What is Facial Processing

- One of the most well studied topics in Computer Vision.
- Tons of applications in this human-centric world:
  - Digital photography
  - Person identification and verification
  - Emotions and expressions

Character recognition from TV shows *[Everingham et. al.]*

Female
Young
Black hair
…

Facial attributes *[Kumar et. al.]*

Human Perception of Faces

- Understanding human perception helps facial processing in Computer Vision.
- Early face recognition systems use the facial landmark distance ratios.
- Who are these people?



Each face is compressed to be 25% of its original width.

Major Milestones

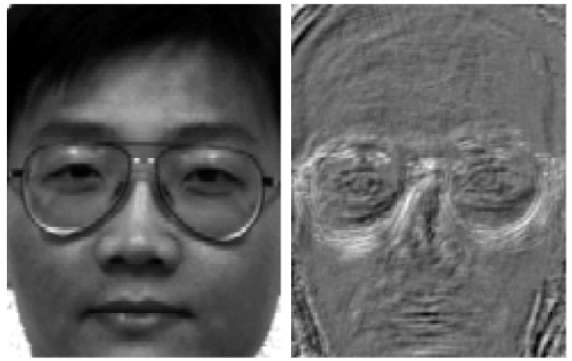

Eigenfaces [1991]

Fisherfaces [1997]

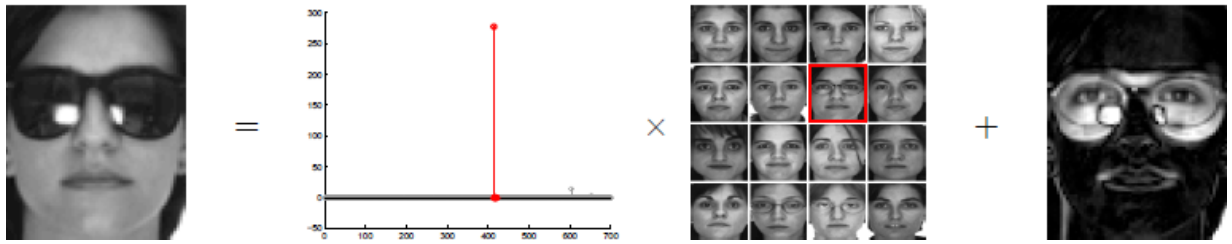

Viola-Jones Face Detector [2001]

Sparse Coding [2009]

