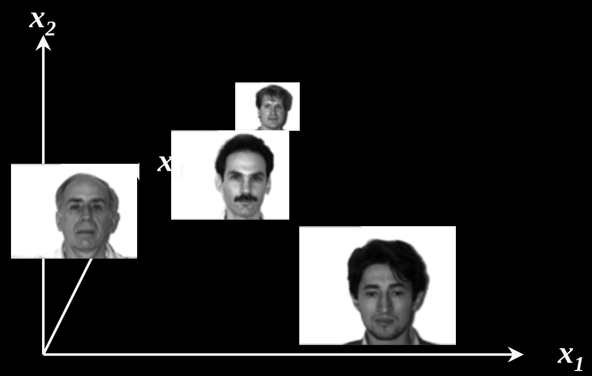

However, even tiny changes in lighting, expression or head orientation cause the location in image space to change dramatically. Plus, large amounts of storage is required.

PCA – Principal Component Analysis

- Principal component analysis is used to calculate the vectors which best represent this small region of image space.
- There are the eigenvectors of the covariance matrix for the training set.
- The eigenvectors are used to define the subspace of face images, known as face space.

In Practice

- Align a set of face images (the training set)
  - Rotate, scale and translate such that the eyes are located at the same coordinates.
- Compute the average face image
- Compute the difference image for each image in the training set
- Compute the covariance matrix of this set of difference images
- Compute the eigenvectors of the covariance matrix

Examples of Eigenfaces

- The eigenvectors of the covariance matrix can be viewed as images.

These are the first 4 eigenvectors, from a training set of 23 images....

Hence the name eigenfaces.

Dimensionality Reduction

- Only selecting the top M eigenfaces, reduces the dimensionality of the data.
- Too few eigenfaces results in too much information loss, and hence less discrimination between faces.

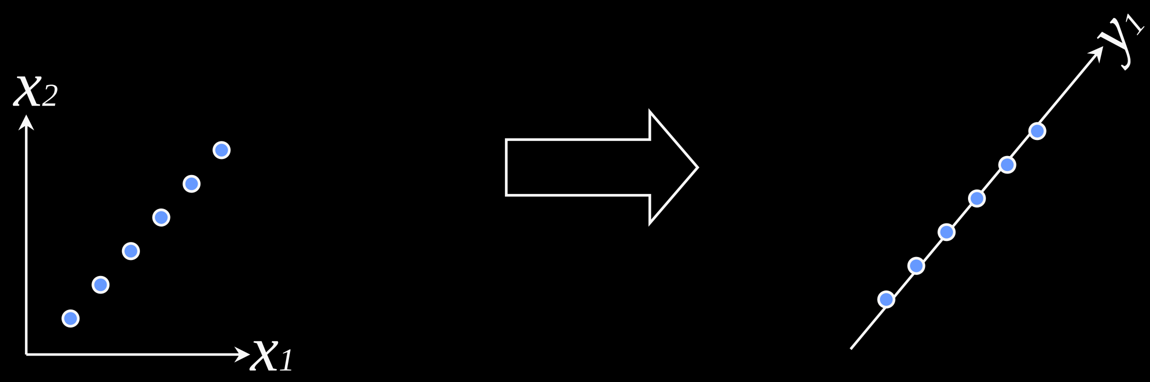

2D data

1D data

The Fisherface method

- *Eigenfaces vs. Fisherfaces: Recognition Using Class Specific Linear Projection*
  - P. Belhumeur, J. Hespanha, D. Kriegman
    - Yale University
- Eigenfaces attempt to maximise the scatter of the training images in face space.
- Fisherfaces attempt to maximise the between class scatter, while minimising the within class scatter.
- In other words, moves images of the same face closer together, while moving images of different faces further apart.

