Introduction to Neural Computation

Prof. Michale Fee
MIT BCS 9.40 — 2017

Lecture 17
Principal Components Analysis

Learning Objectives for Lecture 17

- Eigenvectors and eigenvalues
- Variance and multivariate Gaussian distributions
- Computing a covariance matrix from data
- Principal Components Analysis (PCA)

Learning Objectives for Lecture 17

- Eigenvectors and eigenvalues
- Variance and multivariate Gaussian distributions
- Computing a covariance matrix from data
- Principal Components Analysis (PCA)

Matrix transformations

$$\vec{y} = A\vec{x}$$

- In general A maps the set of vectors in $\mathbb{R}^2$ onto another set of vectors in $\mathbb{R}^2$.

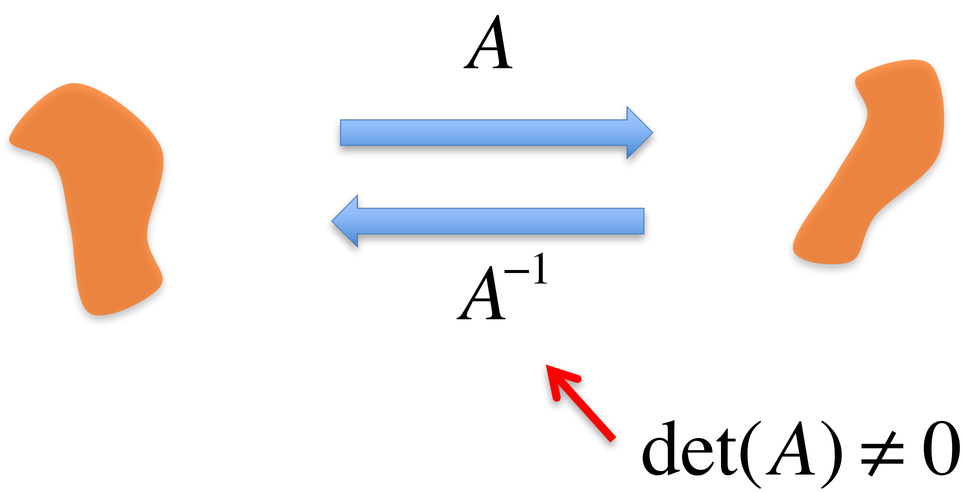

Eigenvectors and eigenvalues

- Matrix transformations have special directions

$$A = \begin{pmatrix} 1+\delta & 0 \\ 0 & 1 \end{pmatrix}$$

$$A = \begin{pmatrix} 1-\delta & 0 \\ 0 & 1+\delta \end{pmatrix}$$

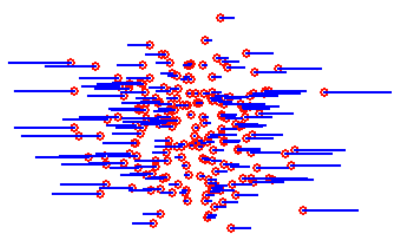

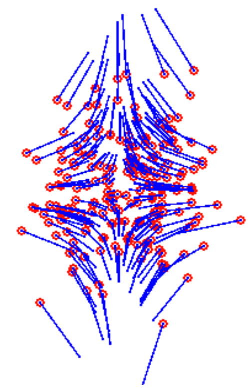

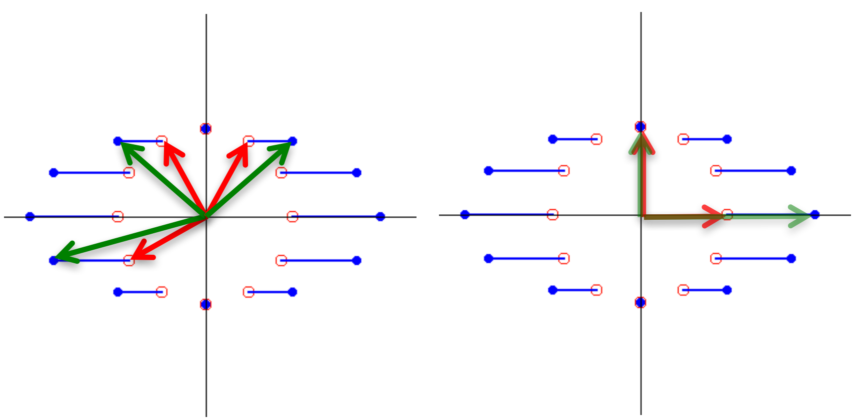

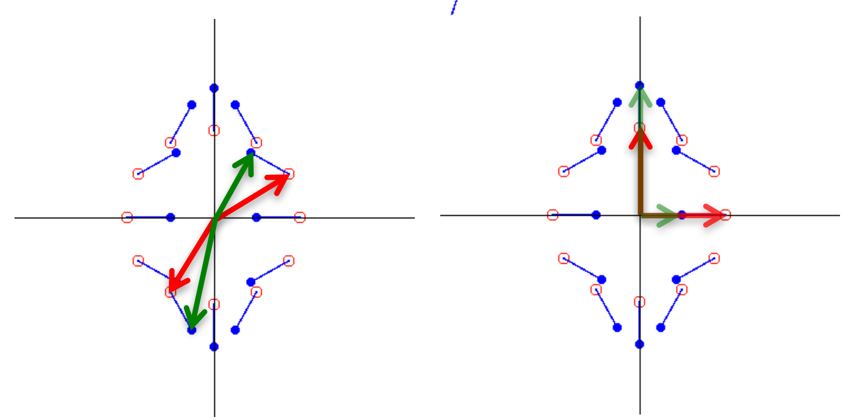

These are all diagonal matrices

$$\Lambda = \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix}$$

Eigenvectors and eigenvalues

- Some vectors are rotated, some are not.

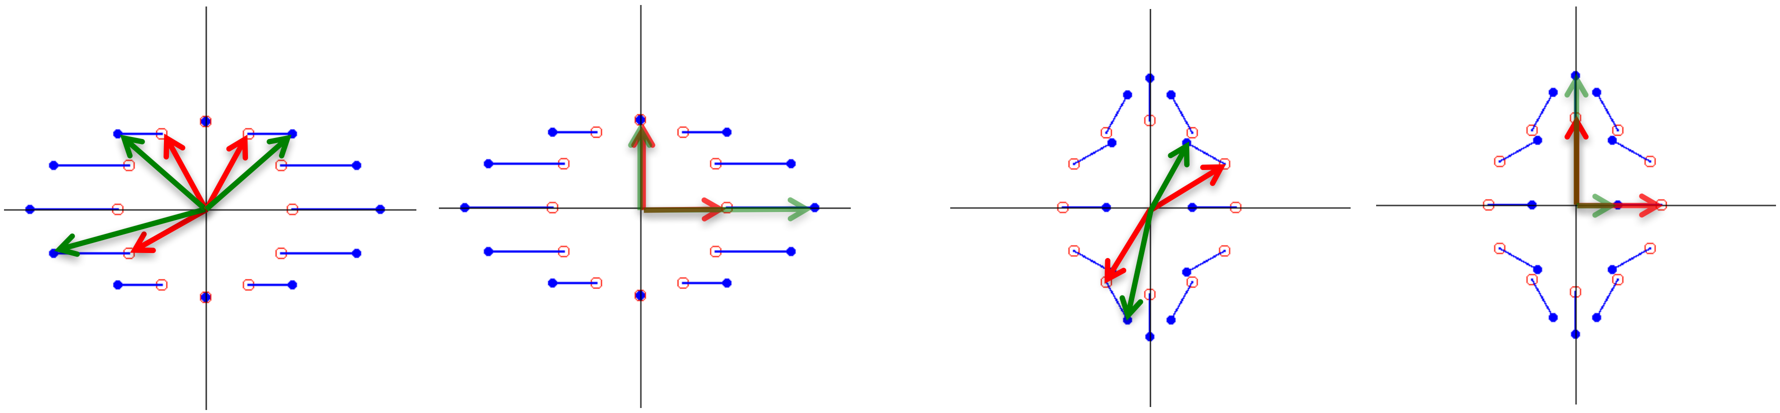

- For a diagonal matrix, vectors along the axes are scaled, but not rotated.

$$\begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix}\begin{pmatrix} 1 \\ 0 \end{pmatrix} = \lambda_1 \begin{pmatrix} 1 \\ 0 \end{pmatrix} \qquad \Lambda \hat{e}_1 = \lambda_1 \hat{e}_1$$

$$\begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix}\begin{pmatrix} 0 \\ 1 \end{pmatrix} = \lambda_2 \begin{pmatrix} 0 \\ 1 \end{pmatrix} \qquad \Lambda \hat{e}_2 = \lambda_2 \hat{e}_2$$

Eigenvectors and eigenvalues

- Diagonal matrices have the property that they map any vector parallel to a standard basis vector into another vector along that standard basis vector.

eigenvalue equation

$$\Lambda = \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix} \qquad \Lambda \hat{e}_i = \lambda_i \hat{e}_i \quad , \quad i = 1,2\ldots.n$$

- Any vector $\vec{v}$ that is mapped by matrix A onto a parallel vector $\lambda\vec{v}$ is called an eigenvector of A. The scale factor λ is called the eigenvalue of vector $\vec{v}$.

- A matrix in $\mathbb{R}^n$ has n eigenvectors and n eigenvalues.

