


# Recognition: Face Recognition

Linda Shapiro

EE/CSE 576

# Face recognition: once you've detected and cropped a face, try to recognize it



Detection



Recognition

"Sally"

# Face recognition: overview

- Typical scenario: few examples per face, identify or verify test example
- What's hard: changes in expression, lighting, age, **occlusion**, **viewpoint**
- Basic approaches (all nearest neighbor)
  1. Project into a new subspace (or kernel space) (e.g., "Eigenfaces"=PCA)
  2. **Measure face features**

# Typical face recognition scenarios

- **Verification:** a person is claiming a particular identity; verify whether that is true
  - E.g., security
- **Closed-world identification:** assign a face to one person from among a known set
- **General identification:** assign a face to a known person or to “unknown”

# What makes face recognition hard?

## Expression





# What makes face recognition hard?

## Lighting





# What makes face recognition hard?

Occlusion






# What makes face recognition hard?

## Viewpoint



# Simple idea for face recognition

1. Treat face image as a vector of intensities



2. Recognize face by nearest neighbor in database



$$k = \underset{k}{\operatorname{argmin}} \|\mathbf{y}_k - \mathbf{x}\|$$


# The space of all face images

- When viewed as vectors of pixel values, face images are extremely high-dimensional
  - 100x100 image = 10,000 dimensions
  - Slow and lots of storage
- But very few 10,000-dimensional vectors are valid face images
- We want to effectively model the subspace of face images



# The space of all face images

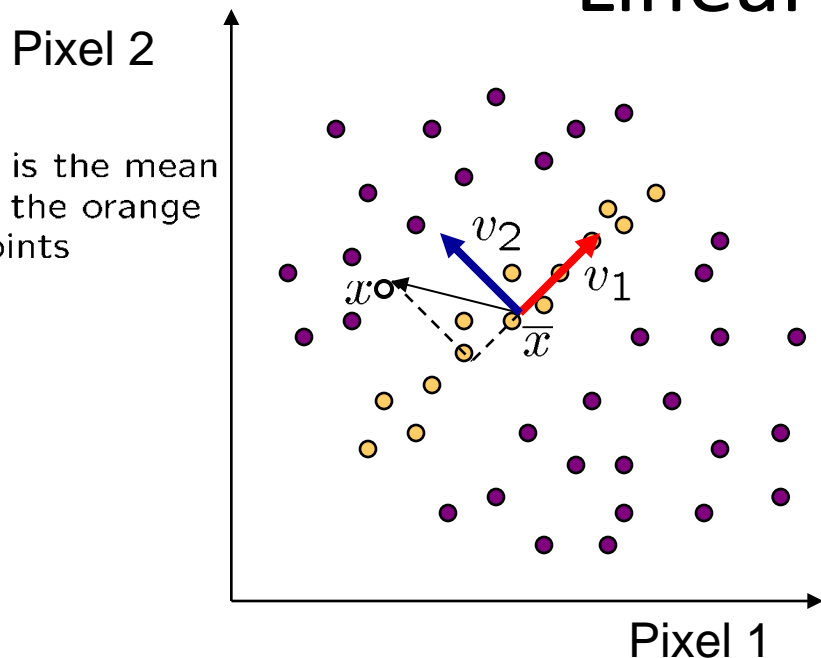
- Eigenface idea: construct a **low-dimensional linear subspace** that best explains the variation in the set of face images



● A face image

● A (non-face) image

# Linear subspaces



$\mathbf{v}_1$  is the major direction of the orange points and  $\mathbf{v}_2$  is perpendicular to  $\mathbf{v}_1$ .

Convert  $\mathbf{x}$  into  $\mathbf{v}_1, \mathbf{v}_2$  coordinates

$$\mathbf{x} \rightarrow ((\mathbf{x} - \bar{\mathbf{x}}) \cdot \mathbf{v}_1, (\mathbf{x} - \bar{\mathbf{x}}) \cdot \mathbf{v}_2)$$

What does the  $\mathbf{v}_2$  coordinate measure?

- distance to line
- use it for classification—near 0 for orange pts

What does the  $\mathbf{v}_1$  coordinate measure?

