Principal Component Analysis

Aryan Mokhtari, Santiago Paternain, and Alejandro Ribeiro
Dept. of Electrical and Systems Engineering
University of Pennsylvania
aribeiro@seas.upenn.edu
`http://www.seas.upenn.edu/users/~aribeiro/`

March 29, 2022

The DFT and iDFT with Unitary Matrices

The Discrete Fourier Transform with Unitary Matrices

Stochastic signals

Principal Component Analysis (PCA) transform

Dimensionality reduction

Principal Components

Face recognition

Signals Represented as Vectors in $\mathbb{R}^N$

- We are ready to write the DFT with better tools ⇒ Which will lead to abstract generalizations

- Write the signal x and the complex exponential e_{kN} as vectors in $\mathbb{R}^N$ ⇒ Call them x and e_{kN}

$$\mathrm{x} = \begin{bmatrix} x(0) \\ x(1) \\ \vdots \\ x(N-1) \end{bmatrix} \quad \Rightarrow \quad \mathrm{x}^H = (\mathrm{x}^*)^T = \begin{bmatrix} x(0), & x(1), & \ldots, & x(N-1) \end{bmatrix}$$

$$\mathrm{e}_{kN} = \frac{1}{\sqrt{N}} \begin{bmatrix} e^{j2\pi k0/N} \\ e^{j2\pi k1/N} \\ \vdots \\ e^{j2\pi k(N-1)/N} \end{bmatrix} \quad \Rightarrow \quad \mathrm{e}_{kN}^H = (\mathrm{e}_{kN}^*)^T = \frac{1}{\sqrt{N}} \begin{bmatrix} e^{-j2\pi k0/N}, & e^{-j2\pi k1/N}, & \ldots, & e^{-j2\pi k(N-1)/N} \end{bmatrix}$$

- We can now rewrite the DFT as ⇒ $y(k) = \langle \mathrm{x}, \mathrm{e}_{kN} \rangle = \mathrm{e}_{kN}^H \mathrm{x} = \frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} x(n) e^{-j2\pi kn/N}$

The DFT in Matrix Notation

- The kth DFT component $\tilde{x}(k)$ is the product of exponential $\mathbf{e}_{kN}^H$ with signal $\mathbf{x}$ $\Rightarrow y(k) = \mathbf{e}_{kN}^H \mathbf{x}$

- Define the DFT vector $\mathbf{y}$ as a stack of all N DFT components and stack individual definitions

$$\mathbf{y} = \begin{bmatrix} y(0) \\ y(1) \\ \vdots \\ y(N-1) \end{bmatrix} = \begin{bmatrix} \mathbf{e}_{0N}^H \mathbf{x} \\ \mathbf{e}_{1N}^H \mathbf{x} \\ \vdots \\ \mathbf{e}_{(N-1)N}^H \mathbf{x} \end{bmatrix} = \begin{bmatrix} \mathbf{e}_{0N}^H \\ \mathbf{e}_{1N}^H \\ \vdots \\ \mathbf{e}_{(N-1)N}^H \end{bmatrix} \mathbf{x} = \mathbf{F}^H \mathbf{x}$$

- Where in the last equality we defined the Fourier matrix Hermitian $\mathbf{F}^H$ to write $\mathbf{y} = \mathbf{F}^H \mathbf{x}$

The DFT Hermitian Matrix

- Each row of the DFT Hermitian is the Hermitian of a complex exponential of a different frequency

$$\mathsf{F}^H = \begin{bmatrix} \mathsf{e}_{0N}^H \\ \mathsf{e}_{1N}^H \\ \vdots \\ \mathsf{e}_{kN}^H \\ \vdots \\ \mathsf{e}_{(N-1)N}^H \end{bmatrix} = \frac{1}{\sqrt{N}} \begin{bmatrix} 1 & 1 & \cdots & 1 & & \\ 1 & e^{-j2\pi(1)(1)/N} & \cdots & e^{-j2\pi(1)(n)/N} & \cdots & e^{-j2\pi(1)(N-1)/N} \\ \vdots & \vdots & & \vdots & & \\ 1 & e^{-j2\pi(k)(1)/N} & \cdots & e^{-j2\pi(k)(n)/N} & \cdots & e^{-j2\pi(k)(N-1)/N} \\ \vdots & \vdots & & \vdots & & \\ 1 & e^{-j2\pi(N-1)(1)/N} & \cdots & e^{-j2\pi(N-1)(n)/N} & \cdots & e^{-j2\pi(N-1)(N-1)/N} \end{bmatrix}$$

- Each row of DFT Hermitian corresponds to a given frequency. We sweep different time indexes
- Each column of DFT Hermitian corresponds to a given time index. We sweep different frequencies

Visualizing the DFT as a Matrix Product

- Each row of a matrix-vector product entails sweeping the corresponding row of the matrix

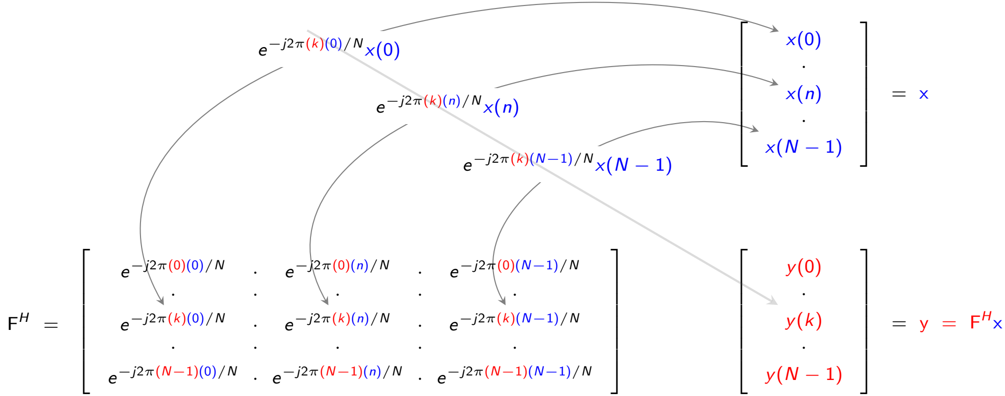

$$e^{-j2\pi(k)(0)/N}x(0)$$

$$e^{-j2\pi(k)(n)/N}x(n)$$

$$e^{-j2\pi(k)(N-1)/N}x(N-1)$$

$$\begin{bmatrix} x(0) \\ \cdot \\ x(n) \\ \cdot \\ x(N-1) \end{bmatrix} = \mathsf{x}$$

$$\mathsf{F}^H = \begin{bmatrix} e^{-j2\pi(0)(0)/N} & \cdot & e^{-j2\pi(0)(n)/N} & \cdot & e^{-j2\pi(0)(N-1)/N} \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ e^{-j2\pi(k)(0)/N} & \cdot & e^{-j2\pi(k)(n)/N} & \cdot & e^{-j2\pi(k)(N-1)/N} \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ e^{-j2\pi(N-1)(0)/N} & \cdot & e^{-j2\pi(N-1)(n)/N} & \cdot & e^{-j2\pi(N-1)(N-1)/N} \end{bmatrix}$$

$$\begin{bmatrix} y(0) \\ \cdot \\ y(k) \\ \cdot \\ y(N-1) \end{bmatrix} = \mathsf{y} = \mathsf{F}^H\mathsf{x}$$

The DFT Matrix

- The Hermitian of the DFT Hermitian matrix is the DFT matrix $\Rightarrow \mathsf{F} = (\mathsf{F}^H)^H$

$$\mathsf{F}^H = \begin{bmatrix} \mathsf{e}_{0N}^H \\ \mathsf{e}_{1N}^H \\ \vdots \\ \mathsf{e}_{kN}^H \\ \vdots \\ \mathsf{e}_{(N-1)N}^H \end{bmatrix} \quad \Rightarrow \quad \mathsf{F} = \begin{bmatrix} \mathsf{e}_{0N} & \mathsf{e}_{1N} & \cdots & \mathsf{e}_{kN} & \cdots & \mathsf{e}_{(N-1)N} \end{bmatrix}$$

- The conjugate of the kth row of the Hermitian F^H becomes the kth column of the DFT matrix F

The DFT Matrix

- Each column of the DFT matrix is a complex exponential of a different frequency

$$\mathsf{F} = \begin{bmatrix} \mathsf{e}_{0N} & \mathsf{e}_{1N} & \cdots & \mathsf{e}_{kN} & \cdots & \mathsf{e}_{(N-1)N} \end{bmatrix}$$

$$\mathsf{F} = \frac{1}{\sqrt{N}} \begin{bmatrix} 1 & 1 & \cdots & 1 & \cdots & 1 \\ 1 & e^{-j2\pi(1)(1)/N} & \cdots & e^{-j2\pi(k)(1)/N} & \cdots & e^{-j2\pi(N-1)(1)/N} \\ \vdots & \vdots & & \vdots & & \vdots \\ 1 & e^{-j2\pi(1)(n)/N} & \cdots & e^{-j2\pi(k)(n)/N} & \cdots & e^{-j2\pi(N-1)(n)/N} \\ \vdots & \vdots & & \vdots & & \vdots \\ 1 & e^{-j2\pi(1)(N-1)/N} & \cdots & e^{-j2\pi(k)(N-1)/N} & \cdots & e^{-j2\pi(N-1)(N-1)/N} \end{bmatrix}$$

- Each column corresponds to a given frequency. Each row corresponds to a given time index.

The product of the Hermitian Fourier Matrix and the Fourier Matrix

- The product of the Hermitian Fourier matrix F^H and the Fourier matrix F is given by

$$\begin{array}{cc} & \begin{bmatrix} \mathsf{e}_{0N} & \mathsf{e}_{1N} & \cdots & \mathsf{e}_{kN} & \cdots & \mathsf{e}_{(N-1)N} \end{bmatrix} \\ \begin{bmatrix} \mathsf{e}_{0N}^H \\ \mathsf{e}_{1N}^H \\ \vdots \\ \mathsf{e}_{kN}^H \\ \vdots \\ \mathsf{e}_{(N-1)N}^H \end{bmatrix} & \begin{bmatrix} \mathsf{e}_{0N}^H\,\mathsf{e}_{0N} & \mathsf{e}_{0N}^H\,\mathsf{e}_{1N} & \cdots & \mathsf{e}_{0N}^H\,\mathsf{e}_{kN} & \cdots & \mathsf{e}_{0N}^H\,\mathsf{e}_{(N-1)N} \\ \mathsf{e}_{1N}^H\,\mathsf{e}_{0N} & \mathsf{e}_{1N}^H\,\mathsf{e}_{1N} & \cdots & \mathsf{e}_{1N}^H\,\mathsf{e}_{kN} & \cdots & \mathsf{e}_{1N}^H\,\mathsf{e}_{(N-1)N} \\ \vdots & \vdots & & \vdots & & \vdots \\ \mathsf{e}_{kN}^H\,\mathsf{e}_{0N} & \mathsf{e}_{kN}^H\,\mathsf{e}_{1N} & \cdots & \mathsf{e}_{kN}^H\,\mathsf{e}_{kN} & \cdots & \mathsf{e}_{kN}^H\,\mathsf{e}_{(N-1)N} \\ \vdots & \vdots & & \vdots & & \vdots \\ \mathsf{e}_{(N-1)N}^H\,\mathsf{e}_{0N} & \mathsf{e}_{(N-1)N}^H\,\mathsf{e}_{1N} & \cdots & \mathsf{e}_{(N-1)N}^H\,\mathsf{e}_{kN} & \cdots & \mathsf{e}_{(N-1)N}^H\,\mathsf{e}_{(N-1)N} \end{bmatrix} = \mathsf{F}^H\mathsf{F} \end{array}$$

- The inner products in the diagonal are one and the inner products outside the diagonal are zero

The product of the Hermitian Fourier Matrix and the Fourier Matrix

- The product of the Hermitian Fourier matrix F^H and the Fourier matrix F is given by

$$\begin{array}{c|cccccc} & \mathrm{e}_{0N} & \mathrm{e}_{1N} & \cdots & \mathrm{e}_{kN} & \cdots & \mathrm{e}_{(N-1)N} \\ \hline \mathrm{e}^H_{0N} & 1 & 0 & \cdots & 0 & \cdots & 0 \\ \mathrm{e}^H_{1N} & 0 & 1 & \cdots & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & & \vdots & & \vdots \\ \mathrm{e}^H_{kN} & 0 & 0 & \cdots & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & & \vdots & & \vdots \\ \mathrm{e}^H_{(N-1)N} & 0 & 0 & \cdots & 0 & \cdots & 1 \end{array} = \mathrm{F}^H\mathrm{F}$$

- The inner products in the diagonal are one and the inner products outside the diagonal are zero

Orthogonality of the Fourier Basis

- We have therefore proved the following fundamental theorem

Theorem

The DFT matrix F *and its Hermitian are inverses of each other* $\Rightarrow$ *We say they are unitary*

$$\mathsf{F}^H\mathsf{F} \;=\; \mathsf{I} \;=\; \mathsf{F}\mathsf{F}^H$$

- This follows from, and is equivalent to, the orthonormality of complex exponentials

$\Rightarrow$ Orthonormality, the professional way.