Eigenvectors and eigenvalues

- What are the special directions of our rotated transformations?

$$A = \Phi\Lambda\Phi^T = \begin{pmatrix} 1.5 & 0.5 \\ 0.5 & 1.5 \end{pmatrix} \qquad \Lambda = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix} \qquad \Phi(45^0) = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}$$

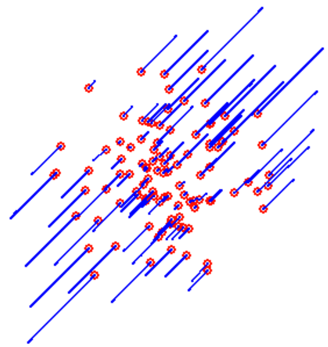

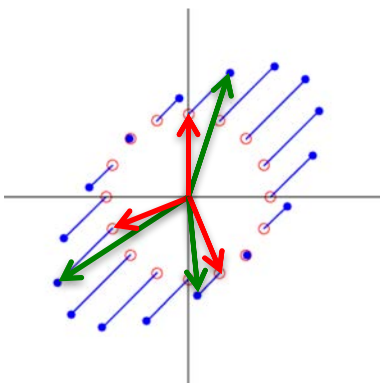

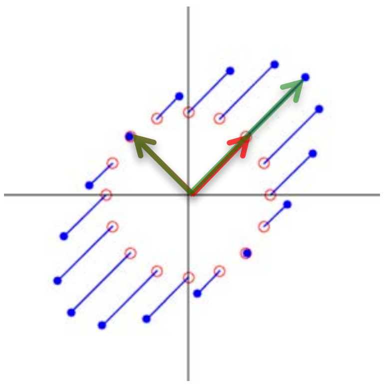

Eigenvectors and eigenvalues

- What are the eigenvectors and eigenvalues of our rotated transformation matrix $A = \Phi\Lambda\Phi^T$?

$$A\vec{x}_i = a_i\,\vec{x}_i$$

$$\Phi\Lambda\Phi^T\ \vec{x}_i = a_i\,\vec{x}_i$$

$$\Phi^T\Phi\Lambda\Phi^T\ \vec{x}_i = \Phi^T(a_i\,\vec{x}_i)$$

$$\Lambda\underbrace{\Phi^T\vec{x}_i}_{\hat{e}_i} = a_i\ \underbrace{\Phi^T\vec{x}_i}_{\hat{e}_i}$$

Remember…

$$\Lambda\hat{e}_i = \lambda_i\,\hat{e}_i \qquad \Lambda = \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix}$$

So we know the solution if $\Phi^T\vec{x}_i = \hat{e}_i$

$$\Lambda\hat{e}_i = a_i\,\hat{e}_i$$

$$\Rightarrow a_i = \lambda_i$$

So the solution to the eigenvalue equation $\Phi\Lambda\Phi^T\ \vec{x}_i = a_i\,\vec{x}_i$ is:

eigenvalues	eigenvectors
$a_i = \lambda_i$	$\vec{x}_i = \Phi\hat{e}_i$

what is this?

Eigenvectors and eigenvalues

- The eigenvectors are just the standard basis vectors rotated by the matrix Φ !

$$\vec{x}_i = \Phi \hat{e}_i$$

$$\Phi(45^0) = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}$$

$$\hat{f}_1 = \frac{1}{\sqrt{2}}\begin{pmatrix} -1 \\ 1 \end{pmatrix}$$

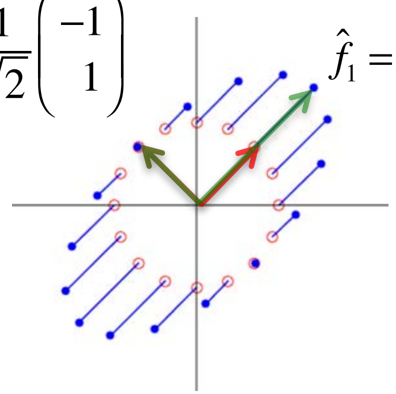

$$\hat{f}_1 = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 \\ 1 \end{pmatrix}$$

$$\vec{x}_1 = \Phi \hat{e}_1 = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}\begin{pmatrix} 1 \\ 0 \end{pmatrix} = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 \\ 1 \end{pmatrix}$$

$$\vec{x}_2 = \Phi \hat{e}_2 = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}\begin{pmatrix} 0 \\ 1 \end{pmatrix} = \frac{1}{\sqrt{2}}\begin{pmatrix} -1 \\ 1 \end{pmatrix}$$

The eigenvectors are just the columns of Φ!

Eigenvectors and eigenvalues

- In summary, a symmetric matrix A can always be written as follows:

$$A = \Phi\Lambda\Phi^T$$

where Φ is a rotation matrix and Λ is a diagonal matrix

- The eigenvectors of A are the columns of Φ (the basis vectors, $\hat{f}_i$)

$$\Phi = \left(\hat{f}_1 \;\middle|\; \hat{f}_2 \right)$$

- The eigenvalue associated with each eigenvector $\hat{f}_i$ is the diagonal element λ_i of Λ .

$$\Lambda = \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix}$$

Eigenvectors and eigenvalues

- Note that eigenvectors are not unique…

If $\vec{x}_i$ is an eigenvector of A then so is $a\vec{x}_i$

$$A\vec{x}_i = \lambda_i \vec{x}_i \qquad\qquad A\left(a\vec{x}_i\right) = \lambda_i \left(a\vec{x}_i\right)$$

… so we usually write eigenvectors as unit vectors

- For a matrix in n-dimensions… there are n different unit eigenvectors

- For a symmetric matrix, the eigenvectors of A are orthogonal (and we write them as unit vectors)…

  … the eigenvectors of A form a complete orthonormal basis set!

Eigenvectors and eigenvalues

- What are the eigenvalues of a general 2-dim matrix A ?

$$A\vec{x} = \lambda\vec{x}$$

we only want solutions where

$$\vec{x} \neq 0$$

$$A\vec{x} = \lambda I\vec{x}$$

$$\left(A - \lambda I\right)\vec{x} = 0$$

$$\det\left(A - \lambda I\right) = 0$$

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \qquad A - \lambda I = \begin{pmatrix} a-\lambda & b \\ c & d-\lambda \end{pmatrix} \qquad \det\left(A - \lambda I\right) = \left(a-\lambda\right)\left(d-\lambda\right) - bc$$

$$\det\left(A - \lambda I\right) = \left(a-\lambda\right)\left(d-\lambda\right) - bc = 0$$

$$ad - \lambda\left(a+d\right) + \lambda^2 - bc = 0$$

Characteristic equation of matrix A

$$\lambda^2 - \left(a+d\right)\lambda + \left(ad - bc\right) = 0$$

Eigenvectors and eigenvalues

- What are the eigenvalues of a general 2-dim matrix?

Characteristic equation of matrix A

$$\lambda^2 - (a+d)\lambda + (ad - bc) = 0$$

- Solutions are given by the quadratic formula

$$\lambda_{\pm} = \frac{1}{2}(a+d) \pm \frac{1}{2}\sqrt{\underbrace{(a-d)^2 + 4bc}}$$

eigenvalues can be real, complex, imaginary

Eigenvectors and eigenvalues

$$\lambda_{\pm} = \frac{1}{2}(a+d) \pm \frac{1}{2}\sqrt{(a-d)^2 + 4bc}$$

- For a symmetric matrix

$$A = \begin{pmatrix} a & b \\ b & d \end{pmatrix}$$

$$\lambda_{\pm} = \frac{1}{2}(a+d) \pm \frac{1}{2}\sqrt{\underbrace{(a-d)^2 + 4b^2}_{\geq 0}}$$

The eigenvalues of a symmetric matrix are always real.

$$A = \begin{pmatrix} 3/2 & 1/2 \\ 1/2 & 1/2 \end{pmatrix}$$

$$\lambda_{\pm} = \frac{1}{2}\left(\frac{3}{2} + \frac{1}{2}\right) \pm \frac{1}{2}\sqrt{\left(\frac{3}{2} - \frac{1}{2}\right)^2 + 4\left(\frac{1}{2}\right)^2}$$

$$\lambda_{\pm} = 1 \pm \frac{\sqrt{2}}{2}$$

Eigenvectors and eigenvalues

- Let's consider a special case of a symmetric matrix

$$A = \begin{pmatrix} a & b \\ b & a \end{pmatrix}$$

$$\lambda_{\pm} = \frac{1}{2}(a+d) \pm \frac{1}{2}\sqrt{(a-d)^2 + 4b^2}$$

$$\lambda_{+} = a+b \qquad \lambda_{-} = a-b$$

$$A\vec{x}_{+} = \lambda_{+}\vec{x}_{+}$$

$$\begin{pmatrix} a & b \\ b & a \end{pmatrix}\vec{x}_{+} = (a+b)\vec{x}_{+}$$

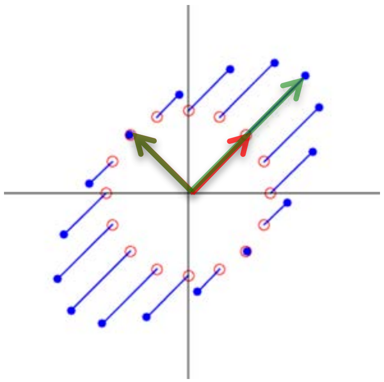

$$\begin{pmatrix} a & b \\ b & a \end{pmatrix}\vec{x}_{+} = \begin{pmatrix} a+b & 0 \\ 0 & a+b \end{pmatrix}\vec{x}_{+}$$

$$\begin{pmatrix} -b & b \\ b & -b \end{pmatrix}\vec{x}_{+} = 0$$

$$\begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix}\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = 0 \qquad \begin{aligned} x_1 &= s \\ x_2 &= s \end{aligned} \qquad \vec{x}_{+} = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 \\ 1 \end{pmatrix} \qquad \vec{x}_{-} = \frac{1}{\sqrt{2}}\begin{pmatrix} -1 \\ 1 \end{pmatrix}$$

Eigen-decomposition

- The process of writing a matrix A as $A = \Phi\Lambda\Phi^T$ is called eigen-decomposition. It works for any symmetric matrix.
    - the eigenvalues λ_i are real numbers
    - the eigenvectors $\hat{f}_i$ form an orthogonal basis set

- Let's rewrite…

$$A = \Phi\Lambda\Phi^T \qquad \Lambda = \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix}$$

$$A\Phi = \Phi\Lambda\Phi^T\Phi$$

- Eigenvalue equation

$$A\Phi = \Phi\Lambda$$

is equivalent to the set of equations $A\,\hat{f}_i = \lambda_i\,\hat{f}_i$

Eigenvectors and eigenvalues

- MATLAB® has a function 'eig' to calculate eigenvectors and eigenvalues

$$A = F\,V\,F^T$$

```
>> A=[1.5 0.5;0.5 1.5]

A =

    1.5000    0.5000
    0.5000    1.5000
```

```
>> [F,V]=eig(A)

F =

   -0.7071    0.7071
    0.7071    0.7071


V =

     1     0
     0     2
```

```
>> F*V*F'

ans =

    1.5000    0.5000
    0.5000    1.5000
```

Learning Objectives for Lecture 17

- Eigenvectors and eigenvalues
- Variance and multivariate Gaussian distributions
- Computing a covariance matrix from data
- Principal Components Analysis (PCA)

Variance

- Let's say we have m observations of a variable x.

