CS 332 Visual Processing in Computer and Biological Vision Systems

The Analysis of Faces in Brains and Machines

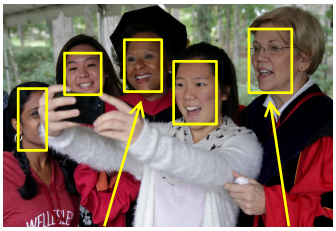

Paula Johnson

Elizabeth Warren

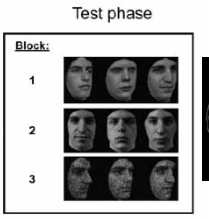

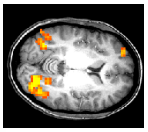

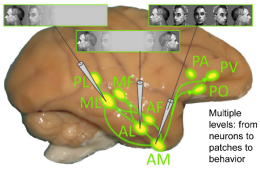

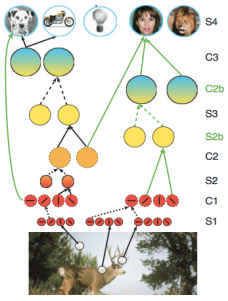

HMAX model

Why is face analysis important?

Remember/recognize people we've seen before

Categorization – e.g. gender, race, age, kinship

Social communication – emotions/mood, intentions, trustworthiness, competence or intelligence, attractiveness

Scene understanding, e.g. direction of gaze suggests focus of attention

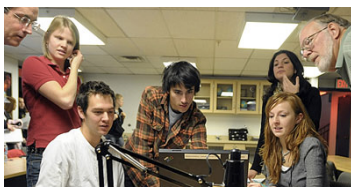

Why is face recognition hard?

changing pose

changing illumination

changing expression

aging

clutter

occlusion

How good are we at face recognition?