The iDFT in Matrix Form

- We can also write the iDFT of y as a matrix product $\Rightarrow$ In this case the product is $\tilde{x} = Fy$

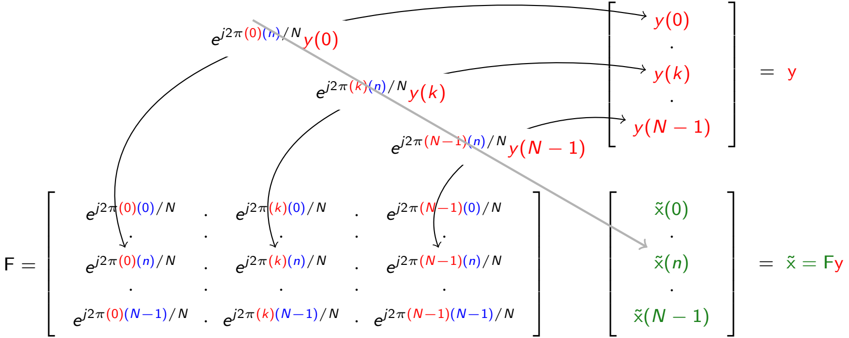

$$e^{j2\pi(0)(n)/N}y(0)$$

$$e^{j2\pi(k)(n)/N}y(k)$$

$$e^{j2\pi(N-1)(n)/N}y(N-1)$$

$$\begin{bmatrix} y(0) \\ \vdots \\ y(k) \\ \vdots \\ y(N-1) \end{bmatrix} = y$$

$$F = \begin{bmatrix} e^{j2\pi(0)(0)/N} & \cdots & e^{j2\pi(k)(0)/N} & \cdots & e^{j2\pi(N-1)(0)/N} \\ \vdots & \cdots & \vdots & \cdots & \vdots \\ e^{j2\pi(0)(n)/N} & \cdots & e^{j2\pi(k)(n)/N} & \cdots & e^{j2\pi(N-1)(n)/N} \\ \vdots & \cdots & \vdots & \cdots & \vdots \\ e^{j2\pi(0)(N-1)/N} & \cdots & e^{j2\pi(k)(N-1)/N} & \cdots & e^{j2\pi(N-1)(N-1)/N} \end{bmatrix}$$

$$\begin{bmatrix} \tilde{x}(0) \\ \vdots \\ \tilde{x}(n) \\ \vdots \\ \tilde{x}(N-1) \end{bmatrix} = \tilde{x} = Fy$$

Inverse theorem, like a pro

- When we proved theorems we had monkey steps and one smart step

$\Rightarrow$ That was orthonormality $\Rightarrow$ matrix F is unitary $\Rightarrow$ $F^H F = I$

Theorem

The iDFT is, indeed, the inverse of the DFT

Proof.

- Write $\tilde{x} = FX$ and $X = F^H x$ and exploit fact that F is unitary

$$\tilde{x} \;=\; FX \;=\; FF^H x \;=\; Ix \;=\; x$$ □

- Actually, this theorem would be true for any transform pair

$$X = T^H x \qquad \Longleftrightarrow \qquad \tilde{x} = TX$$

- As long as the transform matrix T is unitary $\Rightarrow$ $T^H T = I$

Energy conservation (Parseval) theorem, like a pro

Theorem

The DFT preserves energy $\Rightarrow \|\mathbf{x}\|^2 = \mathbf{x}^H\mathbf{x} = \mathbf{X}^H\mathbf{X} = \|\mathbf{X}\|^2$

Proof.

- Use iDFT to write $\mathbf{x} = \mathbf{F}\mathbf{X}$ and exploit fact that $\mathbf{F}$ is unitary

$$\|\mathbf{x}\|^2 \;=\; \mathbf{x}^H\mathbf{x} \;=\; (\mathbf{F}\mathbf{X})^H\,\mathbf{F}\mathbf{X} \;=\; \mathbf{X}^H\mathbf{F}^H\mathbf{F}\mathbf{X} \;=\; \mathbf{X}^H\mathbf{X} \;=\; \|\mathbf{X}\|^2 \qquad \square$$

- This theorem would also be true for any transform pair

$$\mathbf{X} = \mathbf{T}^H\mathbf{x} \qquad \Longleftrightarrow \qquad \tilde{\mathbf{x}} = \mathbf{T}\mathbf{X}$$

- As long as the transform matrix $\mathbf{T}$ is unitary $\Rightarrow \mathbf{T}^H\mathbf{T} = \mathbf{I}$

The discrete cosine transform

- Are there other useful transforms defined by unitary matrices T?

$\Rightarrow$ Many. One we have already found is the DCT

- Define the inverse DCT matrix C to write the iDCT as $\tilde{\mathrm{x}} = \mathrm{CX}$

$$\mathrm{C} = \frac{1}{\sqrt{N}} \begin{bmatrix} 1 & 1 & \cdots & 1 \\ 1 & \sqrt{2}\cos\left[\frac{2\pi(1)((1)+1/2)}{N}\right] & \cdots & \sqrt{2}\cos\left[\frac{2\pi(N-1)((1)+1/2)}{N}\right] \\ \vdots & \vdots & \ddots & \vdots \\ 1 & \sqrt{2}\cos\left[\frac{2\pi(1)((N-1)+1/2)}{N}\right] & \cdots & \sqrt{2}\cos\left[\frac{2\pi(N-1)((N-1)+1/2)}{N}\right] \end{bmatrix}$$

- It is ready to verify that C is unitary (the cosines are orthonormal)
- From where the inverse and energy conservation theorems follow

$\Rightarrow$ Proofs hold for all unitary matrices, C in particular

Designing transformations adapted to signals

- A basic information processing theory can be built for any T
- Then, why do we specifically choose the DFT? Or the DCT?

  ⇒ Oscillations represent different rates of change

  ⇒ Different rates of change represent different aspects of a signal
- Not a panacea, though. E.g., F^H is independent of the signal
- If we know something about signal, we should use it to build better T
- A way of "knowing something" is a stochastic model of the signal
- PCA: Principal component analysis

  ⇒ Use the eigenvectors of the covariance matrix to build T

Stochastic signals

The Discrete Fourier Transform with Unitary Matrices

Stochastic signals

Principal Component Analysis (PCA) transform

Dimensionality reduction

Principal Components

Face recognition

Random Variables

- A random variable X models a random phenomena
  - ⇒ One in which many different outcomes are possible
  - ⇒ And one in which some outcomes may be more likely than others
- Thus, a random variable represents two things
  - ⇒ All possible outcomes and their respective likelihoods

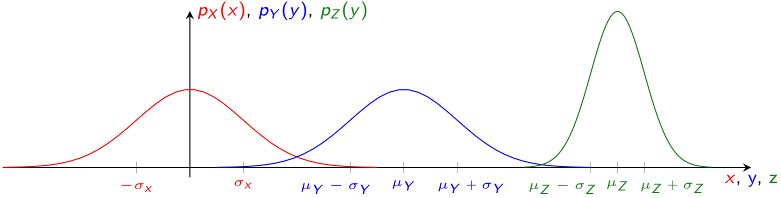

- Random variable X takes values around 0 and Y values around μ_Y
- Z takes values around μ_Z and the values are more concentrated

Probabilities

- Probabilities measure the likelihood of observing different outcomes

  $\Rightarrow$ Larger probability means an outcome that is more likely

  $\Rightarrow$ Or, observed more often when seeing many realizations

- Random variables represented by uppercase $\Rightarrow$ E.g., X
- Values that it can take represented by lowercase $\Rightarrow$ E.g., x
- The probability that X takes values between x and x' is written as

$$\mathrm{P}\left(x < X \leq x'\right)$$

- Here, we describe probabilities with density functions (pdf) $\Rightarrow$ $p_X(x)$

$$\mathrm{P}\left(x < X \leq x'\right) = \int_x^{x'} p_X(u)\, du$$

- $p_X(x) \approx$ How likely random variable X is to take a value around x

Gaussian random variables

- A random variable X is Gaussian (or Normal) if its pdf is of the form

$$p_X(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-(x-\mu)^2/\sigma^2}$$

- The mean μ determines center. The variance σ^2 determines width

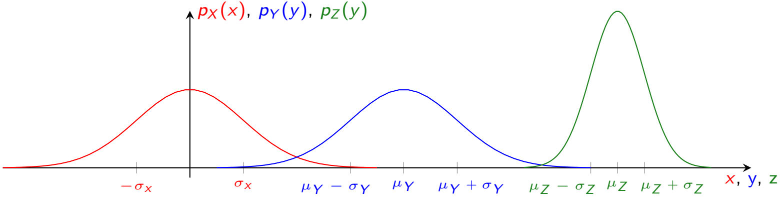

- Means satisfy $0 = \mu_X < \mu_Y < \mu_Z$. Variances are $\sigma_X^2 = \sigma_Y^2 > \sigma_Z^2$

Expectation

- Expectation of random variable is an average weighted by likelihoods

$$\mathbb{E}\left[X\right] = \int_{-\infty}^{\infty} x p_X(x)\, dx$$

- Regular average $\Rightarrow$ Sum all values and divide by number of values

- Expectation $\Rightarrow$ Weight values x by their relative likelihoods $p_X(x)$
- For a Gaussian random variable X the expectation is the mean μ

$$\mathbb{E}\left[X\right] = \int_{-\infty}^{\infty} x \frac{1}{\sqrt{2\pi}\sigma} e^{-(x-\mu)^2/\sigma^2}\, dx = \mu$$

- Not difficult to evaluate integral, but besides the point to do so here

Variance

- Measure of variability around the mean weighted by likelihoods

$$\text{var}\left[X\right] = \mathbb{E}\left[\left(X - \mathbb{E}\left[X\right]\right)^2\right] = \int_{-\infty}^{\infty} \left(x - \mathbb{E}\left[X\right]\right)^2 p_X(x)\, dx$$

- Large variance $\equiv$ likely values are spread out around the mean

- Small variance $\equiv$ likely values are concentrated around the mean
- For a Gaussian random variable X the variance is the variance σ^2

$$\text{var}\left[X\right] = \int_{-\infty}^{\infty} \left(x - \mathbb{E}\left[X\right]\right)^2 \frac{1}{\sqrt{2\pi}\sigma} e^{-(x-\mu)^2/\sigma^2}\, dx = \sigma^2$$

- Not difficult to evaluate either. But also besides the point here

Random signals

- A random signal X is a collection of random variables (length N)

$$\mathrm{X} = \left[X(0),\ X(1),\ \ldots,\ X(N-1)\right]^T$$

- Each of the random variables has its own pdf $\Rightarrow p_{X(n)}(x)$

- This pdf describes the likelihood of $X(n)$ taking a value around x
- This is not a sufficient description. Joint outcomes also important

