Introduction to Computer Vision

Michael J. Black

Oct 2009

PCA and applications

Probability and classification

Admin

Assignment 1 grades back shortly.

Assignment 2, parts 1 and 2 due **October 22 at 11:00am**

Do part 1 this weekend. Don't wait.

Goals

- Applications of PCA.
- Start probability and classification
  - Everything you need for parts 1 and 2

- Monday: finish probability and classification

ImageNet

- May be a useful resource for final projects.
  - http://www.image-net.org/index

“Mocap”

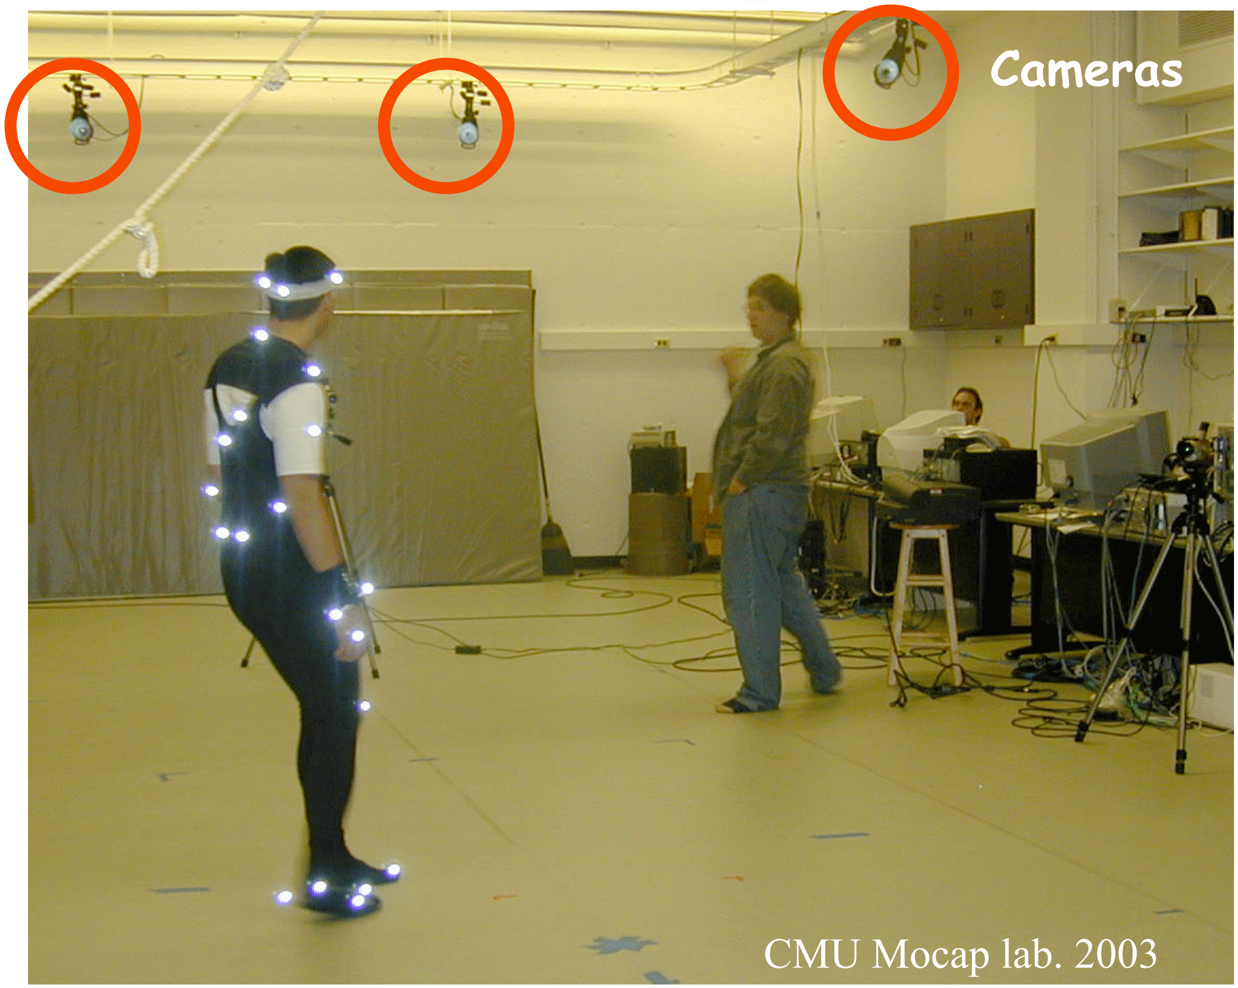

Kinematic Tree

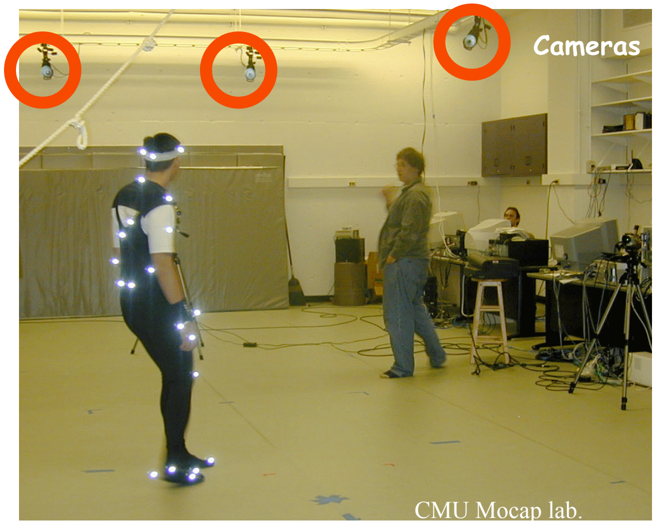

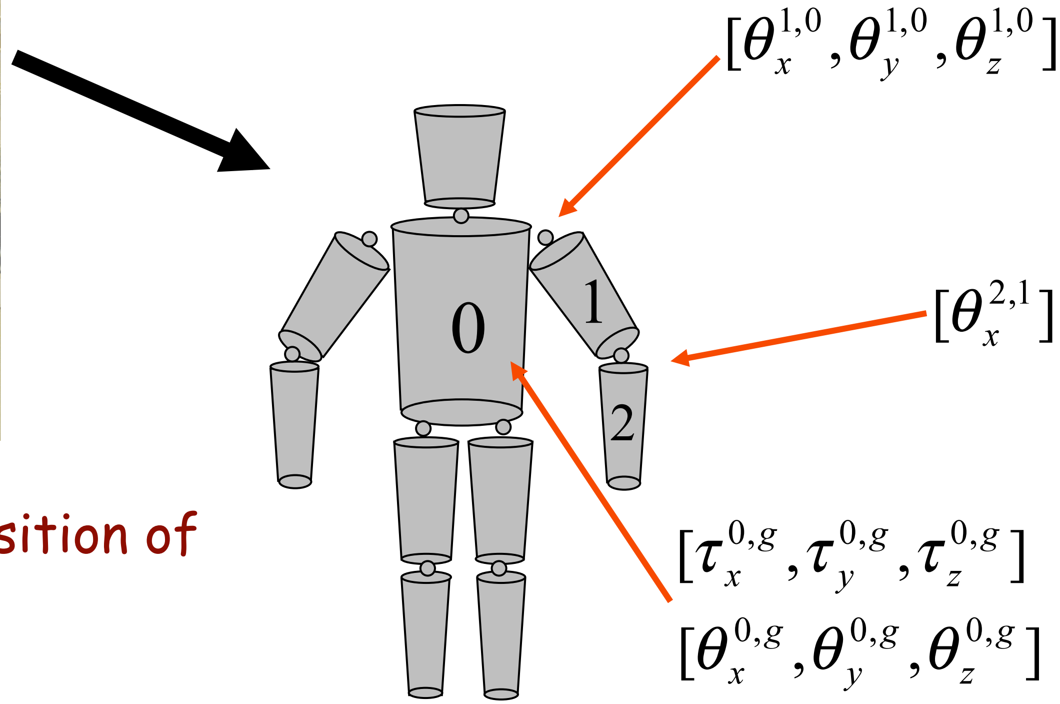

Triangulate to find 3D position of markers.

Fit a human body model.

Compute joint angles.