$$x^{(j)}\ ,\ j = 1,2,3....,m$$

j[th] observation of x

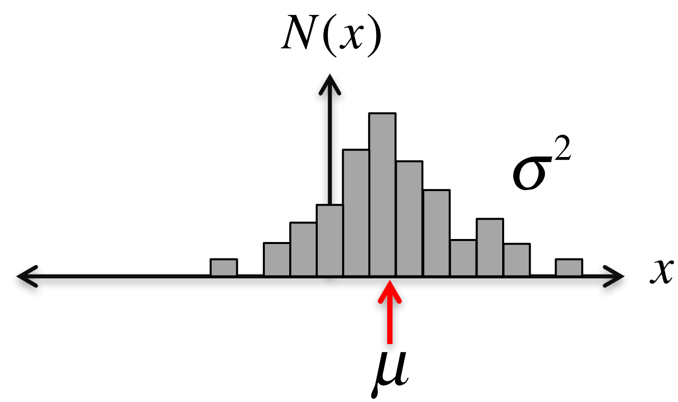

- The mean of these observations is $\mu = \langle x \rangle \ = \frac{1}{m}\sum_{i=1}^{m} x^{(i)}$

- The variance is $\sigma^2 = \frac{1}{m}\sum_{i=1}^{m}\left(x^{(i)} - \mu\right)^2$

Variance

- Now let's say we have m simultaneous observations of variables and . x_1 x_2

$$\begin{pmatrix} x_1^{(j)} \\ x_2^{(j)} \end{pmatrix}, \quad j = 1,2,3\dots,m$$

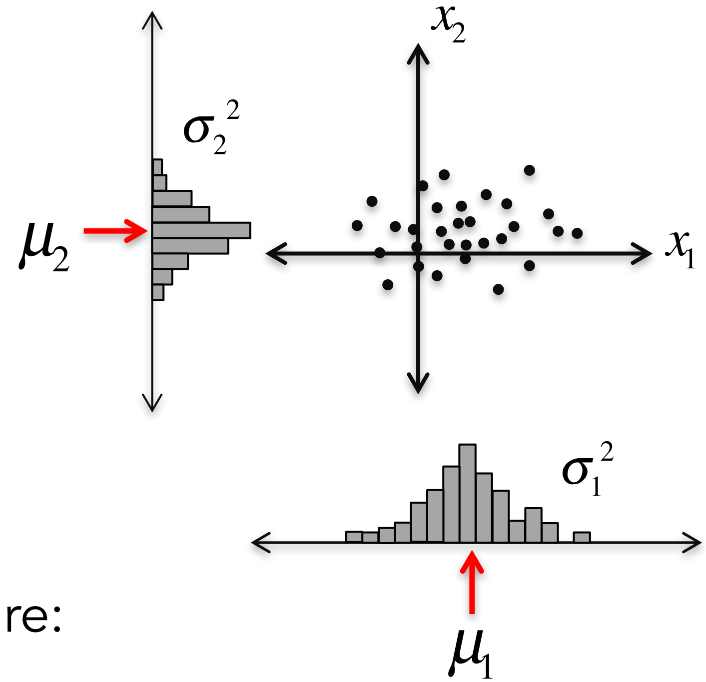

- The mean and variance of x_1 and x_2 are:

$$\mu_1 = \frac{1}{m}\sum_{j=1}^{m} x_1^{(j)} \qquad \mu_2 = \frac{1}{m}\sum_{j=1}^{m} x_2^{(j)}$$

$$\sigma_1^{\ 2} = \frac{1}{m}\sum_{j=1}^{m}\left(x_1^{(j)} - \mu_1\right)^2 \qquad \sigma_2^{\ 2} = \frac{1}{m}\sum_{j=1}^{m}\left(x_2^{(j)} - \mu_2\right)^2$$

Covariance

uncorrelated

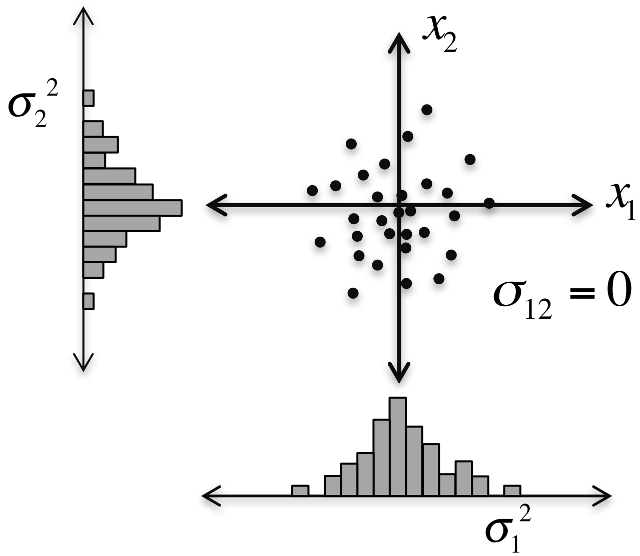

correlated

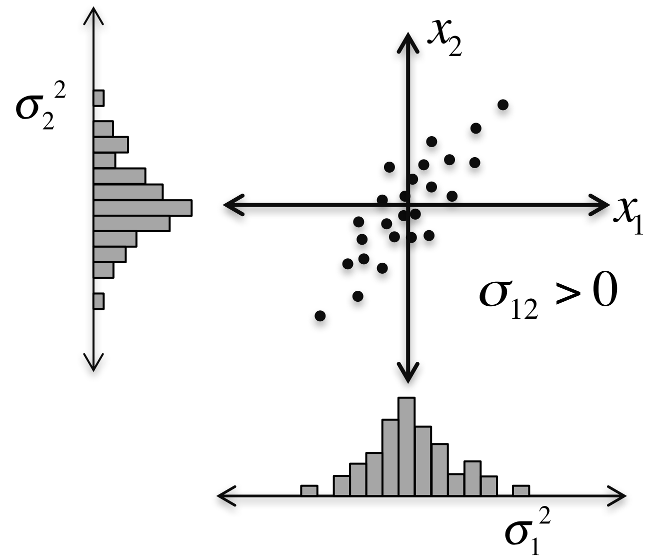

- x_1 has the same variance in both of these cases… also x_2
- So we need another measure to describe the relation between x_1 and x_2 .

covariance

$$\sigma_{12} = \frac{1}{m}\sum_{j=1}^{m}\left(x_1^{(j)} - \mu_1\right)\left(x_2^{(j)} - \mu_2\right)$$

correlation

$$\rho = \frac{\sigma_{12}}{\sqrt{\sigma_1^{\ 2}\sigma_2^{\ 2}}}$$

Gaussian distribution

- Many kinds of data can be fit by a Gaussian distribution

If x is a Gaussian random variable

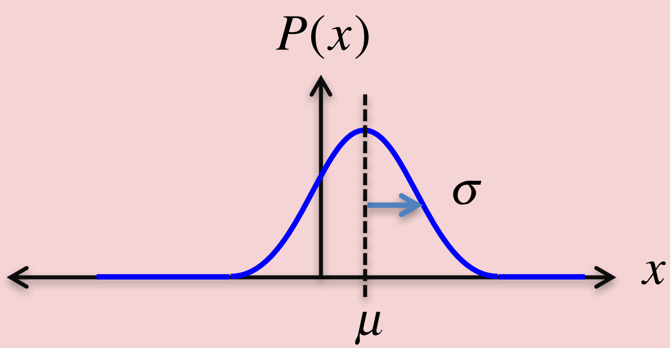

probability 'density'

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

- Gaussian is defined by only its mean and variance.

To find the Gaussian that best fits our data –

Just measure the mean and variance!

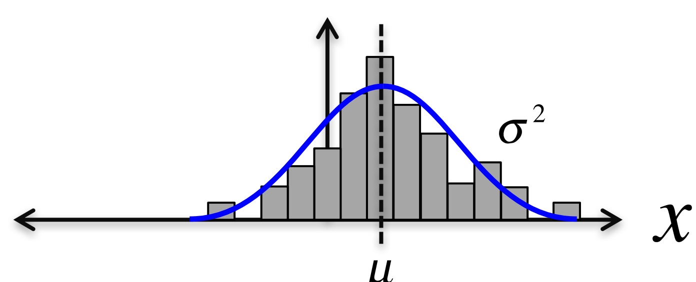

Gaussian distribution

- We are going develop a description of Gaussian distributions in higher dimensions