- Joint pdf $p_{\mathrm{X}}(\mathrm{x})$ says how likely signal X is to be found around x

$$\mathrm{P}\left(\mathrm{x} \in \mathcal{X}\right) = \iint_{\mathcal{X}} p_{\mathrm{X}}(\mathrm{x})\, d\mathrm{x}$$

- The individual pdfs $p_{X(n)}(x)$ are said to be marginal pdfs

Face images

- Random signal X $\Rightarrow$ All possible images of human faces
- More manageable $\Rightarrow$ X is a collection of 400 face images

  $\Rightarrow$ The random variable represents all the images

  $\Rightarrow$ The likelihood of each of them being chosen. E.g., 1/400 each

- Random variable specified by all outcomes and respective probabilities

Vectorization

- Do observe that the dataset consists of images $\equiv$ matrices

- Each image is stored in a matrix of size 112×92

$$\mathrm{M}_i = \begin{bmatrix} m_{1,1} & m_{1,2} & \cdots & m_{1,92} \\ m_{2,1} & m_{2,2} & \cdots & m_{2,92} \\ \vdots & \vdots & \ddots & \vdots \\ m_{112,1} & m_{112,2} & \cdots & m_{112,92} \end{bmatrix}$$

- Stack columns of image M_i into the vector x_i with length $10,304$

$$\mathrm{x}_i = \Big[m_{1,1},\ m_{21},\ \ldots,\ m_{112,1},\ m_{1,2},\ m_{2,2},\ \ldots,\ m_{112,2},\ \vdots,\ m_{1,92},\ m_{2,92},\ \ldots,\ m_{112,92} \Big]^T$$

- Images are matrices $\mathrm{M}_i \in \mathbb{R}^{112 \times 92}$. Signals are vectors $\mathrm{x}_i \in \mathbb{R}^{10,304}$

Realizations

- Realization x is an individual face pulled from set of possible outcomes
- Three possible realizations shown

- Realizations are just regular signals. Nothing random about them

Expectation, variance and covariance

- Signal's expectation is the concatenation of individual expectations

$$\mathbb{E}\left[\mathrm{X}\right] = \Big[\mathbb{E}\left[X(0)\right],\ \mathbb{E}\left[X(1)\right],\ \ldots\ \mathbb{E}\left[X(N-1)\right]\Big]^T = \iint \mathrm{x} p_{\mathrm{X}}(\mathrm{x})\, d\mathrm{x}$$

- Variance of nth element $\Rightarrow \Sigma_{nn} = \mathrm{var}\left[X(n)\right] = \mathbb{E}\left[\big(X(n) - \mathbb{E}\left[X(n)\right]\big)^2\right]$
- Measures variability of nth component

- Covariance between the signal components $X(n)$ and $X(m)$

$$\Sigma_{nm} = \mathbb{E}\Big[\big(X(n) - \mathbb{E}\left[X(n)\right]\big)\big(X(m) - \mathbb{E}\left[X(m)\right]\big)\Big] = \Sigma_{mn}$$

- Measures how much $X(n)$ predicts $X(m)$. Love, hate, and indifference
  - $\Rightarrow \Sigma_{nm} = 0$, components are unrelated. They are orthogonal
  - $\Rightarrow \Sigma_{nm} > 0$ ($\Sigma_{nm} < 0$), move in same (opposite) direction

Covariance matrix

- Assume that $\mathbb{E}\left[\mathrm{X}\right] = 0$ so that covariances are $\Sigma_{nm} = \mathbb{E}\left[X(n)X(m)\right]$

- Consider the expectation $\mathbb{E}\left[\mathrm{x}\mathrm{x}^T\right]$ of the (outer) product $\mathrm{x}\mathrm{x}^T$
- We can write the outer product $\mathrm{x}\mathrm{x}^T$ as

$$\mathrm{x}\mathrm{x}^T = \begin{bmatrix} x(0)x(0) & \cdots & x(0)x(n) & \cdots & x(0)x(N-1) \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ x(n)x(0) & \cdots & x(n)x(n) & \cdots & x(n)x(N-1) \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ x(N-1)x(0) & \cdots & x(N-1)x(n) & \cdots & x(N-1)x(N-1) \end{bmatrix}$$

-

Covariance matrix

- Assume that $\mathbb{E}\left[\mathrm{X}\right] = 0$ so that covariances are $\Sigma_{nm} = \mathbb{E}\left[X(n)X(m)\right]$

- Consider the expectation $\mathbb{E}\left[\mathrm{x}\mathrm{x}^T\right]$ of the (outer) product $\mathrm{x}\mathrm{x}^T$
- Expectation $\mathbb{E}\left[\mathrm{x}\mathrm{x}^T\right]$ implies expectation of each individual element

$$\mathbb{E}\left[\mathrm{x}\mathrm{x}^T\right] = \begin{bmatrix} \mathbb{E}[x(0)x(0)] & \cdots & \mathbb{E}[x(0)x(n)] & \cdots & \mathbb{E}[x(0)x(N-1)] \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \mathbb{E}[x(n)x(0)] & \cdots & \mathbb{E}[x(n)x(n)] & \cdots & \mathbb{E}[x(n)x(N-1)] \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \mathbb{E}[x(N-1)x(0)] & \cdots & \mathbb{E}[x(N-1)x(n)] & \cdots & \mathbb{E}[x(N-1)x(N-1)] \end{bmatrix}$$

-

Covariance matrix

- Assume that $\mathbb{E}[\mathrm{X}] = 0$ so that covariances are $\Sigma_{nm} = \mathbb{E}[X(n)X(m)]$

- Consider the expectation $\mathbb{E}\left[\mathrm{x}\mathrm{x}^T\right]$ of the (outer) product $\mathrm{x}\mathrm{x}^T$
- The (n, m) element of the matrix $\mathbb{E}\left[\mathrm{x}\mathrm{x}^T\right]$ is the covariance Σ_{nm}

$$\mathbb{E}\left[\mathrm{x}\mathrm{x}^T\right] = \begin{bmatrix} \Sigma_{00} & \cdots & \Sigma_{0n} & \cdots & \Sigma_{0(N-1)} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \Sigma_{n0} & \cdots & \Sigma_{nn} & \cdots & \Sigma_{n(N-1)} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \Sigma_{(N-1)0} & \cdots & \Sigma_{(N-1)n} & \cdots & \Sigma_{(N-1)(N-1)} \end{bmatrix}$$

- Define the covariance matrix of random signal X as $\Sigma := \mathbb{E}\left[\mathrm{x}\mathrm{x}^T\right]$

Definition of covariance matrix

- When the mean is not null define the covariance matrix of X as

$$\Sigma := \mathbb{E}\left[\left(\mathrm{x} - \mathbb{E}\left[\mathrm{x}\right]\right)\left(\mathrm{x} - \mathbb{E}\left[\mathrm{x}\right]\right)^T\right]$$

- As before, the (n, m) element of Σ is the covariance Σ_{nm}

$$((\Sigma))_{nm} = \mathbb{E}\left[\left(X(n) - \mathbb{E}\left[X(n)\right]\right)\left(X(m) - \mathbb{E}\left[X(m)\right]\right)\right] = \Sigma_{nm}$$

- The covariance matrix Σ is an arrangement of the covariances Σ_{nm}

- The diagonal of Σ contains the (auto)variances $\Sigma_{nn} = \mathrm{var}\left[X(n)\right]$

- Covariance matrix is symmetric $\Rightarrow ((\Sigma))_{nm} = \Sigma_{nm} = \Sigma_{mn} = ((\Sigma))_{mn}$

Mean of face images

- All images are equally likely $\Rightarrow$ probability 1/400 for each image
- The mean face is the regular average $\Rightarrow \mathbb{E}[\mathbf{x}] = \frac{1}{\ } \sum^{400} \mathbf{x}_i$

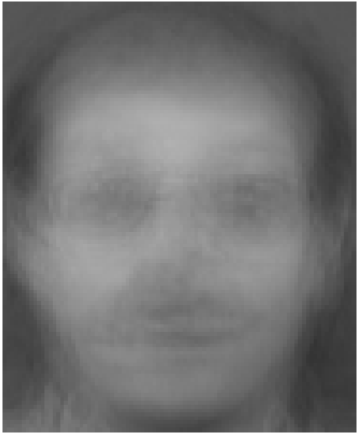

- Average image looks something, sort of, like an average face

Covariance matrix of face images

- Covariance matrix $\Rightarrow \Sigma = \frac{1}{400}\sum_{i=1}^{400}\left(\mathrm{x}_i - \mathbb{E}\left[\mathrm{x}\right]\right)\left(\mathrm{x}_i - \mathbb{E}\left[\mathrm{x}\right]\right)^T$

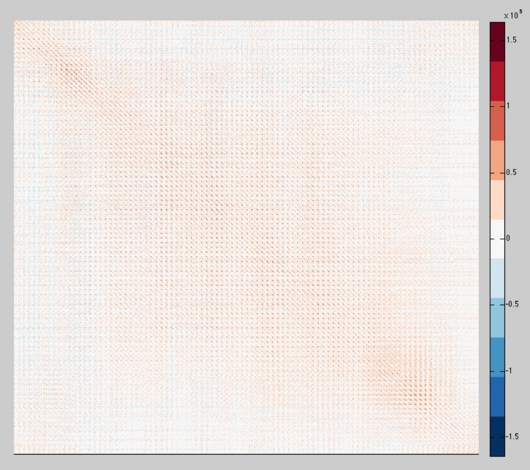

- Heat map of covariance matrix Σ shown on left
- Large correlation values around diagonal
- Large correlation values every 112 elements (jump a row on matrix)

Principal Component Analysis (PCA) transform

The Discrete Fourier Transform with Unitary Matrices

Stochastic signals

Principal Component Analysis (PCA) transform

Dimensionality reduction

Principal Components

Face recognition

Eigenvectors and eigenvalues of covariance matrix

- Consider a vector with N elements $\Rightarrow \mathrm{v} = [v(0), v(1), \ldots, v(N-1)]$
- We say that v is an eigenvector of Σ if for some scalar $\lambda \in \mathbb{R}$

$$\Sigma \mathrm{v} = \lambda \mathrm{v}$$

- We say that λ is the eigenvalue associated to v

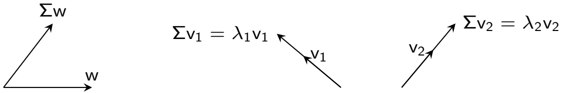

- In general, non-eigenvectors w and Σw point in different directions
- But for eigenvectors v, the product vector Σv is collinear with v

Normalization

- If v is an eigenvector, αv is also an eigenvector for any scalar $\alpha \in \mathbb{R}$,

$$\Sigma(\alpha \mathrm{v}) = \alpha(\Sigma \mathrm{v}) = \alpha\lambda \mathrm{v} = \lambda(\alpha \mathrm{v})$$

- Eigenvectors are defined up to a constant
- We use normalized eigenvectors with unit energy $\Rightarrow \|\mathrm{v}\|^2 = 1$

- If we compute v with $\|\mathrm{v}\|^2 \neq 1$ replace v with $\mathrm{v}/\|\mathrm{v}\|$
- There are N eigenvalues and distinct associated eigenvectors

$\Rightarrow$ Some technical qualifications are needed in this statement

Ordering

Theorem

The eigenvalues of Σ *are real and nonnegative* $\Rightarrow \lambda \in \mathbb{R}$ *and* $\lambda \geq 0$

Proof.