Modeling Cyclic Motion

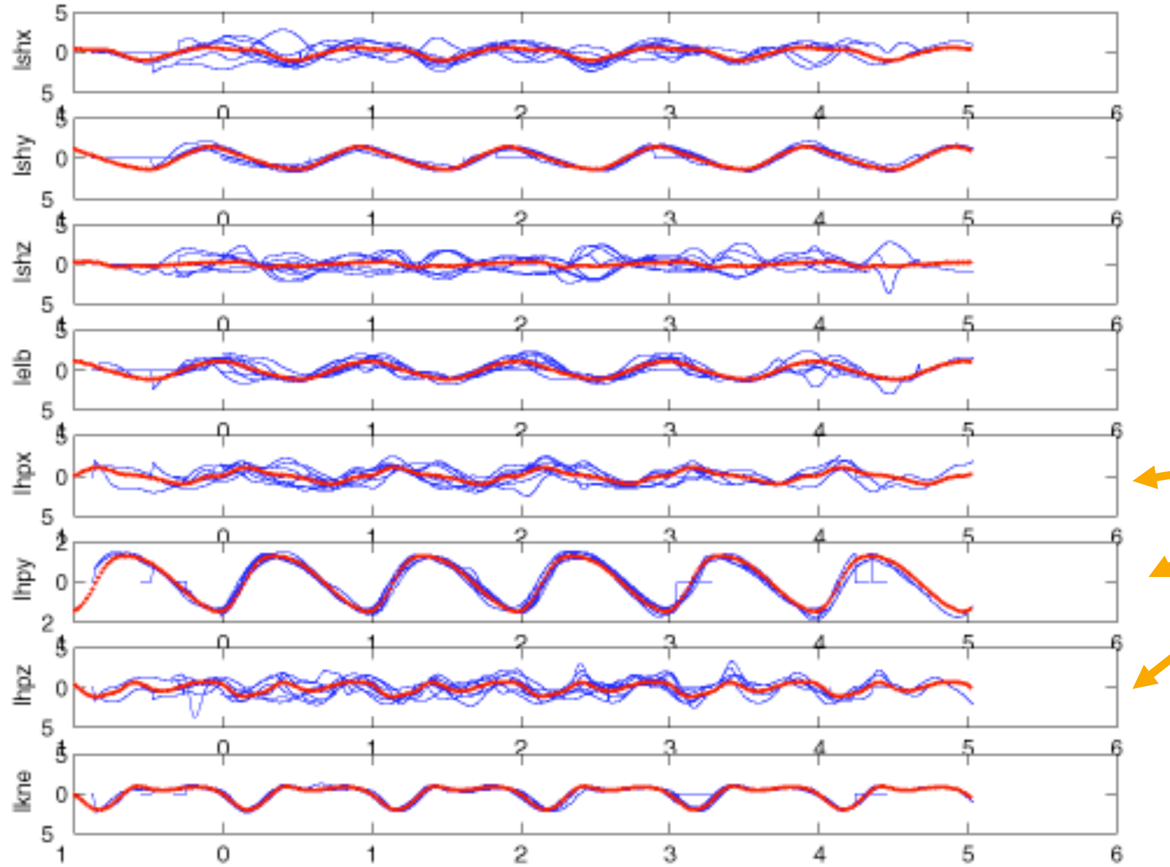

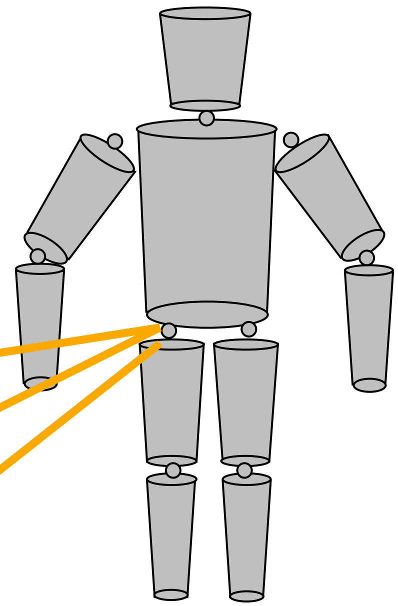

Modeling Cyclic Motion

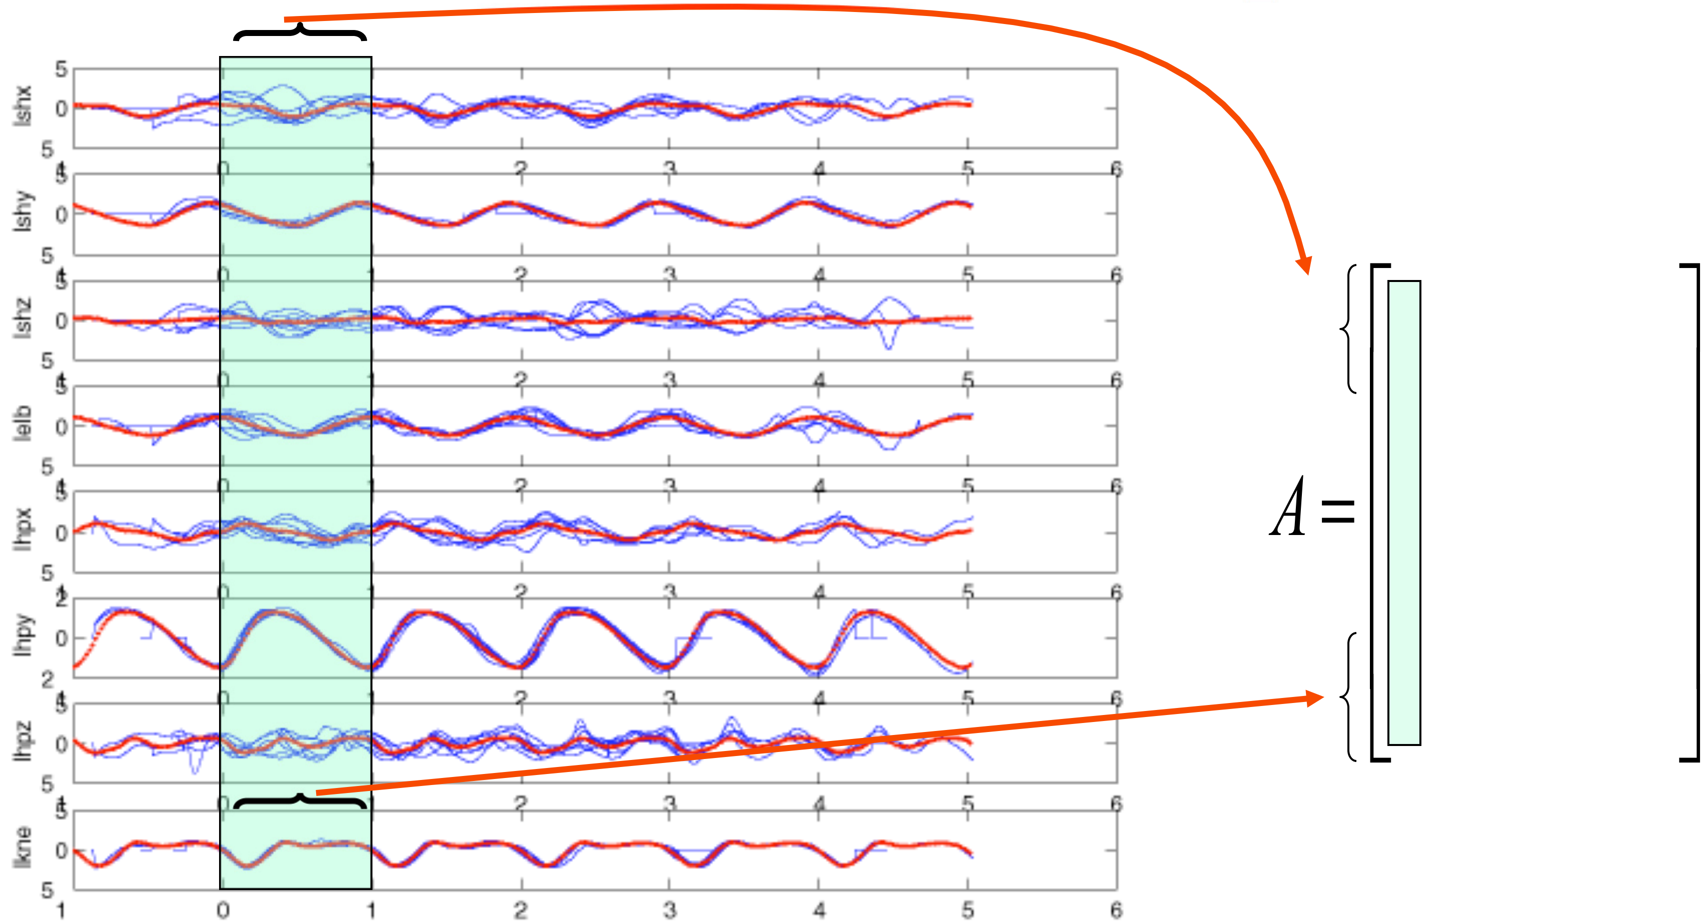

Modeling Cyclic Motion

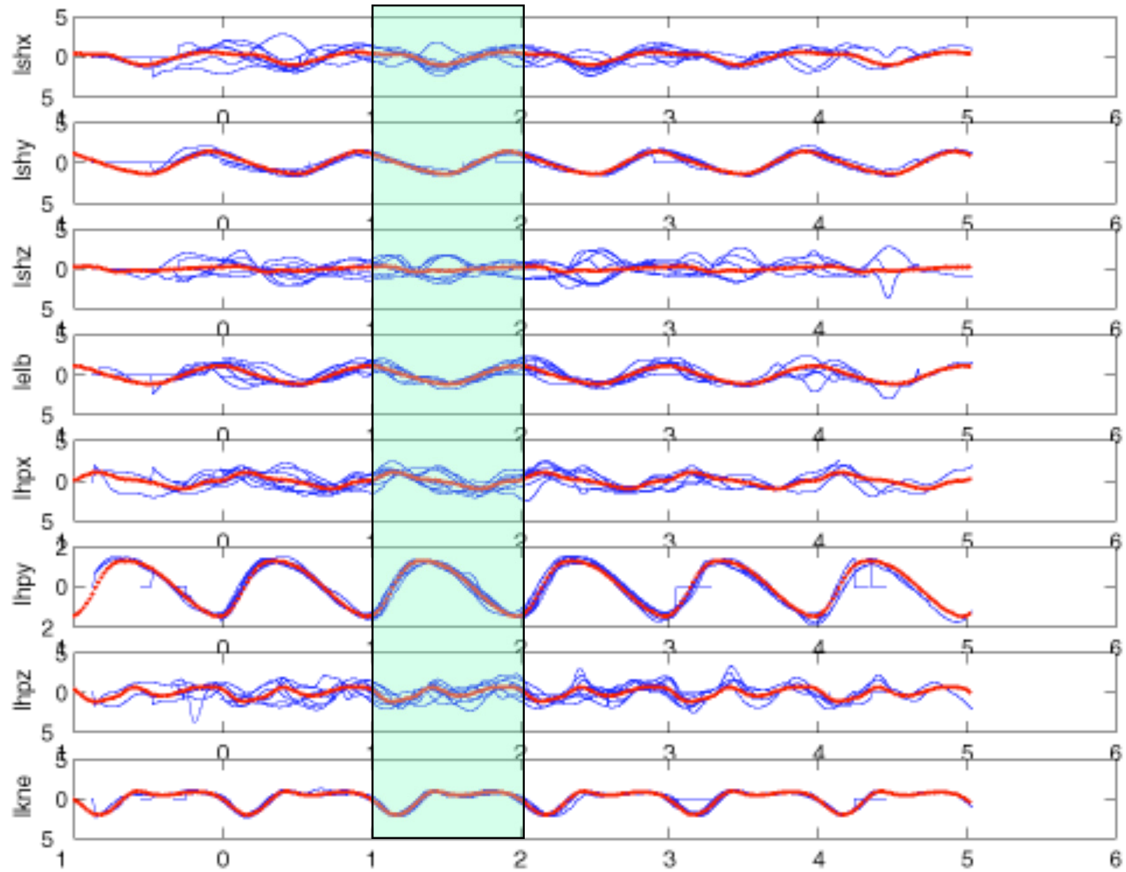

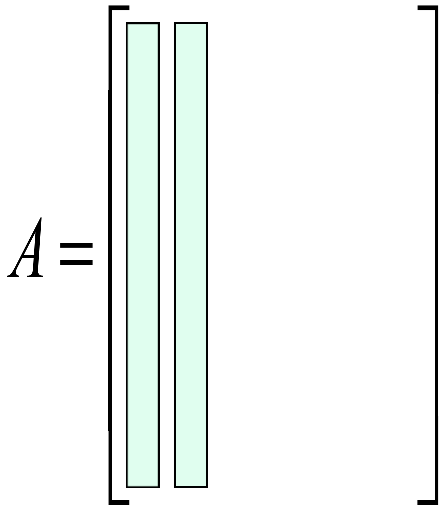

Modeling Cyclic Motion

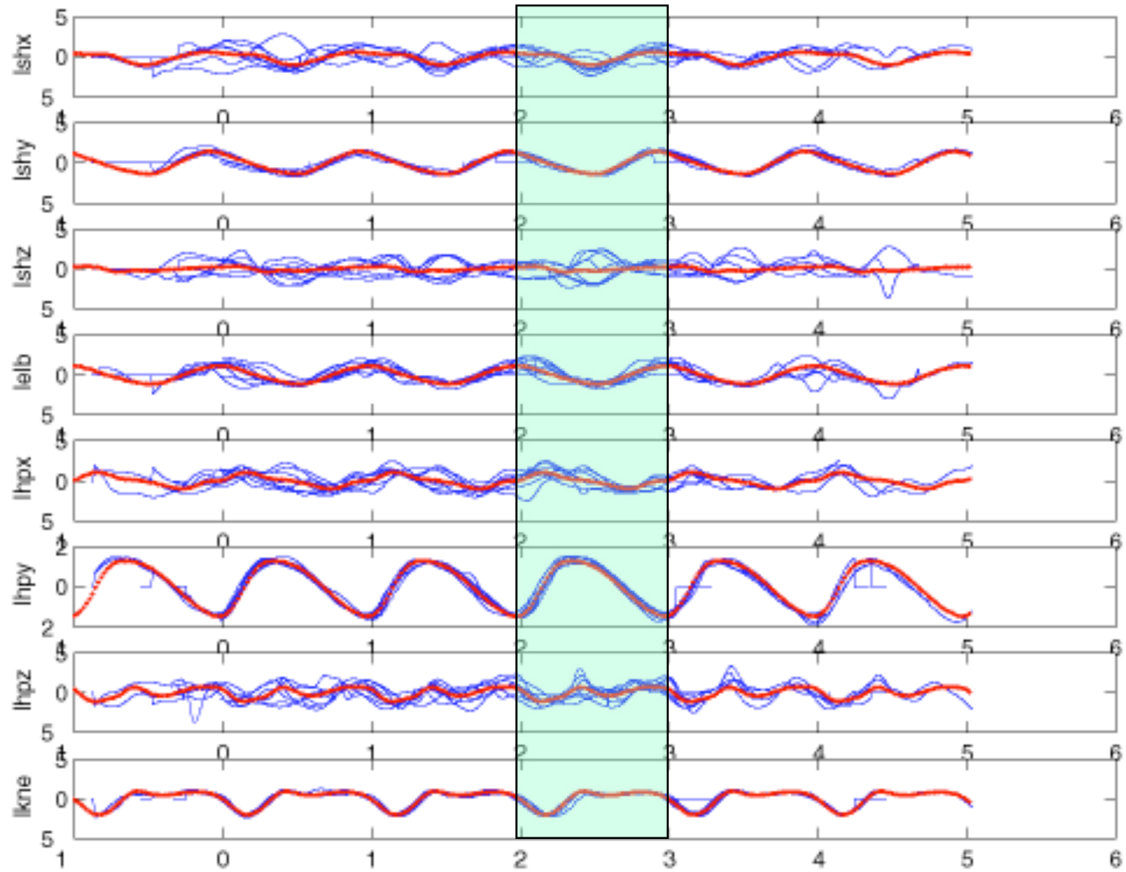

$$A = \begin{bmatrix} \; & \; & \; & \cdots \end{bmatrix}$$

EigenWalking

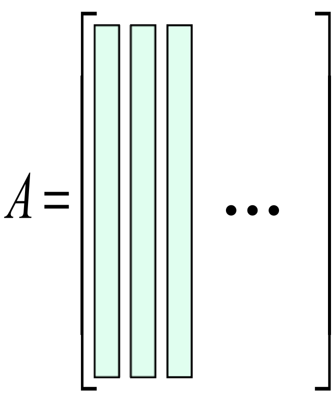

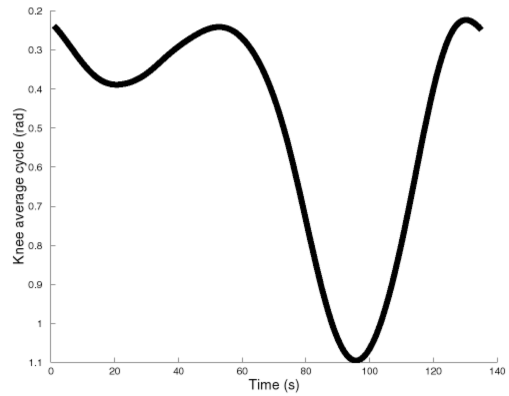

Mean knee motion

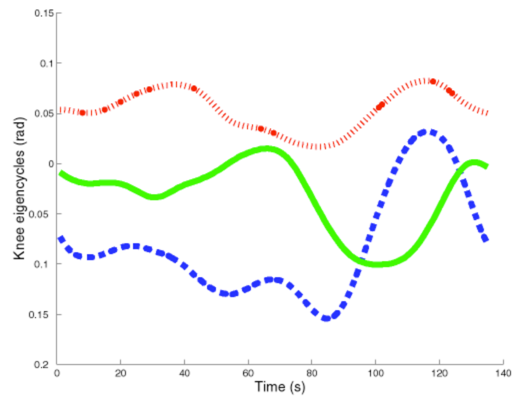

First 3 principal components of knee motion

Action-Specific Model

The joint angles at time *t* are a linear combination of the basis motions evaluated at *phase* ***y***

$$\varphi_t = \tilde{\mu}(\psi_t) + \sum_{k=1}^{q} c_k v_k(\psi_t)$$

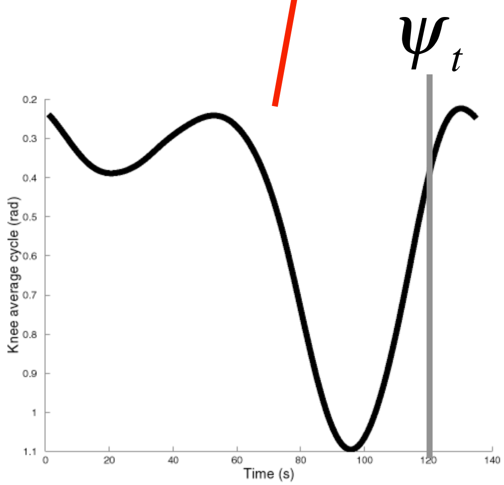

Mean curve

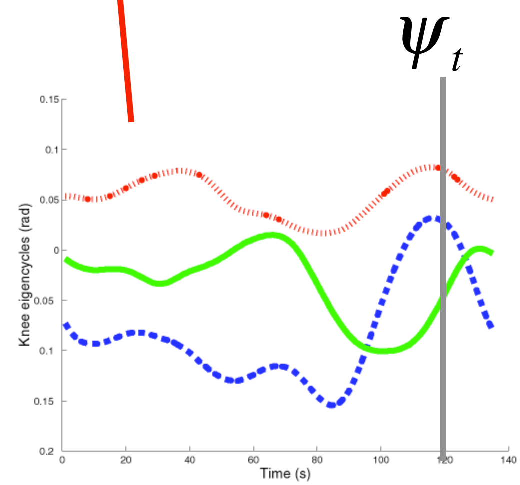

Basis curves

Learned Walking Model

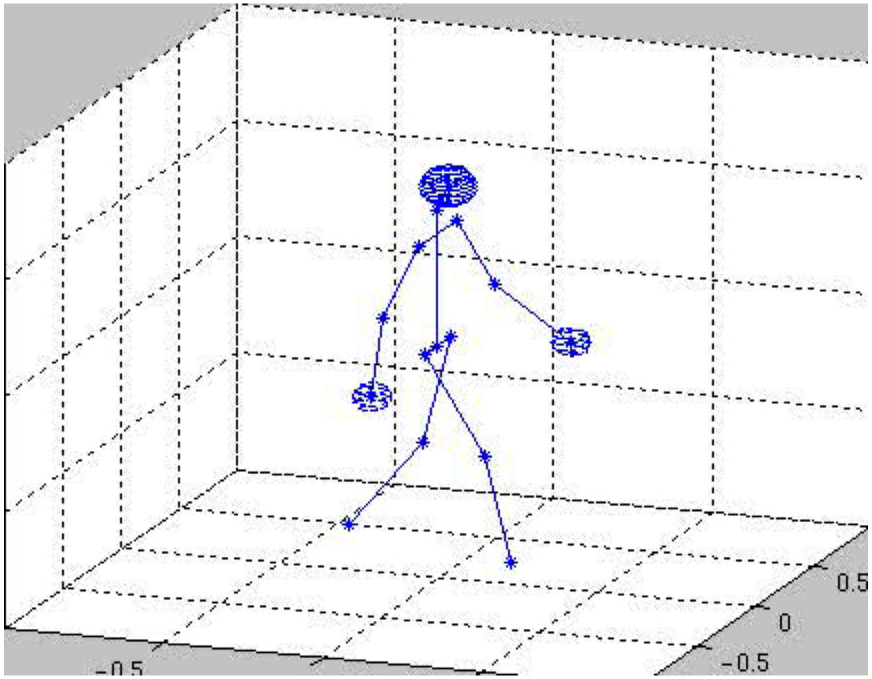

* *mean walker*

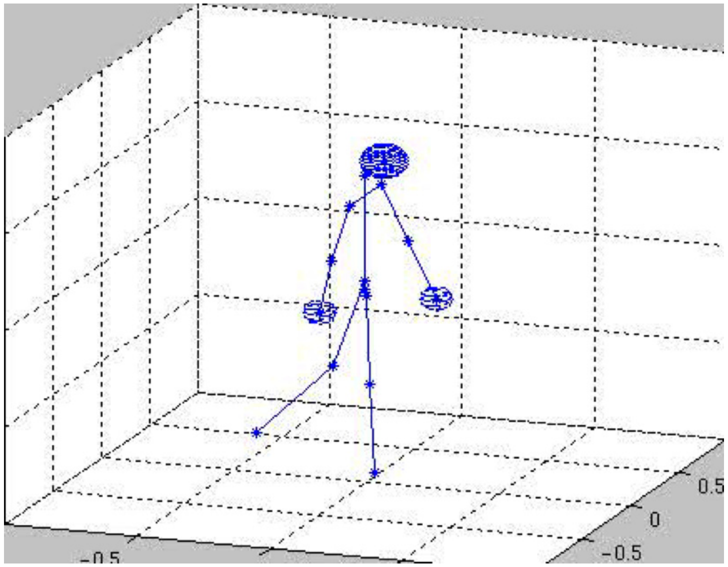

1st Eigenvector

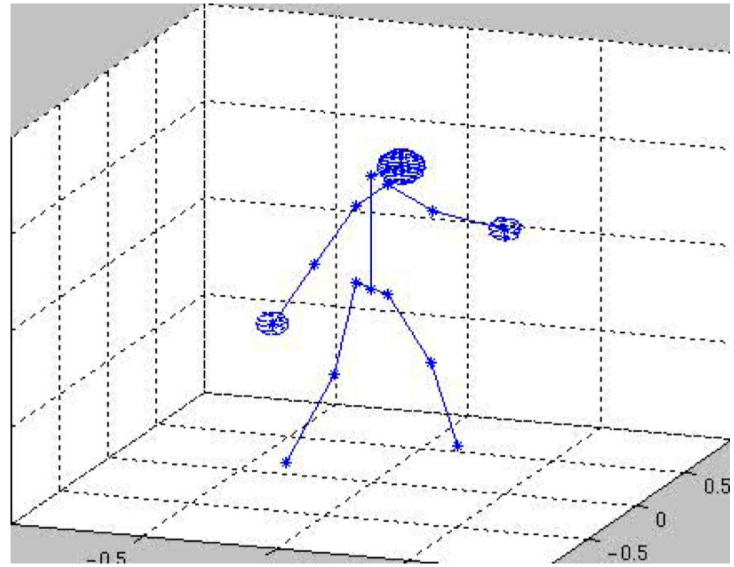

2nd Eigenvector

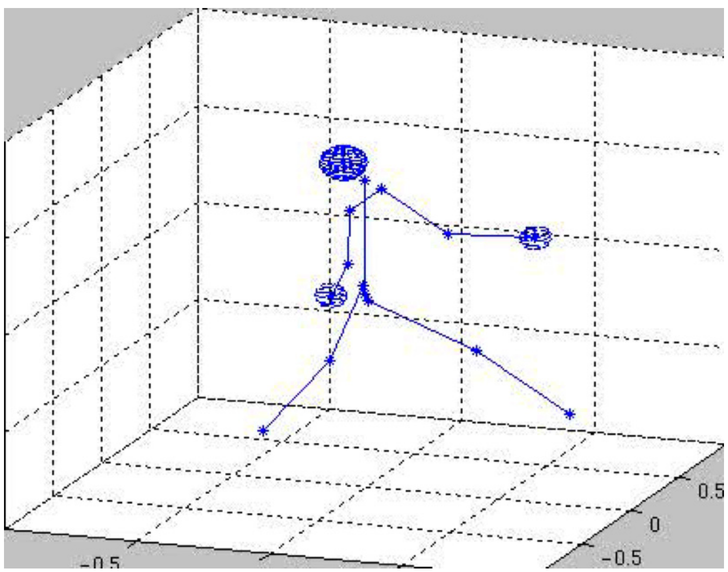

3rd Eigenvector

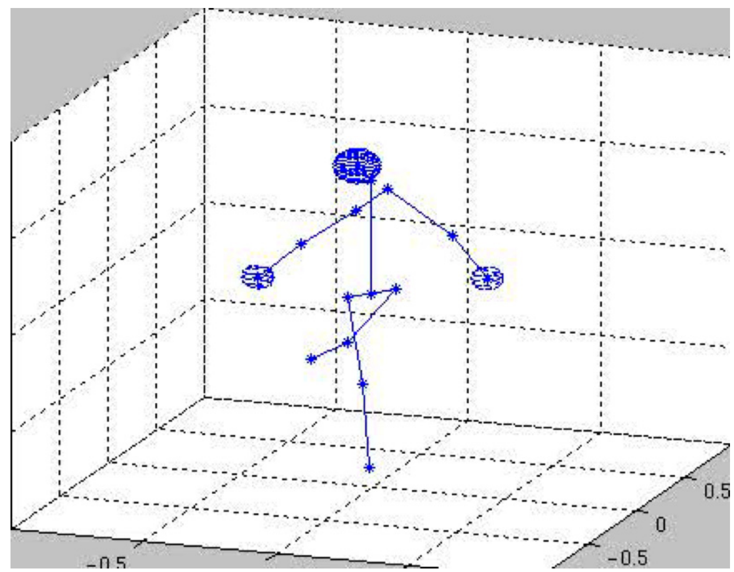

4th Eigenvector

Learned Walking Model

Build a probabilistic model and draw a sample from it.