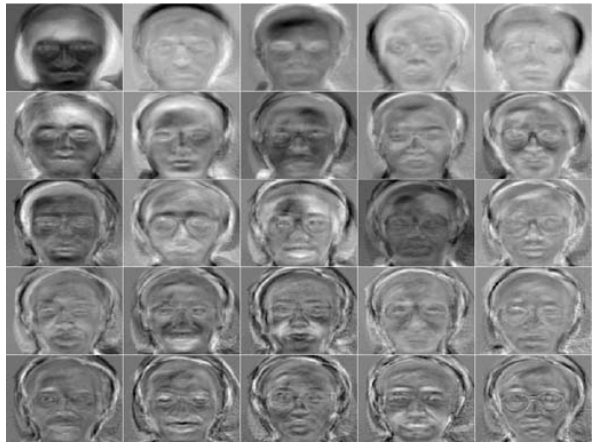

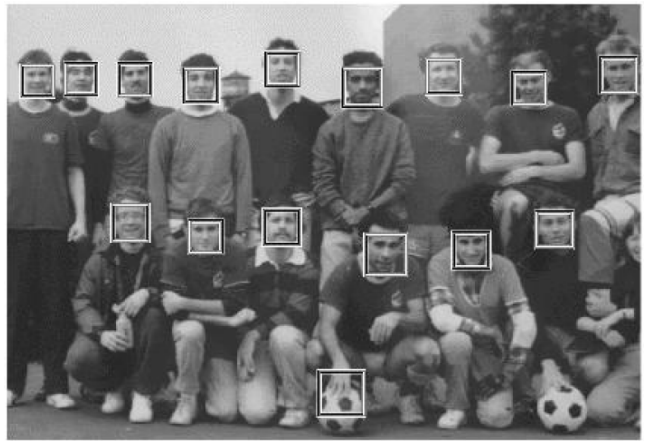

A Typical Facial Processing Pipeline

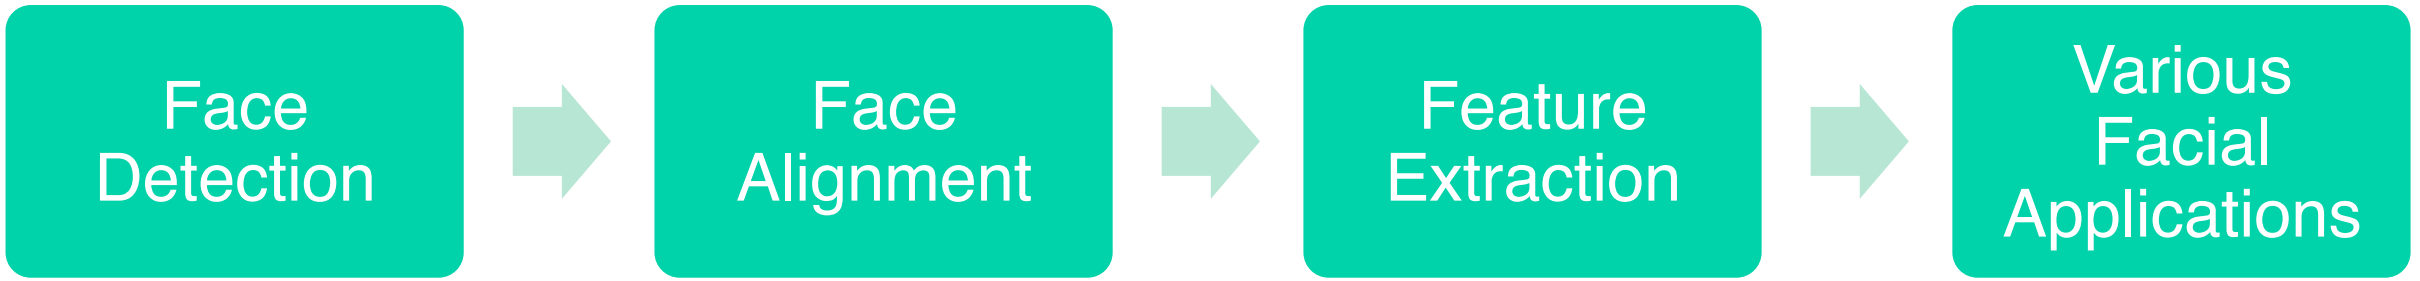

Face Detection

- Human can detect face effortlessly. But also susceptible to occlusions.

EE368 Class Project 2003 – Face Detection

- Most groups used skin detector, followed by connected component analysis.

Face Detection Overview (Viola & Jones, 2001)

- Slide a multi-scale detection window through the image.
- Classify each detection window as face / non-face.

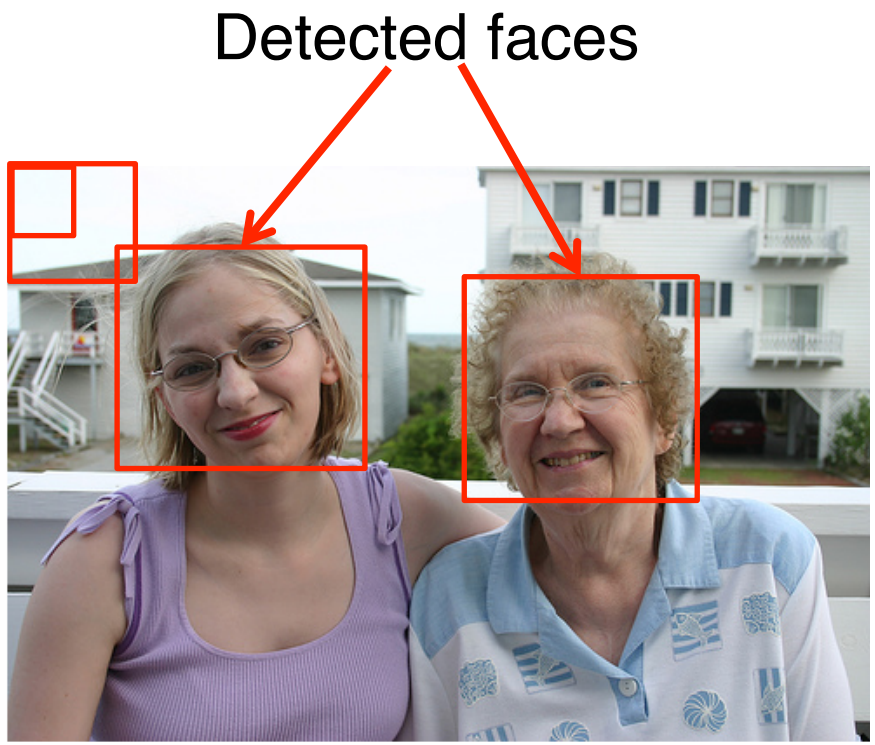

Face Detection (Viola & Jones, 2001)

- Features: difference between the sum of the pixels within two rectangular regions.
  - Captures prior knowledge about faces.
  - Efficient computation.

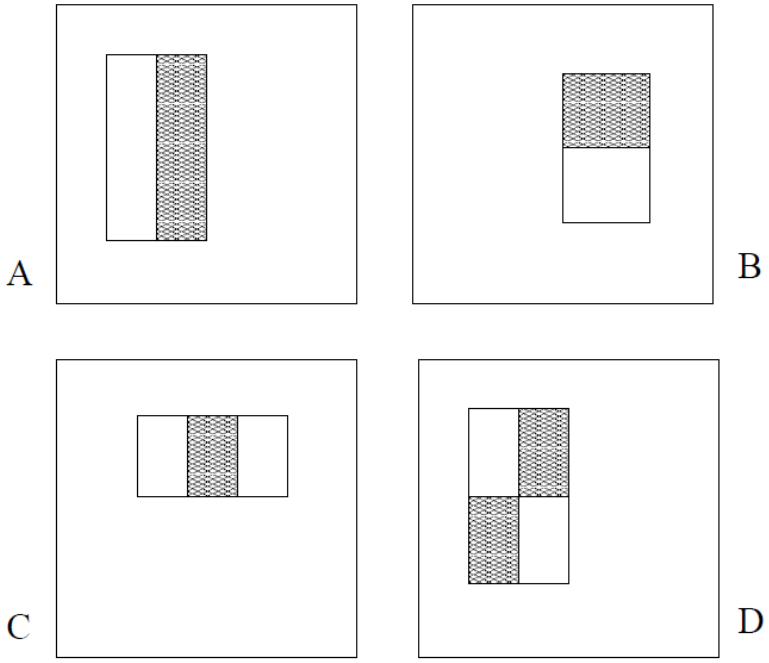

Integral Image

- The value at a point (x, y) in the integral image is the sum of all the pixels above and to the left of (x,y)