- Begin by observing that we can write $\lambda = \mathrm{v}^H \Sigma \mathrm{v} / \|\mathrm{v}\|^2$. Indeed

$$\mathrm{v}^H \Sigma \mathrm{v} = \mathrm{v}^H (\Sigma \mathrm{v}) = \mathrm{v}^H (\lambda \mathrm{v}) = \lambda \mathrm{v}^H \mathrm{v} = \lambda \|\mathrm{v}\|^2$$

- Complete by showing that $\mathrm{v}^T \Sigma \mathrm{v}$ is nonnegative. Indeed (assume $\mathbb{E}[\mathrm{x}] = 0$)

$$\mathrm{v}^H \Sigma \mathrm{v} = \mathrm{v}^H \mathbb{E}\left[\mathrm{x}\mathrm{x}^H\right] \mathrm{v} = \mathbb{E}\left[\mathrm{v}^H \mathrm{x}\mathrm{x}^H \mathrm{v}\right] = \mathbb{E}\left[\left(\mathrm{v}^H \mathrm{x}\right)\left(\mathrm{x}^H \mathrm{v}\right)\right] = \mathbb{E}\left[\left(\mathrm{v}^H \mathrm{x}\right)^2\right] \geq 0$$

□

- Order eigenvalues from largest to smallest $\Rightarrow \lambda_0 \geq \lambda_1 \geq \ldots \geq \lambda_{N-1}$
- Eigenvectors inherit order $\Rightarrow \mathrm{v}_0, \mathrm{v}_1, \ldots, \mathrm{v}_{N-1}$
- The *n*th eigenvector of Σ is associated with its *n*th largest eigenvalue

Eigenvectors are orthonormal

Theorem

Eigenvectors of Σ associated with different eigenvalues are orthogonal

Proof.

- Normalized eigenvectors v and u associated with eigenvalues $\lambda \neq \mu$

$$\Sigma \mathrm{v} = \lambda \mathrm{v}, \qquad \Sigma \mathrm{u} = \mu \mathrm{u}$$

- Since the matrix Σ is symmetric we have $\Sigma^H = \Sigma$, and it follows

$$\mathrm{u}^H \Sigma \mathrm{v} = \left(\mathrm{u}^H \Sigma \mathrm{v}\right)^H = \mathrm{v}^H \Sigma^H \mathrm{u} = \mathrm{v}^H \Sigma \mathrm{u}$$

- Make $\Sigma \mathrm{v} = \lambda \mathrm{v}$ on the leftmost side and $\Sigma \mathrm{u} = \mu \mathrm{u}$ on the rightmost

$$\mathrm{u}^H \lambda \mathrm{v} = \lambda \mathrm{u}^H \mathrm{v} = \mu \mathrm{v}^H \mathrm{u} = \mathrm{v}^H \mu \mathrm{u}$$

- Eigenvalues are different $\Rightarrow$ Relationship can only be true if $\mathrm{v}^H \mathrm{u} = 0$

□

Eigenvectors of face images (1D)

- One dimensional representation of first four eigenvectors v_0, v_1, v_2, v_3

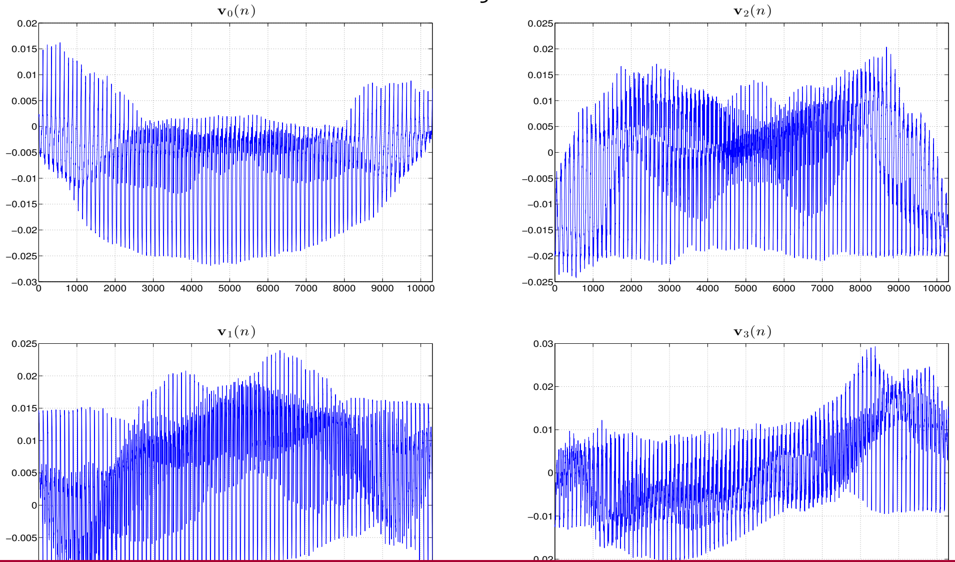

Eigenvectors of face images (2D)

- Two dimensional representation of first four eigenvectors v_0, v_1, v_2, v_3

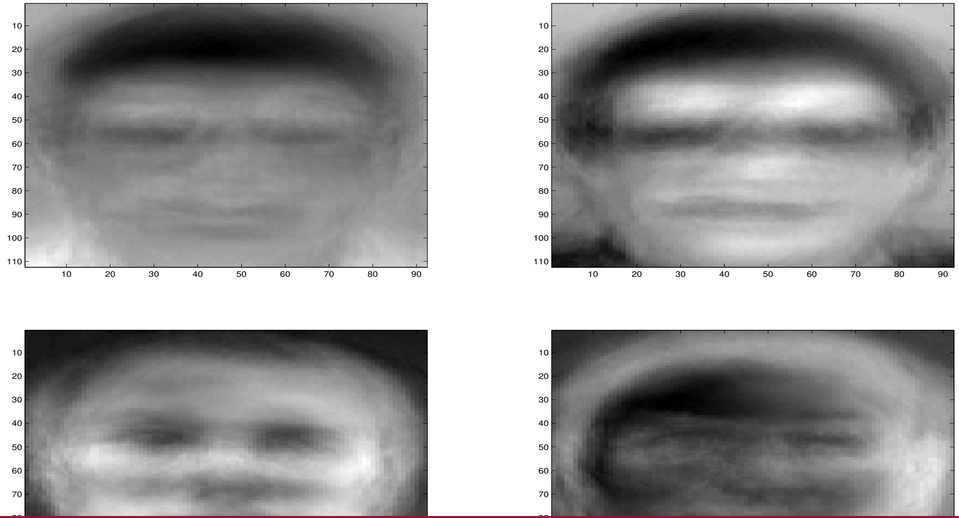

Eigenvector matrix

- Define the matrix T whose kth column is the kth eigenvector of Σ

$$\mathsf{T} = [\mathsf{v}_0, \mathsf{v}_1, \ldots, \mathsf{v}_{N-1}]$$

- Since the eigenvectors v_k are orthonormal, the product $\mathsf{T}^H\mathsf{T}$ is

$$\begin{bmatrix} \mathsf{v}_0 & \cdots & \mathsf{v}_k & \cdots & \mathsf{v}_{N-1} \end{bmatrix}$$

$$\mathsf{T}^H\mathsf{T} = \begin{bmatrix} \mathsf{v}_0^H \\ \vdots \\ \mathsf{v}_k^H \\ \vdots \\ \mathsf{v}_{N-1}^H \end{bmatrix} \begin{bmatrix} \mathsf{v}_0^H\mathsf{v}_0 & \cdots & \mathsf{v}_1^H\mathsf{v}_k & \cdots & \mathsf{v}_0^H\mathsf{v}_{N-1} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \mathsf{v}_k^H\mathsf{v}_0 & \cdots & \mathsf{v}_k^H\mathsf{v}_k & \cdots & \mathsf{v}_k^H\mathsf{v}_{N-1} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \mathsf{v}_{N-1}^H\mathsf{v}_{N-1} & \cdots & \mathsf{v}_{N-1}^H\mathsf{v}_k & \cdots & \mathsf{v}_{N-1}^H\mathsf{v}_{N-1} \end{bmatrix} = \begin{bmatrix} 1 & \cdots & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ 0 & \cdots & 1 & \cdots & 0 \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & \cdots & 1 \end{bmatrix}$$

- The eigenvector matrix T is unitary $\Rightarrow \mathsf{T}^H\mathsf{T} = \mathsf{I}$

Principal component analysis transform

- Any unitary T can be used to define an info processing transform
- Define principal component analysis (PCA) transform $\Rightarrow \mathsf{y} = \mathsf{T}^H \mathsf{x}$

- And the inverse (i)PCA transform $\Rightarrow \tilde{\mathsf{x}} = \mathsf{T}\mathsf{y}$
- Since T is unitary, iPCA is, indeed, the inverse of the PCA

$$\tilde{\mathsf{x}} \;=\; \mathsf{T}\mathsf{y} \;=\; \mathsf{T}\left(\mathsf{T}^H\mathsf{x}\right) \;=\; \mathsf{T}\mathsf{T}^H\mathsf{x} \;=\; \mathsf{I}\mathsf{x} \;=\; \mathsf{x}$$

- Thus y is an equivalent representation of x $\Rightarrow$ Back and forth
- And, also because T is unitary, Parseval's theorem holds

$$\|\mathsf{x}\|^2 \;=\; \mathsf{x}^H\mathsf{x} \;=\; (\mathsf{T}\mathsf{y})^H\,\mathsf{T}\mathsf{y} \;=\; \mathsf{y}^H\mathsf{T}^H\mathsf{T}\mathsf{y} \;=\; \mathsf{y}^H\mathsf{y} \;=\; \|\mathsf{y}\|^2$$

- Modifying elements y_k means altering energy composition of signal

Discussions

- The PCA transform is defined for any signal (vector) x

  $\Rightarrow$ But we expect to work well only when x is a realization of X

- Write the iPCA in expanded form and compare with the iDFT

$$x(n) = \sum_{k=0}^{N-1} y(k) v_k(n) \quad \Leftrightarrow \quad x(n) = \sum_{k=0}^{N-1} X(k) e_{kN}(n)$$

- The same except that they use different bases for the expansion

- Still, like developing a new sense.

- But not one that is generic. Rather, adapted to the random signal X

Coefficients of a projected face image

- PCA transform coefficients for given face image with 10,304 pixels

- Substantial energy in the first 15 PCA coefficients $y(k)$ with $k \leq 15$
- Almost all energy in the first 50 PCA coefficients $y(k)$ with $k \leq 50$

  $\Rightarrow$ This is a compression factor of more than 200

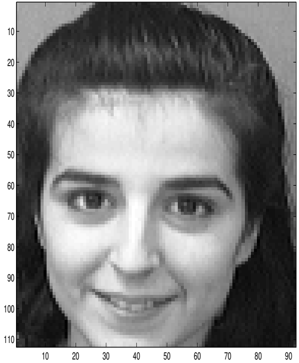

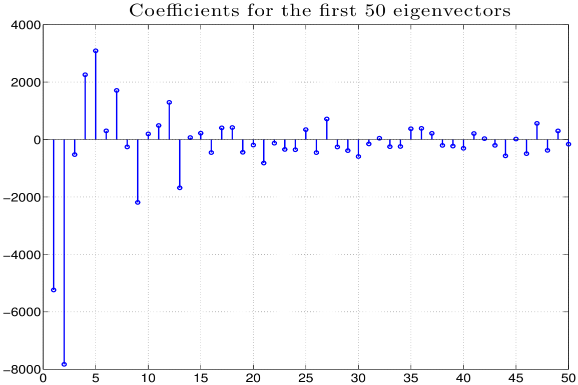

Reconstructed face images

- Reconstructed image for increasing number of PCA coefficients
  ⇒ Increasing number of coefficients increases accuracy.
  ⇒ Using 50 coefficients suffices