Fisher's Linear Discriminant

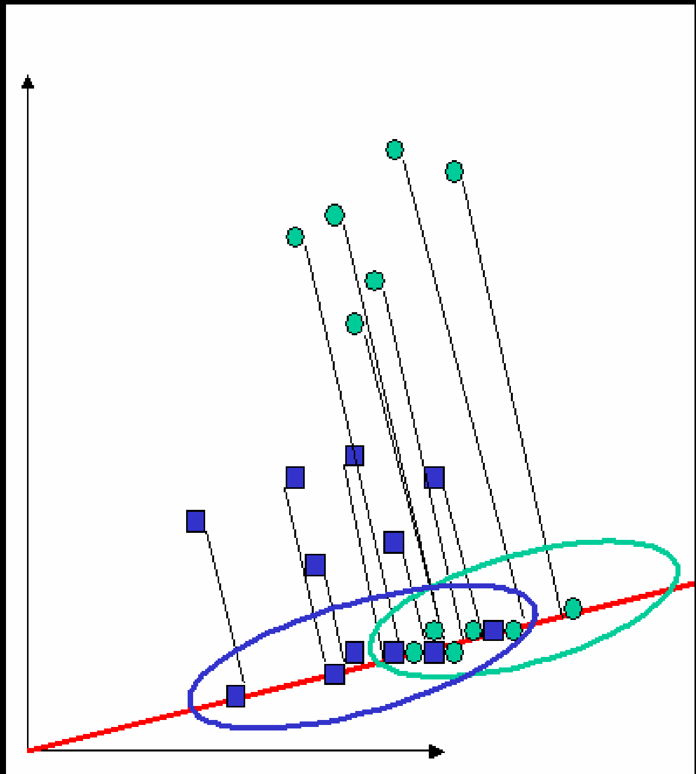

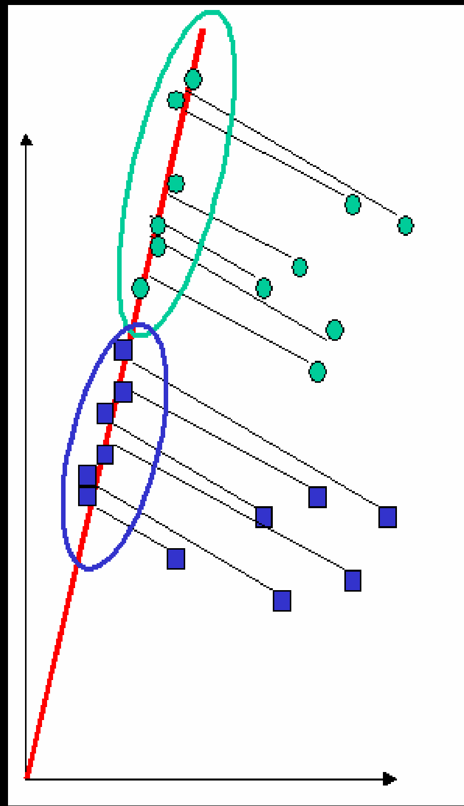

- Attempts to project the data such that the classes are separated.

Disadvantages of Face Recognition

- Not as accurate as other biometrics.
- Large amounts of storage needed.
- Good quality images needed.

Problems:

- Lighting
  - Difference in lighting conditions for enrolment and query.
  - Bright light causing image saturation.
  - Artificial coloured light.
- Pose – Head orientation
  - Difference between enrolment and subsequent images.
- Image quality
  - CCTV etc. is often not good enough for existing systems.

Face Recognition: the Problem of Compensating for Changes in Illumination Direction

- Yael Adini, Yael Moses, Shimon Ullman.

- The Weizmann Institute of Science

- Image representations used:
  - Standard greylevel, edge map, 2D gabor-like filters, first and second derivative.
- Distance measures used:
  - Pointwise, regional, affine-GL, local affine-GL, log distance
- Viewing conditions:
  - Frontal, profile, expressions, lighting.
- Missed-face:
  - If the distance between two images of one face under different conditions is greater than the distance between two different faces under the same conditions.