$$I(x,y) = \sum_{x' \le x, y' \le y} i(x',y')$$

Image pixels

2	5	6	2
3	2	8	4
6	7	9	1
1	9	2	4

Integral image

2	7	13	15
5	12	26	32
11	25	48	55
12	35	60	71

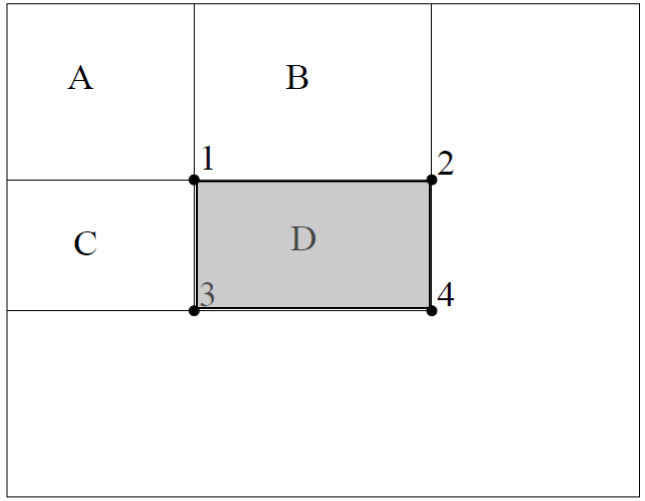

1: A
2: A+B
3: A+C
4: A+B+C+D
Pixel sum in shade: 4+1-2-3

Face/Non-face Classification

- Objective: classify a detection window as face/non-face.
- For each detection window, compute the feature responses.

Feature type

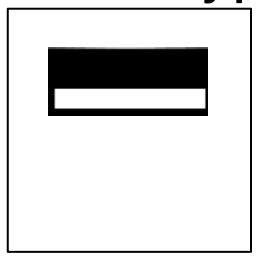

Response scores:

0.9 0.8 0.1 0.4

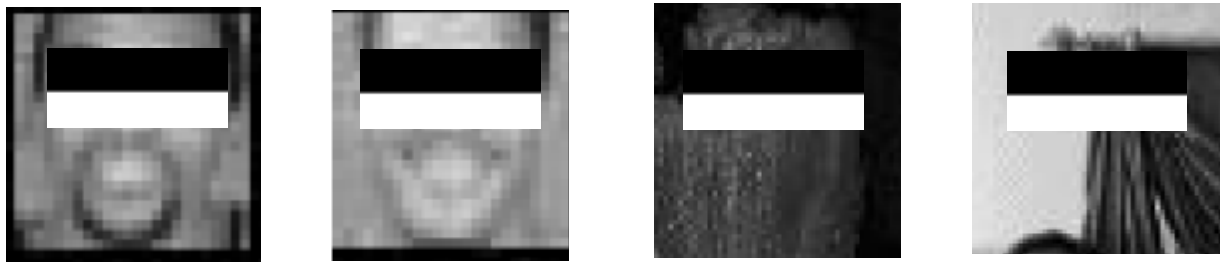

- There are many types of features. Select relevant features for face detection.

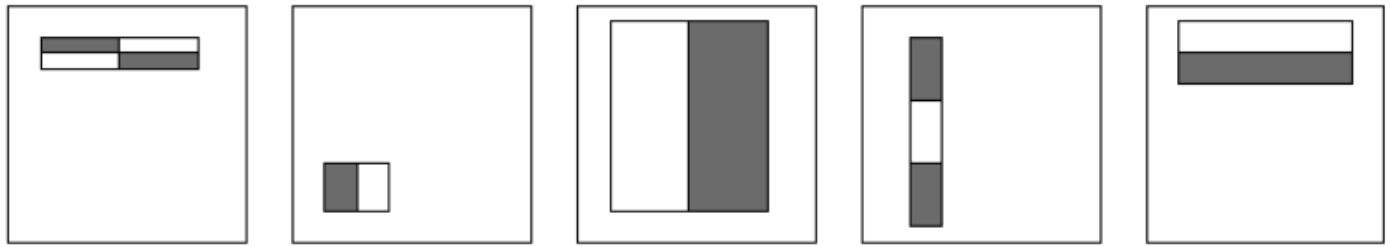

...

[Slide credit: L. Fei-fei]

AdaBoost *[Freund & Schapire 97]*

Feature space

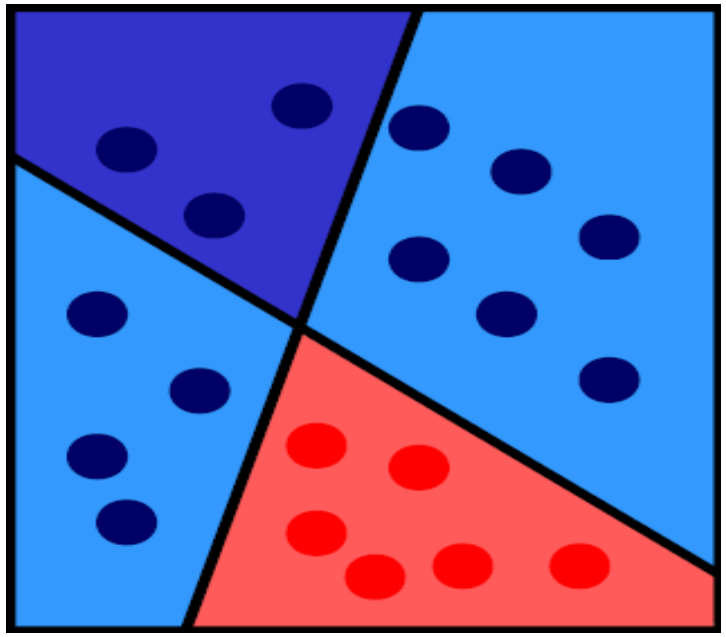

$$= \alpha_1 \times \; h_1 \; + \; \alpha_2 \times \; h_2$$

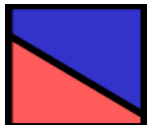

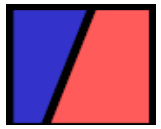

Input

- Set of labeled training samples
- Weight distribution over samples

Algorithm

for n=1 to N // number of classifiers

- train a classifier using samples and weight distribution
- calculate error
- calculate classifier weight
- update sample weights

end

Result

$$H(\mathbf{x}) = \text{sign}\left(\sum_{n=1}^{N} \alpha_n \, h_n(\mathbf{x})\right)$$

