

Data-driven Methods: Faces



Portrait of Piotr Gibas
© Joaquin Rosales Gomez

15-463: Computational Photography
Alexei Efros, CMU, Fall 2008

The Power of Averaging



8-hour exposure



© Atta Kim

Figure-centric averages



Antonio Torralba & Aude Oliva (2002)

Averages: Hundreds of images containing a person are averaged to reveal regularities in the intensity patterns across all the images.

More by Jason Salavon



Homes for Sale



109 Homes for Sale,
Seattle/Tacoma



117 Homes for Sale,
Chicagoland



124 Homes for Sale, The 5
Boroughs



121 Homes for Sale,
LA/Orange County



114 Homes for Sale,
Dallas/Ft. Worth Metroplex



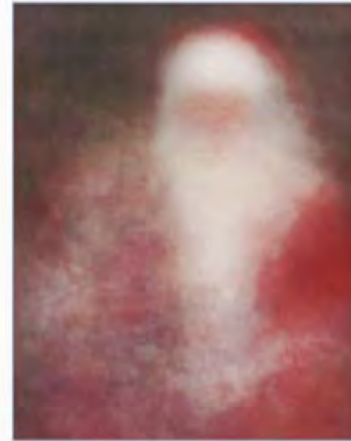
112 Homes for Sale,
Miami-Dade County

More at: <http://www.salavon.com/>

“100 Special Moments” by Jason Salavon



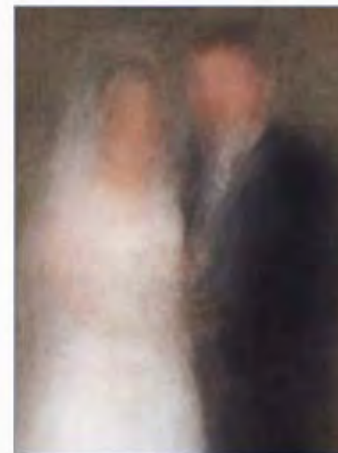
Little Leaguer



Kids with Santa



The Graduate



Newlyweds

Why
blurry?

Computing Means

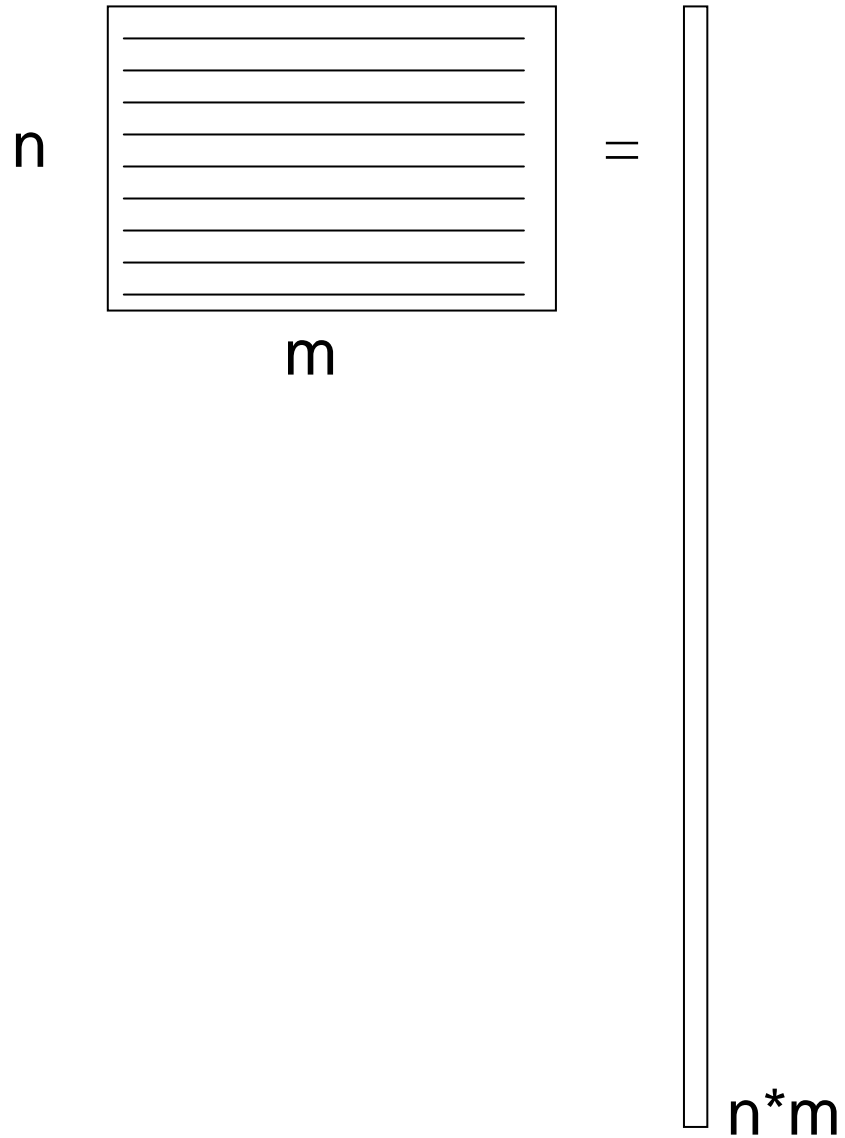
Two Requirements:

- Alignment of objects
- Objects must span a subspace

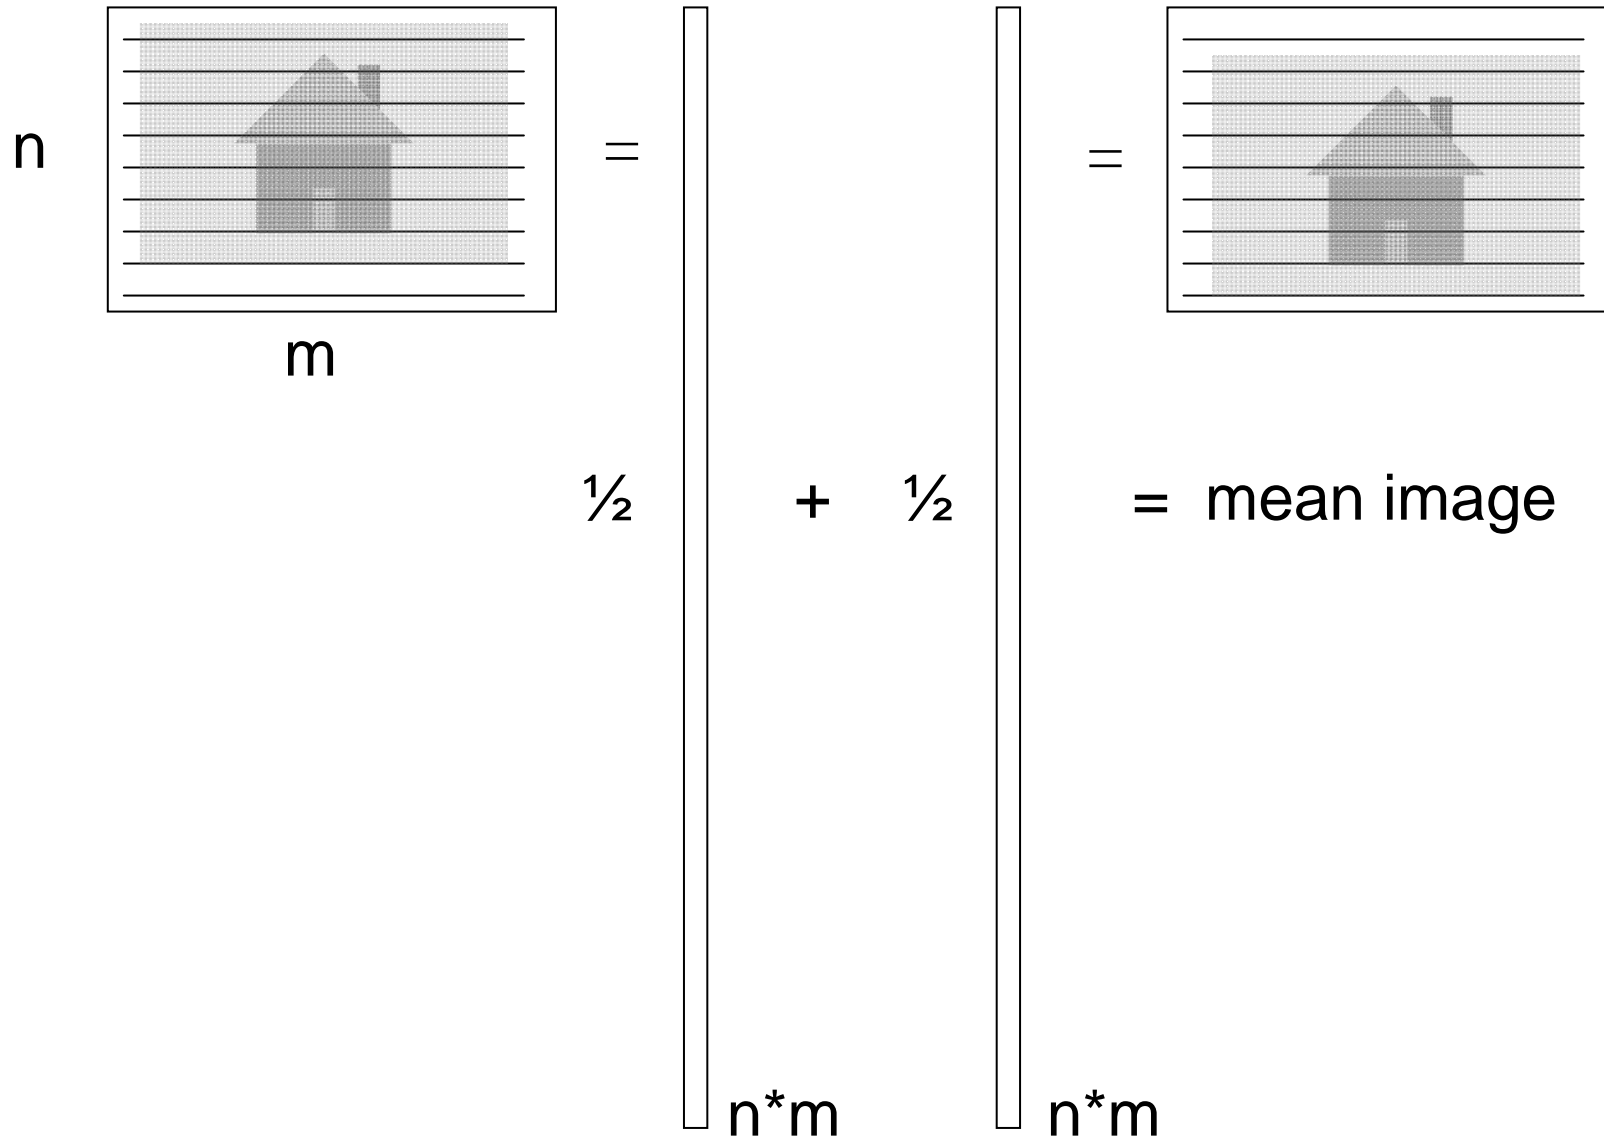
Useful concepts:

- Subpopulation means
- Deviations from the mean

Images as Vectors



Vector Mean: Importance of Alignment




How to align faces?




<http://www2.imm.dtu.dk/~aam/datasets/datasets.html>

Shape Vector



Provides alignment!

=





43

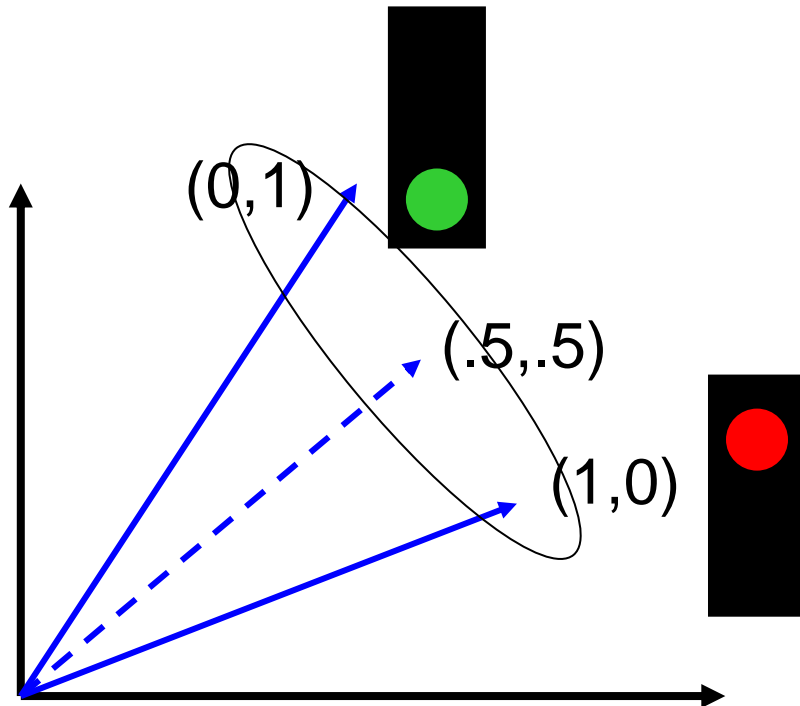
Average Face




1. Warp to mean shape
2. Average pixels



Objects must span a subspace



Example




mean

Does not span a subspace


Subpopulation means

Examples:

- Happy faces
- Young faces
- Asian faces
- Etc.
- Sunny days
- Rainy days
- Etc.
- Etc.