- Develop deep insights into high-dimensional data

- We will develop this description using vector and matrix notation

    - vectors and matrices are the natural format to manipulate data sets

    - very compact notation

    - manipulations are trivial in MATLAB®

Multivariate Gaussian distribution

- We can create a Gaussian distribution in two dimensions

Two independent Gaussian random variables x_1 and x_2

$$P(x_1, x_2) = P(x_1)P(x_2)$$

$$= \beta e^{-\frac{1}{2}x_1^2} e^{-\frac{1}{2}x_2^2}$$

normalization to make total prob = 1

$$= \beta e^{-\frac{1}{2}\left(x_1^2 + x_2^2\right)}$$

$$P(\vec{x}) = \beta e^{-\frac{1}{2}d^2}$$

$$d^2 = |\vec{x}|^2 = \vec{x}^T \vec{x}$$

d = Mahalanobis Distance

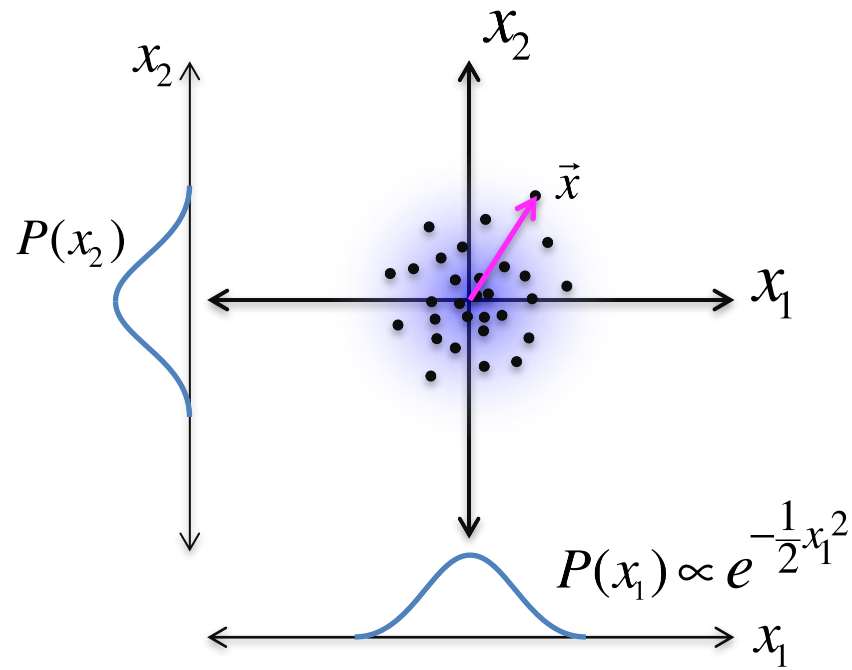

Isotropic multivariate Gaussian distribution

Multivariate Gaussian distribution

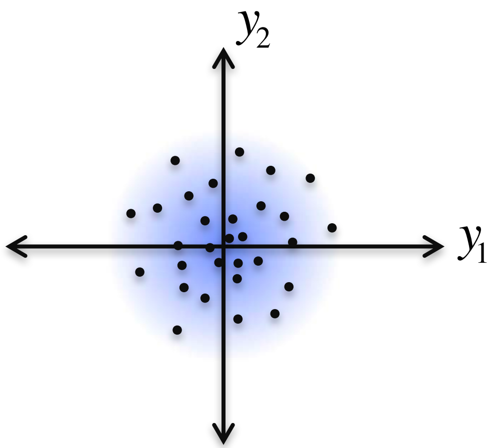

Isotropic

$$P(\vec{y}) = \beta e^{-\frac{1}{2\sigma^2}\vec{y}^T\vec{y}}$$

$$\sigma^2$$

variance

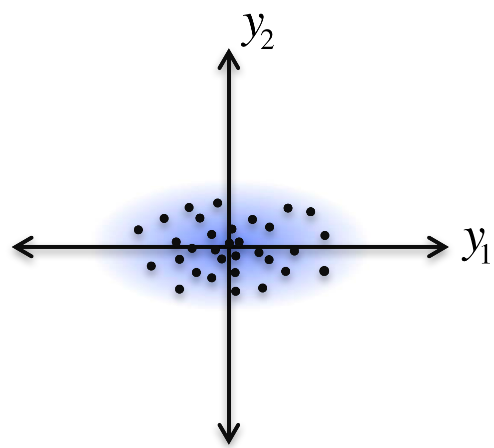

Non-isotropic:
no correlation

$$P(\vec{y}) = \beta e^{-\frac{1}{2}\vec{y}^T\Lambda^{-1}\vec{y}}$$

$$\Lambda = \begin{pmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{pmatrix}$$

variance matrix

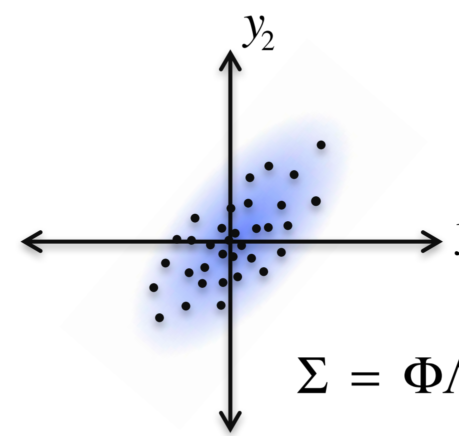

$$\Sigma = \Phi\Lambda\Phi^T$$

Non-isotropic:
with correlation

$$P(\vec{y}) = \beta e^{-\frac{1}{2}\vec{y}^T\Sigma^{-1}\vec{y}}$$

$$\Sigma = \begin{pmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{12} & \sigma_2^2 \end{pmatrix}$$

covariance matrix

Multivariate Gaussian distribution

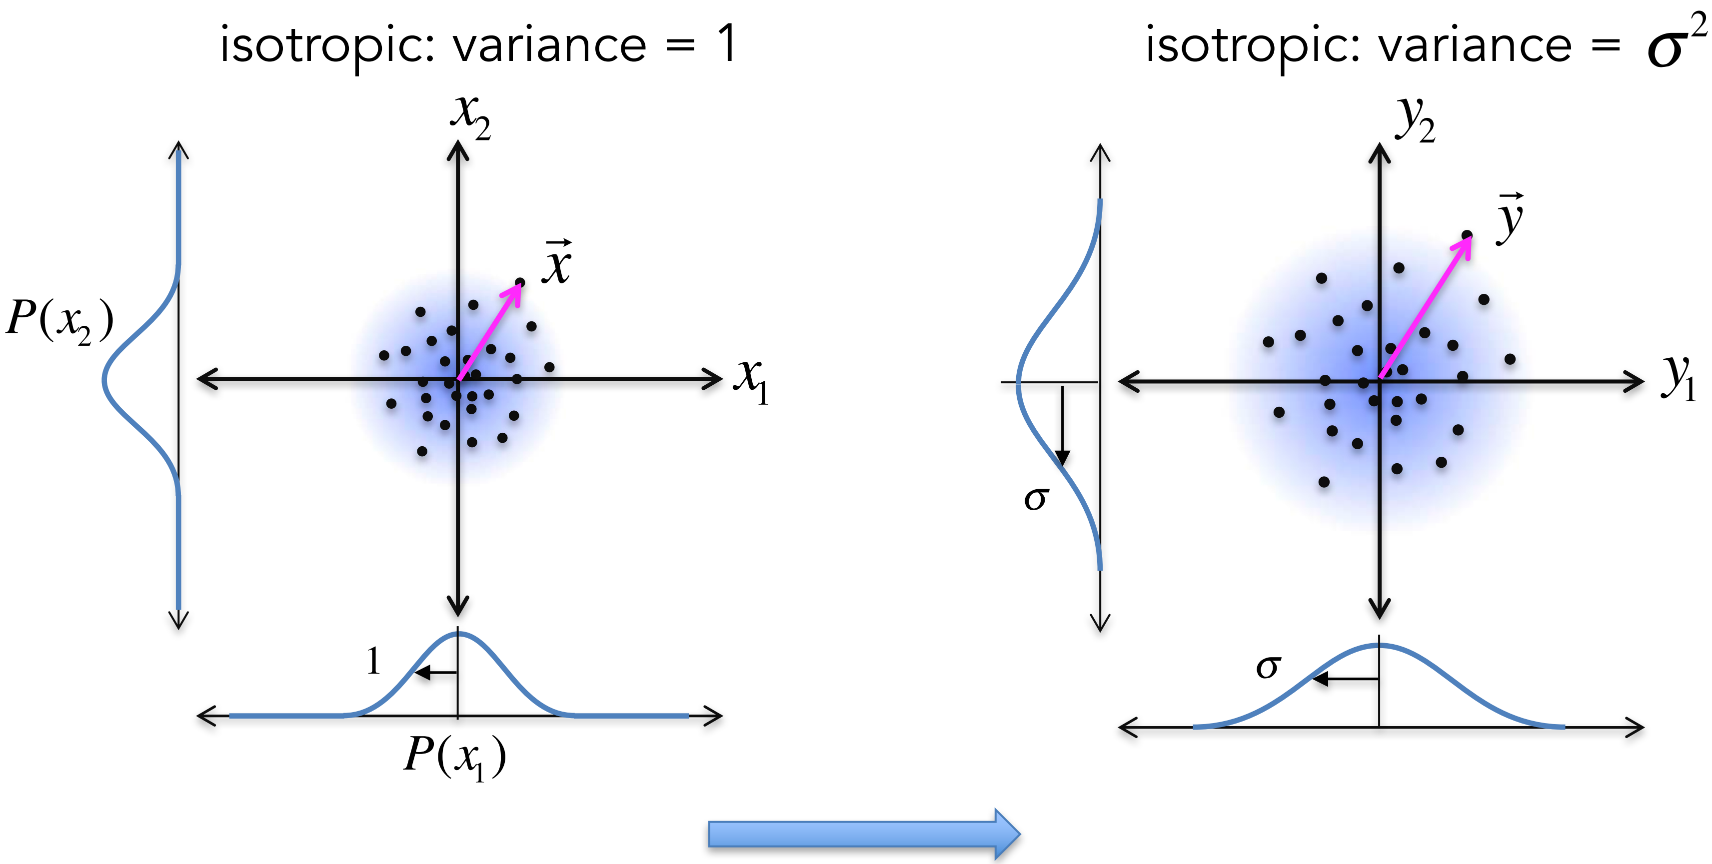

$$\vec{y} = \sigma\,\vec{x}$$

Multivariate Gaussian distribution

$P(\vec{y}) = ?$

$$\vec{y} = \sigma \vec{x}$$

isotropic: variance $= \sigma^2$

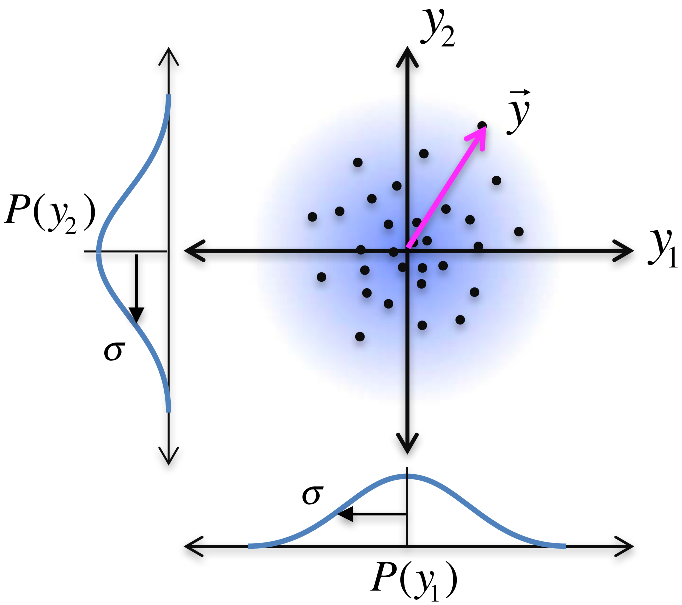

$$P(\vec{x}) = \beta e^{-\frac{1}{2}\vec{x}^T\vec{x}} \qquad \vec{x} = \sigma^{-1}\vec{y}$$

What is the Mahalanobis distance?

$$d^2 = \vec{x}^T\vec{x} = \left(\sigma^{-1}\vec{y}\right)^T\left(\sigma^{-1}\vec{y}\right)$$

$$= \vec{y}^T\sigma^{-1}\sigma^{-1}\vec{y}$$

$$d^2 = \vec{y}^T\sigma^{-2}\vec{y}$$

$$P(\vec{y}) = \beta e^{-\frac{1}{2}\left(\frac{\vec{y}^T\vec{y}}{\sigma^2}\right)} \qquad = \beta e^{-\frac{1}{2}\frac{|\vec{y}|^2}{\sigma^2}}$$