Jenkins, White, Van Montfort & Burton, *Cognition*, 2011

Face recognition performance in humans

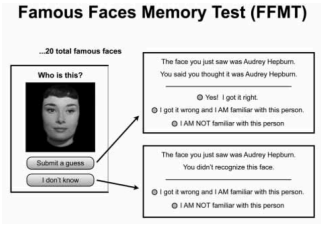

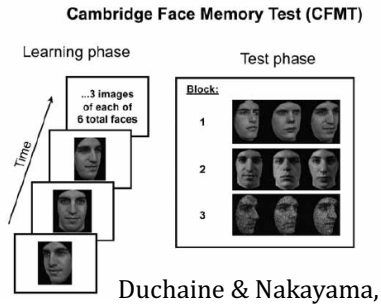

Duchaine & Nakayama, 2006

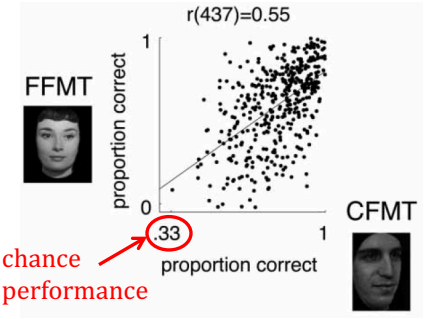

testmybrain.org

Wilmer et al., 2012

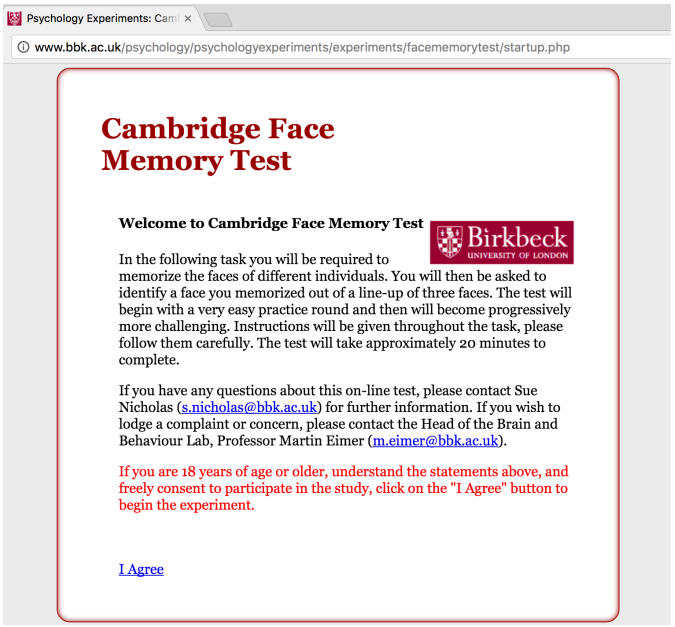

Face recognition performance in humans

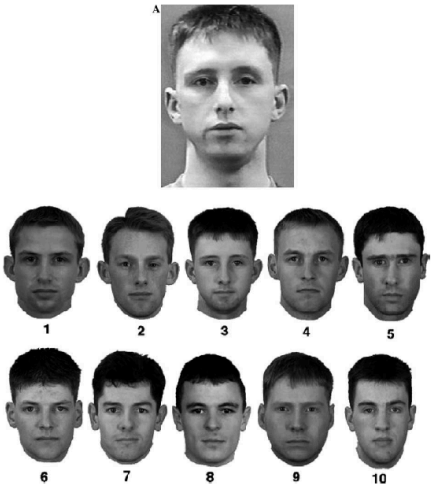

Bruce et al., 1999

Which of the 10 photos on the bottom depicts the target face?

Viewers are ~ 70% correct

Performance degrades with changes in pose, expression

Only slight improvement with short video clip of target

Importance of familiar vs. unfamiliar face recognition!

How good are the best machines?

Public databases of face images serve as benchmarks:

Labeled Faces in the Wild (LFW, http://vis-www.cs.umass.edu/lfw)
> 13,000 images of celebrities, 5,749 different identities

YouTube Faces Database (YTF, http://www.cs.tau.ac.il/~wolf/ytfaces)
3,425 videos, 1,595 different identities

Private face image datasets:

(Facebook) Social Face Classification dataset
4.4 million face photos, 4,030 different identities
(Google) 100-200 million face images, ~ 8 million different identities

	LFW	YTF
Facebook DeepFace	97.4%	91.4%
Google FaceNet	99.6%	95.1%
Human performance	97.5%	89.7%

Machine vision applications of face recognition

security, forensics

access control

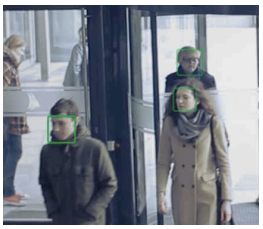

surveillance

More applications of face recognition

Facebook's New Face-Recognition Software Is Scary Good

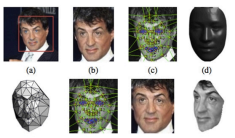

content-based image retrieval

Facial recognition app matches strangers to online profiles

social media

Google Glass Face Recognition App Coming This Month, Whether Google Likes It Or Not

Since Google Glass first appeared, its potential for facial recognition has been seen either as a privacy nightmare or as one of the headset's first truly intriguing uses. Google has declared itself in the first camp. Stephen Balaban is in the second, and he's about to share his vision with Glassheads everywhere, whether Google likes it or not.

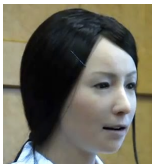

graphics, HCI
humanoid robots

Aspects of face processing

Face detection – find image regions that contain faces

Face identification – who is the person?

Categorization – gender, age, race

Facial expression – mood, emotion

Non-verbal social perception and communication

It all began with Takeo Kanade (1973)…

PhD thesis, *Picture Processing System by Computer Complex and Recognition of Human Faces*

- Special purpose algorithms to locate eyes, nose, mouth, boundaries of face
- ~ 40 geometric features, e.g. ratios of distances and angles between features