Average female



Average male

Deviations from the mean



Image X



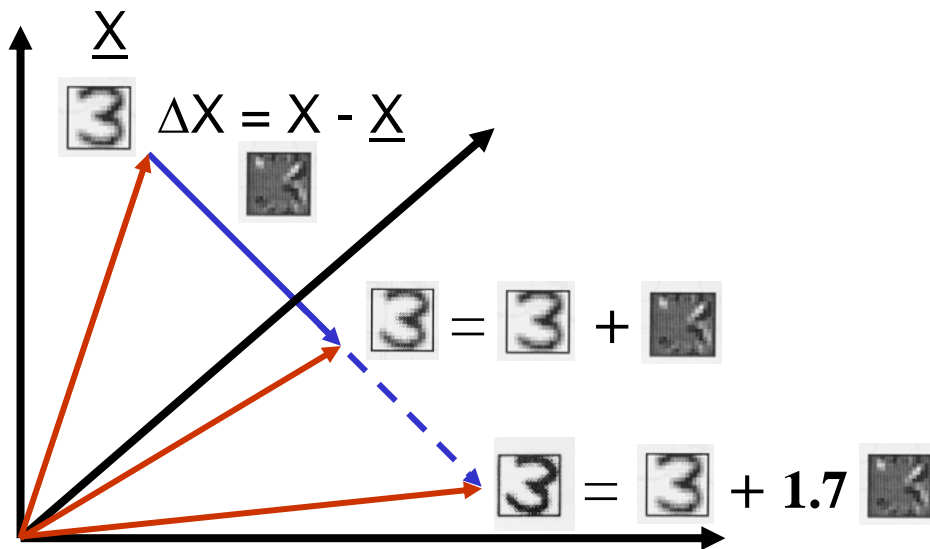
Mean \underline{X}

=



$$\Delta X = X - \underline{X}$$

Deviations from the mean



Manipulating Facial Appearance through Shape and Color


Duncan A. Rowland and David I. Perrett

St Andrews University

IEEE CG&A, September 1995

Face Modeling

Compute *average* faces
(color and shape)




Compute *deviations*
between male and
female (vector and color
differences)

Changing gender

Deform shape and/or color of an input face in the direction of “more female”