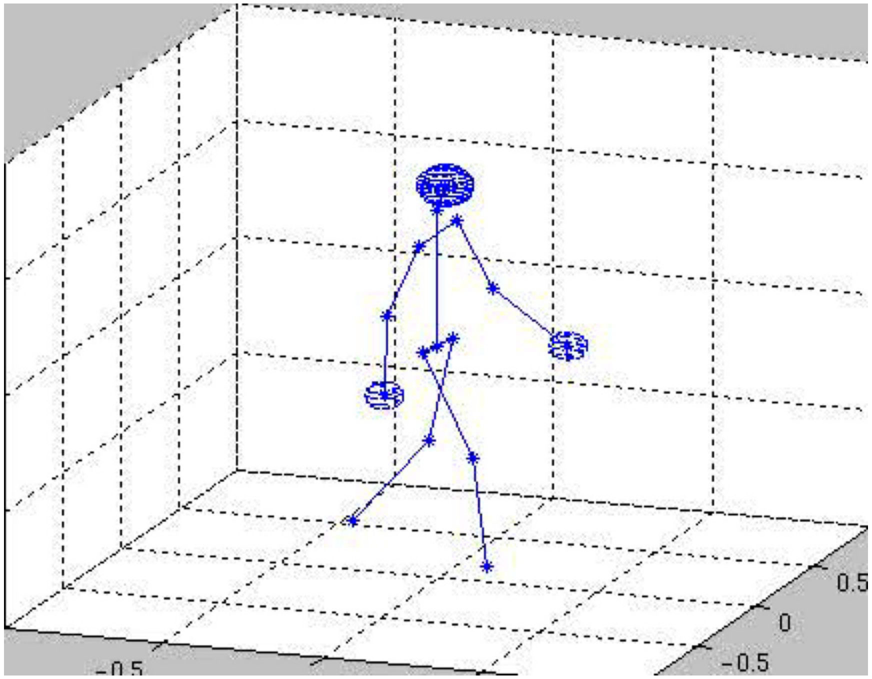

* *sample with small* ε

Learned Walking Model

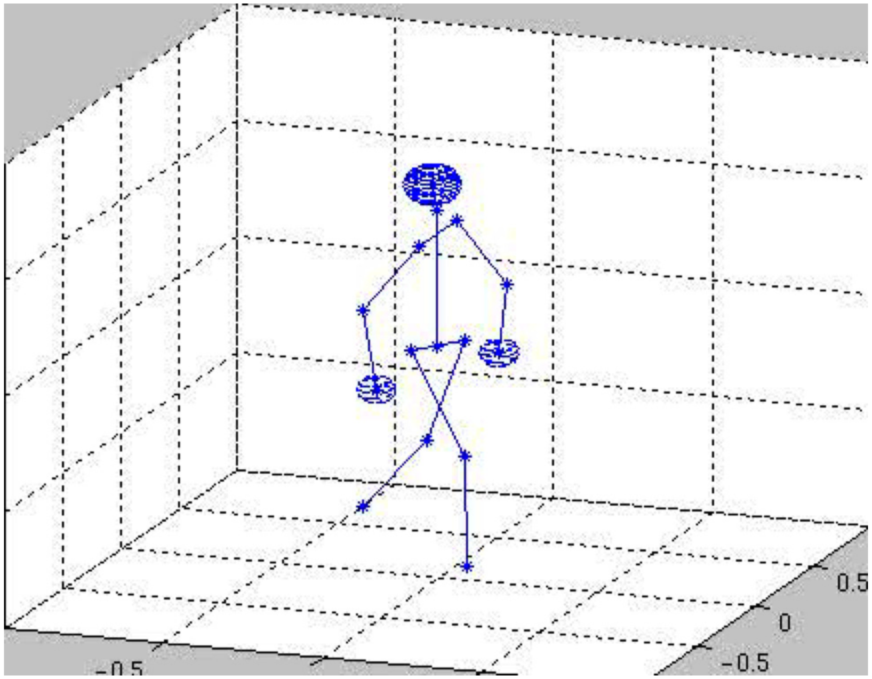

* *sample with moderate* ε

Learned Walking Model

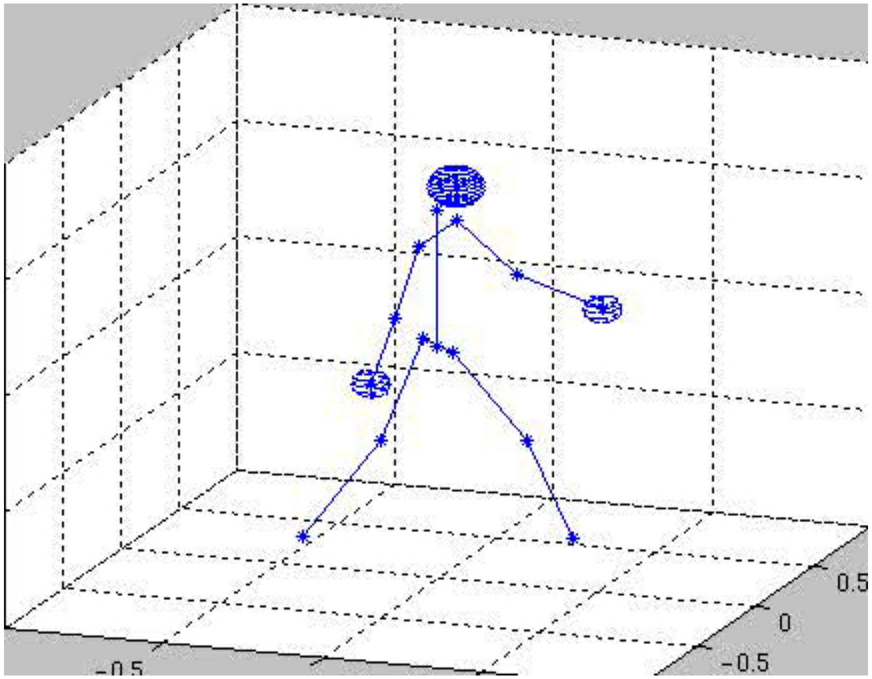

* *sample with large* ε

Learned Walking Model

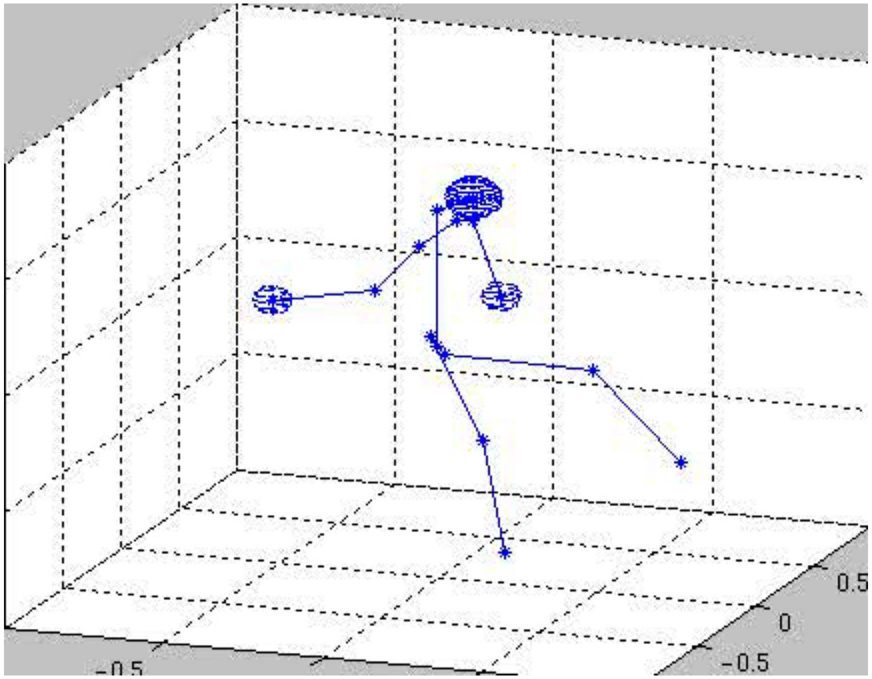

* *sample with very large* ε

Appearance Manifolds

Many objects do not have convex subspaces when one considers different poses and lighting variations.

Fleet & Szeliski

Appearance Manifolds

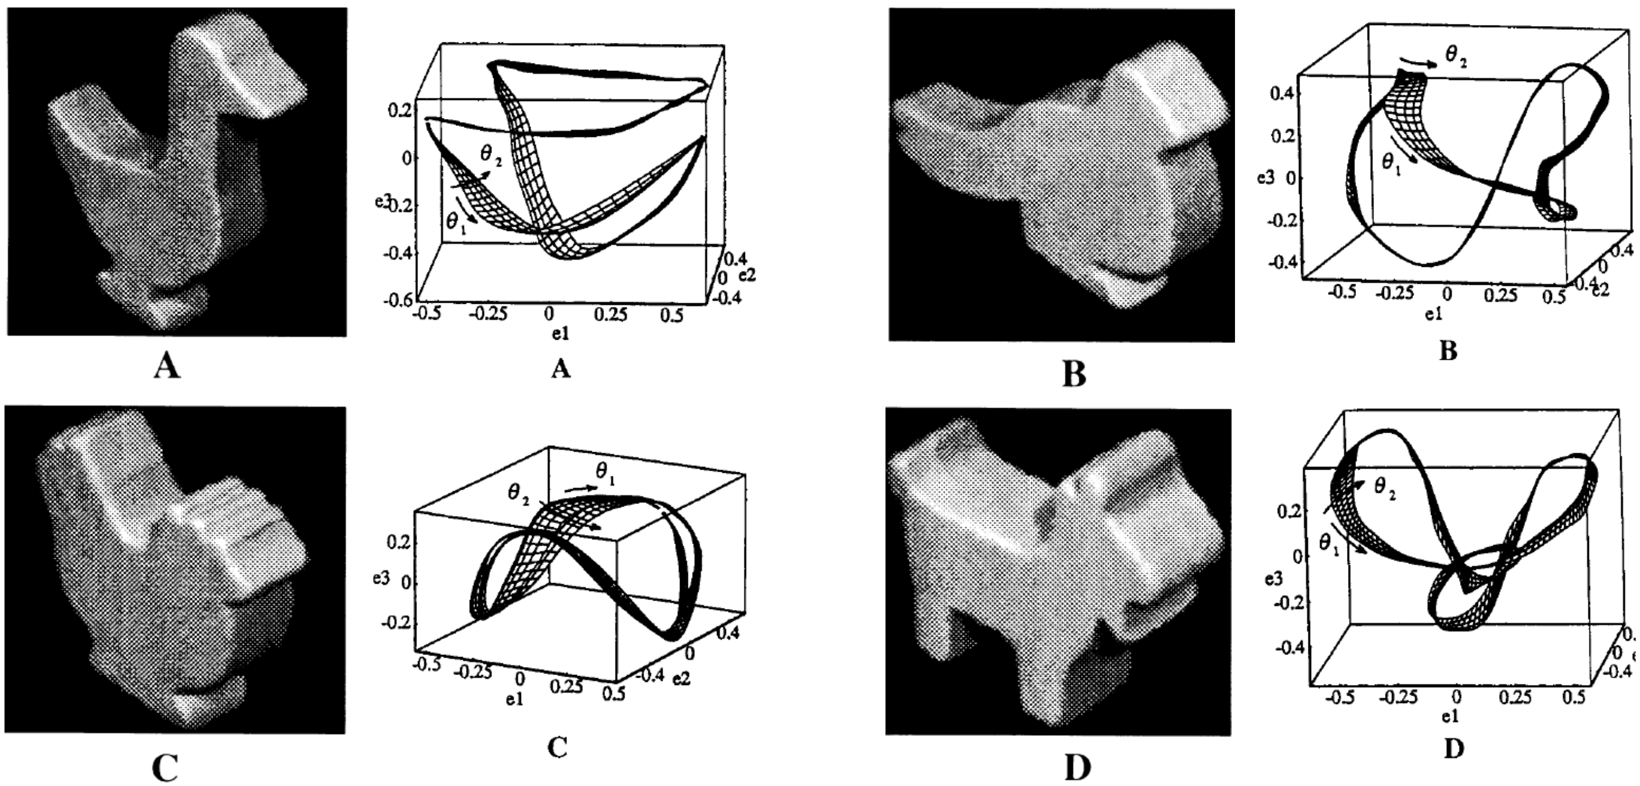

[Murase & Nayar, 1996]

Naïve View-Based Approach

Database of mouth “templates”

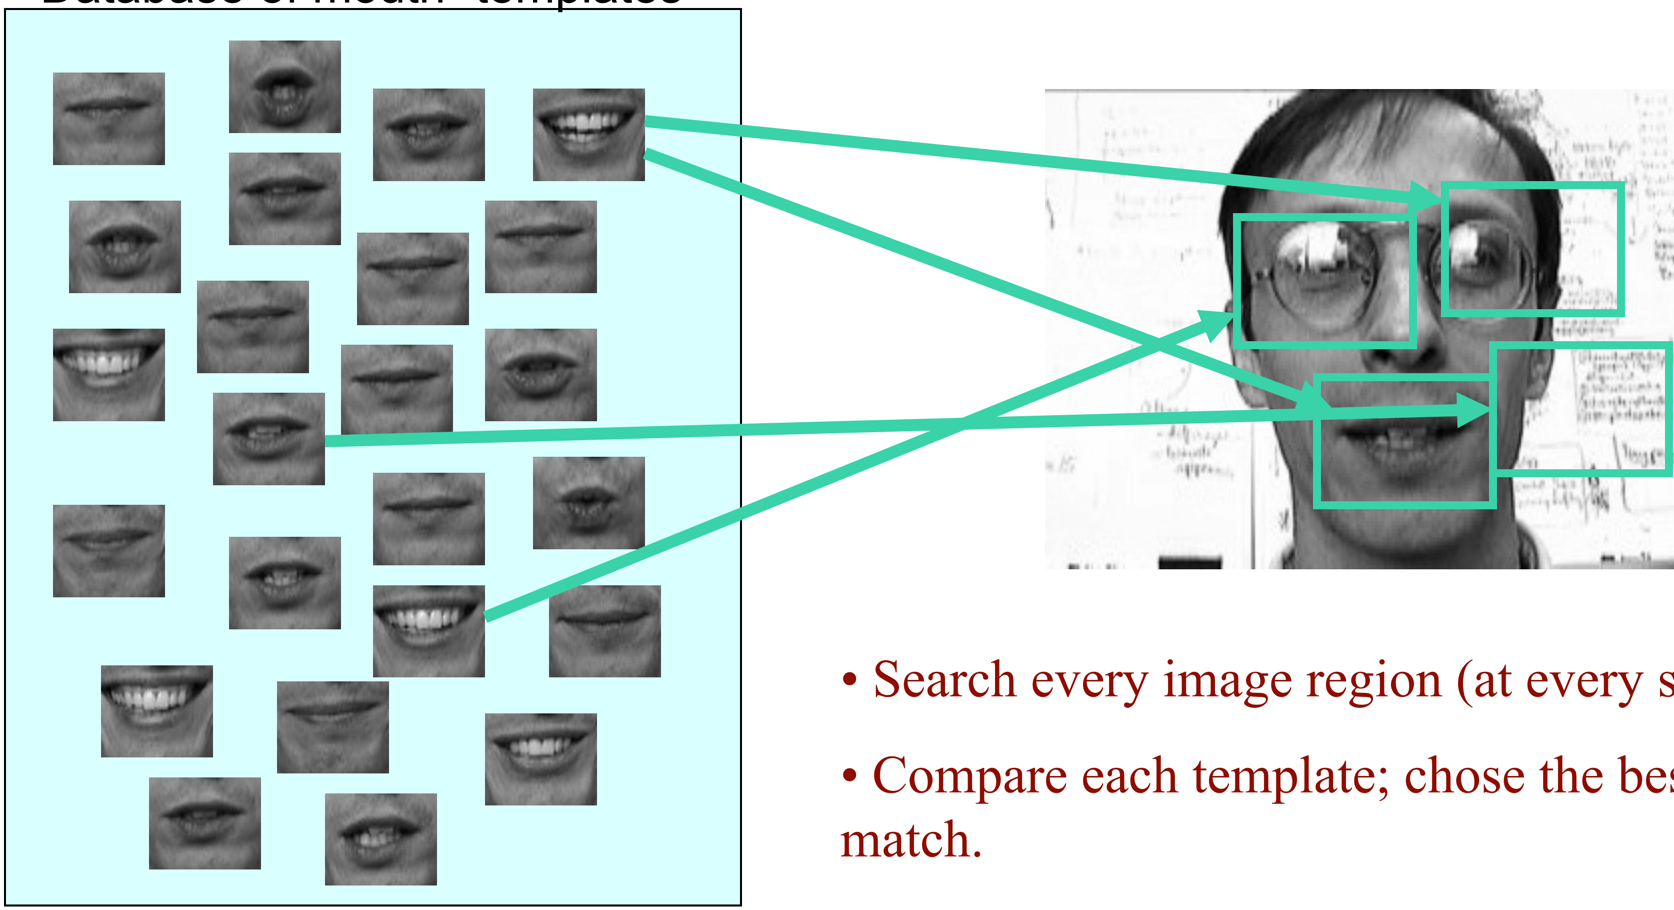

- Search every image region (at every scale).
- Compare each template; chose the best match.

View-Based Approach

Database of mouth “templates”

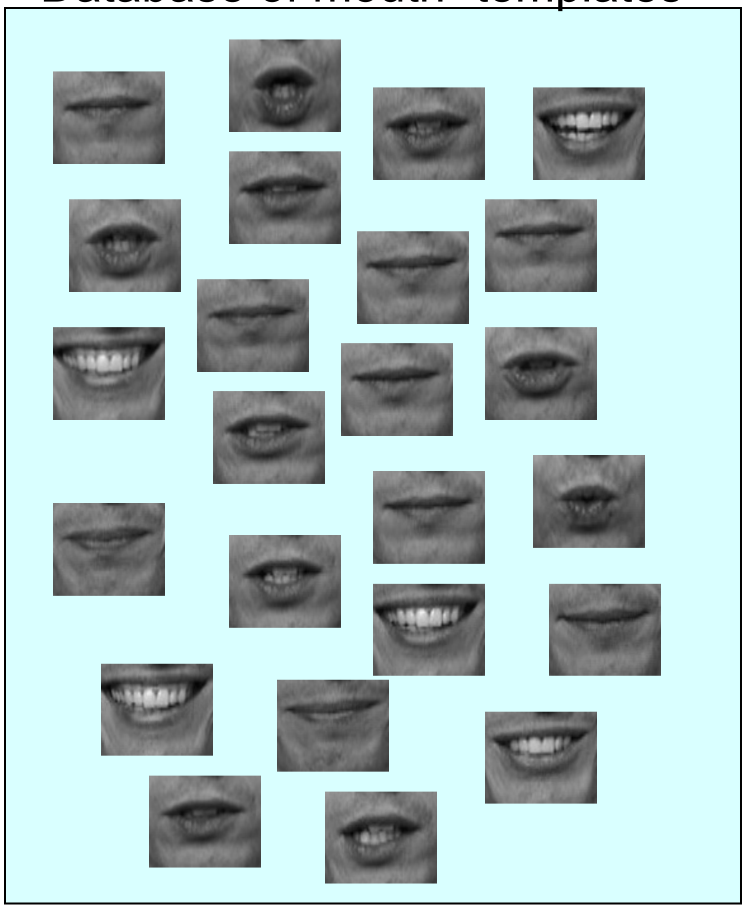

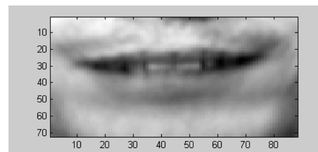

Mean

First three eigenvectors:

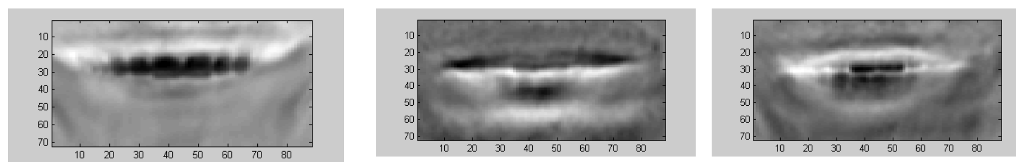

View-Based Approach

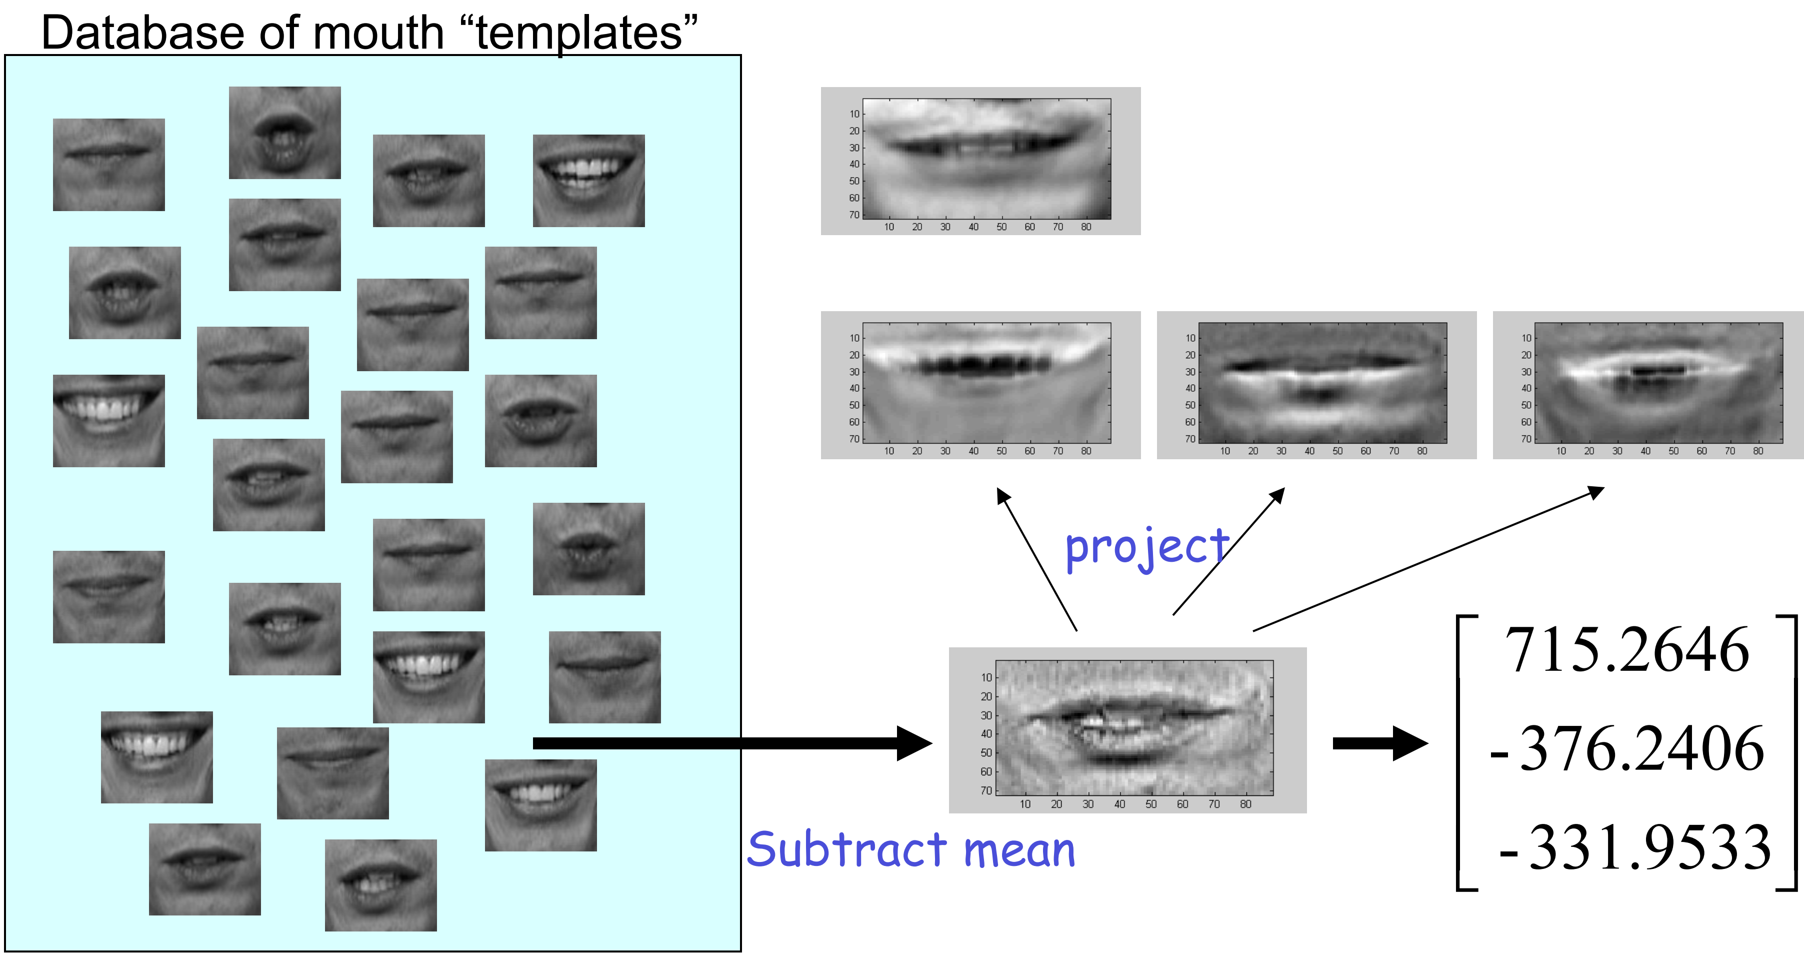

Mouth Space

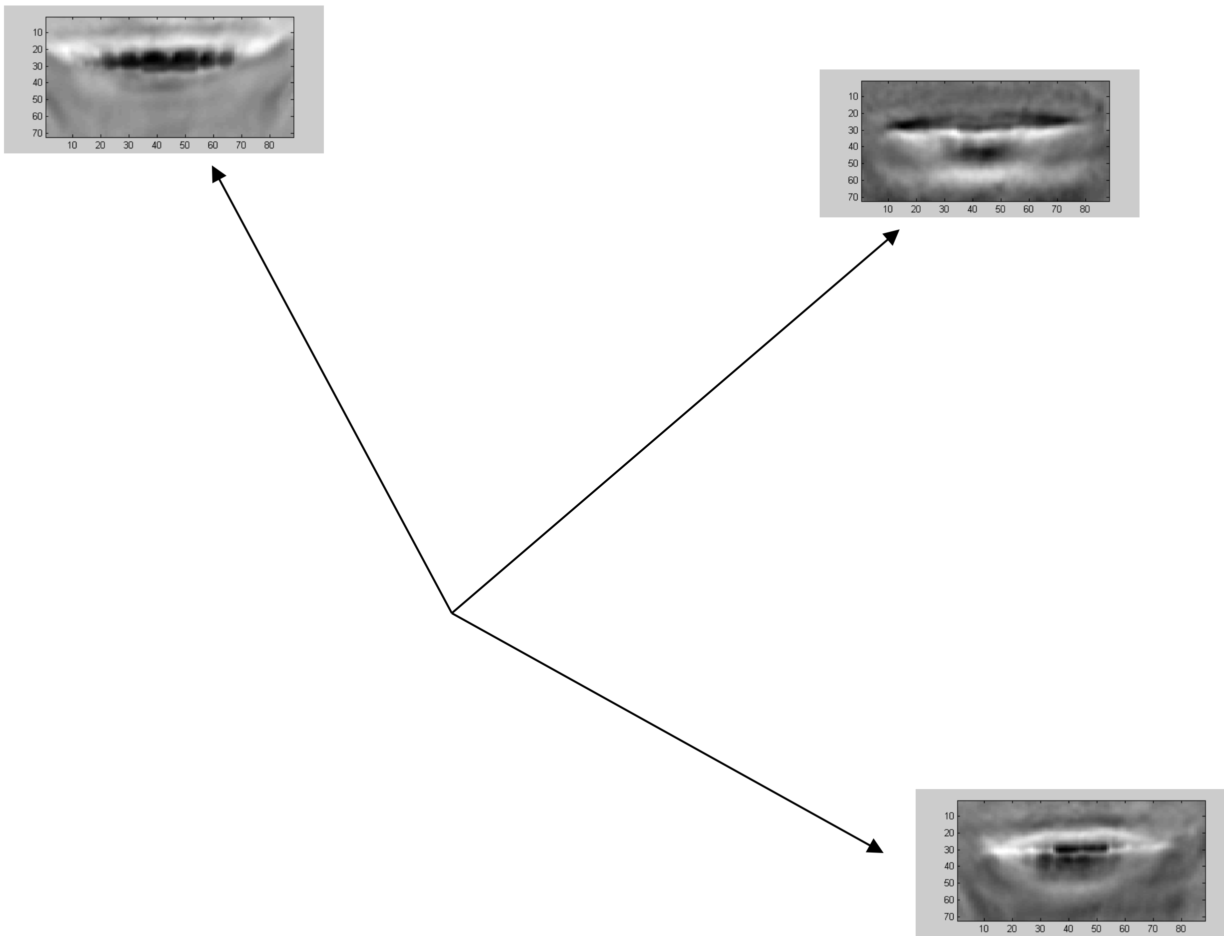

Mouth Space

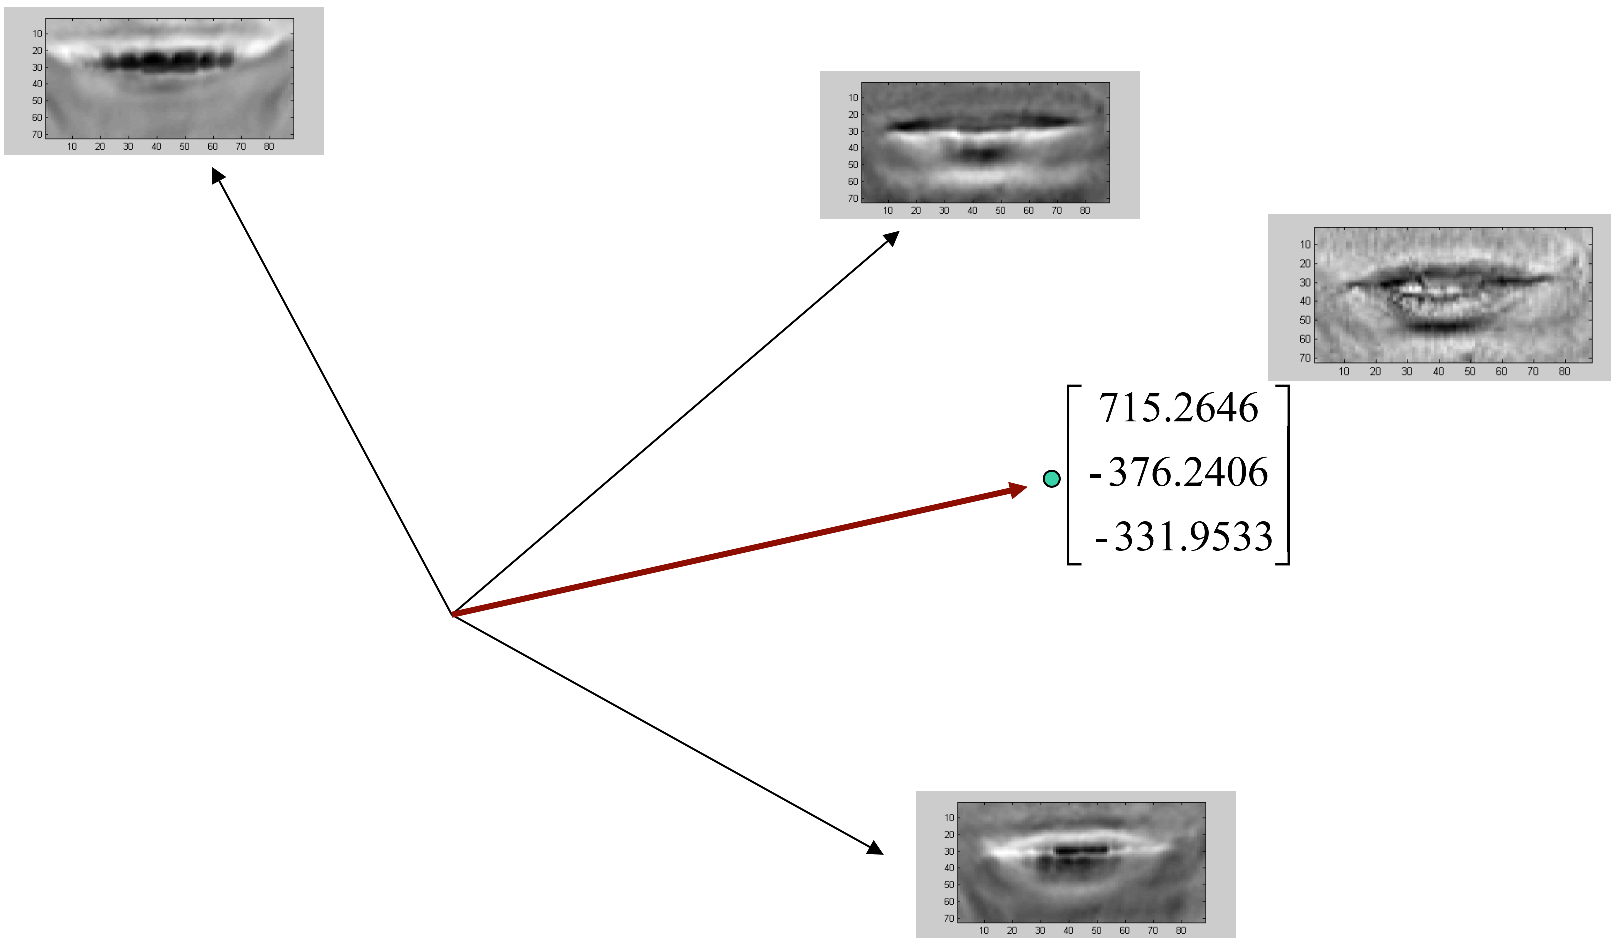

Mouth Space

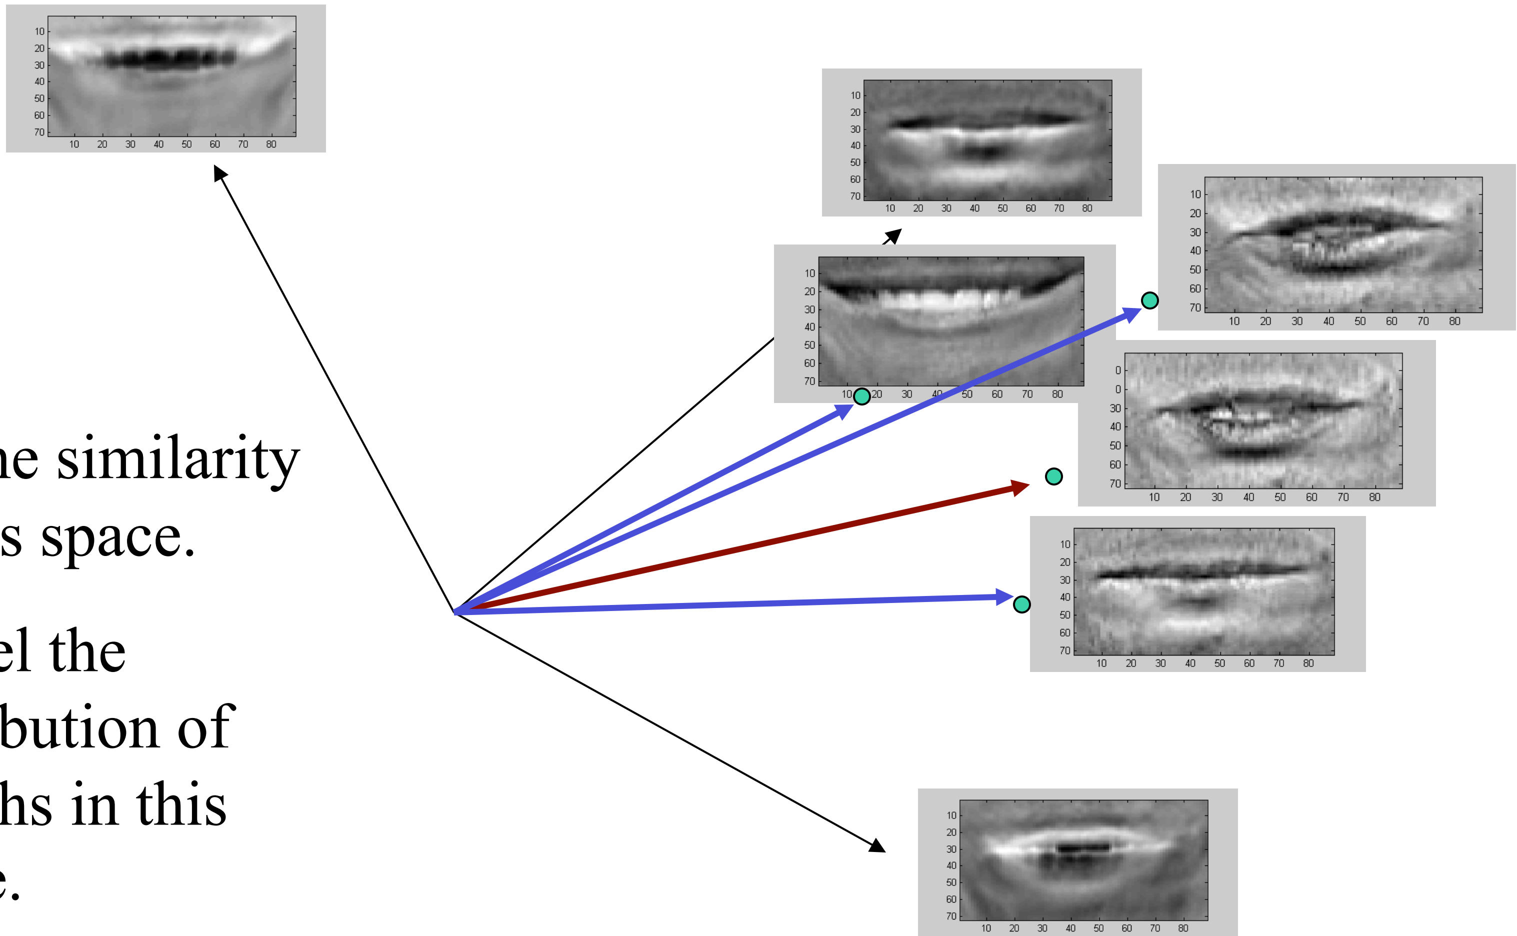

Define similarity in this space.

Model the distribution of mouths in this space.

Images as Vectors

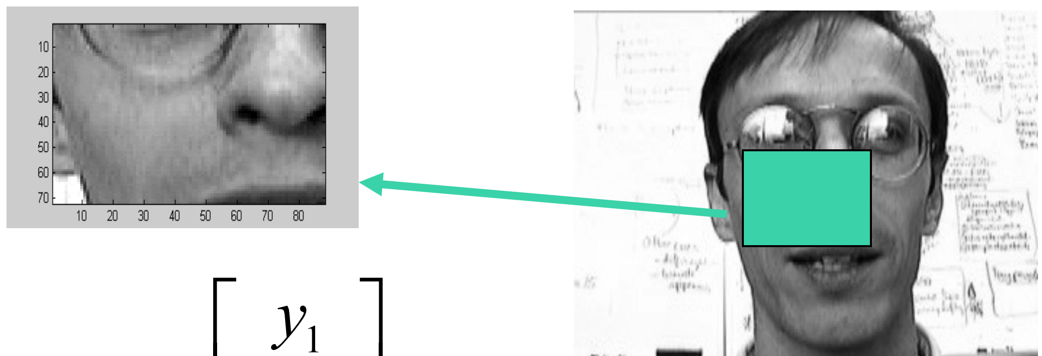

$$= \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_{n \times m} \end{bmatrix}$$

Is it a mouth?