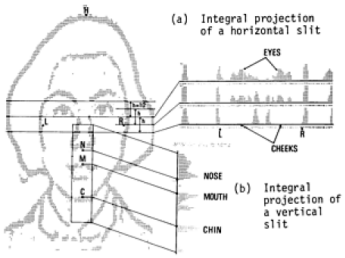

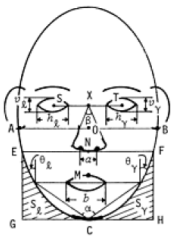

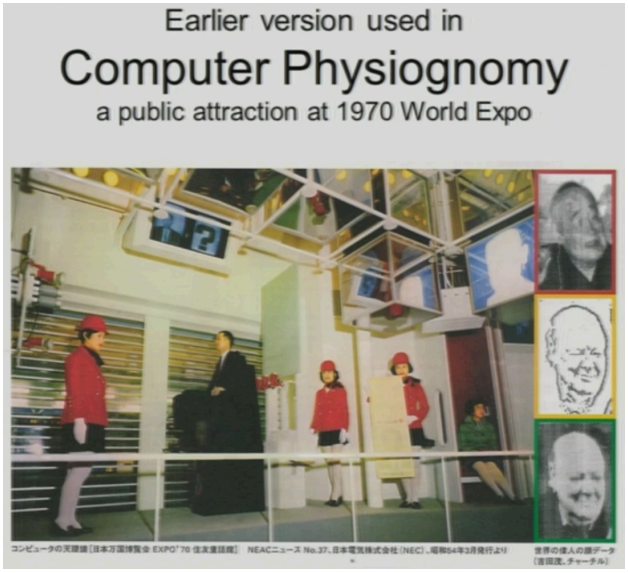

- From talk by Takeo Kanade, *CBMM Face ID Challenge Workshop*

Eigenfaces for recognition (Turk & Pentland)

Principal Components Analysis (PCA)

Goal: reduce the dimensionality of the data while retaining as much information as possible in the original dataset

PCA allows us to compute a linear transformation that maps data from a high dimensional space to a lower dimensional subspace

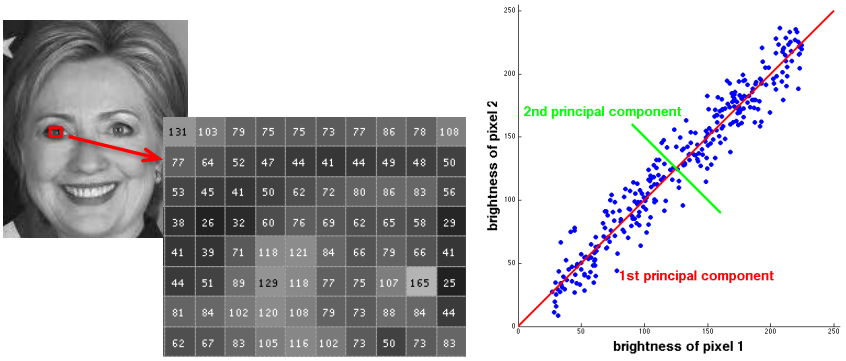

Typical sample training set…

Sample images from the Yale face database

- One or more images per person
- Aligned & cropped to common pose, size
- Simple background

Eigenfaces for recognition (Turk & Pentland)

$\Psi(x,y)$

Perform ***PCA*** on a large set of training images, to create a set of *eigenfaces*, $E_i(x,y)$, that span the data set

First components capture most of the variation across the data set, later components capture subtle variations

$\Psi(x,y)$: average face (across all faces)

http://vismod.media.mit.edu/vismod/demos/facerec/basic.html

Each face image $F(x,y)$ can be expressed as a weighted combination of the eigenfaces $E_i(x,y)$:

$$F(x,y) = \Psi(x,y) + \Sigma_i \, w_i * E_i(x,y)$$

1-16

Representing individual faces

Each face image F(x,y) can be expressed as a weighted combination of the eigenfaces $E_i(x,y)$:

$$F(x,y) = \Psi(x,y) + \Sigma_i \, w_i * E_i(x,y)$$

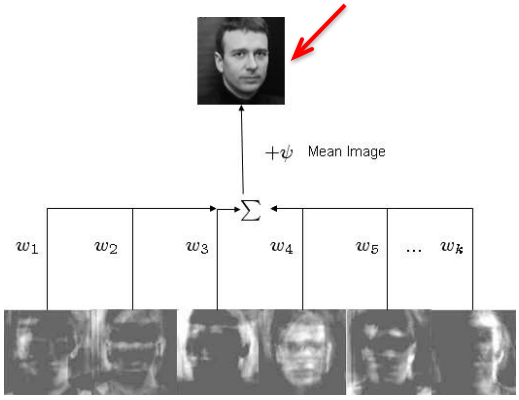

Recognition process:

(1) Compute weights w_i for novel face image

(2) Find image m in face database with most similar weights, e.g.

$$\min \sum_{i=1}^{k} (w_i - w_i^m)^2$$

Faces everywhere…