original



shape

color

both

Enhancing gender



more same **original** androgynous more opposite

Changing age

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

original



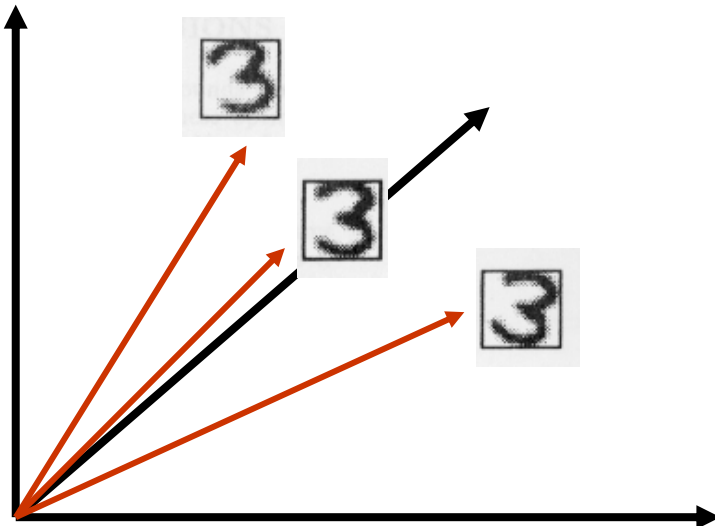
shape

color

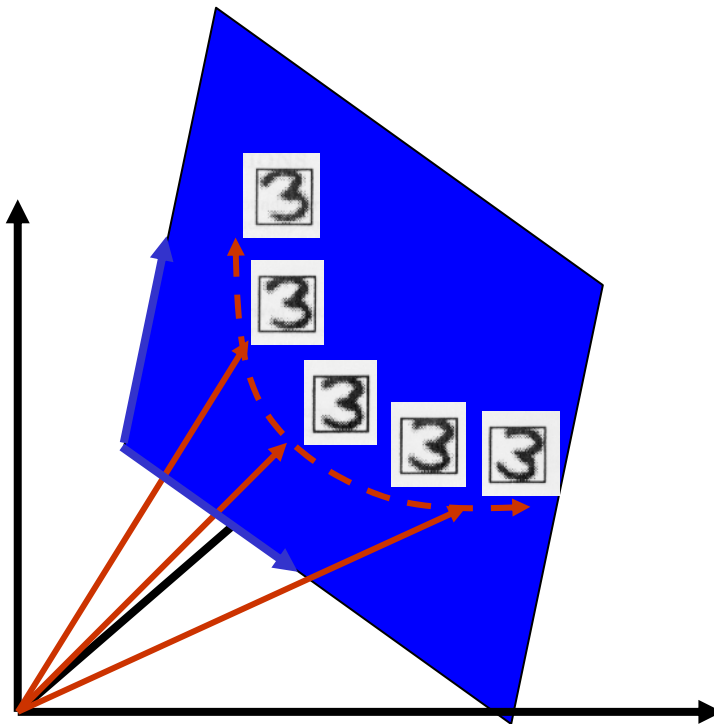


both

Back to the Subspace



Linear Subspace: convex combinations



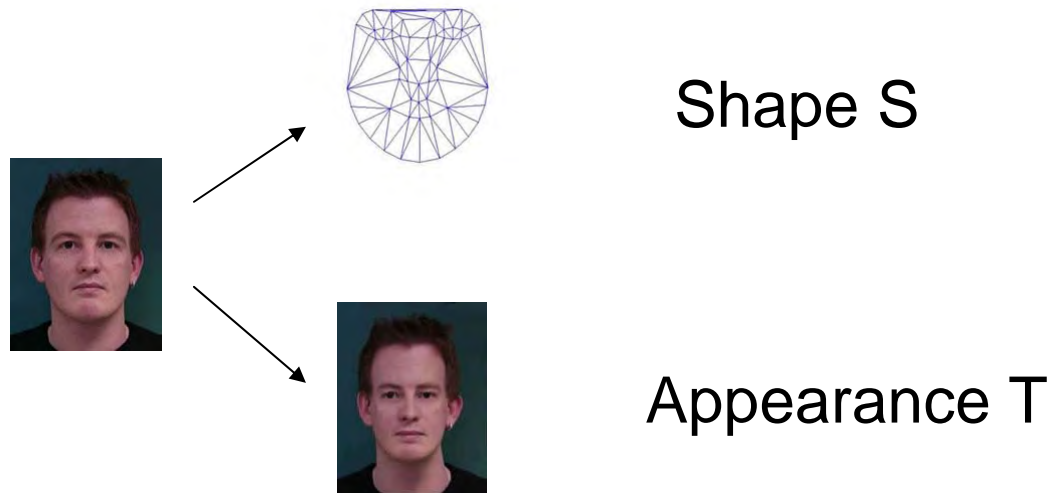
Any new image X can be obtained as weighted sum of stored “basis” images.

$$X = \sum_{i=1}^m a_i X_i$$

Our old friend, change of basis!
What are the new coordinates of X ?

The Morphable Face Model

The actual structure of a face is captured in the shape vector $\mathbf{S} = (x_1, y_1, x_2, \dots, y_n)^T$, containing the (x, y) coordinates of the n vertices of a face, and the appearance (texture) vector $\mathbf{T} = (R_1, G_1, B_1, R_2, \dots, G_n, B_n)^T$, containing the color values of the mean-warped face image.



The Morphable face model

Again, assuming that we have m such vector pairs in full correspondence, we can form new shapes \mathbf{S}_{model} and new appearances \mathbf{T}_{model} as:

$$\mathbf{S}_{model} = \sum_{i=1}^m a_i \mathbf{S}_i \quad \mathbf{T}_{model} = \sum_{i=1}^m b_i \mathbf{T}_i$$

$$s = \alpha_1 \cdot \text{img}_1 + \alpha_2 \cdot \text{img}_2 + \alpha_3 \cdot \text{img}_3 + \alpha_4 \cdot \text{img}_4 + \dots = \mathbf{S} \cdot \mathbf{a}$$

$$t = \beta_1 \cdot \text{img}_1 + \beta_2 \cdot \text{img}_2 + \beta_3 \cdot \text{img}_3 + \beta_4 \cdot \text{img}_4 + \dots = \mathbf{T} \cdot \mathbf{b}$$

If number of basis faces m is large enough to span the face subspace then:

Any new face can be represented as a pair of vectors

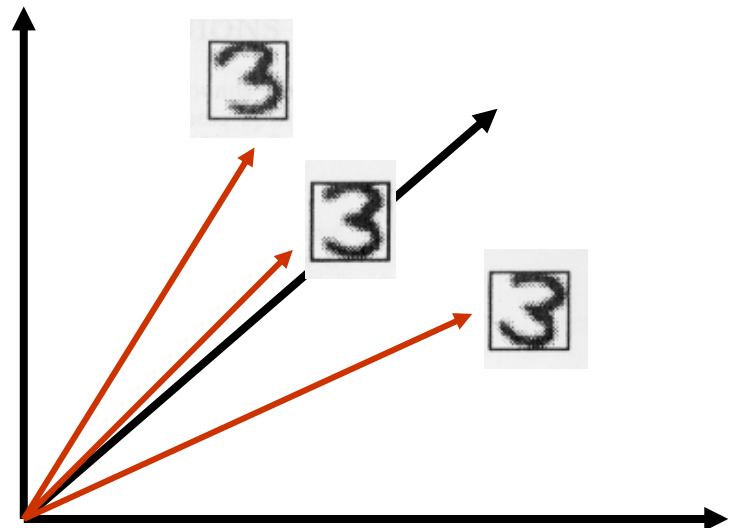
$$(\alpha_1, \alpha_2, \dots, \alpha_m)^T \text{ and } (\beta_1, \beta_2, \dots, \beta_m)^T !$$

Issues:

1. How many basis images is enough?
2. Which ones should they be?
3. What if some variations are more important than others?
 - E.g. corners of mouth carry much more information than haircut

Need a way to obtain basis images automatically, in order of importance!