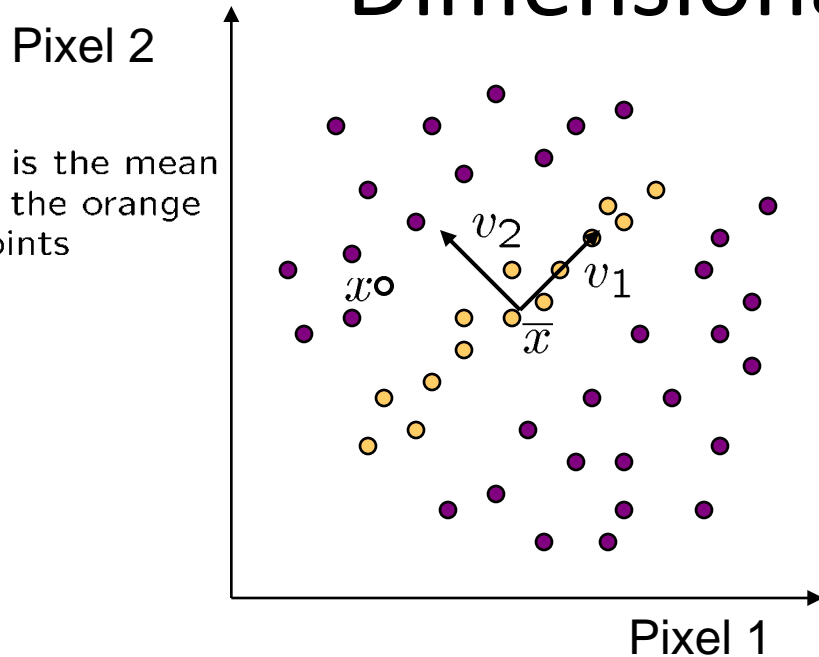
- position along line
- use it to specify which orange point it is

- Classification (to what class does  $\mathbf{x}$  belong) can be expensive
  - Big search problem

Suppose the data points are arranged as above

- Idea—fit a line, classifier measures distance to line

# Dimensionality reduction



## Dimensionality reduction


- We can represent the orange points with *only* their  $\mathbf{v}_1$  coordinates
  - since  $\mathbf{v}_2$  coordinates are all essentially 0
- This makes it much cheaper to store and compare points
- A bigger deal for higher dimensional problems (like images!)



# Eigenvectors and Eigenvalues

Pixel 2

$\bar{x}$  is the mean of the orange points



Consider the variation along a direction  $\mathbf{v}$  among all of the orange points:

$$var(\mathbf{v}) = \sum_{\text{orange point } \mathbf{x}} \|(\mathbf{x} - \bar{\mathbf{x}})^T \cdot \mathbf{v}\|^2$$

What unit vector  $\mathbf{v}$  minimizes  $var$ ?

$$\mathbf{v}_2 = \min_{\mathbf{v}} \{var(\mathbf{v})\}$$

What unit vector  $\mathbf{v}$  maximizes  $var$ ?

$$\mathbf{v}_1 = \max_{\mathbf{v}} \{var(\mathbf{v})\}$$

$$\begin{aligned} var(\mathbf{v}) &= \sum_{\mathbf{x}} \|(\mathbf{x} - \bar{\mathbf{x}})^T \cdot \mathbf{v}\|^2 \\ &= \sum_{\mathbf{x}} \mathbf{v}^T (\mathbf{x} - \bar{\mathbf{x}}) (\mathbf{x} - \bar{\mathbf{x}})^T \mathbf{v} \\ &= \mathbf{v}^T \left[ \sum_{\mathbf{x}} (\mathbf{x} - \bar{\mathbf{x}}) (\mathbf{x} - \bar{\mathbf{x}})^T \right] \mathbf{v} \\ &= \mathbf{v}^T \mathbf{A} \mathbf{v} \quad \text{where } \mathbf{A} = \sum_{\mathbf{x}} (\mathbf{x} - \bar{\mathbf{x}}) (\mathbf{x} - \bar{\mathbf{x}})^T \end{aligned}$$

$\mathbf{A}$  = **covariance matrix** of data points (if divided by no. of points)

**Solution:  $\mathbf{v}_1$  is eigenvector of  $\mathbf{A}$  with *largest* eigenvalue  
 $\mathbf{v}_2$  is eigenvector of  $\mathbf{A}$  with *smallest* eigenvalue**

# Principal component analysis (PCA)


- Suppose each data point is N-dimensional
  - Same procedure applies:

$$\begin{aligned} \text{var}(\mathbf{v}) &= \sum_{\mathbf{x}} \|(\mathbf{x} - \bar{\mathbf{x}})^T \cdot \mathbf{v}\|^2 \\ &= \mathbf{v}^T \mathbf{A} \mathbf{v} \quad \text{where } \mathbf{A} = \sum_{\mathbf{x}} (\mathbf{x} - \bar{\mathbf{x}})(\mathbf{x} - \bar{\mathbf{x}})^T \end{aligned}$$