[slide credit H. Grabner]

Feature Selection by Adaboost

- The top two features selected by Adaboost:

- To achieve good performance, Adaboost needs to select many features.
  - Computationally expensive!
  - An idea: Quickly rule out the windows that are highly unlikely to be faces. Gradually refine detection.

Cascaded Detection

- Each node is a layer of detection with a few number of features.
- At each detection layer, the threshold is tuned such that no face is missed.
- Detection windows that are highly non-face will be rejected at early stage.

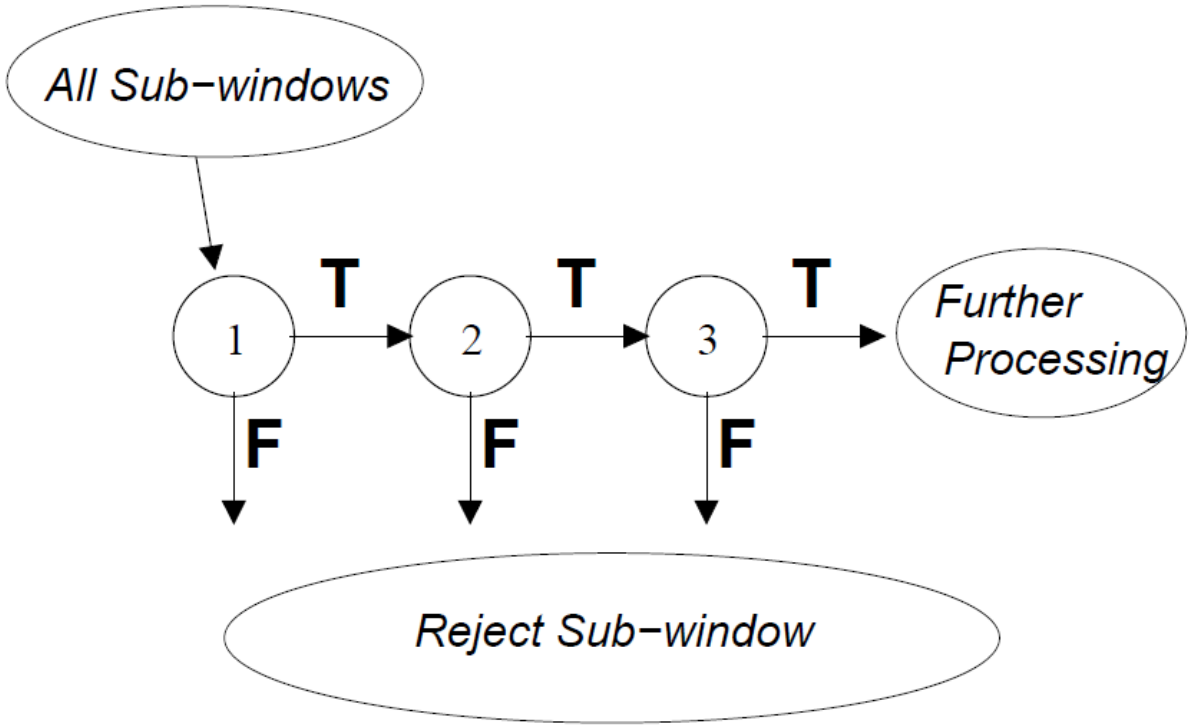

Face Detection Examples

- 700 MHz Intel Pentium III
  - Training: a 38 layer cascaded classifier was trained in “weeks”.
  - Testing: 384×288 pixel image in about 0.067 seconds.

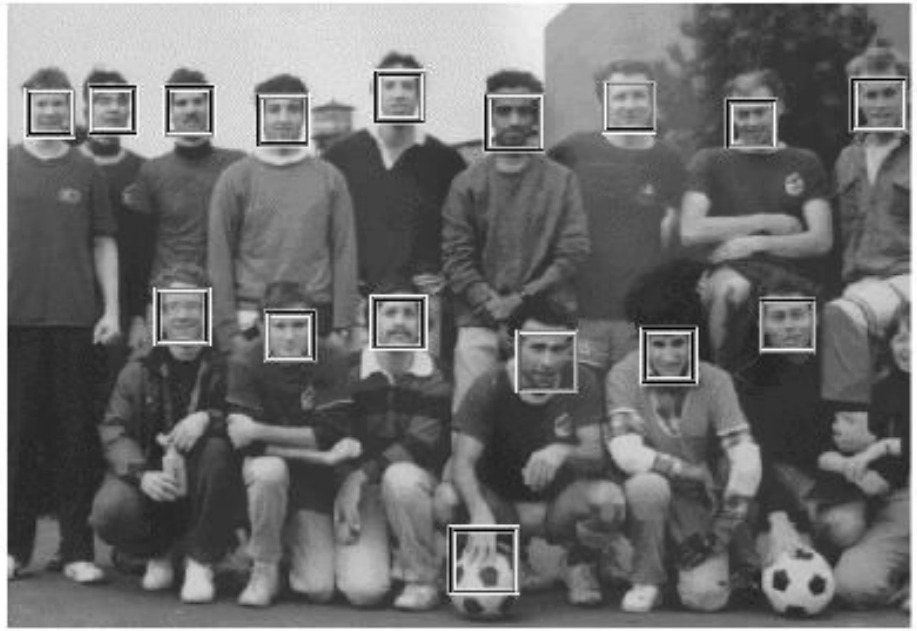

Facial Landmark Detection

- The Viola-Jones face detector can also be used for generic object detection.
- Finding the facial landmarks (eyes, nose, mouth, etc.).