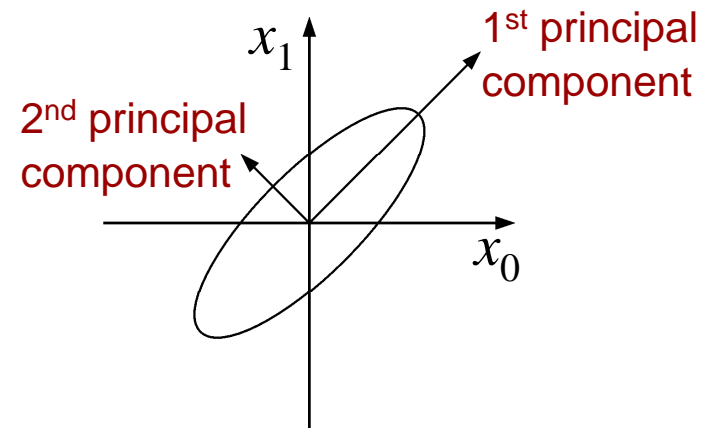
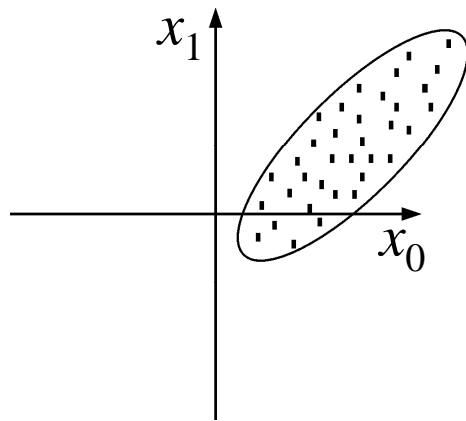
But what's important?