Results

- Changes in lighting direction:
  - Grey-level comparison 100% missed-faces
  - Other representations 20%~100% missed-faces
- Changes in viewing angle:
  - Grey-level comparison 100% missed-faces
  - Missed-faces of all representations above 50%
- Changes in expression
  - Smile
    - Grey-level comparison 0% missed-faces
    - Gabor-like filters reduced the accuracy to 34% even though it was good for the changes illumination
  - Drastic
    - Grey-level comparison 60% missed-faces
    - Other representations decreased accuracy

Lighting: Potential Solutions

- Controlled lighting
  - Dominant light source
  - Infrared images
    - *Face recognition using infrared images and Eigenfaces*. Ross cutler, Uni of Maryland
- Colour normalisation
  - Intensity normalisation
  - Grey-world normalisation
  - Comprehensive normalisation
  - HSV – hue representation
  - Brightness and gamma invariant hue
- Filters
  - Edge detection
  - 2D gabor-like filters
  - First and second derivatives

Comprehensive Colour Image Normalisation

- Graham Finlayson, The University of Derby
- Bernt Schiele, MIT
- James Crowley, INRIA Rhones Alpes

•Apply intensity normalisation, followed by Grey World.

•Repeat until a stable state is reached.

Hue that is invariant to brightness and gamma

- **Graham Finlayson, Gerald Schaefer, University of East Anglia**

•Apply a log transform to the RGBs.

•Gamma becomes mutliplicative scalars and cancel.

•Taking the difference between colour channel cancels the brightness.

•The angle of the resulting vector is analogous to the standard HSV Hue definition.

Examples of Lighting Correction

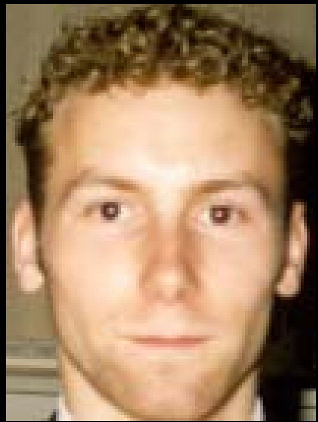

Original Image

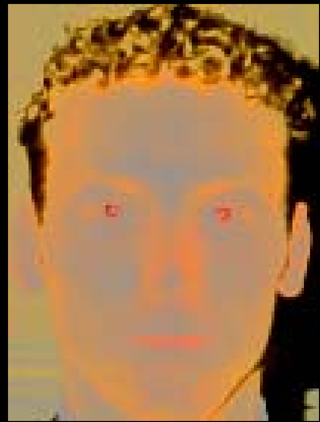

Intensity:
Invariant to light direction

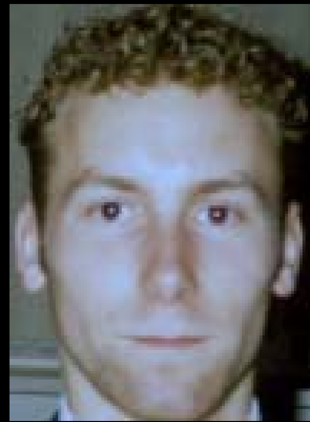

Grey world:
Invariant to coloured light

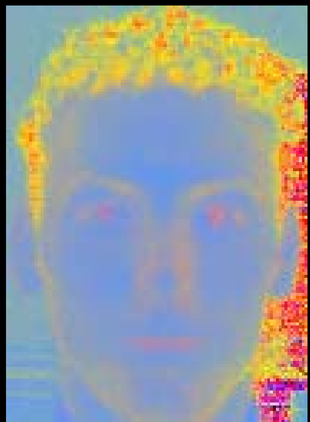

Comprehensive:
Invariant to light colour and direction

HSV Hue:
‘Colour’ representation

BGi Hue:
brightness and gamma invariant

Pose: Potential Solutions

- Multiple enrolment at various orientations
  - Increases FAR
  - Increases required storage space
- Image representations that are invariant to pose
  - Colour histograms
- 3D model enhancement
- View-based Eigenfaces

3D Model Enhanced Face Recognition

- Wen Zhao, Sarnoff Corporation, Princeton
- Rama Chellappa, University of Maryland

- Use a generic 3D shape to estimate light source and pose affect in the 2D image.
- Compensate for the above to render a prototype image.
- Perform face recognition on the prototype image.