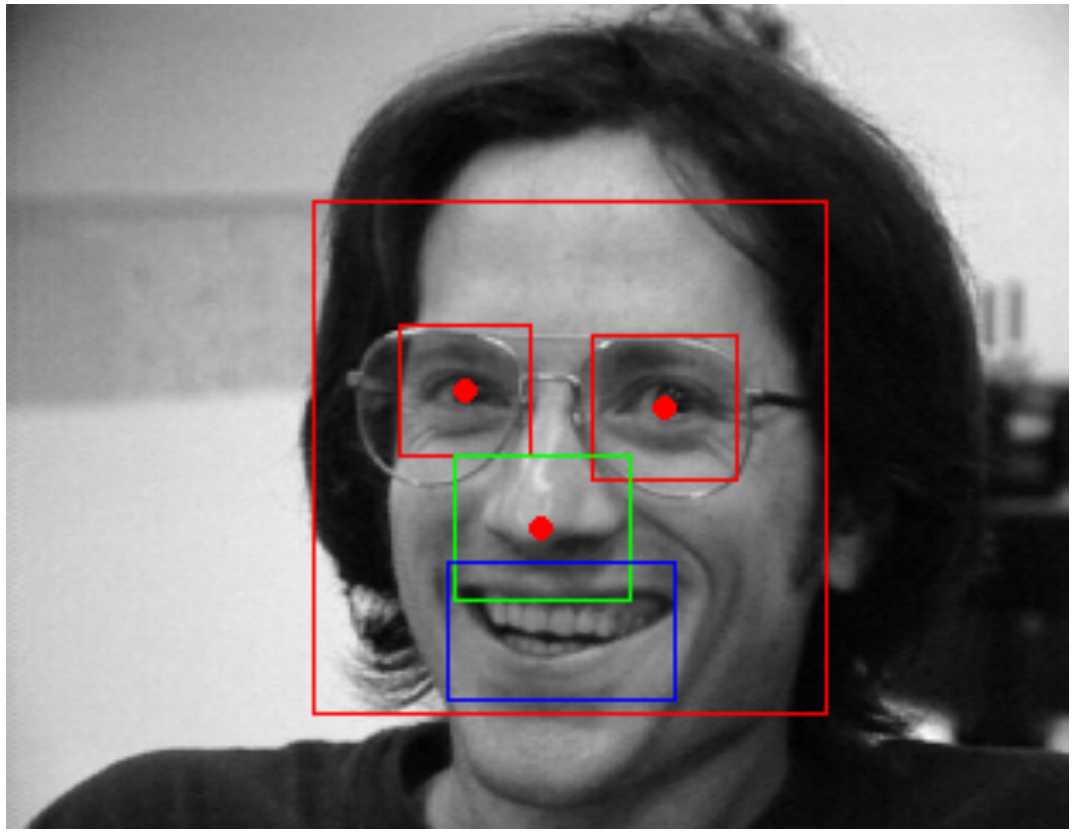

Face Alignment

- Face alignment:
  - Detect face landmarks.
  - Apply similarity transform to align the faces (usually done by aligning eye locations).

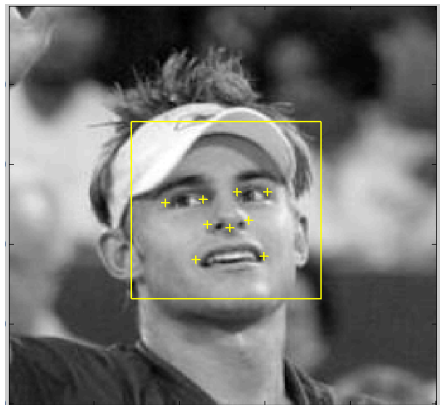
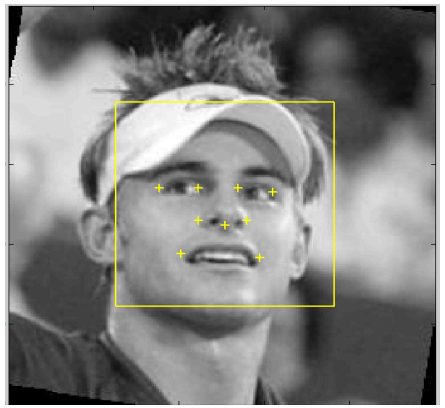
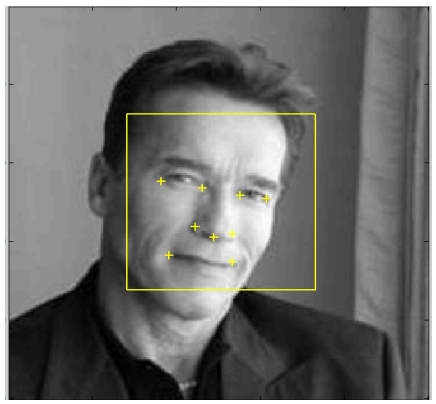
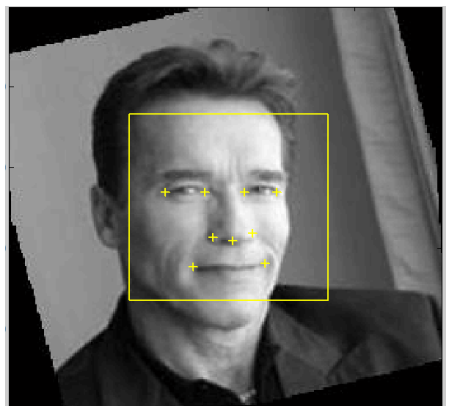