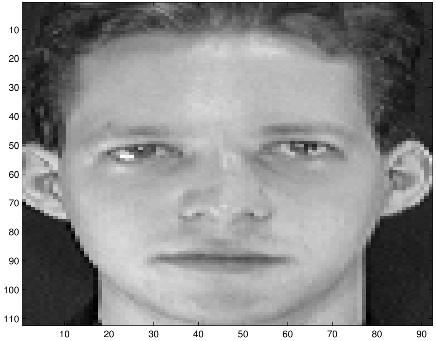

Figure: image

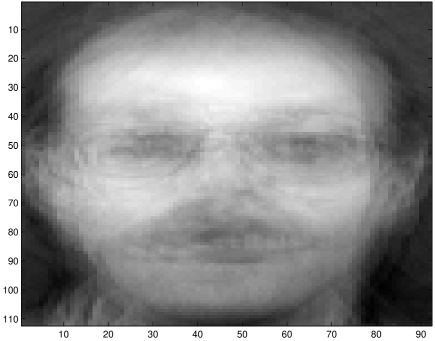

Figure: No. P.C.s = 1

Reconstructed face images

- Reconstructed image for increasing number of PCA coefficients
  $\Rightarrow$ Increasing number of coefficients increases accuracy.
  $\Rightarrow$ Using 50 coefficients suffices

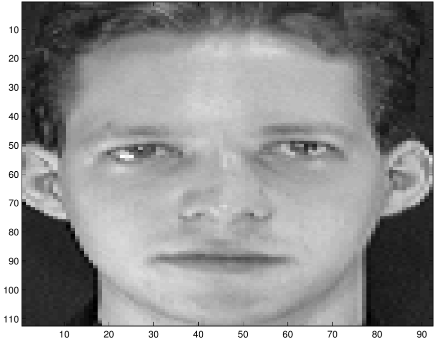

Figure: image

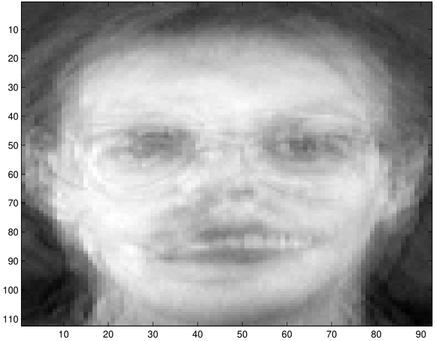

Figure: No. P.C.s = 5

Reconstructed face images

- Reconstructed image for increasing number of PCA coefficients
  - ⇒ Increasing number of coefficients increases accuracy.
  - ⇒ Using 50 coefficients suffices

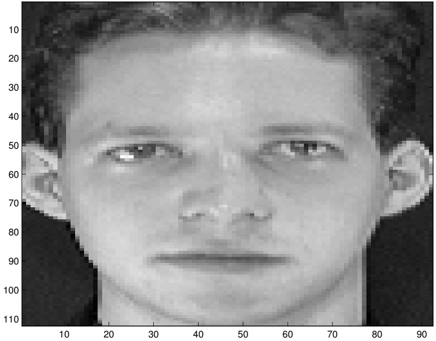

Figure: image

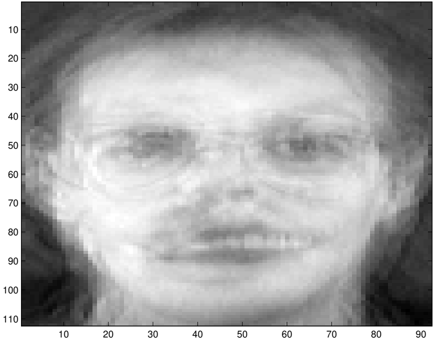

Figure: No. P.C.s $= 10$

Reconstructed face images

- Reconstructed image for increasing number of PCA coefficients
  - $\Rightarrow$ Increasing number of coefficients increases accuracy.
  - $\Rightarrow$ Using 50 coefficients suffices

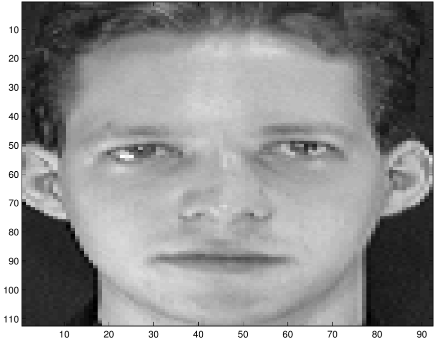

Figure: image

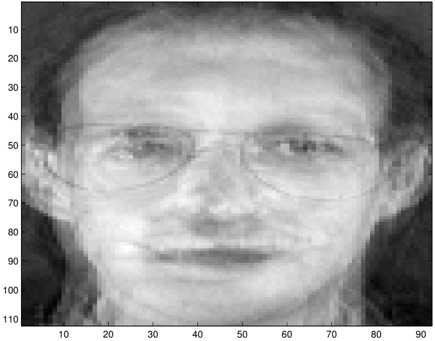

Figure: No. P.C.s = 20

Reconstructed face images

- Reconstructed image for increasing number of PCA coefficients
  - ⇒ Increasing number of coefficients increases accuracy.
  - ⇒ Using 50 coefficients suffices

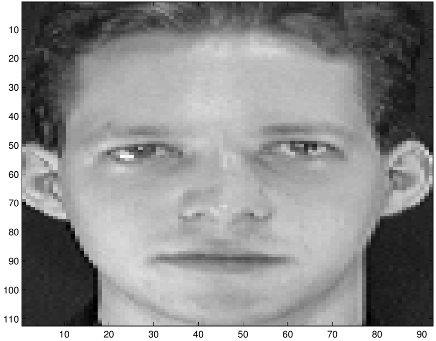

Figure: image

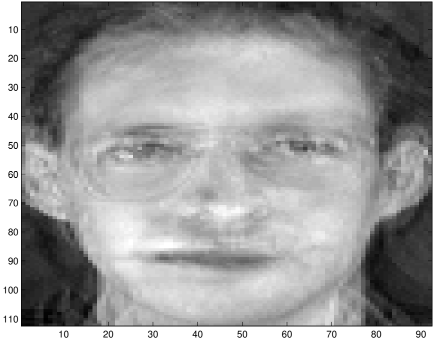

Figure: No. P.C.s = 30

Reconstructed face images

- Reconstructed image for increasing number of PCA coefficients
  - ⇒ Increasing number of coefficients increases accuracy.
  - ⇒ Using 50 coefficients suffices

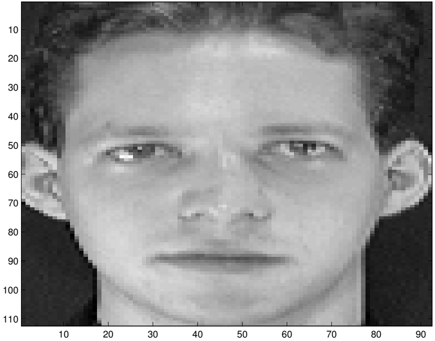

Figure: image

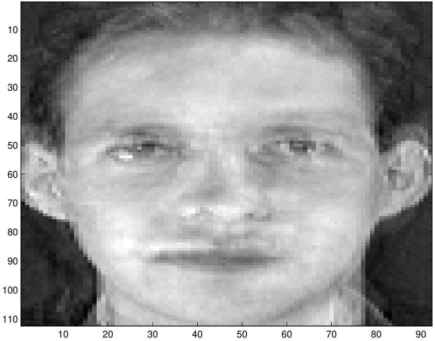

Figure: No. P.C.s = 40

Reconstructed face images

- Reconstructed image for increasing number of PCA coefficients
  - ⇒ Increasing number of coefficients increases accuracy.
  - ⇒ Using 50 coefficients suffices

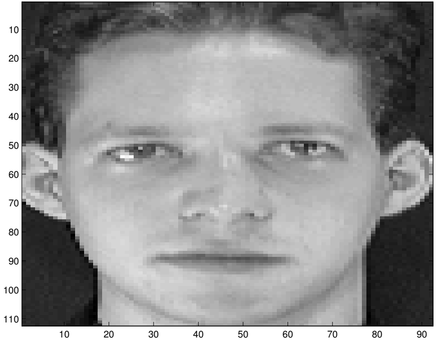

Figure: image

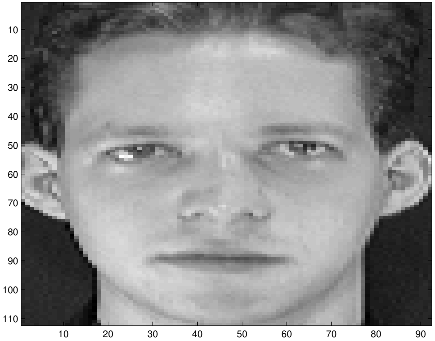

Figure: No. P.C.s = 50

Coefficients of the same person

- PCA transform y for two different pictures of the same person
- Coefficients are similar, even if pose and attitude are different
  - $\Rightarrow$ E.g., first two coefficients almost identical

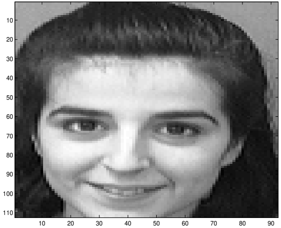

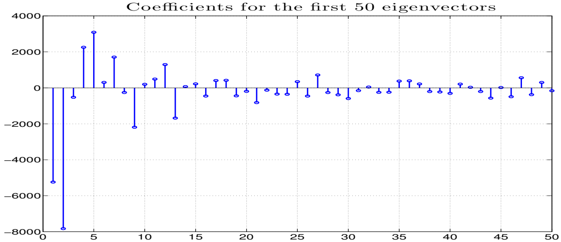
Coefficients for the first 50 eigenvectors

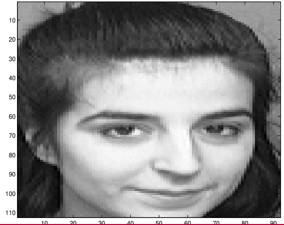

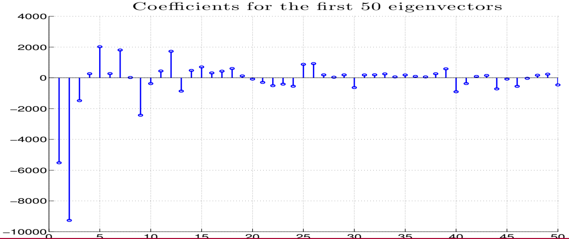
Coefficients for the first 50 eigenvectors

Coefficients of different persons

- PCA transform y for pictures of different persons
- Similar pose and attitude, but PCA coefficients are still different

$\Rightarrow$ Can be used to perform face recognition. More later

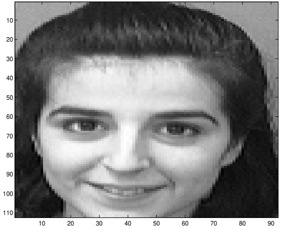

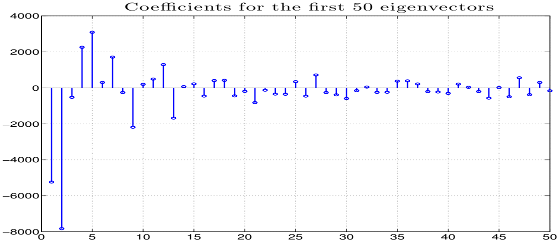

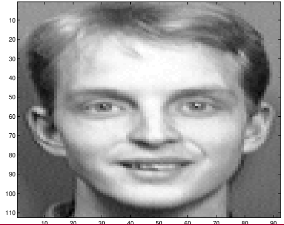

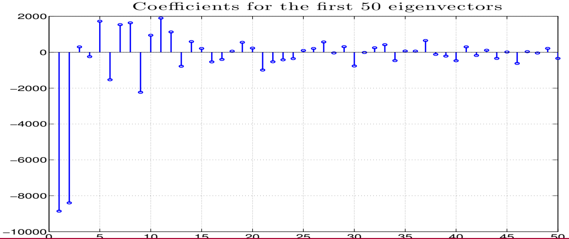

Dimensionality reduction

The Discrete Fourier Transform with Unitary Matrices

Stochastic signals

Principal Component Analysis (PCA) transform

Dimensionality reduction

Principal Components

Face recognition

Compression with the DFT

- Transform signal x into frequency domain with DFT $\mathrm{X} = \mathrm{F}^H \mathrm{x}$