Multivariate Gaussian distribution

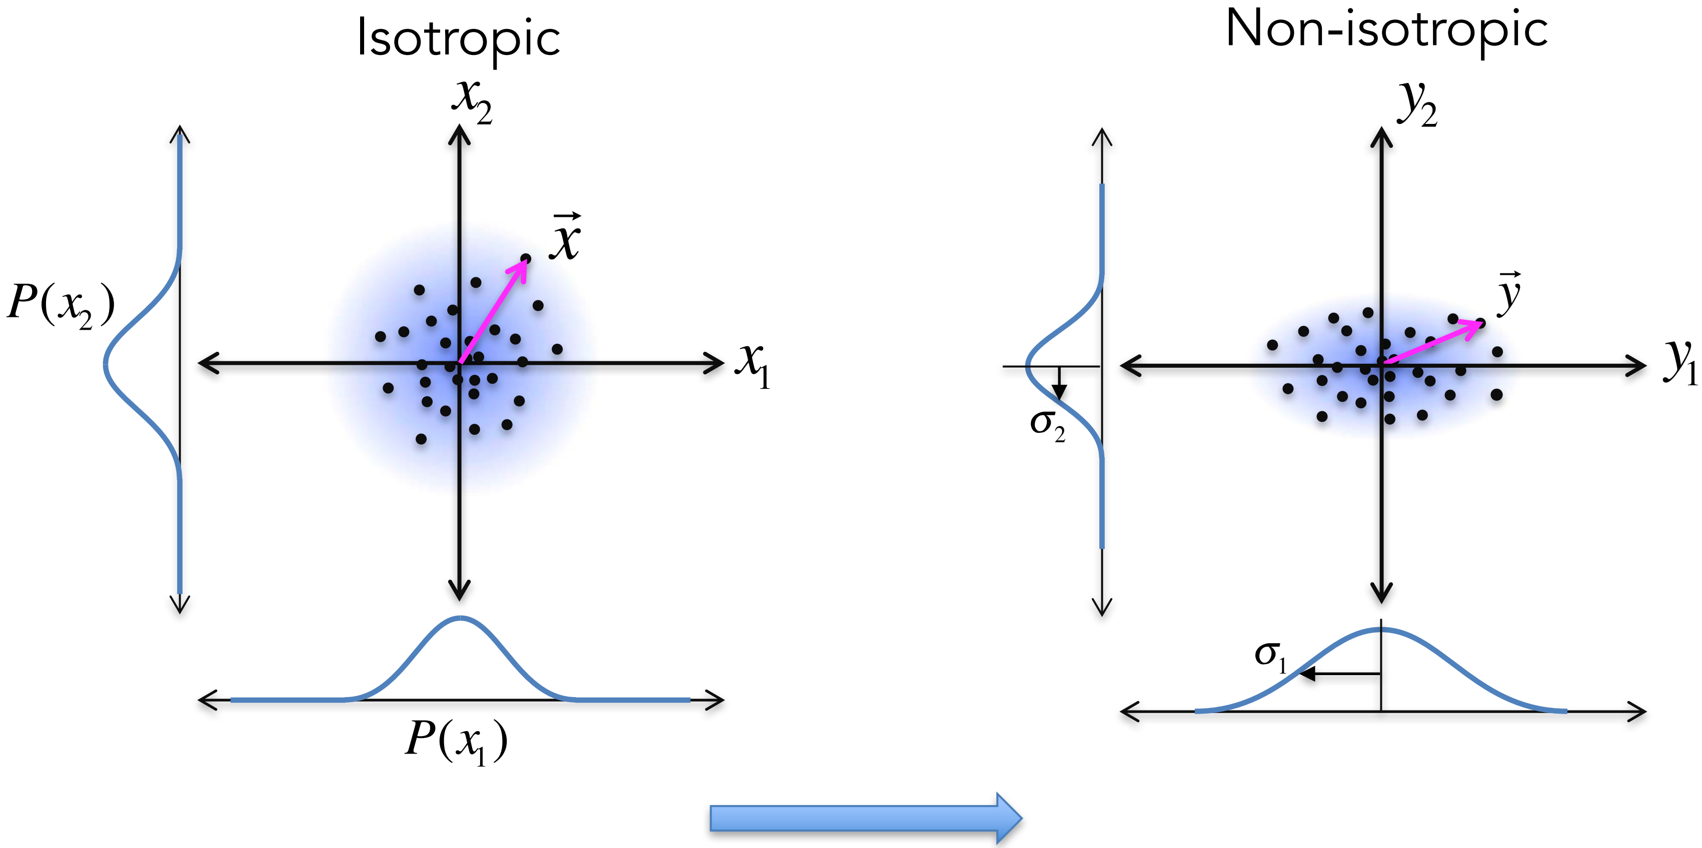

$$\vec{y} = S\,\vec{x}$$

$$S = \begin{pmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{pmatrix}$$

Multivariate Gaussian distribution

$$\vec{y} = S\vec{x} \qquad \vec{x} = S^{-1}\vec{y}$$

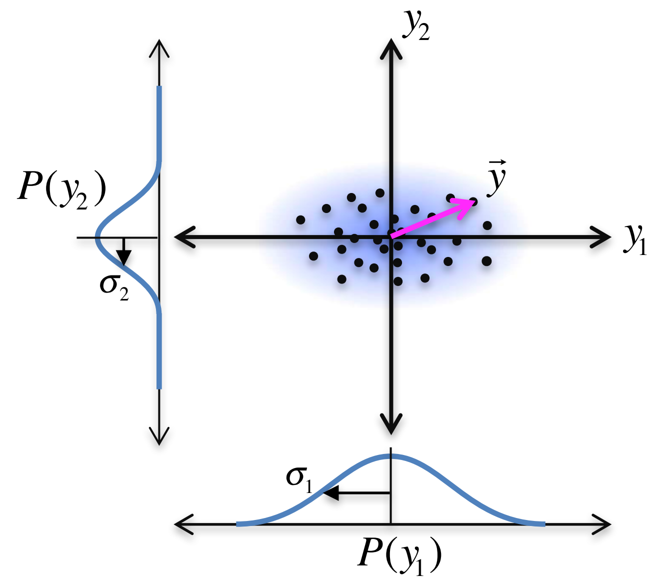

What is the Mahalanobis distance?

$$d^2 = \vec{x}^T\vec{x} = \left(S^{-1}\vec{y}\right)^T\left(S^{-1}\vec{y}\right)$$

$$= \vec{y}^T S^{-1} S^{-1}\vec{y}$$

$$= \vec{y}^T S^{-2}\vec{y}$$

$$d^2 = \vec{y}^T \Lambda^{-1}\vec{y}$$

$$\Lambda^{-1} = S^{-2} = \begin{pmatrix} \sigma_1^{-2} & 0 \\ 0 & \sigma_2^{-2} \end{pmatrix}$$

$$\Lambda = \begin{pmatrix} \sigma_1^{2} & 0 \\ 0 & \sigma_2^{2} \end{pmatrix}$$

matrix of variances

$$P(\vec{y}) = \beta e^{-\frac{1}{2}\vec{y}^T\Lambda^{-1}\vec{y}} = \beta e^{-\frac{1}{2}\left(\frac{y_1^2}{\sigma_1^2}+\frac{y_2^2}{\sigma_2^2}\right)}$$