Images as Vectors

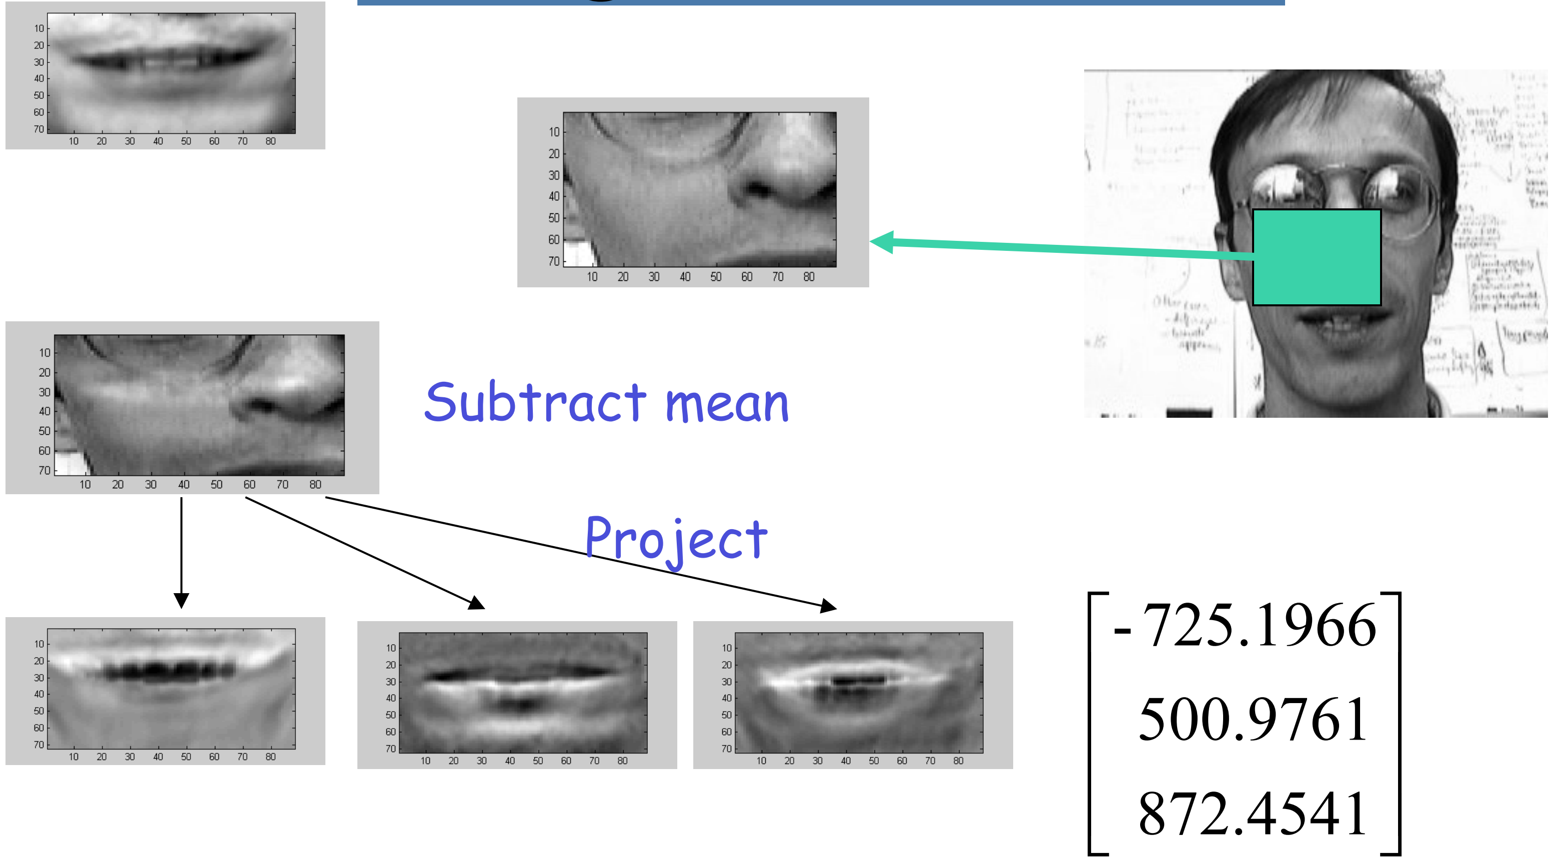

Mouth Space

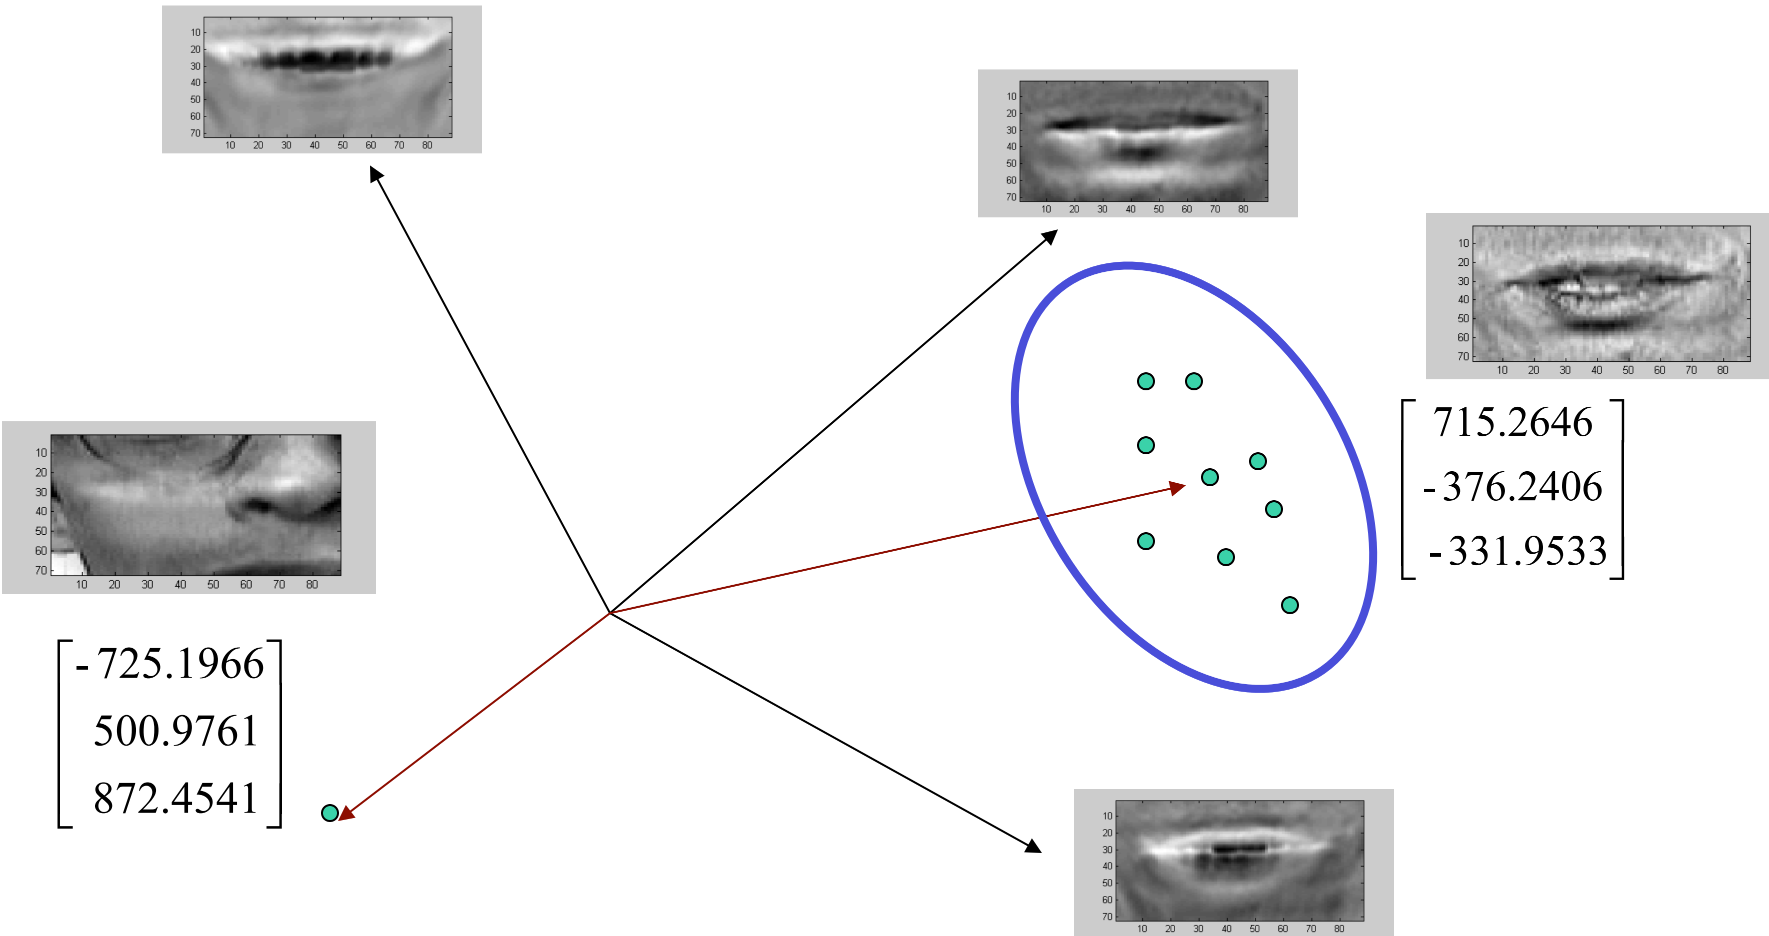

Simple Search Strategy

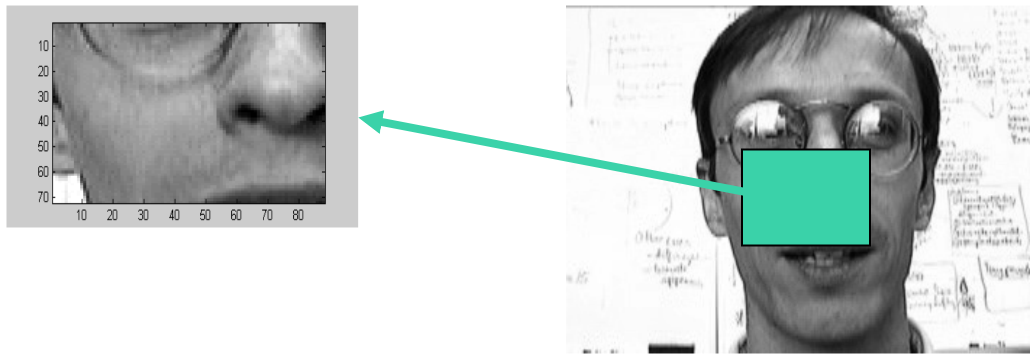

Project each training image onto the low-dimensional subspace. Store the vectors of coefficients

For each image region

1 project it onto the low-dimensional subspace

2 compare this to each stored coefficient vector (cheap)

3 if the smallest distance is less than some threshold, then it is a mouth

Probabilistic Model

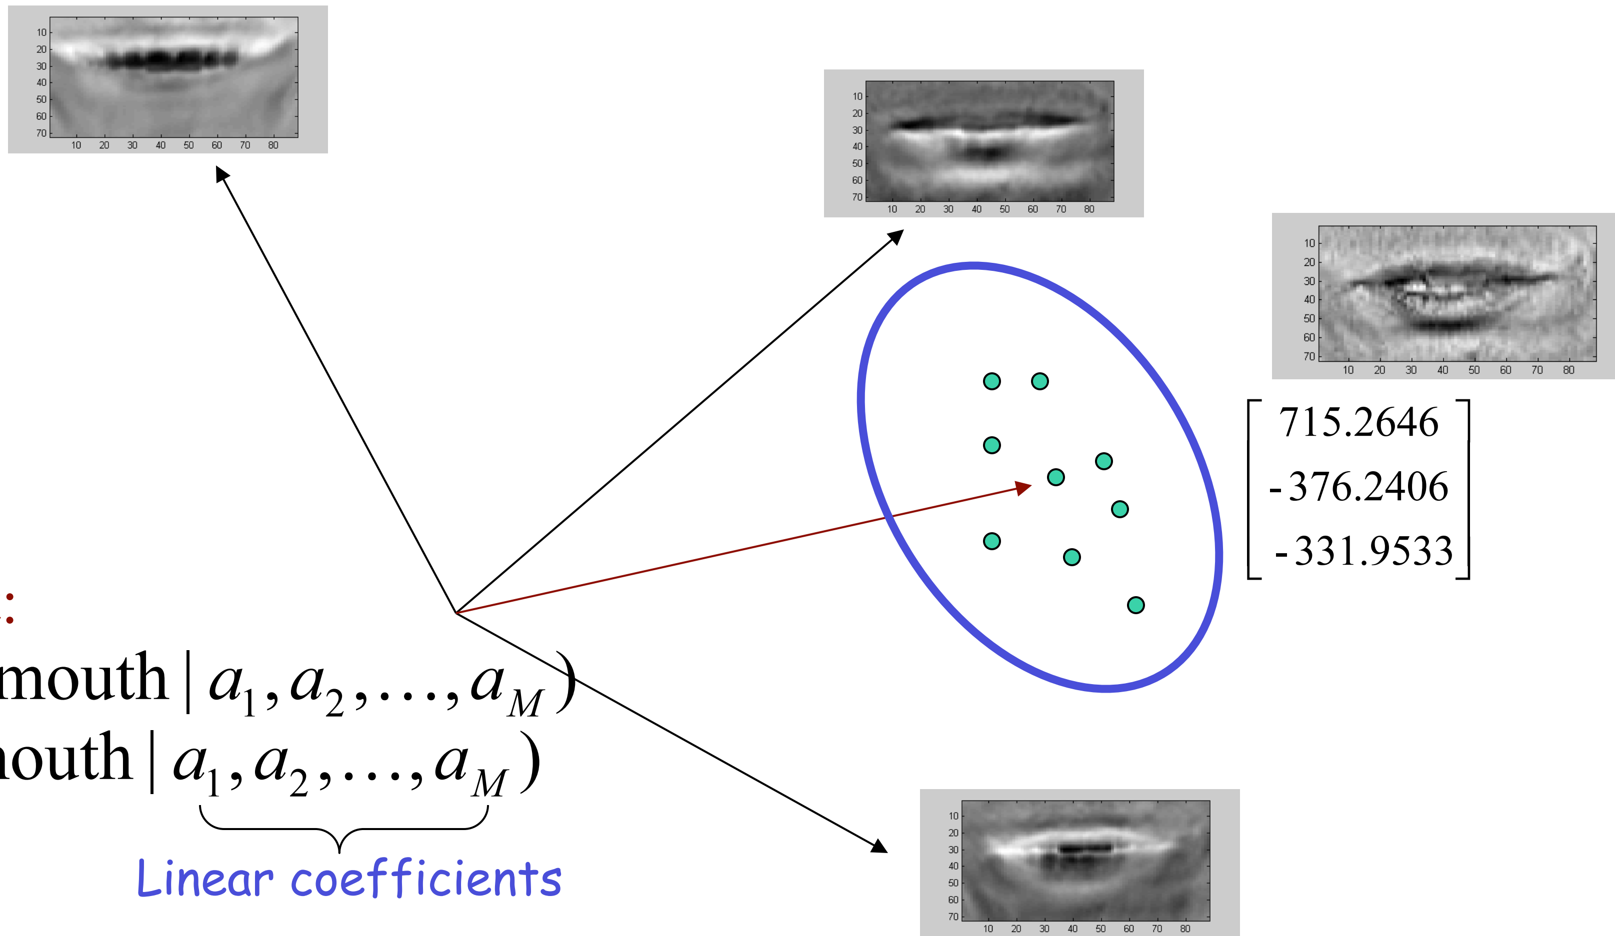

Want:

$$p(\neg\text{mouth} \mid a_1, a_2, \ldots, a_M)$$
$$p(\text{mouth} \mid \underbrace{a_1, a_2, \ldots, a_M}_{\text{Linear coefficients}})$$