- Recover x from X through iDFT matrix multiplication $\mathrm{x} = \mathrm{F}\mathrm{X}$
- We compress by retaining $K < N$ DFT coefficients to write

$$\tilde{\mathrm{x}}(n) = \sum_{k=0}^{K-1} X(k) e^{j2\pi kn/N}$$

- Equivalently, we define the compressed DFT as

$$\tilde{\mathrm{X}}(k) = X(k) \quad \text{for} \quad k < K, \qquad \tilde{\mathrm{X}}(k) = 0 \text{ otherwise}$$

- Reconstructed signal is obtained with iDFT $\Rightarrow \tilde{\mathrm{x}} = \mathrm{F}\tilde{\mathrm{X}}$

Compression with the PCA

- Transform signal x into eigenvector domain with PCA $\mathrm{y} = \mathrm{T}^H \mathrm{x}$

- Recover x from y through iPCA matrix multiplication $\mathrm{x} = \mathrm{T}\mathrm{y}$
- We compress by retaining $K < N$ PCA coefficients to write

$$\tilde{\mathrm{x}}(n) = \sum_{k=0}^{K-1} y(k) \mathrm{v}_k(n)$$

- Equivalently, we define the compressed PCA as

$$\tilde{\mathrm{y}}(k) = y(k) \quad \text{for} \quad k < K, \qquad \tilde{\mathrm{y}}(k) = 0 \text{ otherwise}$$

- Reconstructed signal is obtained with iPCA $\Rightarrow \tilde{\mathrm{x}} = \mathrm{T}\tilde{\mathrm{y}}$

Why keeping the first K coefficients?

- Why do we keep the first K DFT coefficients?

  ⇒ Because faster oscillations tend to represent faster variation

  ⇒ Also, not always, sometimes we keep the largest coefficients

- Why do we keep the first K PCA coefficients?

  ⇒ Eigenvectors with lower ordinality have larger eigenvalues

  ⇒ Larger eigenvalues entail more variability

  ⇒ And more variability signifies more dominant features

- Eigenvectors with large ordinality represent finer signal features

  ⇒ And can often be omitted

Dimensionality reduction

- PCA compression is (more accurately) called dimensionality reduction
  $\Rightarrow$ Do not compress signal. Reduce number of dimensions

$$\Sigma = \begin{bmatrix} 3/2 & 1/2 \\ 1/2 & 3/2 \end{bmatrix}$$

- Covariance eigenvectors mix coordinates

$$\mathrm{v}_0 = \begin{bmatrix} 1 \\ 1 \end{bmatrix} \quad \mathrm{v}_1 = \begin{bmatrix} 1 \\ -1 \end{bmatrix}$$

- Eigenvalues are $\lambda_0 = 2$ and $\lambda_1 = 1$

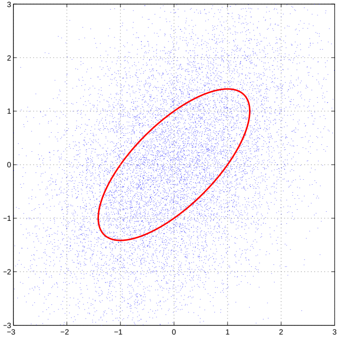

- Signal varies more in $\mathrm{v}_0 = [1, 1]^T$ direction than in $\mathrm{v}_1 = [1, -1]^T$
  $\Rightarrow$ Study one dimensional signal $\tilde{\mathrm{x}} = y(0)\mathrm{v}_0$
  $\Rightarrow$ instead of the original two dimensional signal x

Expected reconstruction error

- PCA dimensionality reduction minimizes the expected error energy
- To see that this is true, define the error signal as ⇒ $\mathbf{e} := \mathbf{x} - \tilde{\mathbf{x}}$
- The energy of the error signal is ⇒ $\|\mathbf{e}\|^2 = \|\mathbf{x} - \tilde{\mathbf{x}}\|^2$
- The expected value of the energy of the error signal is

$$\mathbb{E}\left[\|\mathbf{e}\|^2\right] = \mathbb{E}\left[\|\mathbf{x} - \tilde{\mathbf{x}}\|^2\right]$$

- Keeping the first K PCA coefficients minimizes $\mathbb{E}\left[\|\mathbf{e}\|^2\right]$

⇒ Among all reconstructions that use, at most, K coefficients

Dimensionality reduction expected error

Theorem

The expectation of the reconstruction error is the sum of the eigenvalues corresponding to the eigenvectors of the coefficients that are discarded

$$\mathbb{E}\left[\|\mathsf{e}\|^2\right] = \sum_{k=K}^{N-1} \lambda_k$$

- It follows that keeping the first K PCA coefficients is optimal

  $\Rightarrow$ In the sense that it minimizes the Expected error energy
- Good on average. Across realizations of the stochastic signal X
- Need not be good for given realization (but we expect it to be good)

Proof of expected error expression

Proof.

- Error signal signal is $\mathrm{e} := \mathrm{x} - \tilde{\mathrm{x}}$. Define error PCA transform as $\mathrm{f} = \mathrm{T}^H \mathrm{x}$
- Using Parseval's (energy conservation) we can write the energy of e as

$$\|\mathrm{e}\|^2 = \|\mathrm{f}\|^2 = \sum_{k=K}^{N-1} y^2(k)$$

- In the last equality we used that $\mathrm{f} = \mathrm{y} - \tilde{\mathrm{y}} = [0, \ldots, 0, y(K), \ldots, y(N-1)]$
- Here, we are interested in the expected value of the error's energy
- Take expectation on both sides of equality $\Rightarrow \mathbb{E}\left[\|\mathrm{e}\|^2\right] = \sum_{k=K}^{N-1} \mathbb{E}\left[y^2(k)\right]$
- Used the fact that expectations are linear operators

Proof of expected error expression

Proof.

- Compute expected value $\mathbb{E}\left[y^2(k)\right]$ of the squared PCA coefficient $y(k)$
- As per PCA transform definition $y(k) = \mathrm{v}^H\mathrm{x}$, which implies

$$\mathbb{E}\left[y^2(k)\right] = \mathbb{E}\left[(\mathrm{v}_k^H\mathrm{x})^2\right] = \mathbb{E}\left[\mathrm{v}_k^H\mathrm{x}\mathrm{x}^T\mathrm{v}_k\right] = \mathrm{v}_k^H\mathbb{E}\left[\mathrm{x}\mathrm{x}^T\right]\mathrm{v}_k$$

- Covariance matrix: $\Sigma := \mathbb{E}\left[\mathrm{x}\mathrm{x}^T\right]$. Eigenvector definition $\Sigma\mathrm{v}_k = \lambda_k$. Thus

$$\mathbb{E}\left[y^2(k)\right] = \mathrm{v}_k^H\Sigma\mathrm{v}_k = \mathrm{v}_k^H\lambda_k\mathrm{v}_k = \lambda_k\|\mathrm{v}_k\|^2$$

- Substitute into expression for $\mathbb{E}\left[\|\mathrm{e}\|^2\right]$ to write $\Rightarrow \mathbb{E}\left[\|\mathrm{e}\|^2\right] = \sum_{k=K}^{N-1}\lambda_k$

□

Principal eigenvalues for face dataset

- Covariance matrix eigenvalues for faces dataset.
- Expected approximation error $\Rightarrow$ Tail sum of eigenvalue distribution
  $\Rightarrow$ Average across all realizations. Not the same as actual error

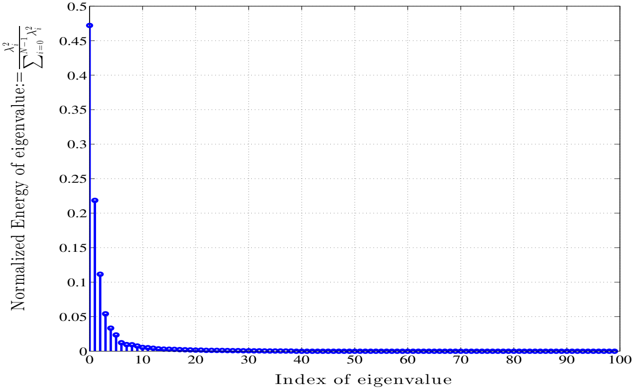

- First 10 coefficients have 98% of energy.
- Eigenvectors with index $k > 50$ have 10^{-3}% of energy on average

Reconstructed face images

- Increasing number of coefficients reduces reconstruction error
- Average and actual reconstruction not the same (although “close”)
- Keep 1 coefficient ⇒ Reconstruction error ⇒ 0.06
  ⇒ Sum of removed eigenvalues ⇒ 0.52

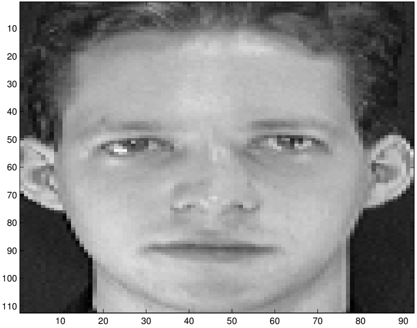

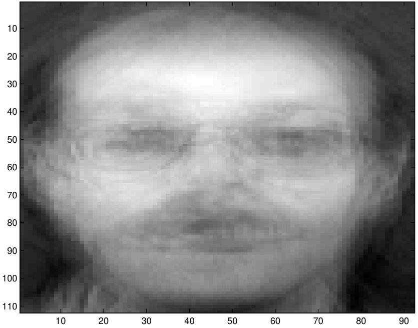

Reconstructed face images

- Increasing number of coefficients reduces reconstruction error
- Average and actual reconstruction not the same (although "close")
- Keep 5 coefficients $\Rightarrow$ Reconstruction error $\Rightarrow$ 0.03
  $\Rightarrow$ Sum of removed eigenvalues $\Rightarrow$ 0.11

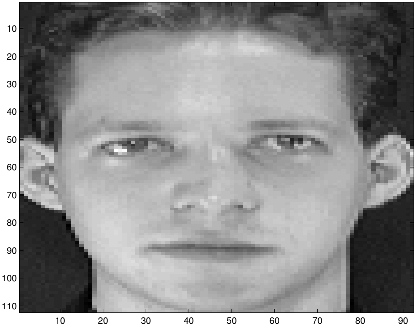

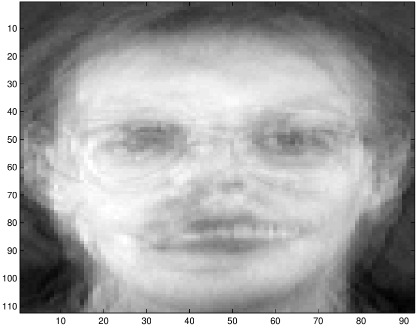

Reconstructed face images

- Increasing number of coefficients reduces reconstruction error
- Average and actual reconstruction not the same (although “close”)
- Keep 10 coefficients $\Rightarrow$ Reconstruction error $\Rightarrow$ 0.02
  $\Rightarrow$ Sum of removed eigenvalues $\Rightarrow$ 0.04

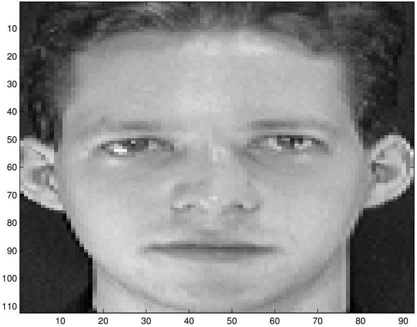

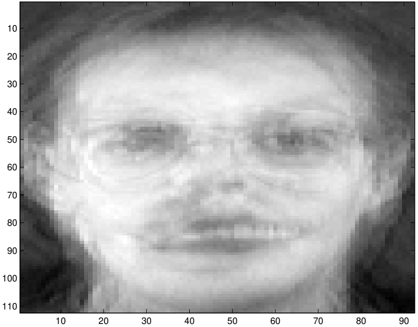

Reconstructed face images

- Increasing number of coefficients reduces reconstruction error
- Average and actual reconstruction not the same (although “close”)
- Keep 20 coefficients $\Rightarrow$ Reconstruction error $\Rightarrow$ 0.01
  $\Rightarrow$ Sum of removed eigenvalues $\Rightarrow$ 0.01