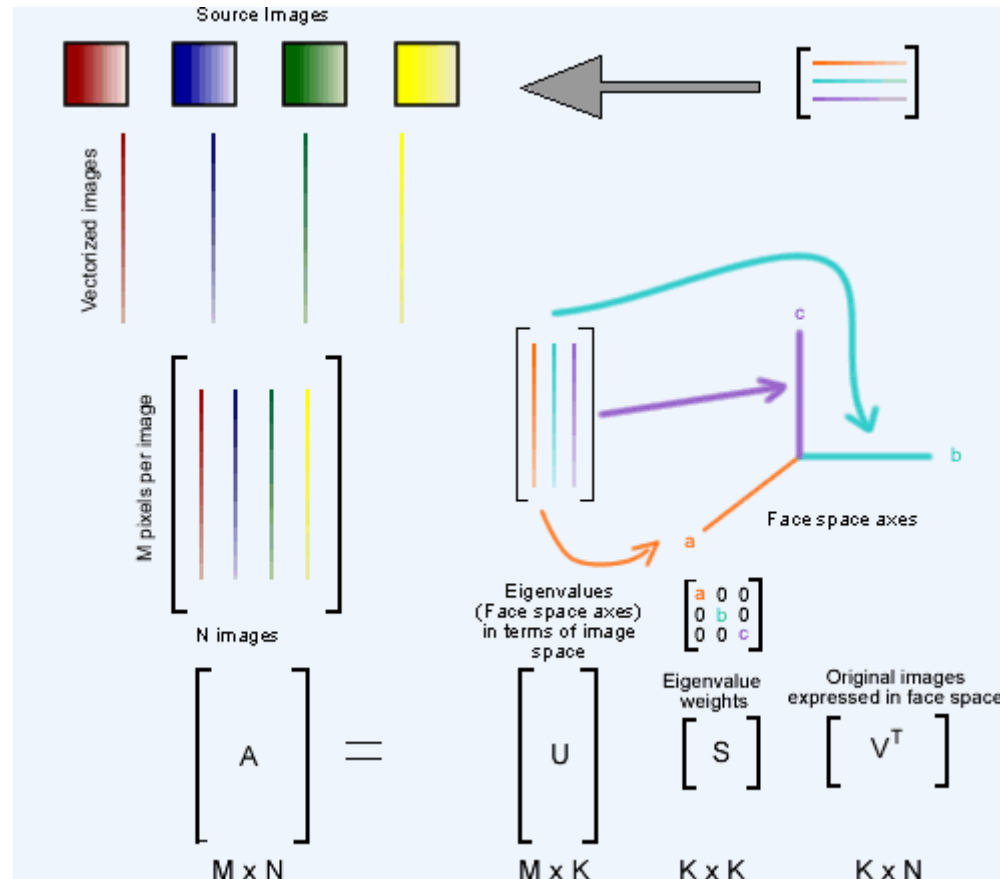
Principal Component Analysis

Given a point set $\{\vec{p}_j\}_{j=1\dots P}$, in an M -dim space, PCA finds a basis such that

- coefficients of the point set in that basis are uncorrelated
- first $r < M$ basis vectors provide an approximate basis that minimizes the mean-squared-error (MSE) in the approximation (over all bases with dimension r)



PCA via Singular Value Decomposition



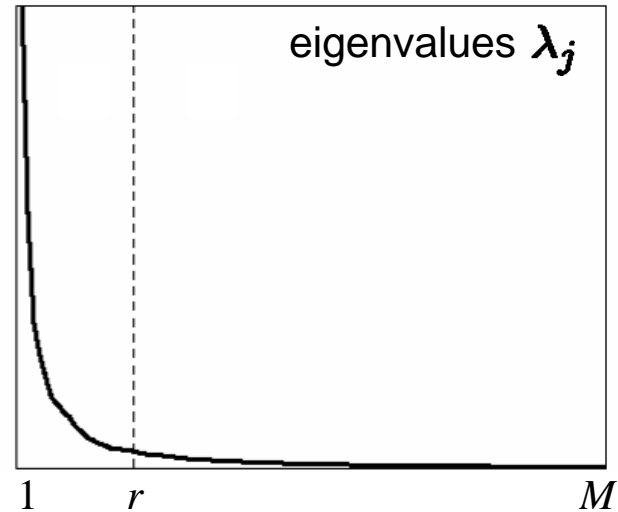
$$[u,s,v] = \text{svd}(A);$$

Principal Component Analysis

Choosing subspace dimension

r :

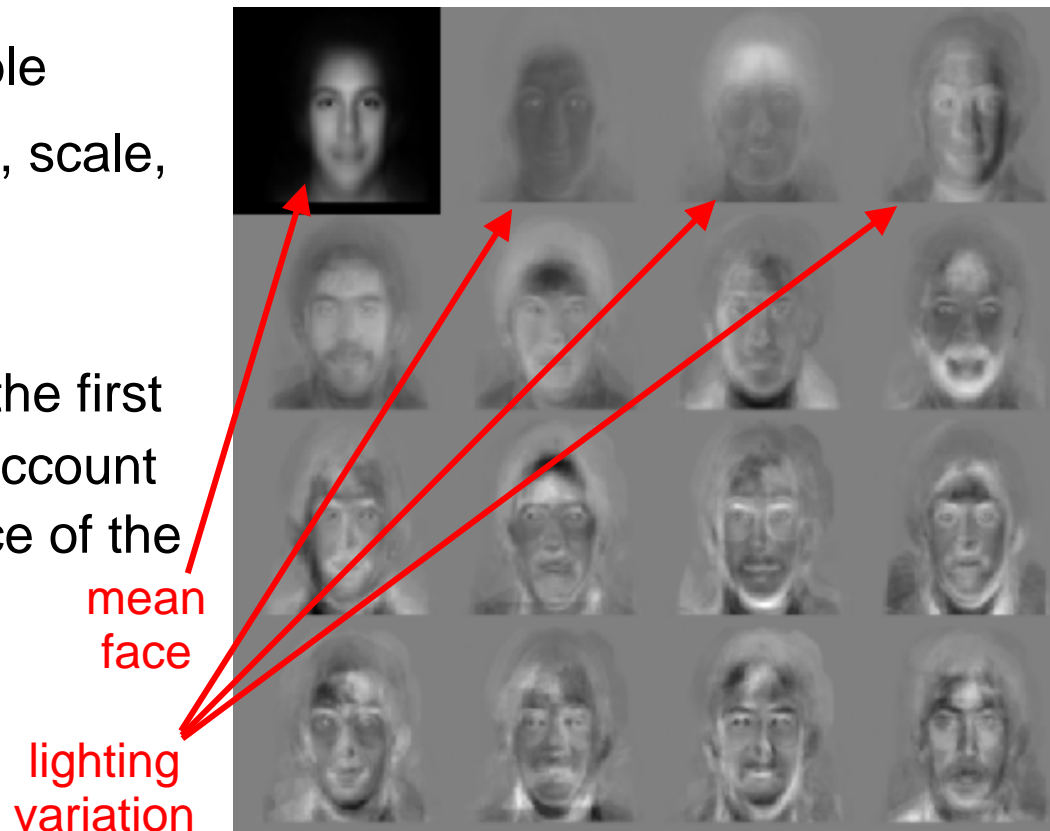
- look at decay of the eigenvalues as a function of r
- Larger r means lower expected error in the subspace data approximation