Multivariate Gaussian distribution

Isotropic

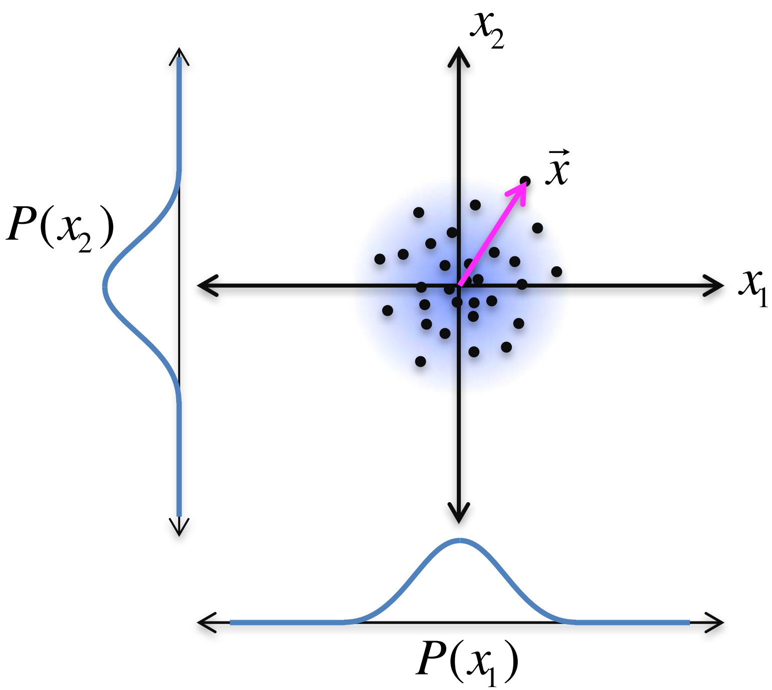

Non-isotropic:
with correlation

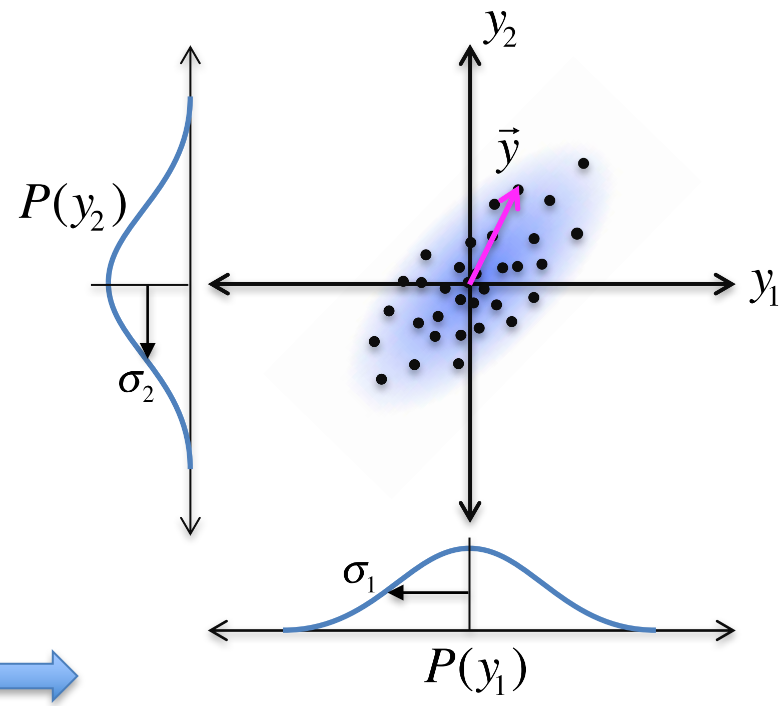

$$\vec{y} = \Phi S \Phi^T \vec{x}$$

stretch matrix

$$S = \begin{pmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{pmatrix}$$

rotation matrix

$$\Phi$$

Covariance matrix

$$\vec{y} = \Phi S \Phi^T \vec{x} \qquad \vec{x} = \Phi S^{-1} \Phi^T \vec{y}$$

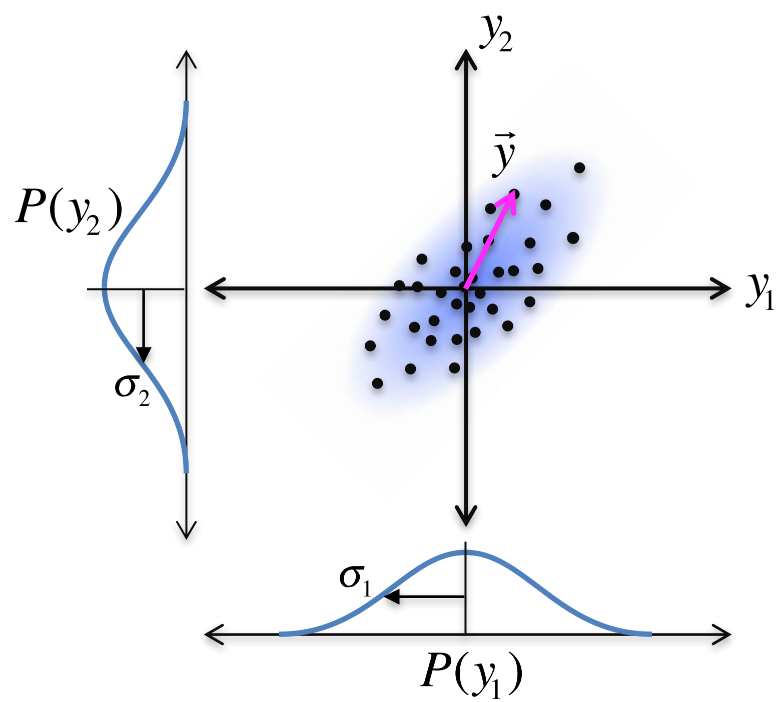

What is the Mahalanobis distance?

$$d^2 = \vec{x}^T \vec{x} = \left(\Phi S^{-1} \Phi^T \vec{y}\right)^T \left(\Phi S^{-1} \Phi^T \vec{y}\right)$$

$$= \vec{y}^T \Phi S^{-1} \Phi^T \Phi S^{-1} \Phi^T \vec{y}$$

$$= \vec{y}^T \Phi S^{-2} \Phi^T \vec{y}$$

$$= \vec{y}^T \underbrace{\Phi \Lambda^{-1} \Phi^T}_{} \vec{y}$$

$$d^2 = \vec{y}^T \Sigma^{-1} \vec{y}$$

$$\Sigma^{-1} = \Phi \Lambda^{-1} \Phi^T$$

$$\Sigma = \begin{pmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{12} & \sigma_2^2 \end{pmatrix}$$

covariance matrix

$$P(\vec{y}) = \beta e^{-\frac{1}{2}\vec{y}^T \Sigma^{-1} \vec{y}}$$

$$\Sigma = \Phi \Lambda \Phi^T$$

Multivariate Gaussian distribution

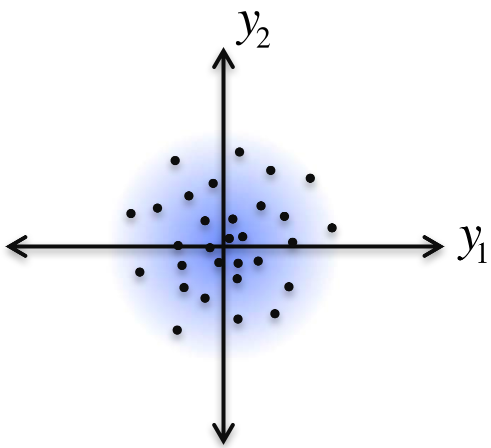

Isotropic

$$P(\vec{y}) = \beta e^{-\frac{1}{2\sigma^2}\vec{y}^T\vec{y}}$$

$$\sigma^2$$

variance

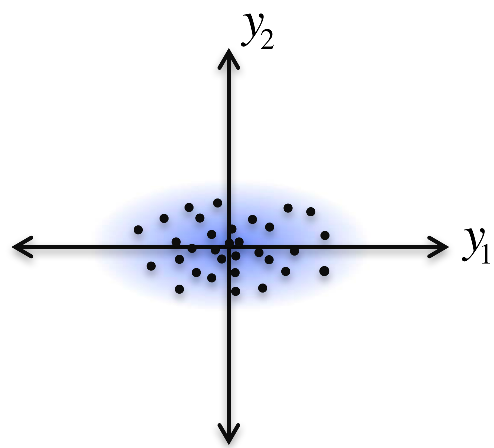

Non-isotropic:
no correlation

$$P(\vec{y}) = \beta e^{-\frac{1}{2}\vec{y}^T\Lambda^{-1}\vec{y}}$$

$$\Lambda = \begin{pmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{pmatrix}$$

variance matrix

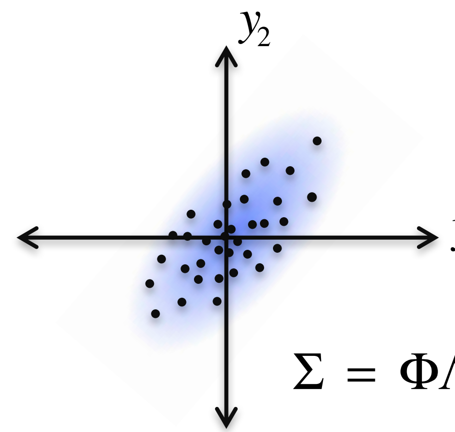

$$\Sigma = \Phi\Lambda\Phi^T$$

Non-isotropic:
with correlation

$$P(\vec{y}) = \beta e^{-\frac{1}{2}\vec{y}^T\Sigma^{-1}\vec{y}}$$

$$\Sigma = \begin{pmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{12} & \sigma_2^2 \end{pmatrix}$$

covariance matrix

Eigen-decomposition of the covariance matrix

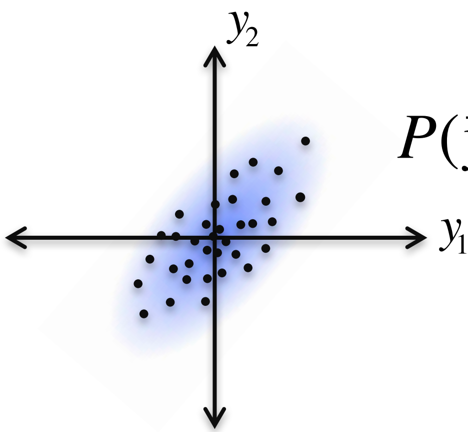

$$P(\vec{y}) = \beta e^{-\frac{1}{2}\vec{y}^T \Sigma^{-1} \vec{y}}$$

$$\Sigma^{-1} = \Phi \Lambda^{-1} \Phi^T$$

$$\Sigma = \Phi \Lambda \Phi^T$$

- Thus, our covariance matrix is just a transformation matrix that turns an isotropic Gaussian distribution (of variance 1) into non-isotropic multivariate Gaussian.

- The eigenvectors of the covariance matrix are just the basis vectors of the rotated transformation.

- And the eigenvalues of the covariance matrix are are the variances of the Gaussian in the directions of these basis vectors.

Learning Objectives for Lecture 17

- Eigenvectors and eigenvalues
- Variance and multivariate Gaussian distributions
- Computing a covariance matrix from data
- Principal Components Analysis (PCA)

How do we fit a Gaussian to multivariate data?

- In 1-dimension

  To find the Gaussian that best fits our data –

  Just measure the mean and variance!

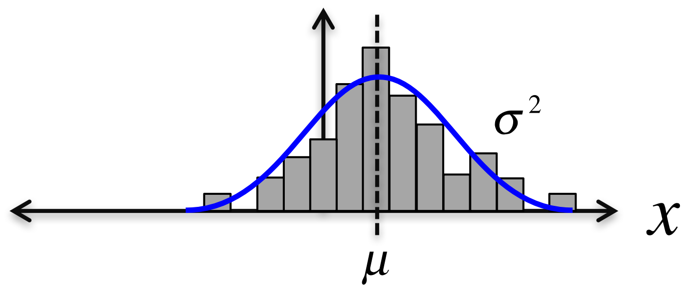

- Compute the covariance matrix!

$$\vec{x}^{(j)}\ ,\ j = 1,2,3....,m$$

- First we subtract the mean

$$\vec{z}^{(j)} = \vec{x}^{(j)} - \vec{\mu} \qquad \vec{\mu} = \frac{1}{m}\sum_{j=1}^{m}\vec{x}^{(j)}$$