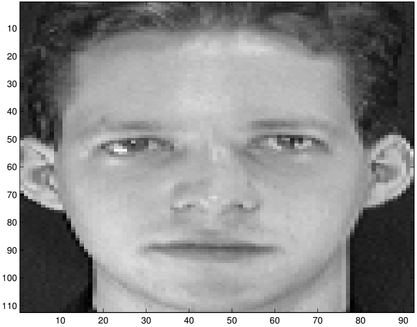

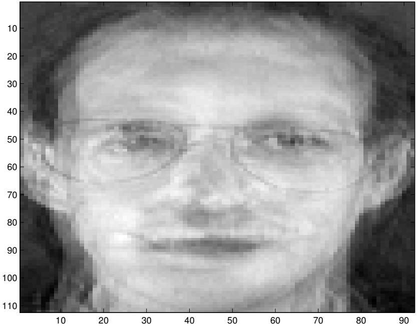

Reconstructed face images

- Increasing number of coefficients reduces reconstruction error
- Average and actual reconstruction not the same (although “close”)
- Keep 30 coefficients $\Rightarrow$ Reconstruction error $\Rightarrow$ 0.006
  $\Rightarrow$ Sum of removed eigenvalues $\Rightarrow$ 0.003

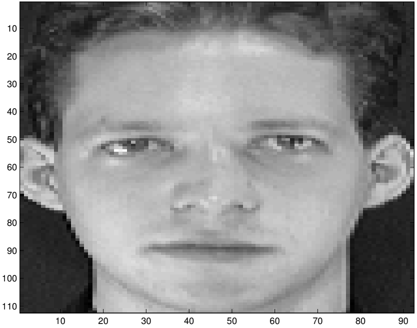

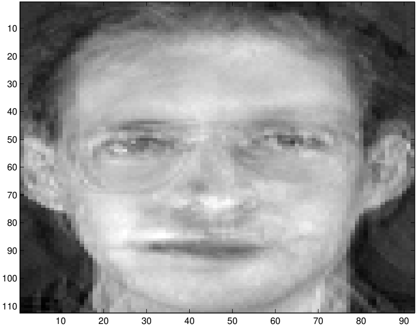

Reconstructed face images

- Increasing number of coefficients reduces reconstruction error
- Average and actual reconstruction not the same (although “close”)
- Keep 40 coefficients ⇒ Reconstruction error ⇒ 0
  ⇒ Sum of removed eigenvalues ⇒ 0

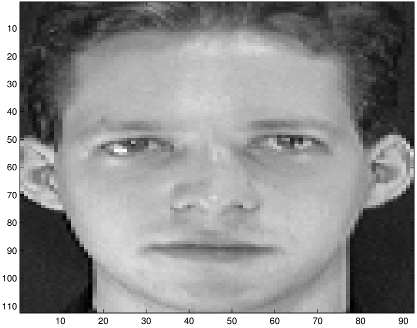

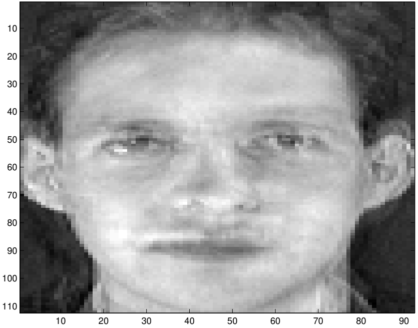

Reconstructed face images

- Increasing number of coefficients reduces reconstruction error
- Average and actual reconstruction not the same (although “close”)
- Keep 50 coefficients $\Rightarrow$ Reconstruction error $\Rightarrow$ 0
  $\Rightarrow$ Sum of removed eigenvalues $\Rightarrow$ 0

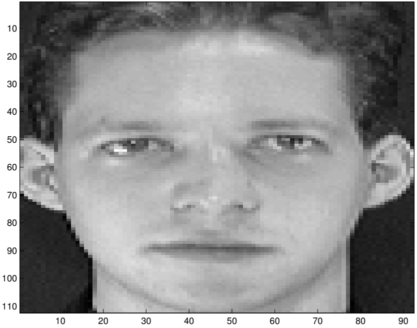

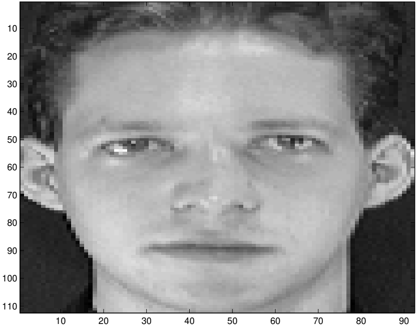

Evolution of reconstruction error

- Error for reconstruction process
- one realization (red), energy of removed eigenvalues (blue)

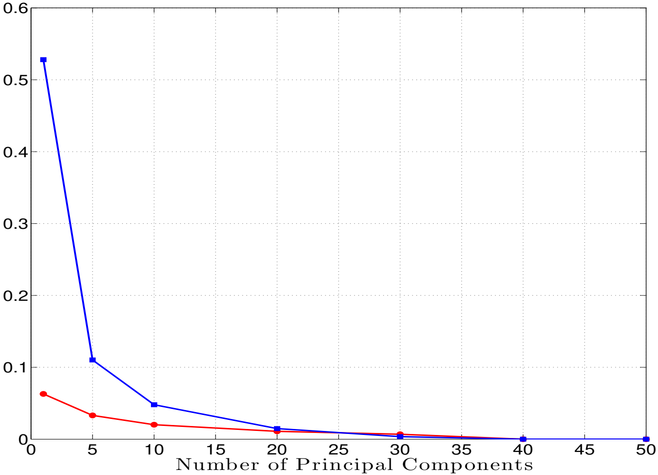

Principal Components

The Discrete Fourier Transform with Unitary Matrices

Stochastic signals

Principal Component Analysis (PCA) transform

Dimensionality reduction

Principal Components

Face recognition

Signals with uncorrelated components

- A random signal X with uncorrelated components is one with

$$\Sigma_{nm} = \mathbb{E}\left[\left(X(n) - \mathbb{E}\left[X(n)\right]\right)\left(X(m) - \mathbb{E}\left[X(m)\right]\right)\right] = 0$$

- Different components are unrelated to each other.
- They represent different (orthogonal) aspects of signal

- Components uncorrelated $\Rightarrow$ The covariance matrix is diagonal

$$\Sigma = \mathbb{E}\left[\left(\mathrm{x} - \mathbb{E}\left[\mathrm{x}\right]\right)\left(\mathrm{x} - \mathbb{E}\left[\mathrm{x}\right]\right)^T\right] = \begin{bmatrix} \Sigma_{00} & \cdots & \Sigma_{0n} & \cdots & \Sigma_{0(N-1)} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \Sigma_{n0} & \cdots & \Sigma_{nn} & \cdots & \Sigma_{n(N-1)} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \Sigma_{(N-1)0} & \cdots & \Sigma_{(N-1)n} & \cdots & \Sigma_{(N-1)(N-1)} \end{bmatrix}$$

- How do eigenvectors (principal components) of uncorrelated signals look?

Uncorrelated signal with 2 components

- Signal $\mathsf{X} = [X(0), X(1)]^T$ with 2 components and diagonal covariance

$$\Sigma = \begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix}$$

- Covariance eigenvectors are

$$\mathsf{v}_0 = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \quad \mathsf{v}_1 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

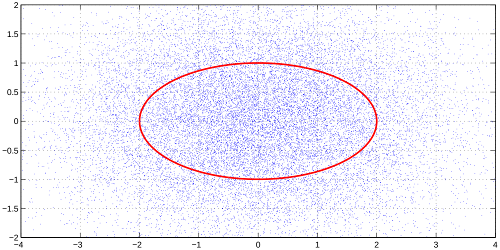

- The respective associated eigenvalues are $\lambda_0 = 2$ and $\lambda_1 = 1$

- Eigenvectors are orthogonal, as they should.
  $\Rightarrow$ Represent directions of separate signal variability
  $\Rightarrow$ Rate of variability given by associated eigenvalue

Another uncorrelated signal with 2 components

- Signal $\mathrm{X} = [X(0), X(1)]^T$ with 2 components and diagonal covariance

$$\Sigma = \begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix}$$

- Covariance eigenvectors reverse order

$$\mathrm{v}_0 = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \quad \mathrm{v}_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$$

- Associated eigenvalues are $\lambda_0 = 2$ and $\lambda_1 = 1$
- Eigenvectors still orthogonal, as they should.
  - ⇒ Directions of separate signal variability
  - ⇒ Rate given by associated eigenvalue

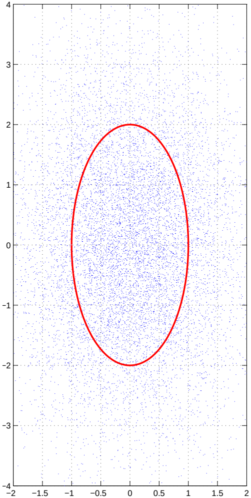

Signal with correlated components

- Signal $\mathrm{X} = [X(0), X(1)]^T$ with 2 components and diagonal covariance

$$\Sigma = \left[\begin{array}{cc} 3/2 & 1/2 \\ 1/2 & 3/2 \end{array}\right]$$

- Covariance eigenvectors mix coordinates

$$\mathrm{v}_0 = \left[\begin{array}{c} 1 \\ 1 \end{array}\right] \quad \mathrm{v}_1 = \left[\begin{array}{r} 1 \\ -1 \end{array}\right]$$

- Eigenvalues are $\lambda_0 = 2$ and $\lambda_1 = 1$

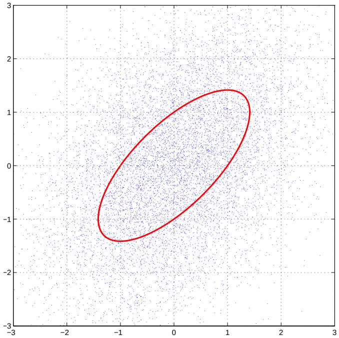

- The eigenvalues are orthogonal. This is true for any covariance matrix
  $\Rightarrow$ Mix coordinates but still represent directions of separate variability
  $\Rightarrow$ Rate of change also given by associated eigenvalue

Eigenvectors in uncorrelated signals

- Uncorrelated components means diagonal covariance matrix

$$\Sigma = \begin{bmatrix} \Sigma_{00} & \cdots & \Sigma_{0n} & \cdots & \Sigma_{0(N-1)} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \Sigma_{n0} & \cdots & \Sigma_{nn} & \cdots & \Sigma_{n(N-1)} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \Sigma_{(N-1)0} & \cdots & \Sigma_{(N-1)n} & \cdots & \Sigma_{(N-1)(N-1)} \end{bmatrix}$$

- If variances are ordered, kth eigenvector is k-shifted delta $\delta(n-k)$

- The corresponding variance Σ_{kk} is the associated eigenvalue
- Eigenvectors represent directions of orthogonal variability

- Rate of variability given by associated eigenvalue

Eigenvectors in correlated signals

- Correlated components means a full covariance matrix

$$\Sigma = \begin{bmatrix} \Sigma_{00} & \cdots & \Sigma_{0n} & \cdots & \Sigma_{0(N-1)} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \Sigma_{n0} & \cdots & \Sigma_{nn} & \cdots & \Sigma_{n(N-1)} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \Sigma_{(N-1)0} & \cdots & \Sigma_{(N-1)n} & \cdots & \Sigma_{(N-1)(N-1)} \end{bmatrix}$$