Probabilistic Model

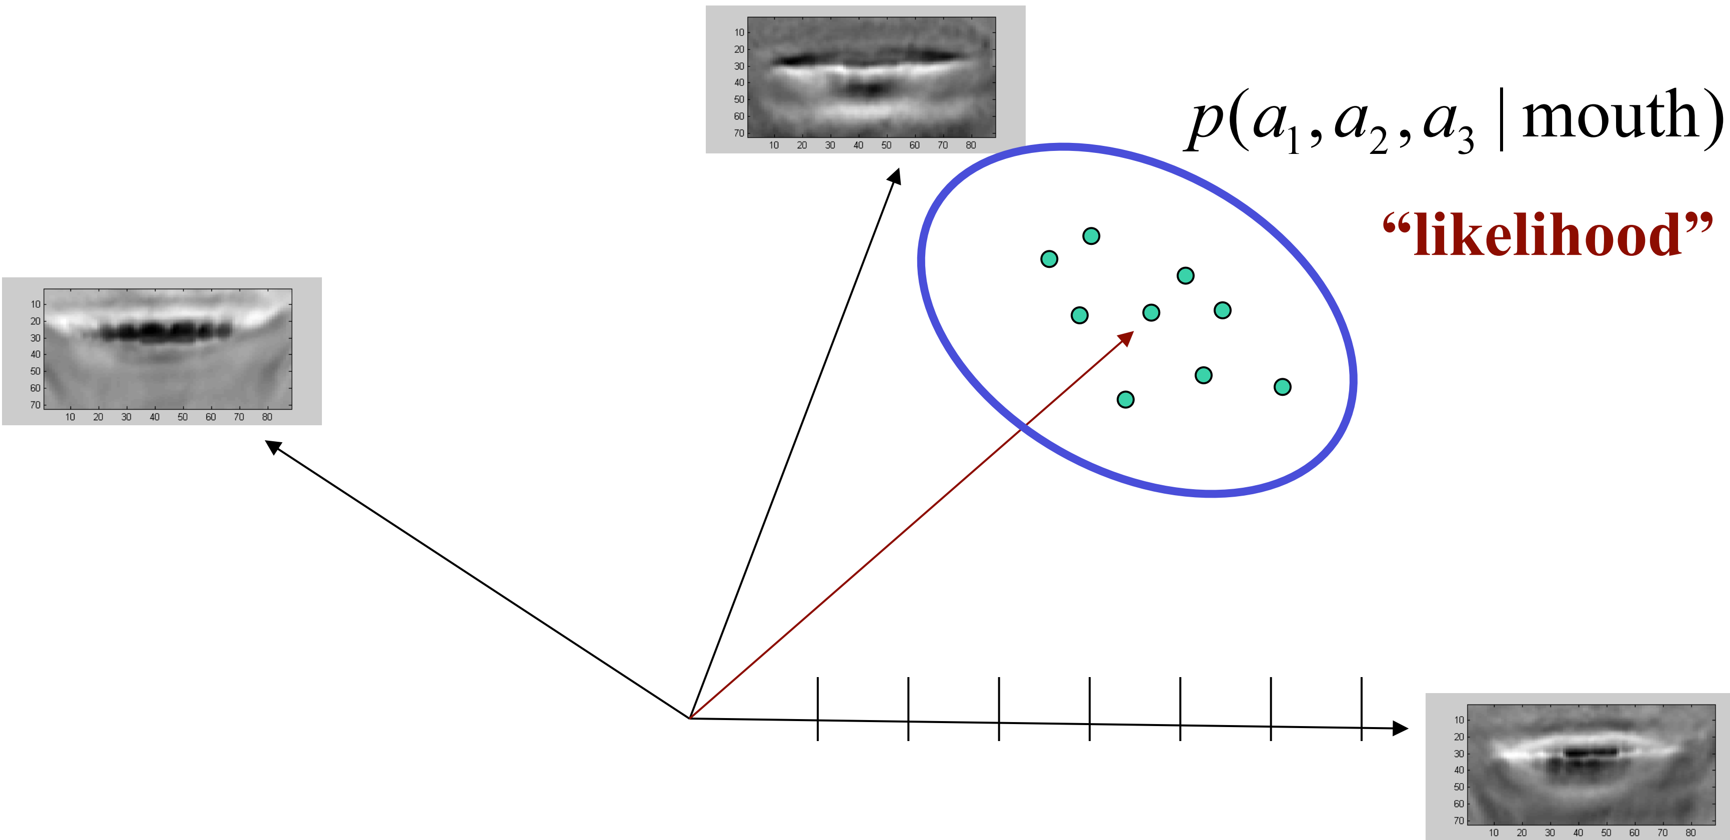

Probabilistic Model

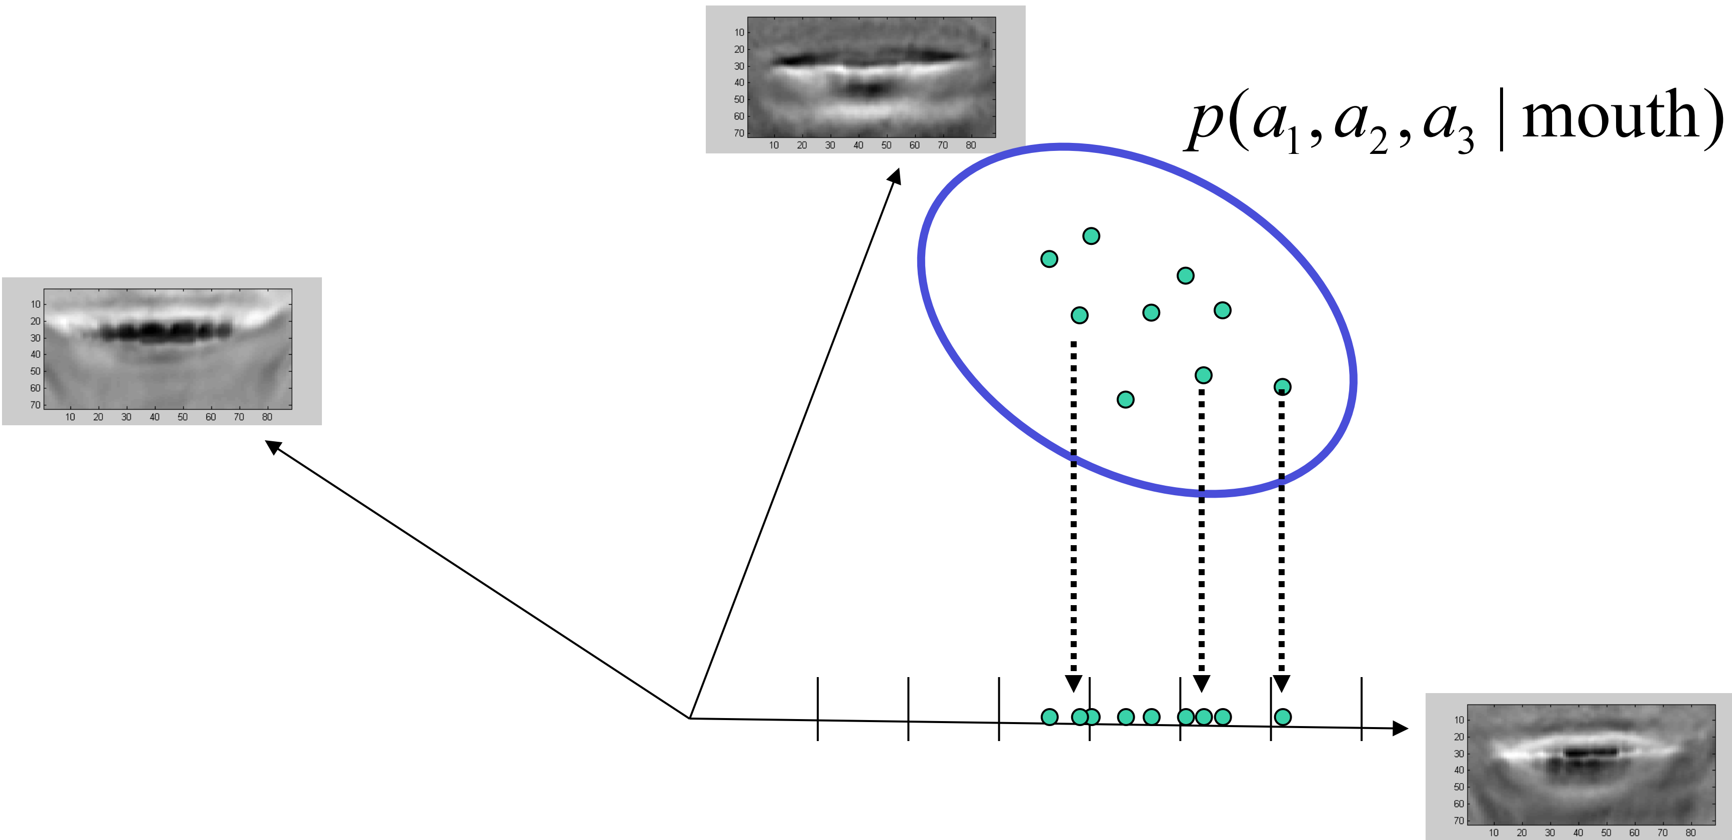

Probabilistic Model

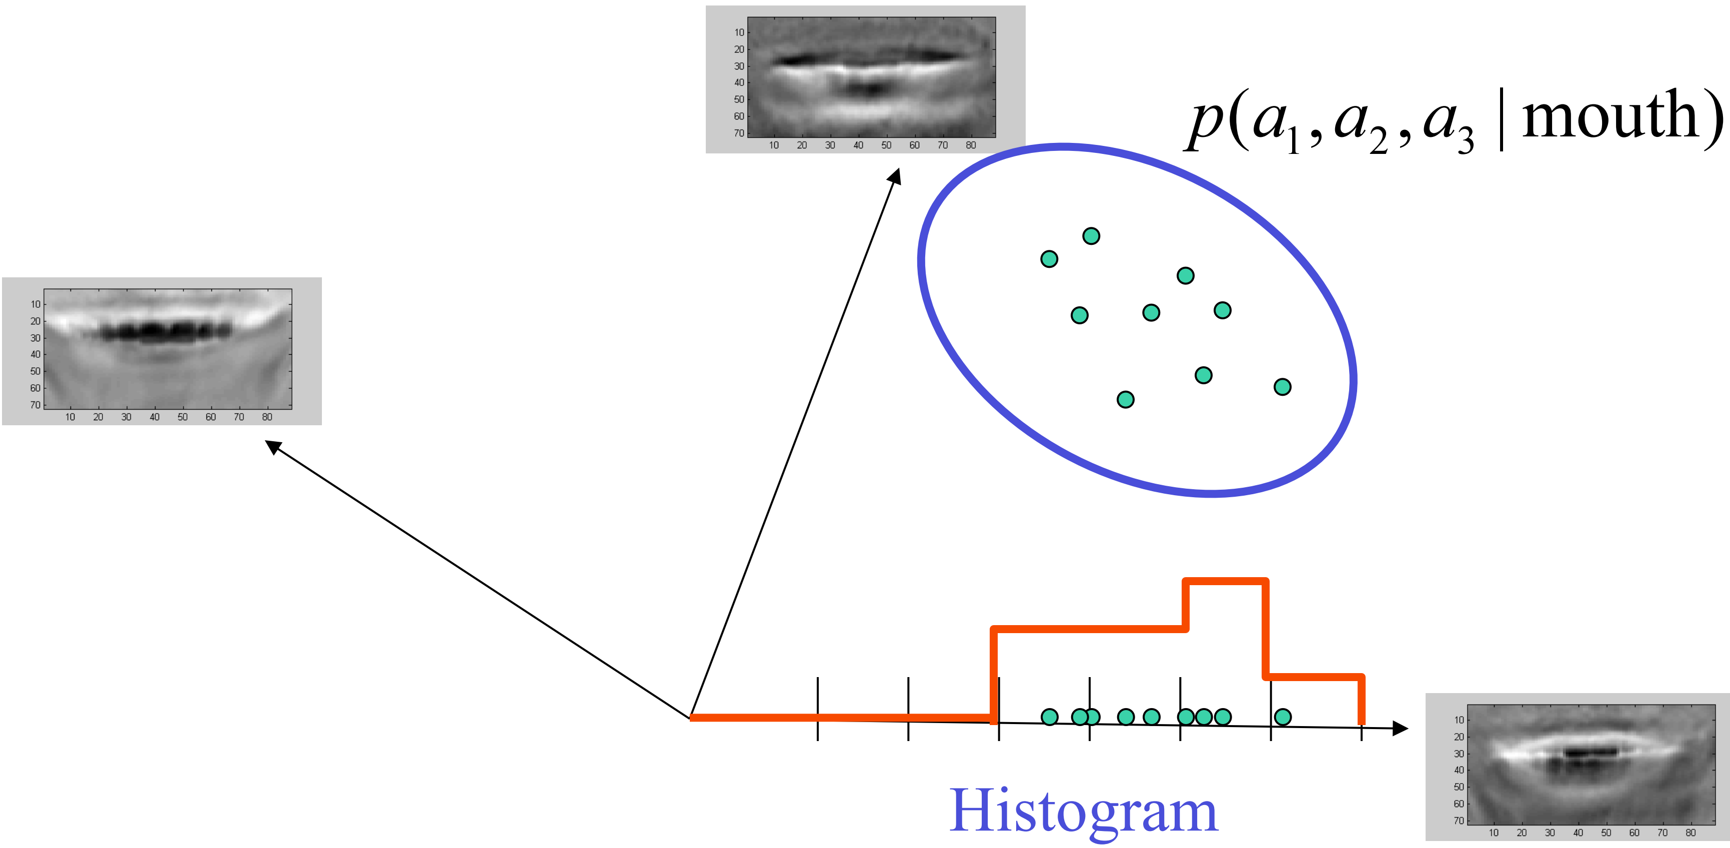

Probabilistic Model

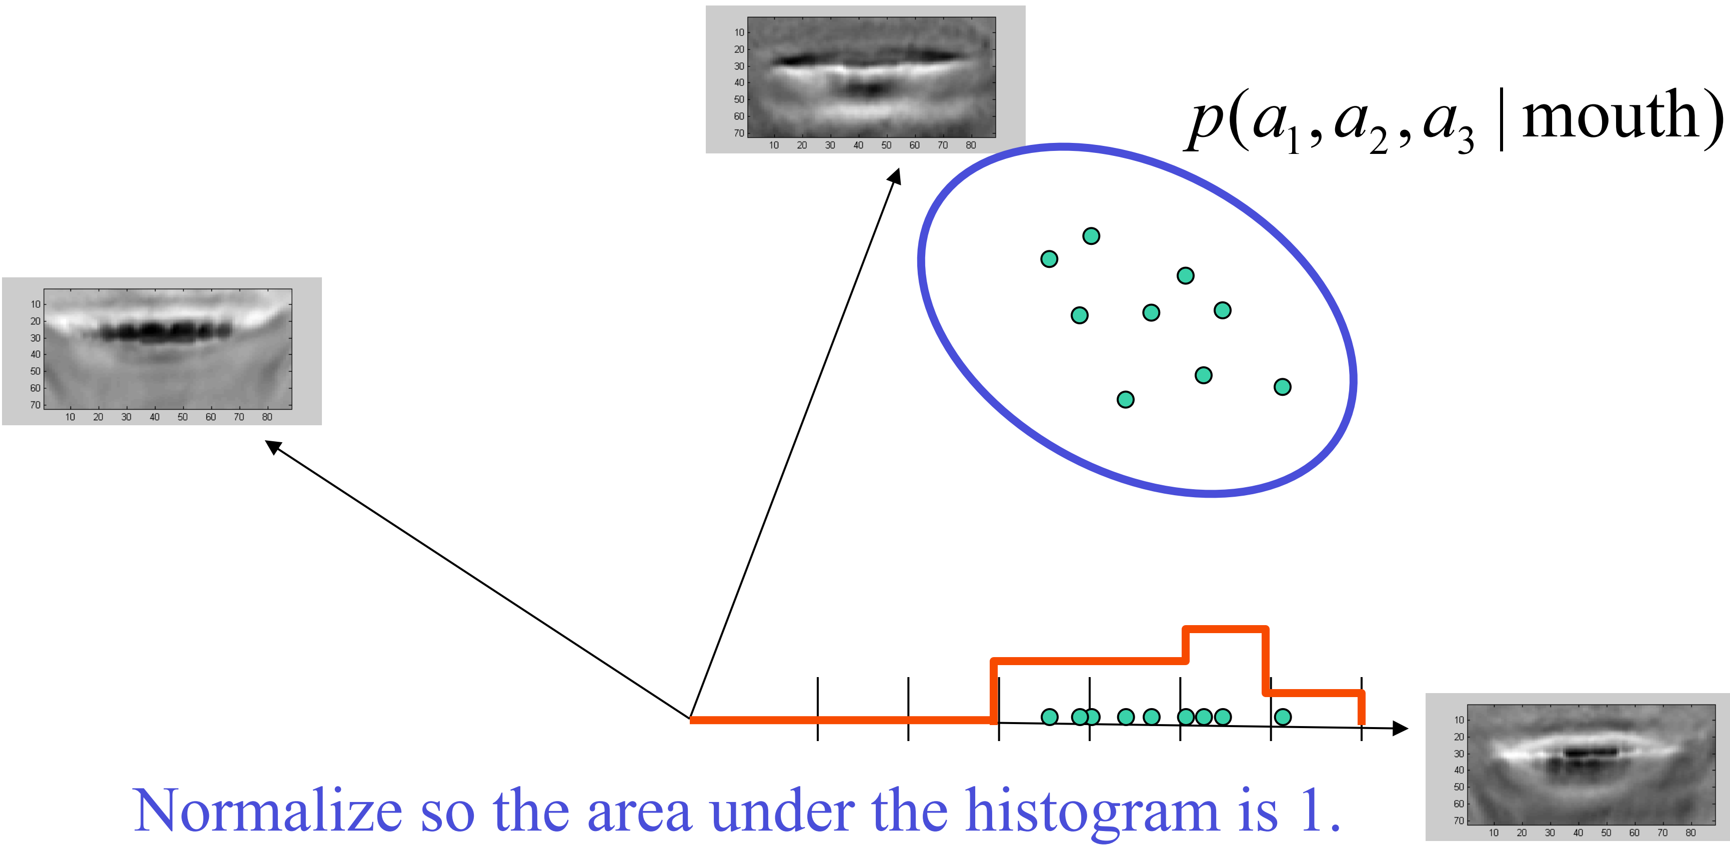

Normalize so the area under the histogram is 1.
(empirical probability)

Probabilistic Model

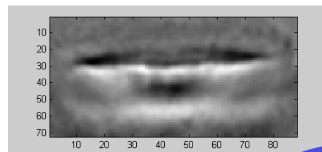

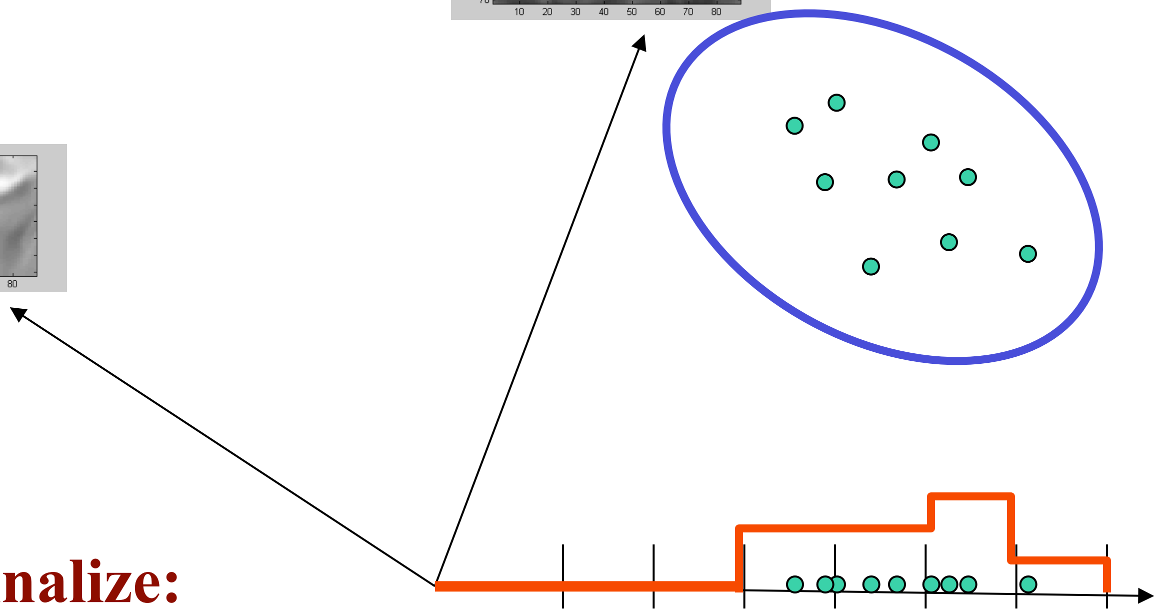

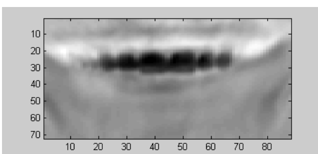

Marginalize:

$$\sum_{a_1, a_2} p(a_1, a_2, a_3 \mid \text{mouth}) = p(a_3 \mid \text{mouth})$$

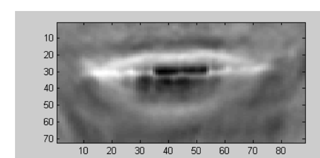

Random variables

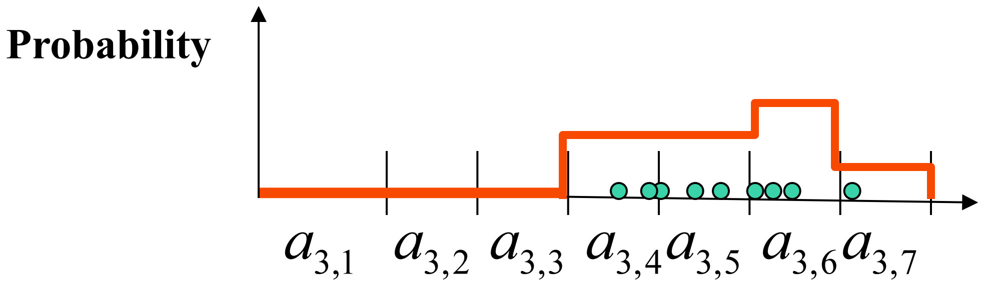

Let X be a *random variable* that can take on one of the discrete histogram bins

$$X \in \{a_{3,1},\ldots,a_{3,7}\}$$

Basic facts

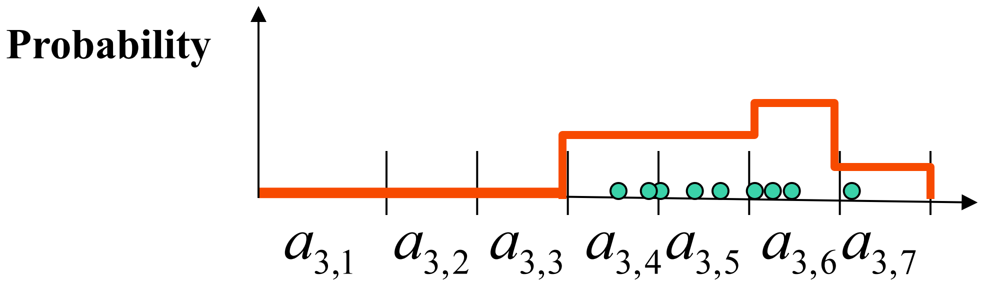

$$p(X = a_{3,i}) \text{ or just } p(a_{3,i})$$

$$0 \leq p(X = a_{3,i}) \leq 1$$

$$\sum_{i=1}^{7} p(a_{3,i}) = 1$$

Basic facts

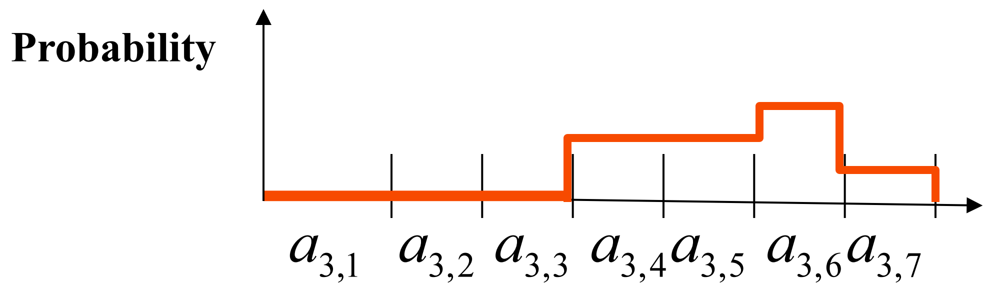

Expected value or expectation of a random variable

$$\mu = E[x] = ?$$

$$\sigma^2 = \mathrm{var}[x] = E[(x - E(x))^2] = \sum_x (x - \mu)^2 p(x)$$

Joint Probability

$$p(X_1 = a_{1,i}, X_2 = a_{2,j}) = p(a_{1,i}, a_{2,j})$$

$$\sum_{a_{1,i}} \sum_{a_{2,j}} p(a_{1,i}, a_{2,j}) = 1$$

Statistical independence

If: $$p(x, y) = p(x)p(y)$$

- knowing *y* tells you nothing about *x*

Conditional Probability

Dependence - Knowing the value of one random variable tells us something about the other.

$$p(A \mid B) = \frac{p(A,B)}{p(B)}$$

$$p(A \mid B)\,p(B) = p(A,B)$$

Statistical Independence

$$p(A \mid B) = \text{?}$$

If A and B are statistically independent?