- The eigenvectors of **A** define a new coordinate system
  - eigenvector with largest eigenvalue captures the most variation among training vectors **x**
  - eigenvector with smallest eigenvalue has least variation
- We can compress the data by only using the top few eigenvectors
  - corresponds to choosing a “linear subspace”
    - represent points on a line, plane, or “hyper-plane”
  - these eigenvectors are known as the **principal components**

# The space of faces


---



- An image is a point in a high dimensional space
  - An  $N \times M$  image is a point in  $\mathbb{R}^{NM}$
  - We can define vectors in this space as we did in the 2D case

# Dimensionality reduction

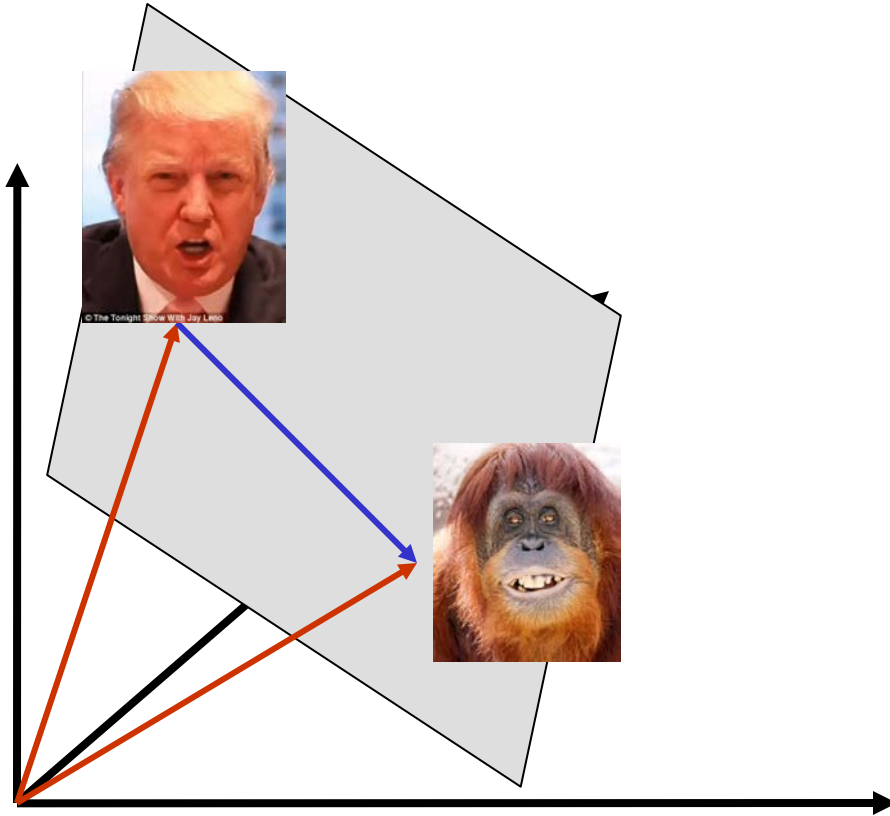
---



- The set of faces is a “subspace” of the set of images
  - Suppose it is  $K$  dimensional
  - We can find the best subspace using PCA
  - This is like fitting a “hyper-plane” to the set of faces
    - spanned by vectors  $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_K$
    - any face  $\mathbf{x} \approx \bar{\mathbf{x}} + a_1\mathbf{v}_1 + a_2\mathbf{v}_2 + \dots + a_k\mathbf{v}_k$

# Dimensionality reduction

---




- The set of faces is a “subspace” of the set of images
  - Suppose it is  $K$  dimensional
  - We can find the best subspace using PCA
  - This is like fitting a “hyper-plane” to the set of faces
    - spanned by vectors  $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_K$
    - any face  $\mathbf{x} \approx \bar{\mathbf{x}} + a_1\mathbf{v}_1 + a_2\mathbf{v}_2 + \dots + a_k\mathbf{v}_k$

# Eigenfaces

---

- PCA extracts the eigenvectors of  $\mathbf{A}$ 
  - Gives a set of vectors  $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \dots$
  - Each one of these vectors is a direction in face space
    - what do these look like?