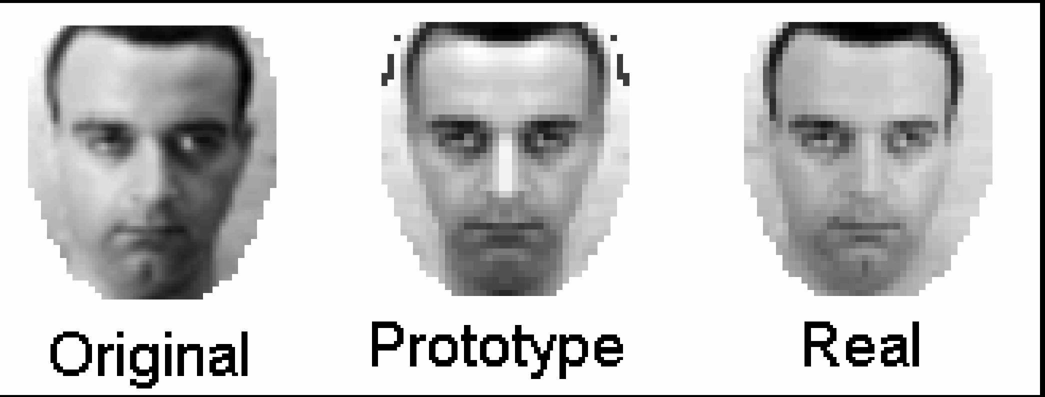

View-based and Modular Eigenfaces for Face Recognition

- Alex Pentland, Baback Moghadden, Thad Starner, MIT

-Use several projections into face space.

-Each projection represents a different viewing angle.

-When comparing faces use all projections.

-Use the nearest to face space$_{angle}$ or just identify as the nearest known face across all projections.

3D Facial Recognition

- Increase accuracy.
- Removes pose and lighting problems.
- Enough invariant information to cope with changes in expression, beards, glasses etc.

Existing Approaches:

- Profile matching.
- Surface segmentation matching.
- Point signature.
- Self-organising matching.
- PCA.
- AURA – coming soon.

Automatic 3D Face Authentication

- Charles Beumier, Mark Acheroy, Royal Military Academy, Belgium

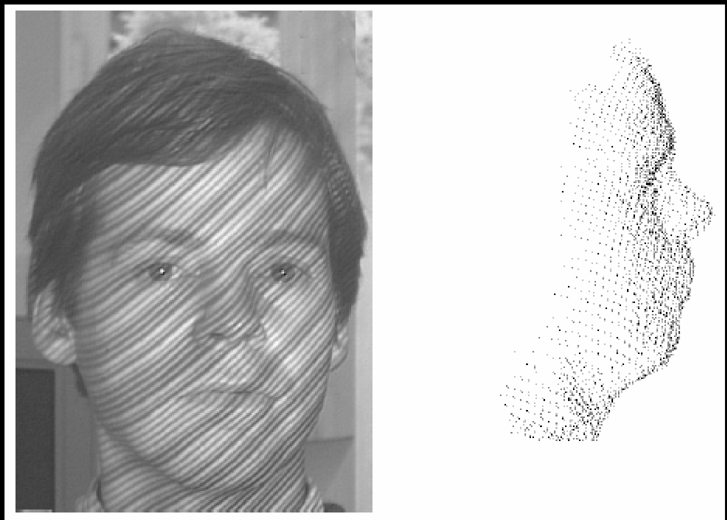

- 3D surface too noisy for global surface matching.
- Take central and lateral profiles from the 3D surface.
- Compare 13 2D profiles.