Recap: Similarity transform

$$\mathbf{x}' = \begin{bmatrix} s\cdot\cos\theta & -s\cdot\sin\theta & t_x \\ s\cdot\sin\theta & s\cdot\cos\theta & t_y \end{bmatrix} \underline{\mathbf{x}}$$

A Typical Facial Processing Pipeline

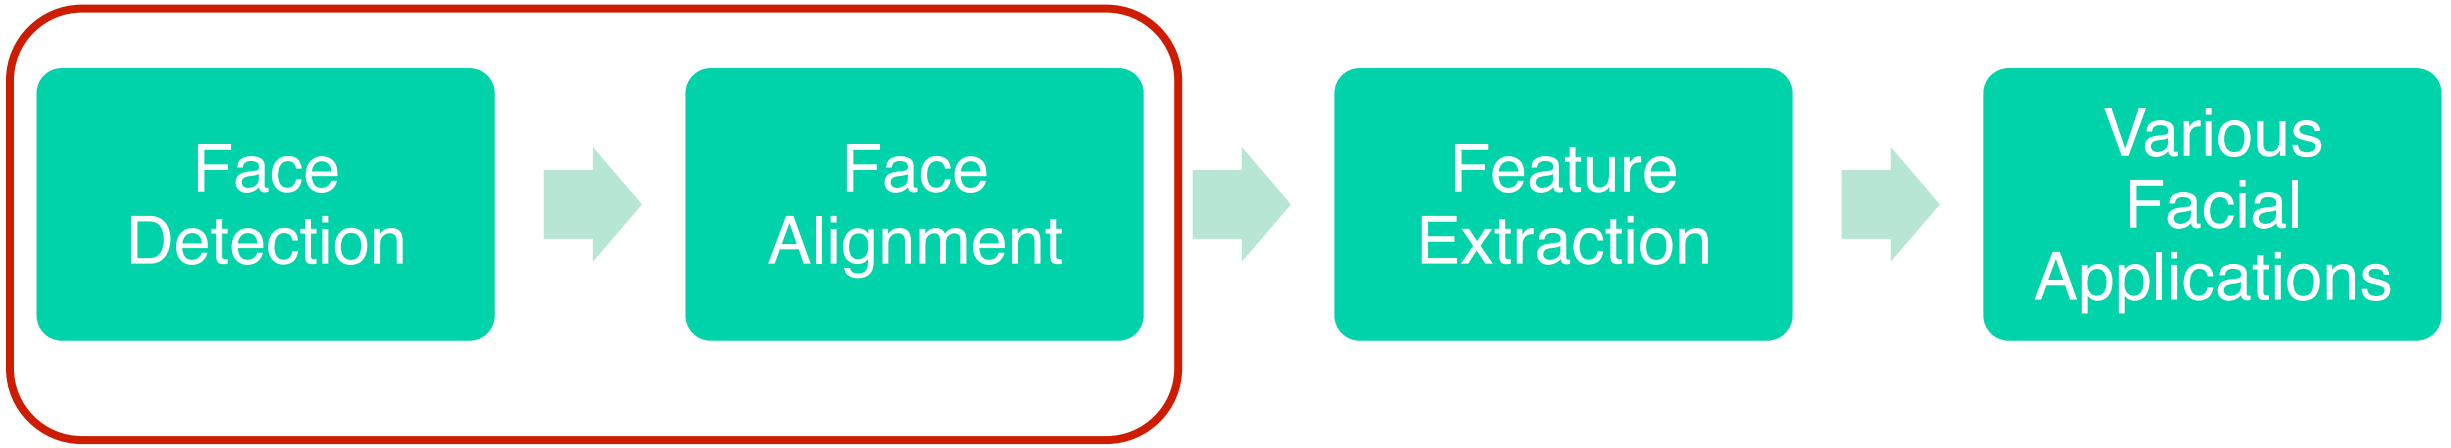

What we have discussed.

A Case Study – The Hidden Sides Names

- Psychology studies: Human has facial stereotypes about names.
- Questions:
  - Is there a relation between a person's first name and facial appearance?
  - Can we use this name-face relation for practical applications?