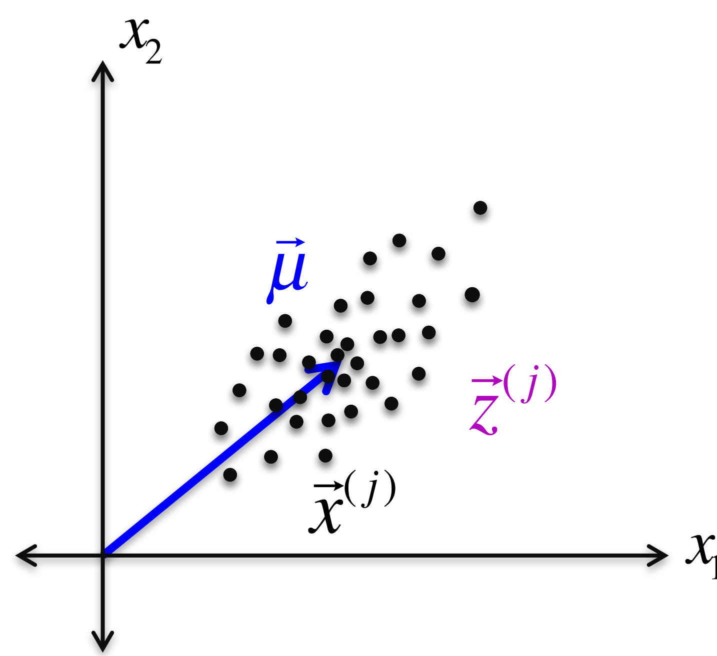

Computing the covariance matrix from data

- Compute the covariance matrix of a set of multivariate observations

$$\vec{x}^{(j)}\ ,\ j = 1,2,3....,m$$

- First we subtract the mean

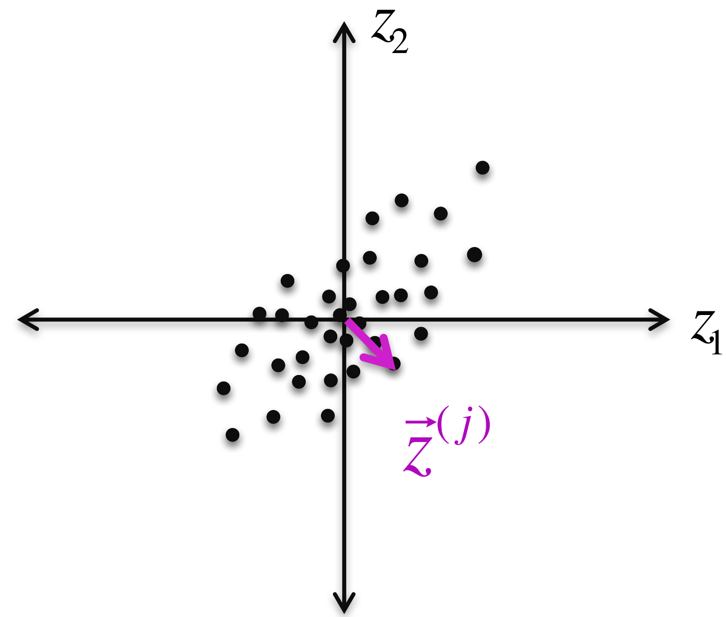

$$\vec{z}^{(j)} = \vec{x}^{(j)} - \vec{\mu} \qquad \vec{\mu} = \frac{1}{m}\sum_{j=1}^{m}\vec{x}^{(j)}$$

$$\Sigma = \begin{pmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{21} & \sigma_2^2 \end{pmatrix} = \begin{pmatrix} \sigma_1^2 = \frac{1}{m}\sum_{j=1}^{m} z_1^{(j)} z_1^{(j)} & \sigma_{12} = \frac{1}{m}\sum_{j=1}^{m} z_1^{(j)} z_2^{(j)} \\ \sigma_{21} = \frac{1}{m}\sum_{j=1}^{m} z_2^{(j)} z_1^{(j)} & \sigma_2^2 = \frac{1}{m}\sum_{j=1}^{m} z_2^{(j)} z_2^{(j)} \end{pmatrix}$$

Outer product

- We are going to implement a useful trick called the vector 'outer product'.

$$\vec{z} = \begin{pmatrix} z_1 \\ z_2 \end{pmatrix}$$

Inner product

$$\vec{z}^T \vec{z} = \begin{pmatrix} z_1 & z_2 \end{pmatrix} \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} = |\vec{z}|^2$$

$1\ x\ 2 \quad 2\ x\ 1 \quad 1\ x\ 1$

Outer product

$$\vec{z}\,\vec{z}^T = \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} \begin{pmatrix} z_1 & z_2 \end{pmatrix} = \begin{pmatrix} z_1 z_1 & z_1 z_2 \\ z_1 z_2 & z_2 z_2 \end{pmatrix}$$

$2\ x\ 1 \quad 1\ x\ 2 \quad 2\ x\ 2$

Computing the covariance matrix

- The covariance matrix has a simpler form using outer product.

$$\vec{z}^{(j)} = \begin{pmatrix} z_1^{(j)} \\ z_2^{(j)} \end{pmatrix}$$

$$\frac{1}{m}\sum_{j=1}^{m} \vec{z}^{(j)} (\vec{z}^{(j)})^T = \frac{1}{m}\sum_{j=1}^{m} \begin{pmatrix} z_1^{(j)} z_1^{(j)} & z_1^{(j)} z_2^{(j)} \\ z_2^{(j)} z_1^{(j)} & z_2^{(j)} z_2^{(j)} \end{pmatrix}$$

$$\Sigma = \begin{pmatrix} \sigma_1^{\ 2} & \sigma_{12} \\ \sigma_{21} & \sigma_{22}^{\ 2} \end{pmatrix}$$

Computing the covariance matrix

- Representing data as a matrix
- We have m observations of vector $\vec{z}$

$$\vec{z}^{(j)} = \begin{pmatrix} z_1^{(j)} \\ z_2^{(j)} \end{pmatrix}, \ j = 1 \ldots m$$

- Put them in matrix form as follows

$$Z = \begin{pmatrix} \vec{z}^{(1)} & \vec{z}^{(2)} & \vec{z}^{(3)} & \cdots & \vec{z}^{(m)} \end{pmatrix}$$

$$\begin{array}{ccccccc} j = & 1 & 2 & 3 & \ldots & m & = \text{number of samples} \\ Z = & \left(z_{11} \right. & z_{12} & z_{13} & \cdots & \left. z_{1m} \right) & n \quad \text{dimension of} \\ & \left(z_{21} \right. & z_{22} & z_{23} & \cdots & \left. z_{2m} \right) & \text{data vector} = 2 \end{array}$$

$$Z = \begin{pmatrix} z_{11} & z_{12} & z_{13} & \cdots & z_{1m} \\ z_{21} & z_{22} & z_{23} & \cdots & z_{2m} \end{pmatrix} \quad n$$

dimension of data vector =2

$n \ x \ m$

Computing the covariance matrix

- Now finding the covariance matrix is trivial!

$$\Sigma = \frac{1}{m} Z Z^T$$

$$= \frac{1}{m} \begin{pmatrix} z_{11} & z_{12} & z_{13} & \cdots & z_{1m} \\ z_{21} & z_{22} & z_{23} & \cdots & z_{2m} \end{pmatrix} \begin{pmatrix} z_{11} & z_{21} \\ z_{12} & z_{22} \\ z_{13} & z_{23} \\ & \\ z_{1m} & z_{2m} \end{pmatrix} = \begin{pmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{21} & \sigma_{22}^2 \end{pmatrix}$$

$2\ x\ m$ $\qquad$ $m\ x\ 2$ $\qquad$ $2\ x\ 2$

Subtracting the mean

- The covariance calculation we just did assumes the data were mean-subtracted. How to subtract the mean?

- We have m observations of vector $\vec{x}$

$$X = \underbrace{\begin{pmatrix} x_{11} & x_{12} & x_{13} & \cdots & x_{1m} \\ x_{21} & x_{22} & x_{23} & \cdots & x_{2m} \end{pmatrix}}_{n \ x \ m} \rightarrow \begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}$$

- First compute the mean in matrix notation and make a matrix with m copies of this column vector.

```
Mu=mean(X,2);
MU=repmat(mu,1,m);
```

$$\mathbf{M} = \underbrace{\begin{pmatrix} \mu_1 & \mu_1 & \mu_1 & \cdots & \mu_1 \\ \mu_2 & \mu_2 & \mu_2 & \cdots & \mu_2 \end{pmatrix}}_{n \ x \ m}$$

- Now subtract this from X to get Z

```
Z=X-MU;
```

Learning Objectives for Lecture 17

- Eigenvectors and eigenvalues
- Variance and multivariate Gaussian distributions
- Computing a covariance matrix from data
- Principal Components Analysis (PCA)

Principal Components Analysis

- A method for finding the directions in high-dimensional data that contain information.

Eigenfaces

Screen shot ©mikedusenberry.com. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.

Genetic profiling

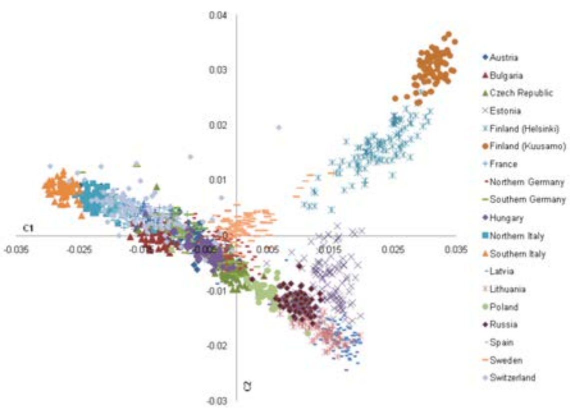

Single-cell transcriptional profiling

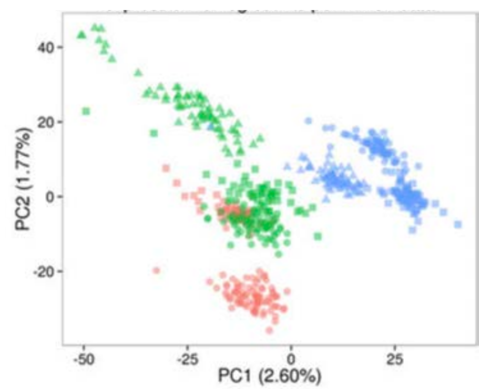

Eigenworm

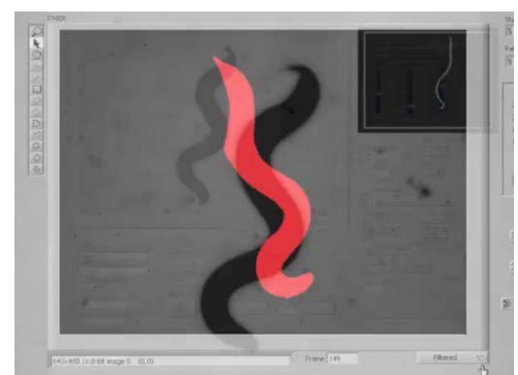

Screen shot ©JoVE. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.