- The eigenvectors v_k now mix different components

  ⇒ But they still represent directions of orthogonal variability

  ⇒ With the rate of variability given by associated eigenvalue

- PCA transform represents a signal as a sum of orthonormal vectors

  ⇒ Each of which represents independent variability

- Principal components (eigenvectors) with larger eigenvalues represent directions in which the signal has more variability

Face recognization

The Discrete Fourier Transform with Unitary Matrices

Stochastic signals

Principal Component Analysis (PCA) transform

Dimensionality reduction

Principal Components

Face recognition

Face Recognition

- Observe faces of known people $\Rightarrow$ Use them to train classifier
- Observe a face of unknown character $\Rightarrow$ Compare and classify
- The dataset we've used contains 10 different images of 40 people

Training set

- Separate the first 9 of each person to construct training set

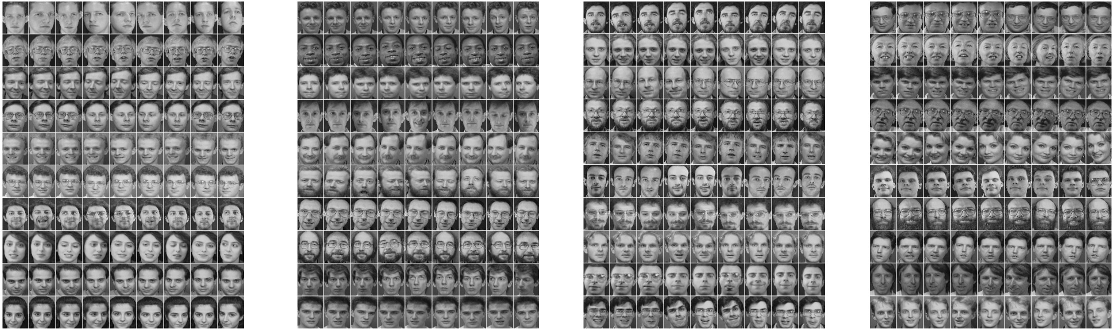

- Interpret these images as know, and use them to train classifier

Test set

- Utilize the last image of each person to construct a test set

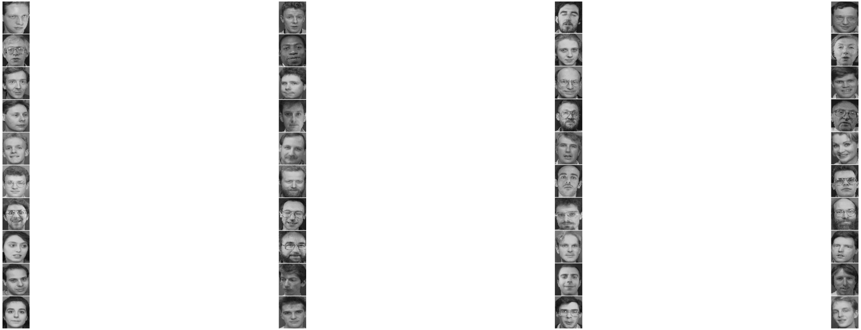

- Interpret these images as unknown, and use them to test classifier

Nearest neighbor classification

- Training set contains (signal, label) pairs $\Rightarrow \mathcal{T} = \{(\mathrm{x}_i, z_i)\}_{i=1}^N$

- Signal x is the face image. Label z is the person's "name"
- Given (unknown) signals x, we want to assign a label
- Nearest neighbor classification rule

$\Rightarrow$ Find nearest neighbor signal in the training set

$$\mathrm{x}_{\mathrm{NN}} := \operatorname*{argmin}_{\mathrm{x}_i \in \mathcal{T}} \|\mathrm{x}_i - \mathrm{x}\|^2$$

$\Rightarrow$ Assign the label associated with the nearest neighbor

$$\mathrm{x}_{\mathrm{NN}} \quad \Rightarrow \quad (\mathrm{x}_i, z_i) \quad \Rightarrow \quad z = z_i$$

- Reasonable enough. It should work. But it doesn't

The signal and the noise

- Image has a part that is inherent to the person ⇒ The actual signal
- But it also contains variability ⇒ Which we model as noise

$$x_i = \tilde{x}_i + w$$

- Problem is, there is more variability (noise) than signal

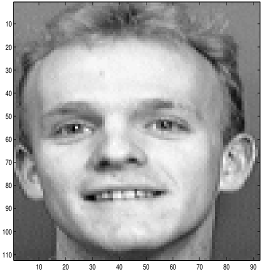

Figure: Test image

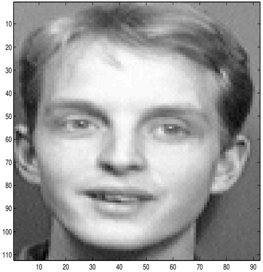

Figure: Nearest neighbor

PCA nearest neighbor classification

- Compute PCA for all elements of training set $\Rightarrow \mathrm{y}_i = \mathrm{T}^H \mathrm{x}_i$
- Redefine training set as one with PCA transforms $\Rightarrow \mathcal{T} = \{(\mathrm{y}_i, z_i)\}_{i=1}^N$
- Compute PCA transform of (unknown) signal x $\Rightarrow \mathrm{y} = \mathrm{T}^H \mathrm{x}$
- PCA nearest neighbor classification rule

  $\Rightarrow$ Find nearest neighbor signal in training set with PCA transforms

$$\mathrm{y}_{\mathrm{NN}} := \underset{\mathrm{y}_i \in \mathcal{T}}{\operatorname{argmin}} \left\| \mathrm{y}_i - \mathrm{y} \right\|^2$$

  $\Rightarrow$ Assign the label associated with the nearest neighbor

$$\mathrm{y}_{\mathrm{NN}} \quad \Rightarrow \quad (\mathrm{y}_i, z_i) \quad \Rightarrow \quad z = z_i$$

- Reasonable enough. It should work. And it does

Why does PCA work for face recognition?

- Recall: image = a part that belongs to the person + noise

$$\mathrm{x}_i = \tilde{\mathrm{x}}_i + \mathrm{w}$$

- PCA transformation $\mathrm{T} = [\mathrm{v}_0^T; \ldots; \mathrm{v}_{N-1}^T]$ leads to

$$\mathrm{y}_i = \mathrm{T}\mathrm{x}_i = \mathrm{T}\tilde{\mathrm{x}}_i + \mathrm{T}\mathrm{w}$$

- PCA concentrates energy of $\tilde{\mathrm{x}}_i$ on a few components

- But it keeps the energy of the noise on all components
- Keeping principal components improves the accuracy of classification

$\Rightarrow$ Because it increases the signal to noise ratio

PCA on the training set

- The training set $D = \{x_1, \ldots, x_{360}\}$ where $x_i \in \mathbb{R}^{10304}$ is given
- Compute the mean vector and the covariance matrix as

$$\bar{x} = \frac{1}{n}\sum_{i=1}^{n} x_i \quad \text{and} \quad \Sigma := \frac{1}{n}\sum_{i=1}^{n}(x_i - \bar{x}_i)(x_i - \bar{x}_i)^T.$$

- Find the k largest eigenvalues of Σ
- Store their corresponding eigenvalues $v_0, \ldots, v_{k-1} \in \mathbb{R}^{10304}$ as P.C.

  $\Rightarrow$ The Principal Components $v_0, \ldots, v_{k-1}$ are called eigenfaces

- Create the PCA transform matrix as $T = [v_0^T; \ldots; v_{k-1}^T]$
- Project the training set into the space of P.C.s $y_i = Tx_i$
- Σ depends training set, but is also a good description of the test set

Average face of the training set

- The average face of the training set

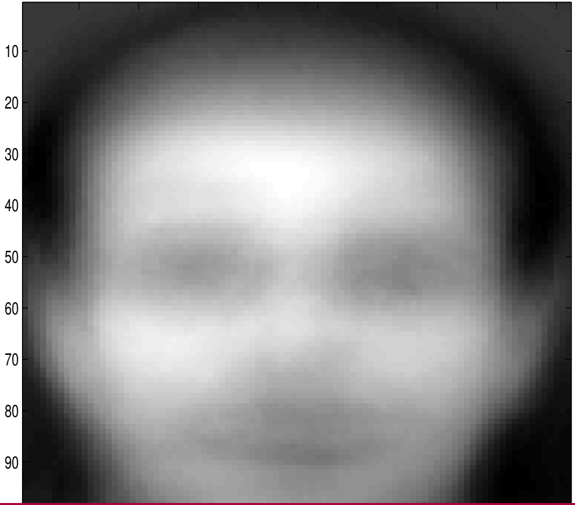

PCA on the training set

- The top 6 eigenfaces of the training set

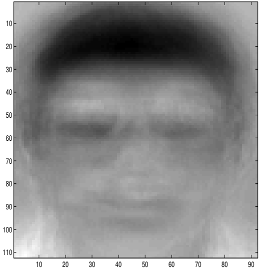

(1)

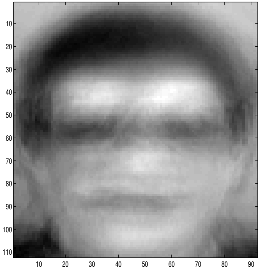

(2)

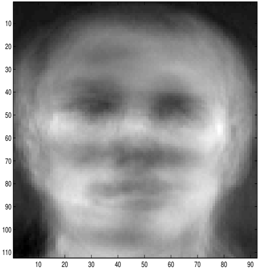

(3)

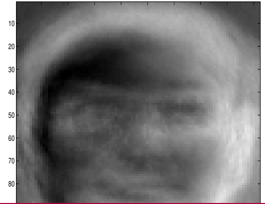

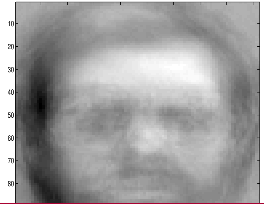

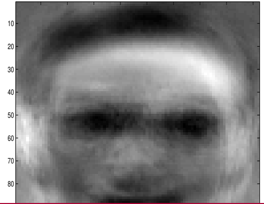

Finding the nearest neighbor

Num. of P.C. test point N.N. in the training set

$k = 1$

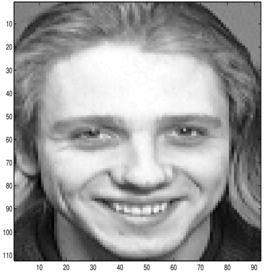

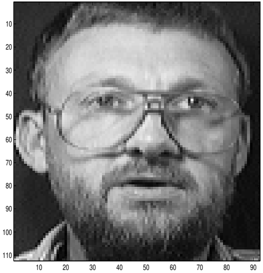

$k = 5$

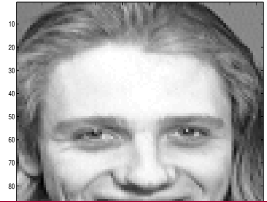

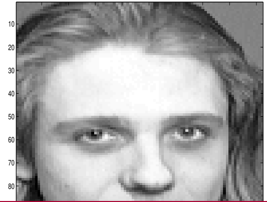

PCA improves classification accuracy

Classification method	test point	result of classification
Naive N.N.	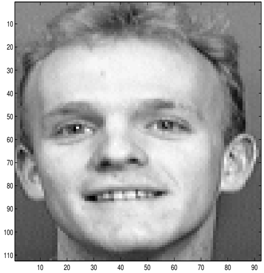	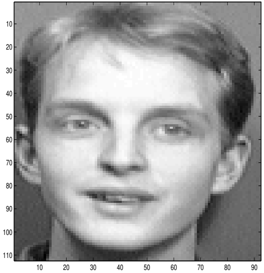
PCA-ed($k = 5$) N.N.