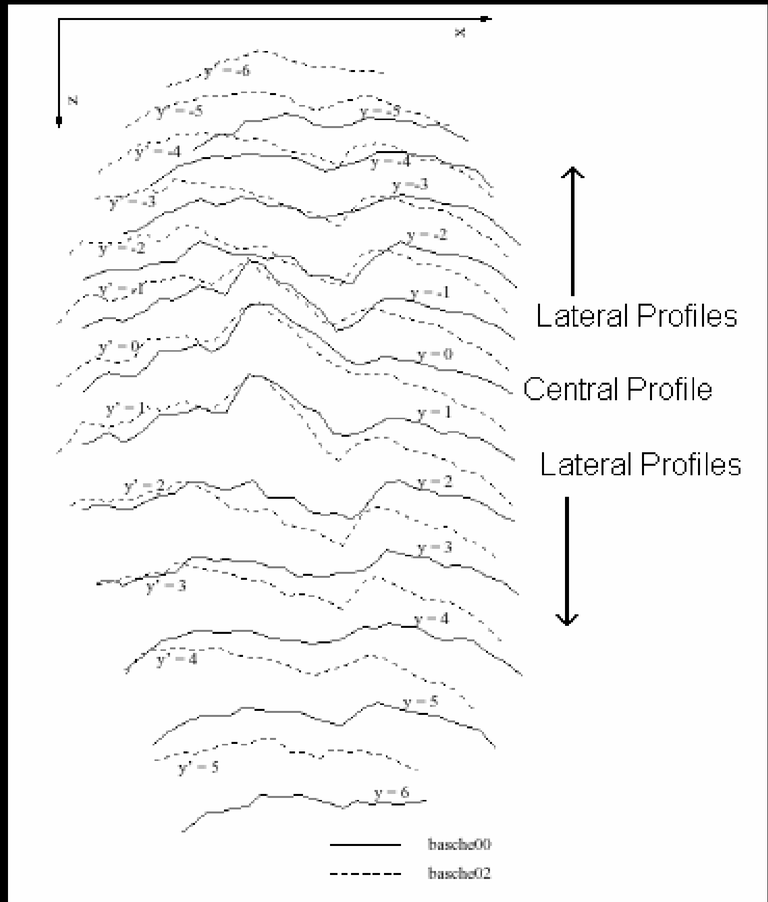

Description and Recognition of Faces from 3D Data

- A. Coombes, R. Richards, A. Linney, University College London
- V. Bruce, University of Nottingham
- R. Fright, Christchurch Hospital, New Zealand

- 3D Data acquired by optical surface scanning.
- Eight fundamental surface types are defined:
  – Peak, pit, ridge, valley, saddle ridge, saddle valley, minimal, flat.
- Facial surface is segmented into surface types.
- Facial features are manually localised by a user.
- Local regions are analysed for the surface type present.
- It is argued that faces can be distinguished by the surface types present in these local regions.
- No results are presented.

3D Human Face Recognition Using Point Signature

- Chin-Seng Chua, Feng Han, Yeong_Khing Ho.
- Nanyang Technological University, Singapore

- Treats the face recognition problem as 3D non-rigid surface recognition problem.
- For each person an analysis over four expressions is carried out to determine the rigid parts of the face.
- A face model of those rigid parts is constructed.
- Each model and test surface is represented by point signatures.

Point Signature

- Each plot in the 3D surface is represented by its point signature.
  - Place a sphere of radius r centred at point p.
  - The intersection of the sphere and surface creates a 3D space curve C.
  - This curve is projected such that its planar approximation is parallel to its normal, to make a new 3D curve C`.
  - A point signature is the set of distances from the points on C to the corresponding points on C`, at intervals of B^{o} around the sphere.

Matching 3D Surfaces

- Point signatures are compared by taking the difference between each distance pairs in the two point signatures.
- All distance must be within a tolerance level for the point signatures to match.
- 100% Accuracy achieved…
- But only tested on 6 people.

Some Other Approaches

- *3-D human face recognition by self-organizing matching approach*
  - S. Gerl, P. Levi
    - Implemented on a massively parallel field computer with 16387 processors.
    - A graph matching approach is used, by minimising a fitting function by simulated annealing.
- *Towards 3-dimensional face recognition*
  - A. Eriksson, D. Weber
    - Face meshes produced from a stereo image pair.
    - Recognition performed by attempting to project meshes onto test images.
- *Face recognition using 3D distance maps and principal component analysis*
  - H. Grecu, V. Buzuloiu, R. Beuran, E. Podaru

The Advanced Uncertain Reasoning Architecture, AURA

- J. Austin, J. Kennedy, K. Lees, University of York

- Correlation Matrix Memories based architecture.
- Simple hardware implementation.
- Able to match incomplete and noisy data at high speeds.
- Graph matcher uses AURA technology.
- Could this be applied to 3D facial surfaces?

Some Existing Aura Applications

- Chemical structure matching.
  - *Performance Evaluation of a Fast Chemical Structure Matching Method using Distributed Neural Relaxation*
    - **A Turner, J Austin, University of York**
- Trade mark matching.
  - *Content-Based Retrieval of Trademark Images*
    - **Sujeewa Alwis, University of York**
- Postal address matching.