EigenFaces

First popular use of PCA on images was for modeling and recognition of faces [Kirby and Sirovich, 1990, Turk and Pentland, 1991]




- Collect a face ensemble
- Normalize for contrast, scale, & orientation.
- Remove backgrounds
- Apply PCA & choose the first N eigen-images that account for most of the variance of the data.



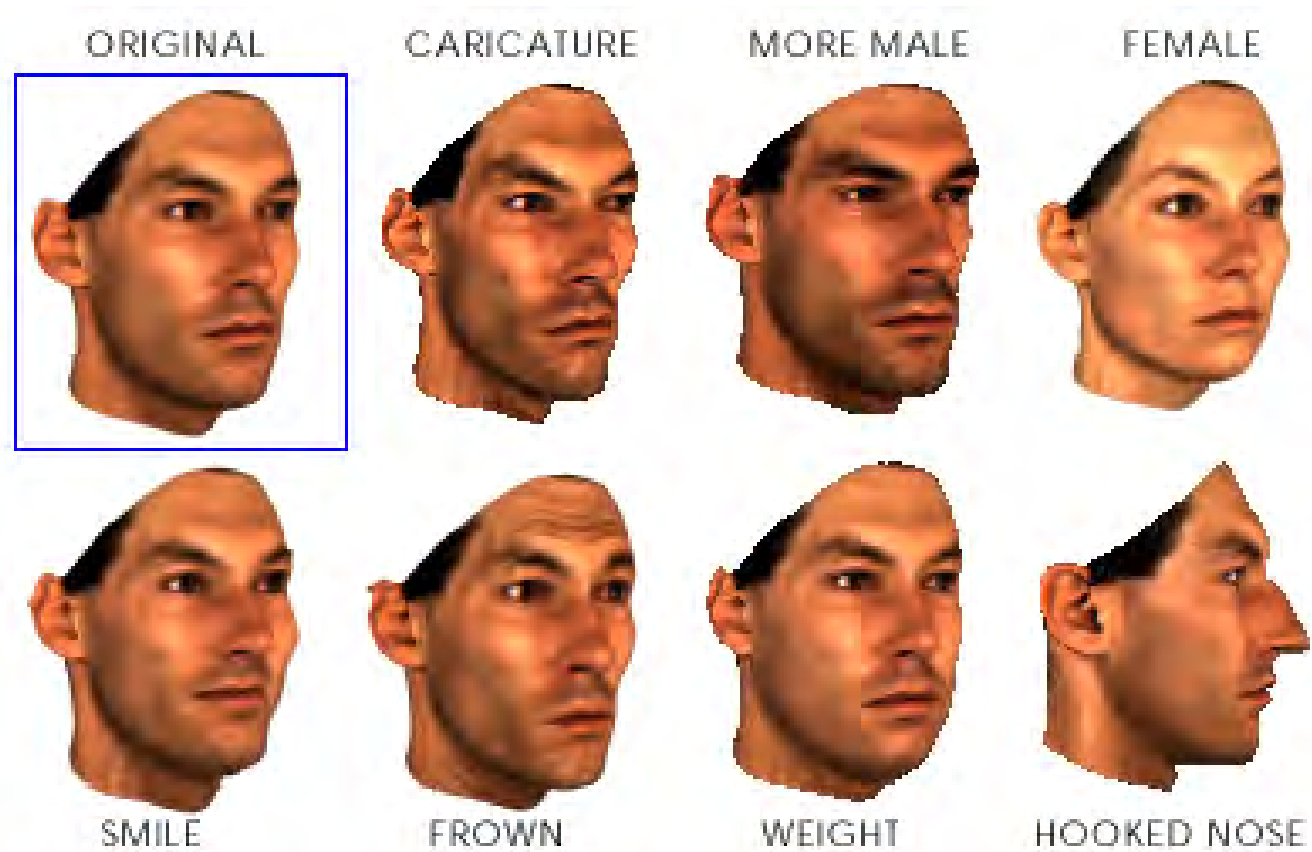
First 3 Shape Basis



Mean appearance



Using 3D Geometry: Blinz & Vetter, 1999



show SIGGRAPH video