Statistical Independence

$$p(A \mid B) = \frac{p(A,B)}{p(B)} = \frac{p(A)p(B)}{p(B)} = p(A)$$

A and B are *statistically independent* if and only if

$$p(A,B) = p(A \mid B)p(B) = p(A)p(B)$$

$$p(A \mid B) = p(A)$$

$$p(B \mid A) = p(B)$$

Conditional Independence

A is independent of B, conditioned on C

$$p(A,B,C) = p(A,B \mid C)\,p(C)$$
$$= p(A \mid C)\,p(B \mid C)\,p(C)$$

If I know *C*, then knowing *B* doesn't give me any more information about *A*.

This does not mean that *A* and *B* are statistically independent

Example: Conditional Independence

$$p(A,B,C) = p(A,B \mid C)\,p(C)$$
$$= p(A \mid C)\,p(B \mid C)\,p(C)$$

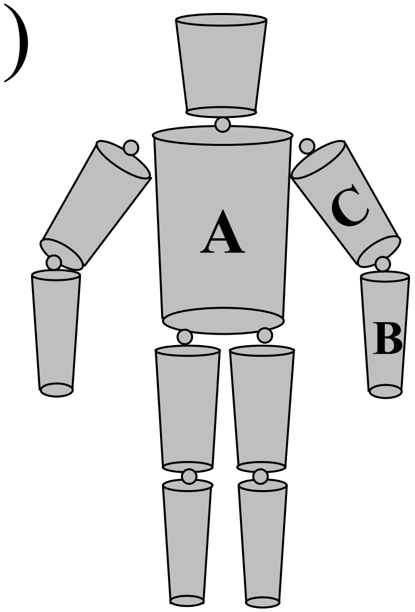

The torso and lower arm poses are not independent.

But if I know the pose of the upper arm then knowing the pose of *B* tells me nothing new about *A*.

Classification

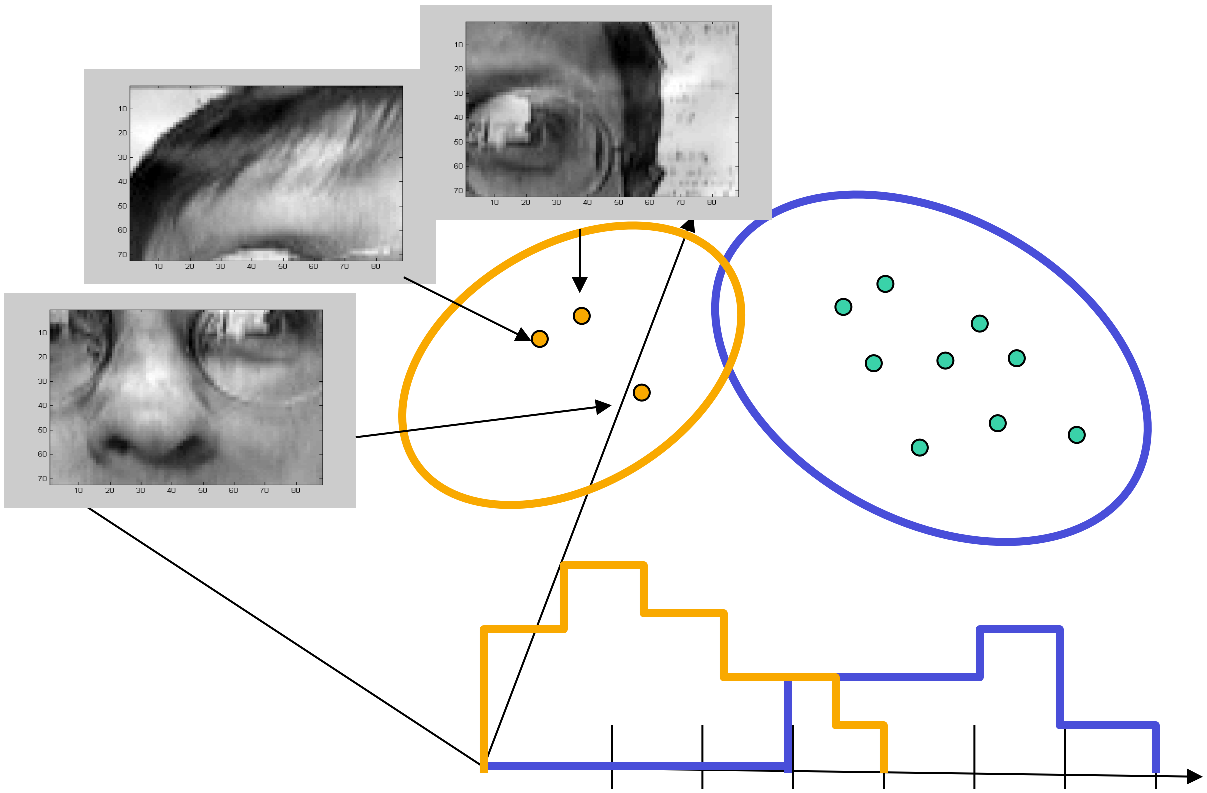

Classification

Imagine we just consider one dimension (one linear coefficient).

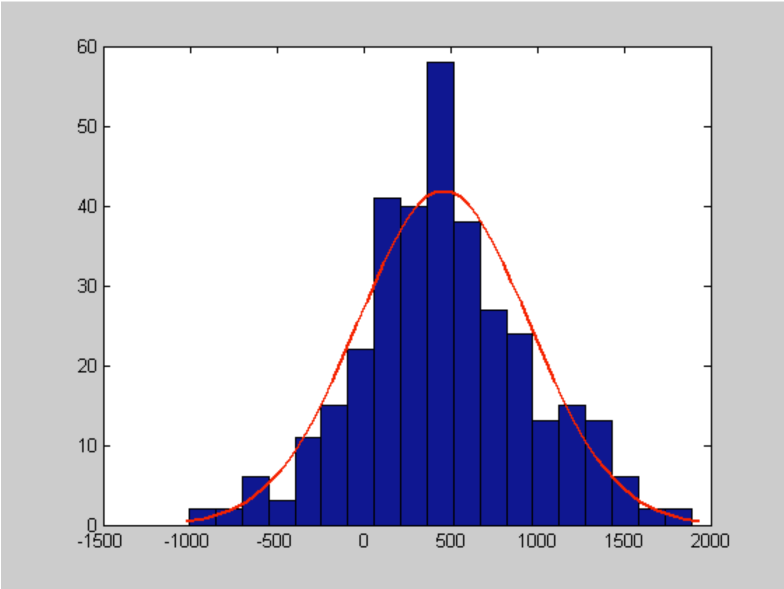

$p(a_3 \mid \neg\text{mouth})$

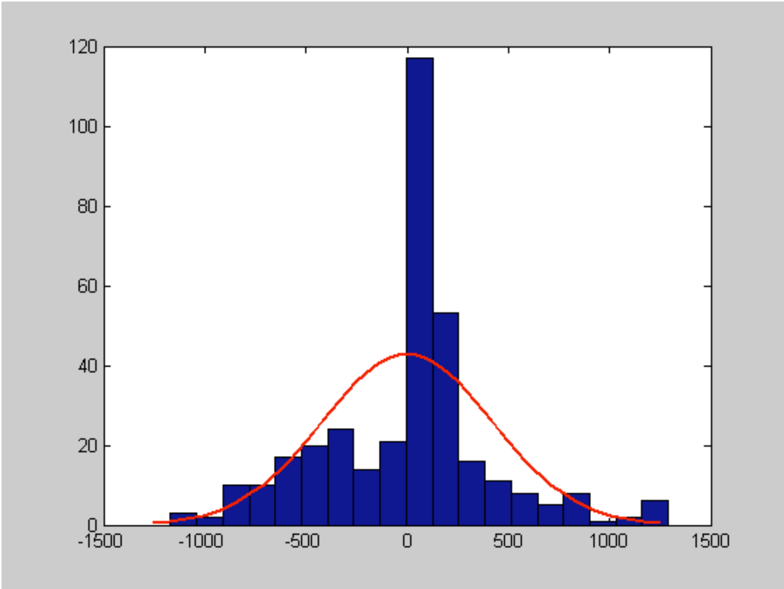

$p(a_3 \mid \text{mouth})$

Parametric models

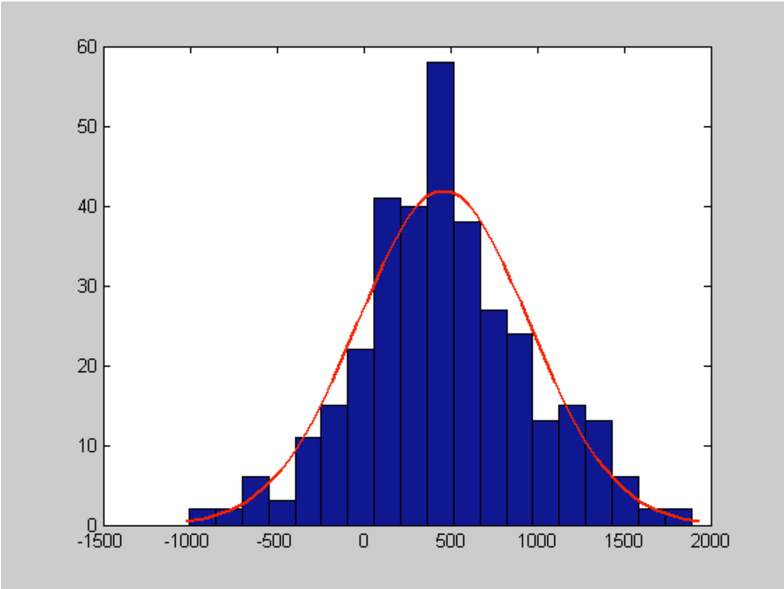

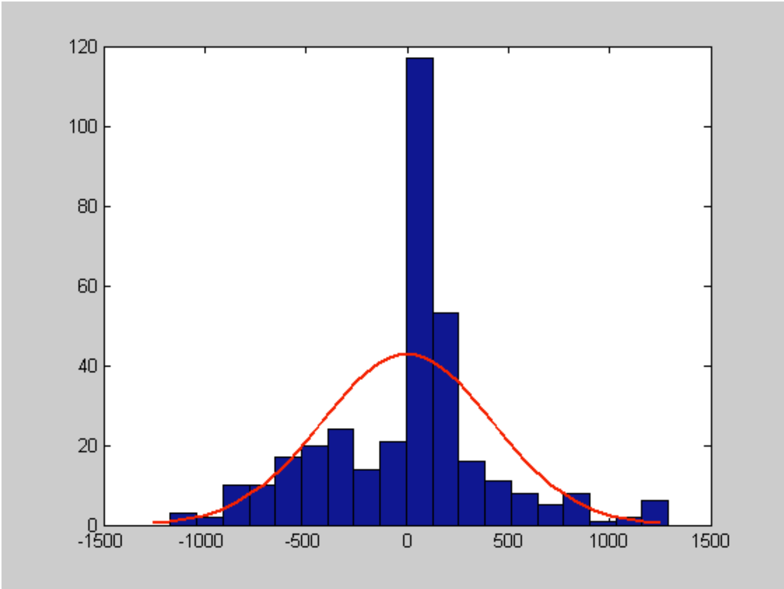

Discrete: normalized histograms.

Parametric (here Gaussian):

$$p(a_{3,i} \mid \text{mouth}) = \frac{1}{\sqrt{2\pi\sigma}} \exp(-\frac{1}{2}(a_{3,i} - \mu)^2 / \sigma^2)$$

Matlab notes

This plot was made using
histfit(*vector of data*)

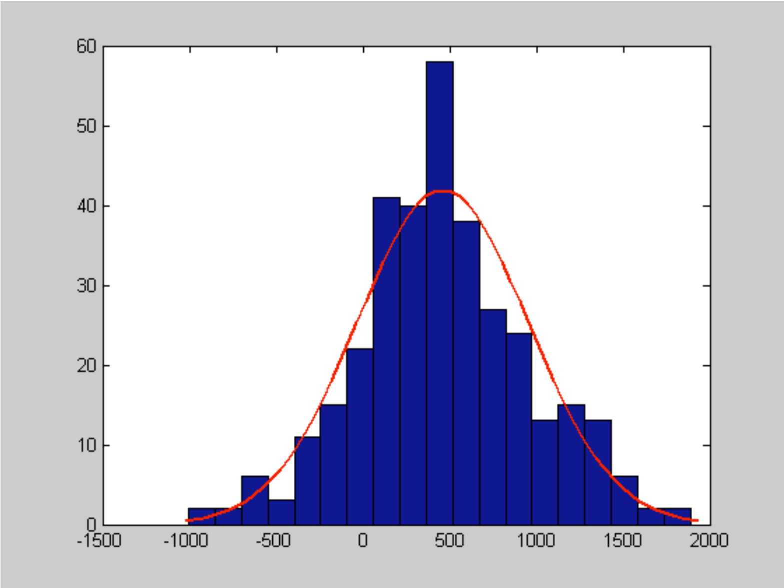

To actually fit the mean and variance:

[mu, sig]=normfit(*vector of data*)

plot(normpdf(*min*:*max*, mu,sig),'r')

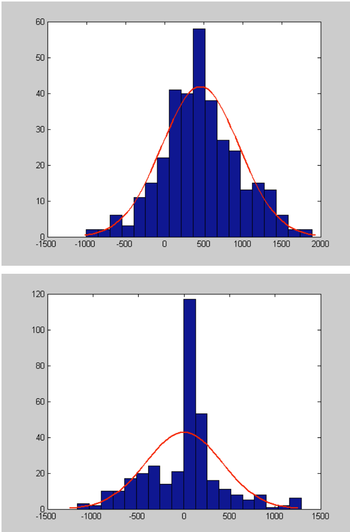

Classification

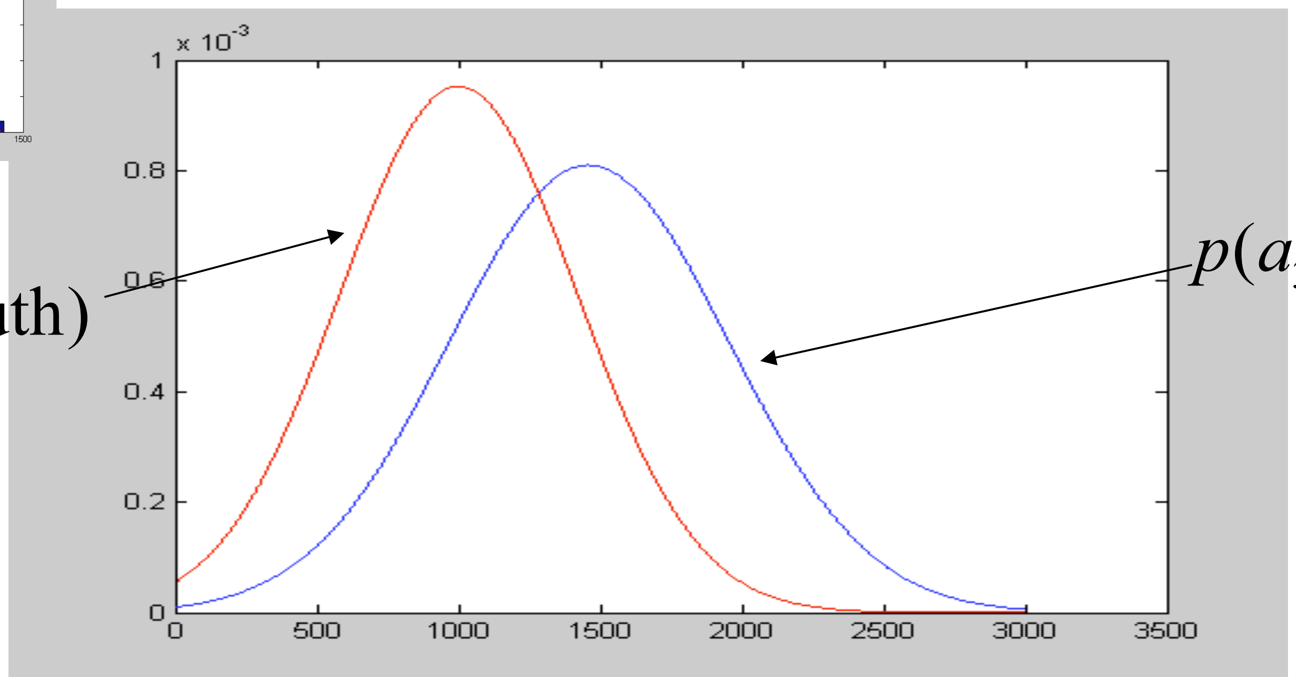

Given a value of a_3, how can I classify it as mouth or not mouth?