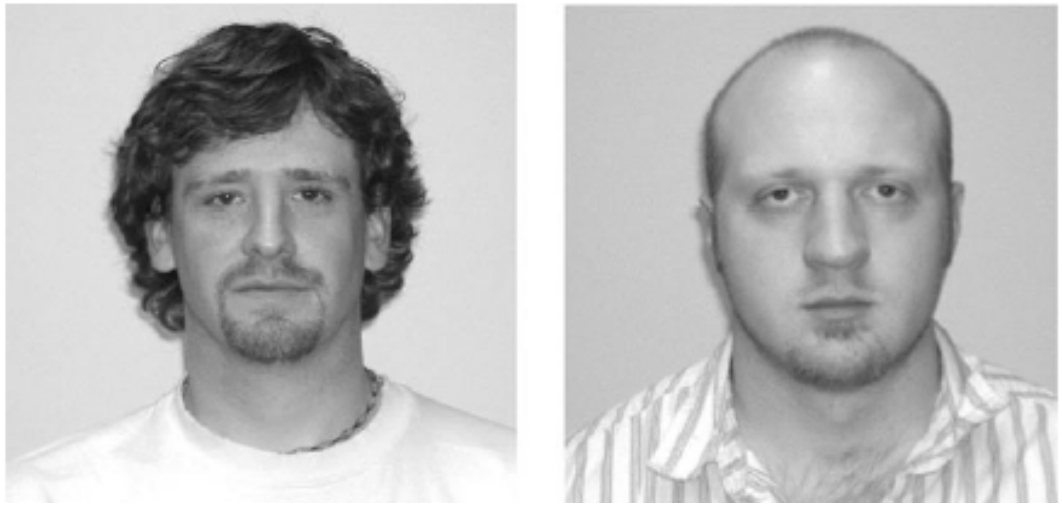

Who is Tim and who is Bob?

[Lea et. al., 07]

What's Behind the Names?

- Gender, age, and possible many other hidden factors.
- Can we guess the name of an unseen person?
  - Surprisingly, machine can do far better than chance.

Alejandra

Heather

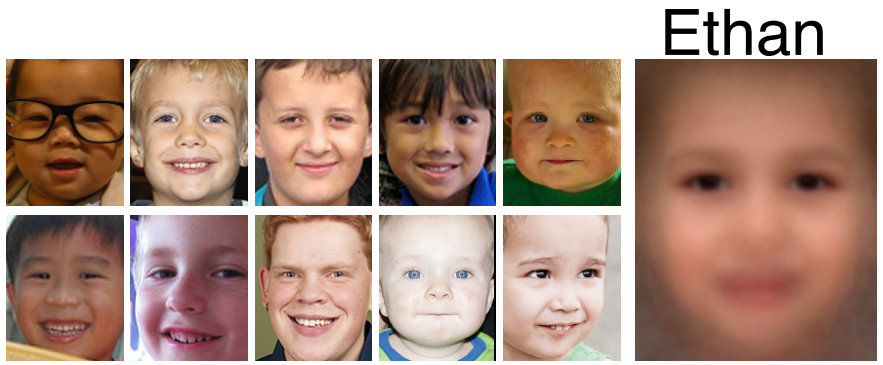

Ethan

System Overview

- The data:
  - 100 popular names in the United States.
  - 800 faces for each name.
  - The data is collected for FREE from image tags with careful filtering to reject noisy tags.
- Use Support Vector Machines to train a set of name classifiers.
- For an input face, generate a ranked list of names.

Multi-Feature SVMs

Ranked Name List

Lisa
Sarah
Christina
Jennifer
⋮

[Chen et. al., 2013]

Name Guessing

- Our findings:
  - Face appearance and first names are correlated.
  - Machine can guess names at an accuracy far better than random.

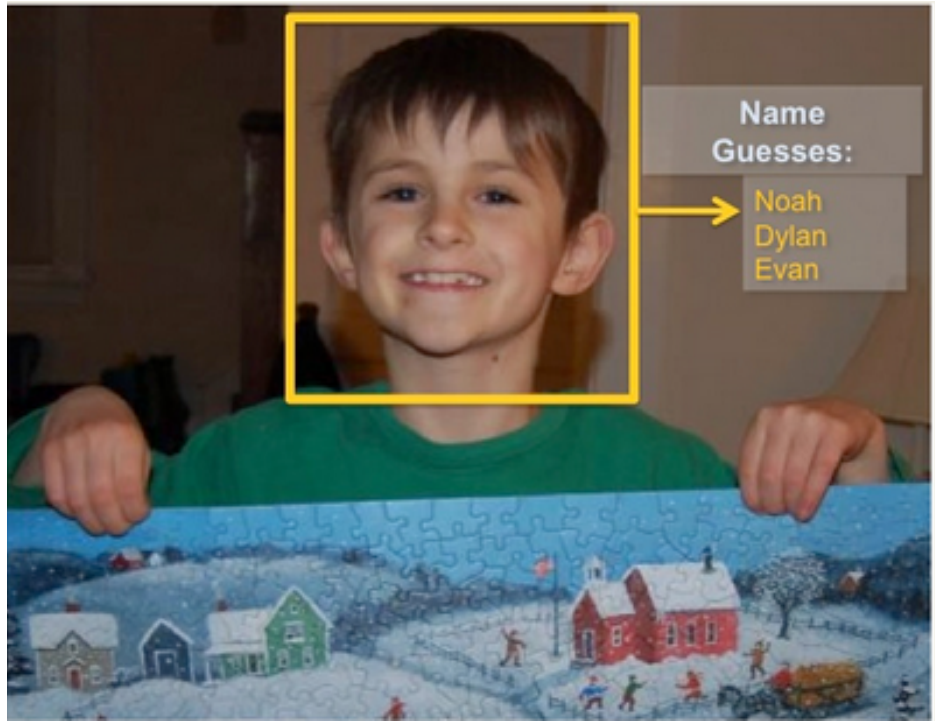

Name guessing among 100 names

Method	Top-1 Accuracy
Our approach	**4.17%**
Random guess	1.00%

Applications of Name Models

- For a given face, generate a ranked list of names that best fit the facial appearance.
- This ranked list of names provide important information about the person, especially gender and age.

Probably female.

Ranked name list
Erin
Emma
Emily
Lauren
Rebecca
⋮

Probably born after 1970.

