Eigenfaces

A Generative Model for Human Faces

Uncompressed Bitmaps

- What is a greyscale image?
- From Wikipedia's entry on Bitmap

Pixel storage [edit]

In typical uncompressed bitmaps, image pixels are generally stored with a color depth of 1, 4, 8, 16, 24, 32, 48, or 64 bits per pixel. Pixels of 8 bits and fewer can represent either grayscale or indexed color. An alpha channel (for transparency) may be stored in a separate bitmap, where it is similar to a greyscale bitmap, or in a fourth channel that, for example, converts 24-bit images to 32 bits per pixel.

- 8 bit: A matrix of numbers, 0 through 255.
- Octave
  - Assign a matrix to A
  - imshow(A)
- Linear Algebra
  - What is a matrix, but a vector? reshape(A, numel(A), 1)

Covariance

- http://en.wikipedia.org/wiki/Covariance
  - Covariance is a measure of how much two variables change together
- http://en.wikipedia.org/wiki/Covariance_matrix
  - Consider all the pixels. We want to know how pixel one changes with respect to pixel two…
  - The covariance matrix, Σ, will be 256x256 for a 16x16 image.
  - For a 256x256 image, the covariance matrix will have 65536 rows!
- How to compute the covariance matrix?
  - For 100 images, data = (65k x 100) matrix, with each image as a row.
  - Subtract the mean image from each to get A. Then…
  - $\Sigma = A A^T$
- But… Do we really need the whole thing?

PCA (Principle Component Analysis)

- http://en.wikipedia.org/wiki/Principal_component_analysis
  - PCA is mathematically defined as an orthogonal linear transformation that transforms the data to a new coordinate system such that the greatest variance by any projection of the data comes to lie on the first coordinate (called the first principal component)
- This amounts to an eigenvalue problem.
- Wanted: Eigenvalues of $\Sigma = A A^T$
  - Really, find eigenvalues of a 65k x 65k matrix?
  - Sort of…

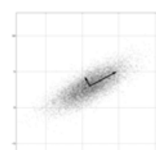

Find Eigenvalues of AA'

- Since AA' is too hard, let's see if we can find a relationship between eigenvalues/eigenvectors of AA' and A'A.
- Suppose **v** is an eigenvector of A'A, corresponding to λ.
- What does that mean about eigenvalues/eigenvectors of AA'?
- How is this helpful? Why is this easier to compute?

Result of PCA

- The higher the eigenvalue, the more variability is explained by that direction. Take the first 20. These are their eigenvectors.

How Well Do the Top 20 Do?

- Project a NEW image of a face onto those first 20 directions. We can display the face (and calculate per pixel error)
- How to project?
  - Observation 1: The eigenvectors will be orthogonal. Why?
  - Observation 2: We must normalize them. How?
  - Observation 3: We are in Euclidean space with the Euclidean inner product, dealing with orthonormal vectors. How can we can calculate the projection?

Result of Projection

- We get faces, and they look different (potentially recognizable)!

How Many Eigenvectors to Take?

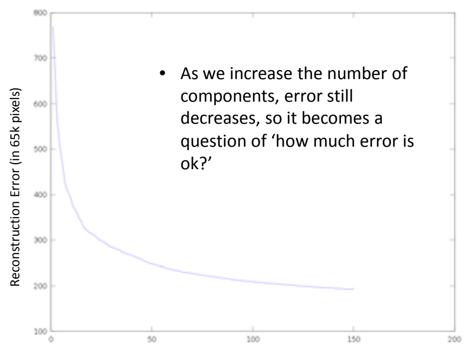

- As we increase the number of components, error still decreases, so it becomes a question of 'how much error is ok?'

Can We Do Better?

- Of course!
- Transform each image so that landmarks are aligned.

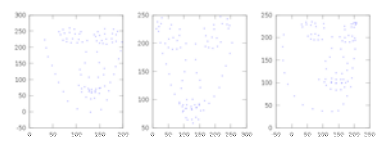

- Projections without Landmark Alignment

- Projections WITH Landmark Alignment

Want More Information?

- As always, wikipedia is a good starting place
  - http://en.wikipedia.org/wiki/Eigenface