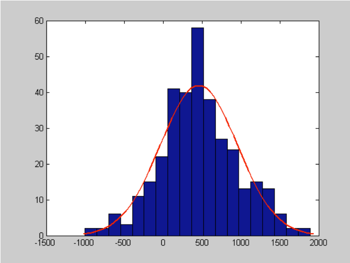

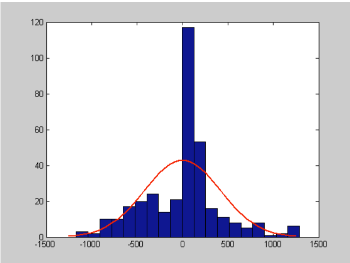

Classification

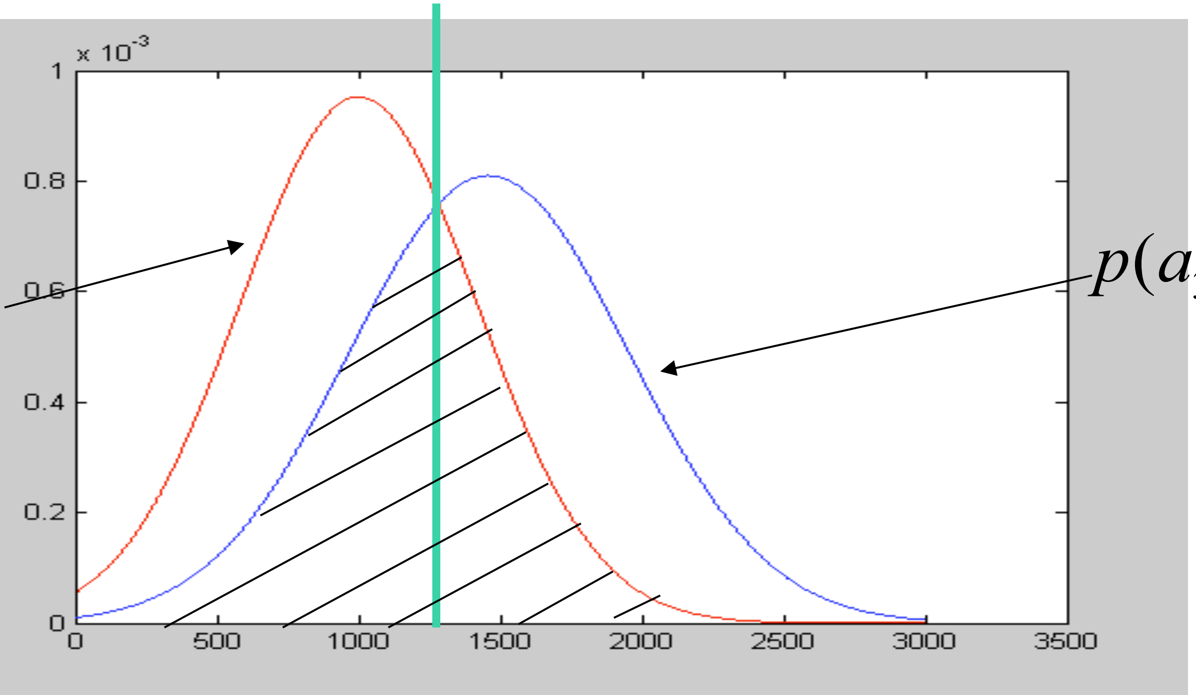

$p(a_3 \mid \text{mouth})$

$p(a_3 \mid \neg\text{mouth})$

Maximum likelihood classification

$$\text{if } p(a_3 \mid \text{mouth}) > p(a_3 \mid \neg\text{mouth}) \text{ then mouth}$$

Classification

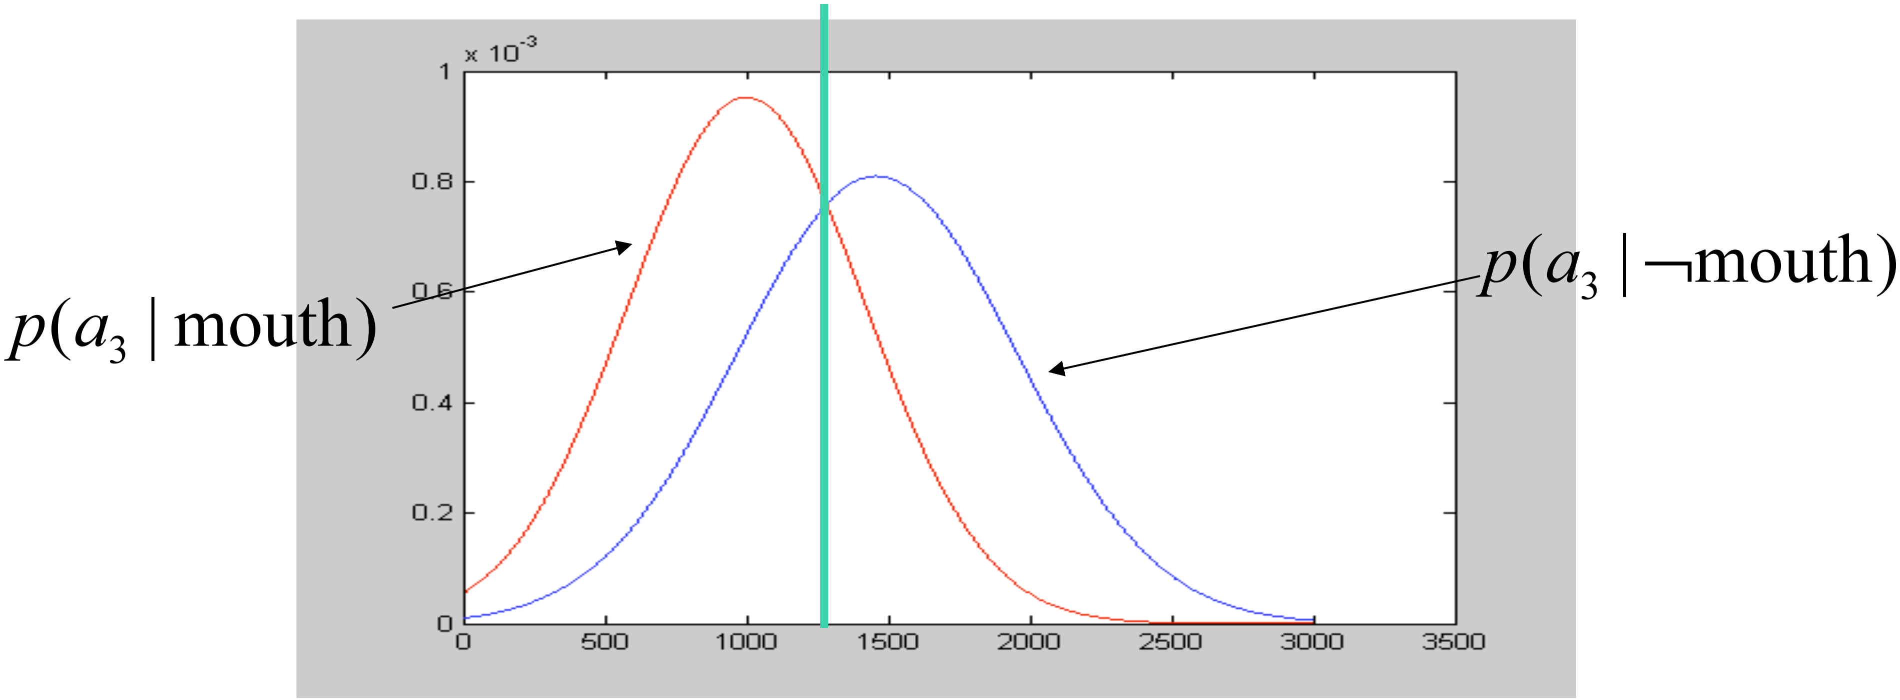

There is an implicit assumption with this approach. What?

$$\text{if } p(a_3 \mid \text{mouth}) > p(a_3 \mid \neg\text{mouth}) \text{ then mouth}$$

Classification

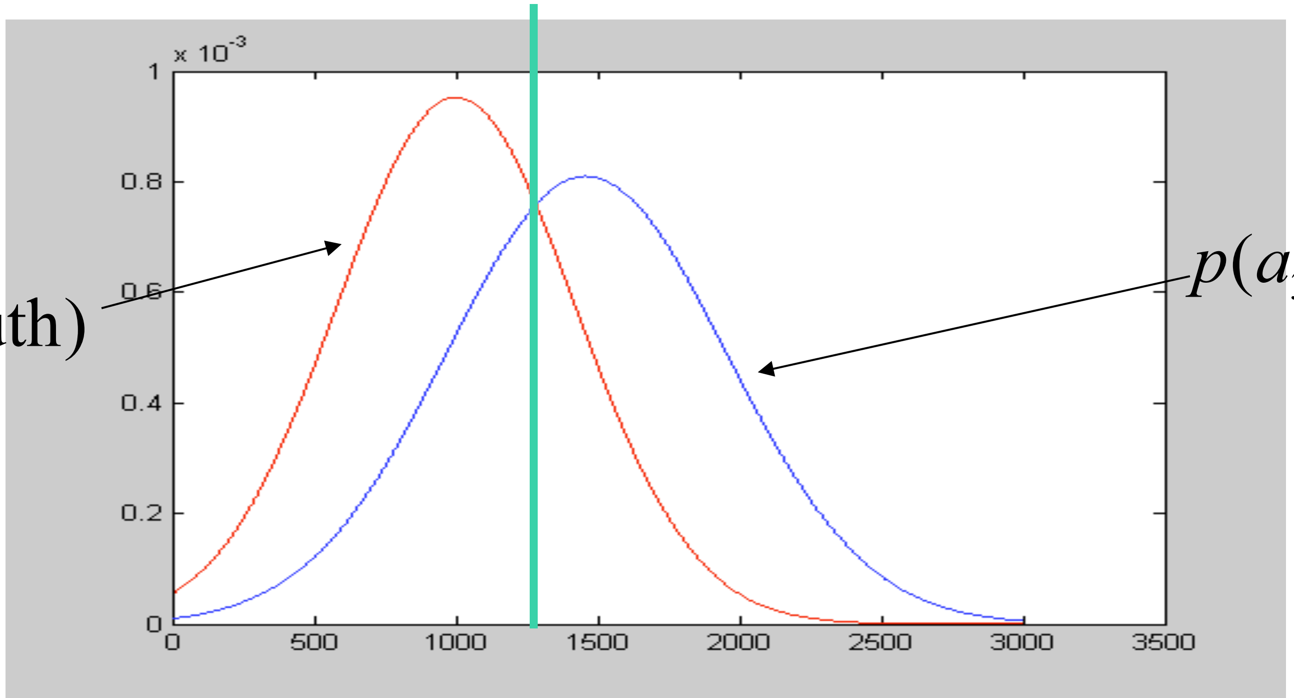

$p(a_3 \mid \text{mouth})$

$p(a_3 \mid \neg\text{mouth})$

There is an implicit assumption with this approach. What?

$$p(\text{mouth}) = p(\neg\text{mouth})$$

Bayes' Theorem

$$p(A,B) = p(A \mid B)\,p(B) = p(B \mid A)\,p(A)$$

Revd. Thomas Bayes, 1701-1761

Posterior Probability

likelihood prior

$$p(\text{mouth} \mid a_3) = \frac{p(a_3 \mid \text{mouth})\, p(\text{mouth})}{p(a_3)}$$

normalization constant
(independent of mouth)

Maximum A Posteriori Classification

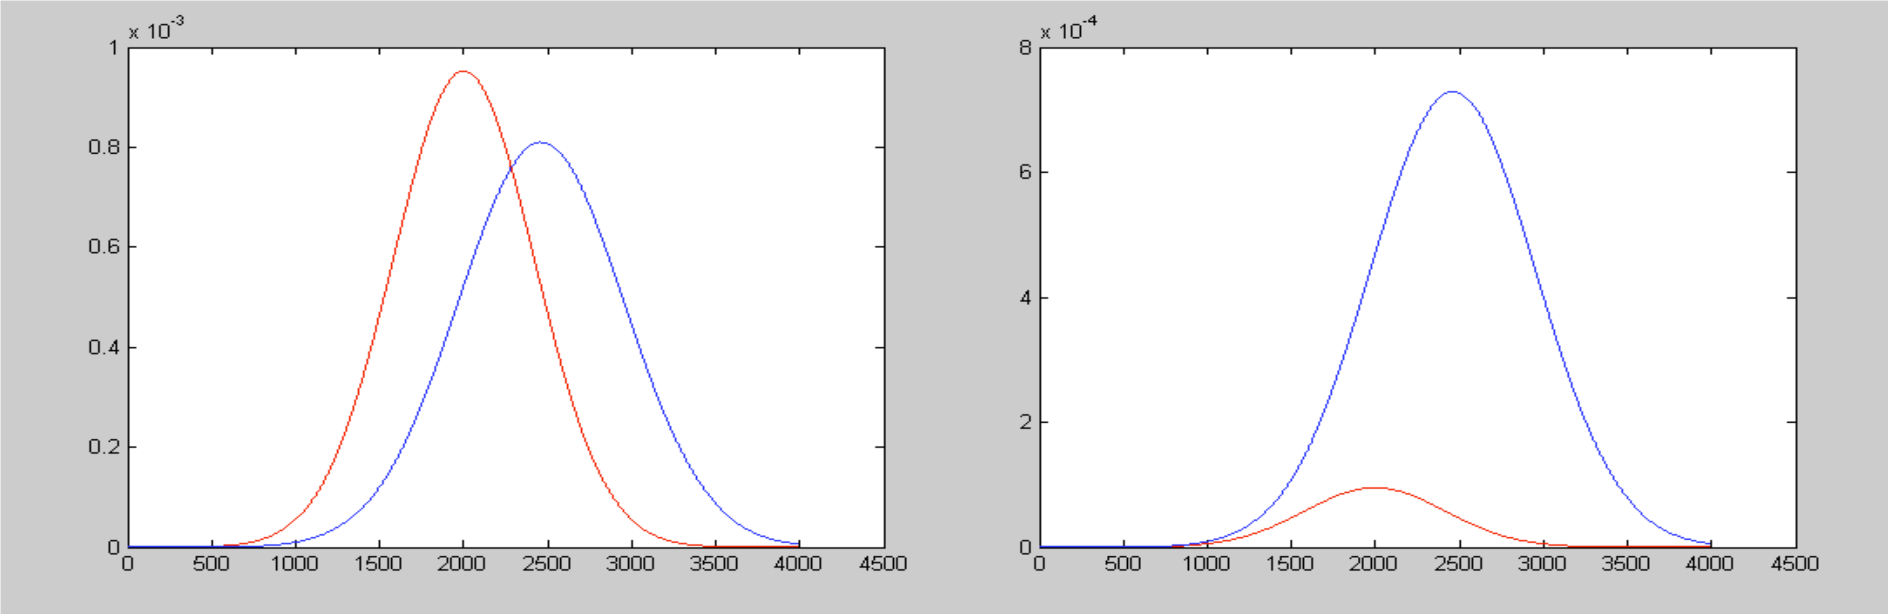

likelihood posterior

From a_3 alone, it looks like MAP classification will always prefer the not-mouth interpretation.

What about the other coefficients?

$p(a_6 \mid \neg \text{mouth})$

$p(a_6 \mid \text{mouth})$

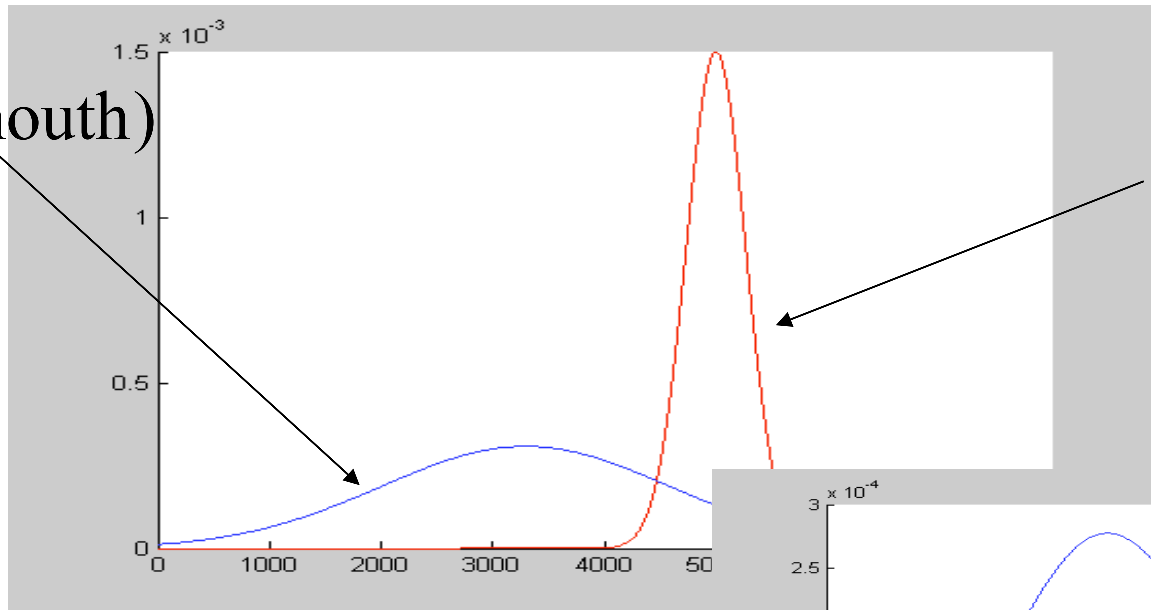

posterior:

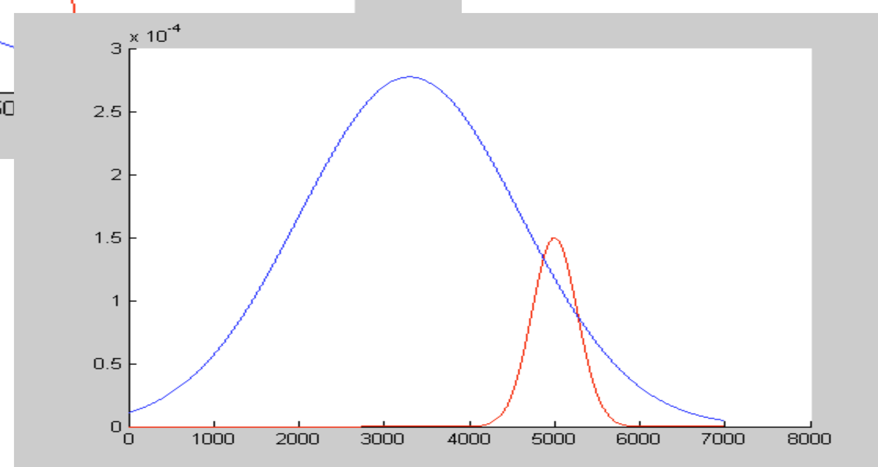

Conditional Independence

$$p(a_1, a_2, \ldots, a_M \mid \text{mouth}) = \prod_{i=1}^{M} p(a_i \mid \text{mouth})$$

Conditional Independence

$$p(a_1, a_2, \ldots, a_M \mid \text{mouth}) = \prod_{i=1}^{M} p(a_i \mid \text{mouth})$$

Where does this break?