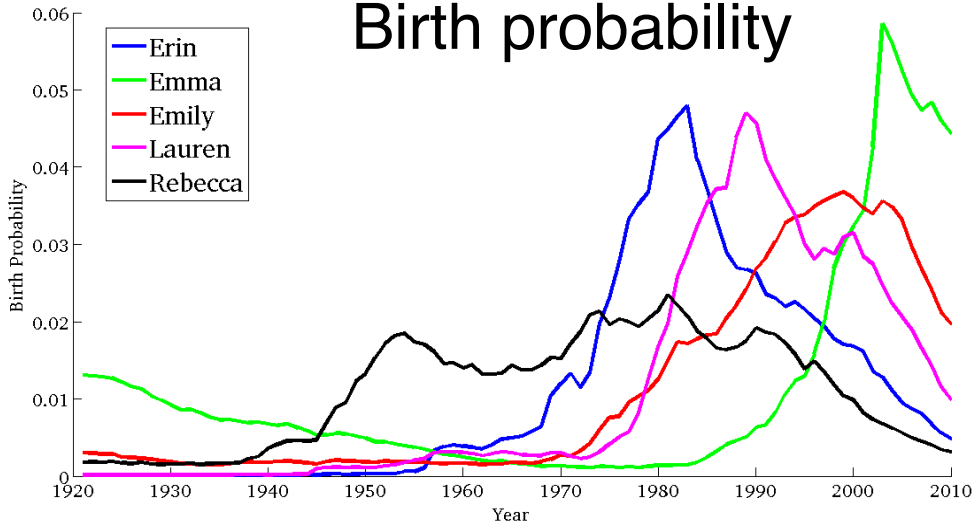

[source: US social security administration]

Gender Recognition From Faces

- We declare female if the majority of the top-5 names are female.
- Extremely simple but achieve state-of-the-art performance.
- Advantage: no gender labels required to perform gender recognition!

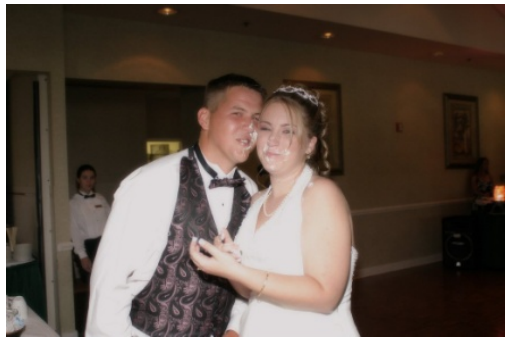

	Classification accuracy
Gender-from-names	**90.4%**
Kumar (PAMI 2011)	86.4%
Random prior	52.4%

Age Estimation From Names

- Given a face, generate a ranked list of 100 names.
- Combine the birth probability curves by weighted product.
- Achieves very competitive result without any age labels!

Ranked name list

Erin
Emma
Emily
Lauren
Rebecca
⋮

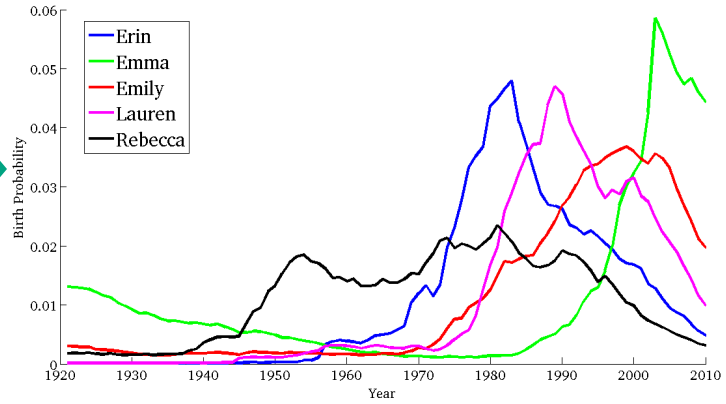

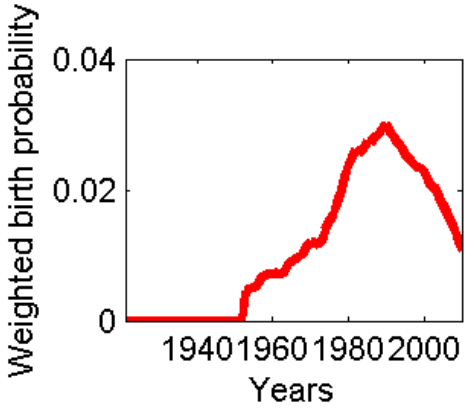

Combined birth probability

Predicted age category: 13-19

Thank you!

hchen2@stanford.edu

Additional Slides

Human Perception of Faces

- Human are good at recognizing faces, but not always!
- The Thatcher Effect

Is Face Recognition Dead?

Yes (on well aligned frontal faces)

Noise corrupted

Reconstructed by sparse representation

30%

50%

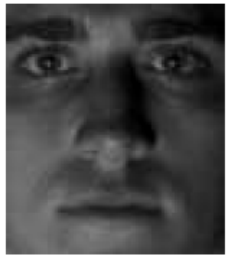

70%

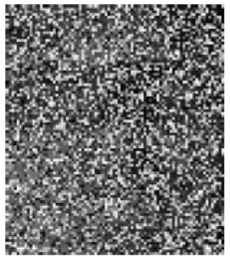

~100% accuracy up to 60% of pixel corruption!

[Wright et. al. 2009]

No (faces in the wild)

Pubfig Dataset

Anna Paquin

Audrey Tautou

Barack Obama

Clive Owen

Cristiano Ronaldo

Daniel Craig

78.65% face verification accuracy
[Kumar et al. 2009]