# Visualization of eigenfaces

Principal component (eigenvector)  $u_k$



$\mu + 3\sigma_k u_k$



$\mu - 3\sigma_k u_k$



# Projecting onto the eigenfaces

- The eigenfaces  $\mathbf{v}_1, \dots, \mathbf{v}_K$  span the space of faces


– A face is converted to eigenface coordinates by

$$\mathbf{x} \rightarrow \left( \underbrace{(\mathbf{x} - \bar{\mathbf{x}}) \cdot \mathbf{v}_1}_{a_1}, \underbrace{(\mathbf{x} - \bar{\mathbf{x}}) \cdot \mathbf{v}_2}_{a_2}, \dots, \underbrace{(\mathbf{x} - \bar{\mathbf{x}}) \cdot \mathbf{v}_K}_{a_K} \right)$$

$$\mathbf{x} \approx \bar{\mathbf{x}} + a_1\mathbf{v}_1 + a_2\mathbf{v}_2 + \dots + a_K\mathbf{v}_K$$



$\mathbf{x}$



$a_1\mathbf{v}_1$   $a_2\mathbf{v}_2$   $a_3\mathbf{v}_3$   $a_4\mathbf{v}_4$   $a_5\mathbf{v}_5$   $a_6\mathbf{v}_6$   $a_7\mathbf{v}_7$   $a_8\mathbf{v}_8$



# Recognition with eigenfaces

---

- Algorithm

1. Process the image database (set of images with labels)

- Run PCA—compute eigenfaces
- Calculate the K coefficients for each image

2. Given a new image (to be recognized)  $\mathbf{x}$ , calculate K coefficients

$$\mathbf{x} \rightarrow (a_1, a_2, \dots, a_K)$$

3. Detect if  $\mathbf{x}$  is a face


$$\|\mathbf{x} - (\bar{\mathbf{x}} + a_1\mathbf{v}_1 + a_2\mathbf{v}_2 + \dots + a_K\mathbf{v}_K)\| < \text{threshold}$$

4. If it is a face, who is it?

- Find closest labeled face in database
- Nearest-neighbor in K-dimensional space

# Choosing the dimension K

---



- How many eigenfaces to use?
- Look at the decay of the eigenvalues
  - the eigenvalue tells you the amount of variance “in the direction” of that eigenface
  - ignore eigenfaces with low variance

# PCA

- General dimensionality reduction technique
- Preserves most of variance with a much more compact representation
  - Lower storage requirements (eigenvectors + a few numbers per face)
  - Faster matching
- What other applications?

# Enhancing gender



more

same

**original**

androgynous

more opposite

D. Rowland and D. Perrett, [“Manipulating Facial Appearance through Shape and Color,”](#) IEEE CG&A, September 1995



# Changing age

- Face becomes “rounder” and “more textured” and “grayer”


- original



- color

D. Rowland and D. Perrett, [“Manipulating Facial Appearance through Shape and Color,”](#) IEEE CG&A, September 1995

# Which face is more attractive?






<http://www.beautycheck.de>

# Use in Cleft Severity Analysis

- We have a large database of normal 3D faces.
- We construct their principal components.
- We can reconstruct any normal face accurately using these components.
- But when we reconstruct a cleft face from the normal components, there is a lot of error.
- This error can be used to measure the severity of the cleft.

# Use of PCA Reconstruction Error to Judge Cleft Severity



Aligned head mesh

Error map from PCA reconstruction

# Extension to 3D Objects

- Murase and Nayar (1994, 1995) extended this idea to 3D objects.
- The training set had **multiple views of each object**, on a dark background.
- The views included **multiple (discrete) rotations** of the object on a turntable and also **multiple (discrete) illuminations**.
- The system could be used first to **identify** the object and then to determine its (approximate) **pose** and illumination.

# Sample Objects

## Columbia Object Recognition Database

COLUMBIA UNIVERSITY IMAGE LIBRARY (COIL-20)



# Significance of this work

- The extension to 3D objects was an important contribution.
- Instead of using brute force search, the authors observed that  
  
All the views of a single object, when transformed into the eigenvector space became points on a manifold in that space.
- Using this, they developed fast algorithms to find the closest object manifold to an unknown input image.
- **Recognition with pose finding took less than a second.**

# Appearance-Based Recognition

- Training images must be representative of the instances of objects to be recognized.
- The object must be well-framed.
- Positions and sizes must be controlled.
- Dimensionality reduction is needed.
- It is not powerful enough to handle general scenes without prior segmentation into relevant objects.
- \* The newer systems that use “parts” from interest operators are an answer to these restrictions.