Spike Sorting

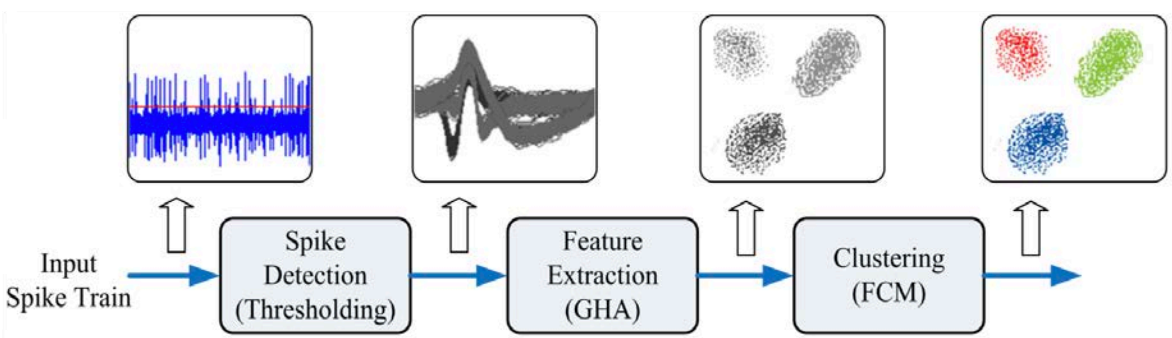

Figure 12 by Hwang, Wen-Jyi, et al., "[Efficient Architecture for Spike Sorting in Reconfigurable Hardware](#)." *Sensors* 13 no. 11 (2013): 14860-14887. MDPI Open Access. License: CC BY.

PCA demo on Gaussian points

$$\Phi = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix} \qquad S = \begin{pmatrix} \sqrt{3} & 0 \\ 0 & \left(\sqrt{3}\right)^{-1} \end{pmatrix}$$

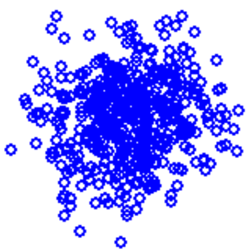

$$\vec{z} = \Phi S \Phi^T \vec{x}$$

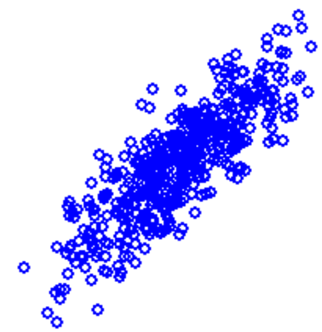

X

Z

```
m=500;
X=randn(2,m);

R=[1 -1;1 1]/sqrt(2);
S=[1.73  0 ; 0  0.577];
%
Z=R*S*R'*X;
```

PCA demo on Gaussian points

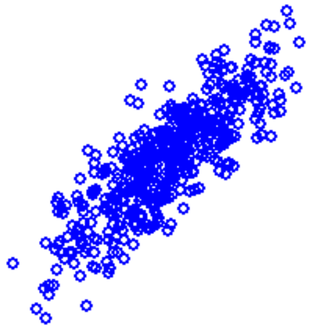

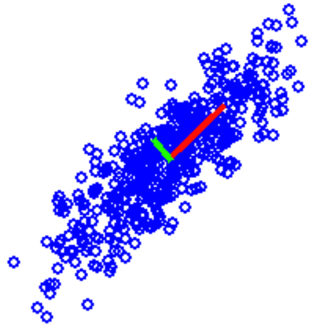

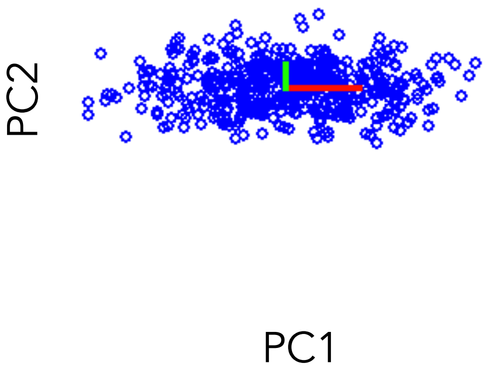

$$\Sigma = \frac{1}{m} Z Z^T$$

$$\Sigma = F V F^T$$

```
Q=Z*Z'/m;
[F,V]=eig(Q);
```

$$F = \begin{pmatrix} 0.72 & -0.70 \\ 0.70 & 0.72 \end{pmatrix}$$

$$V = \begin{pmatrix} 3.28 & 0 \\ 0 & 0.31 \end{pmatrix}$$

```
F=-fliplr(F);
V=flip(sum(V));
```

$$\vec{z}_f = F^T \vec{z}$$

```
Zf=F'*Z;
```

Clustering

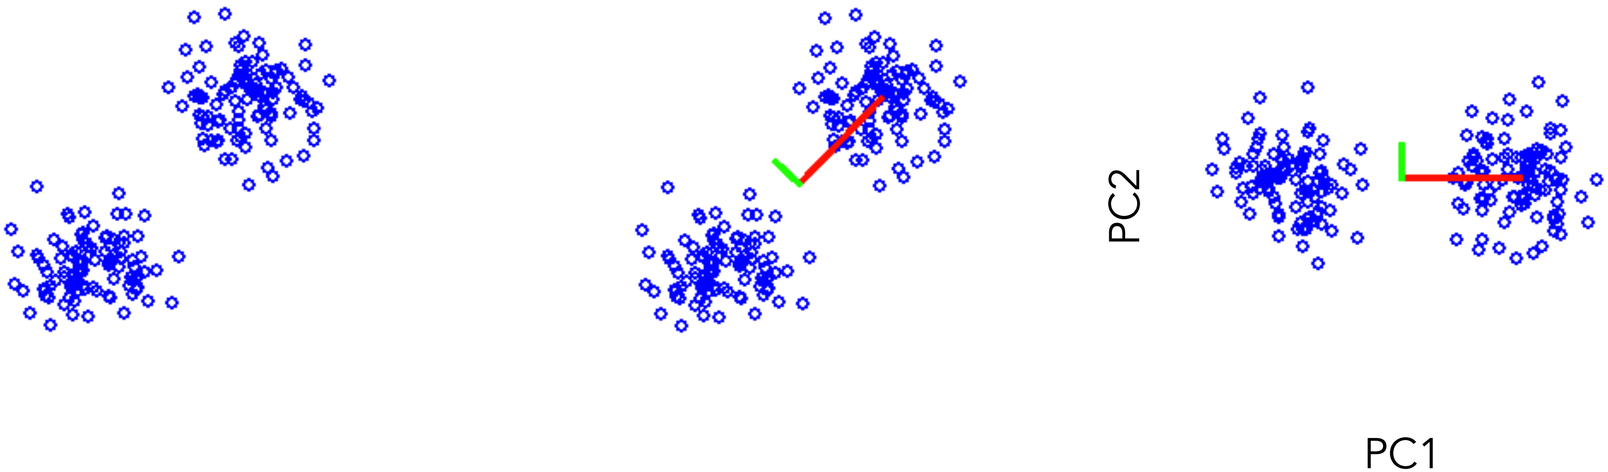

Q=Z*Z'/m;
[F,V]=eig(Q);

Zf=F'*Z;

PCA on time-domain signals

- Let's look at a problem in the time domain.
- Here we have many examples of a noisy signal in time.

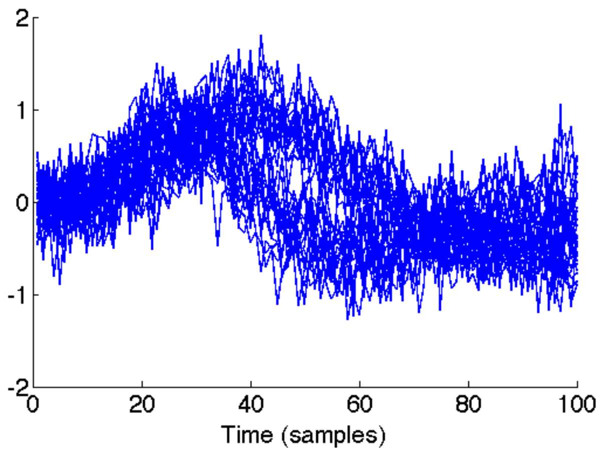

Each example has 100 time points

$$\vec{x}^{j} = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} \qquad n = 100$$

There are 200 different vectors

$$X = \begin{pmatrix} \vec{x}^{(1)} & \vec{x}^{(2)} & \vec{x}^{(3)} & \cdots & \vec{x}^{(m)} \end{pmatrix}$$

$$m = 200$$

$$X = \underbrace{\begin{pmatrix} \\ n \;\; x \;\; m \\ \\ \end{pmatrix}}_{200} 100$$

Covariance matrix

- Do PCA
  - Subtract the mean
  - Compute the covariance matrix
  - Find the eigenvectors and eigenvalues

$$\Sigma = \frac{1}{m} Z Z^T$$

```
Mu=mean(X,2);
MU=repmat(mu,1,m);
Z=X-MU;
Q=Z*Z'/m;


[F,V]=eig(cov);
```

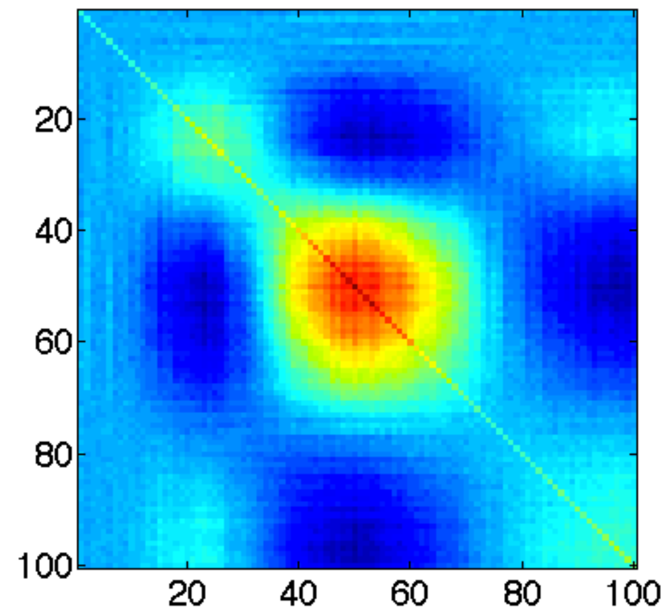

Eigenvalues

```
[F,V]=eig(cov);
var=flip(sum(V));
```

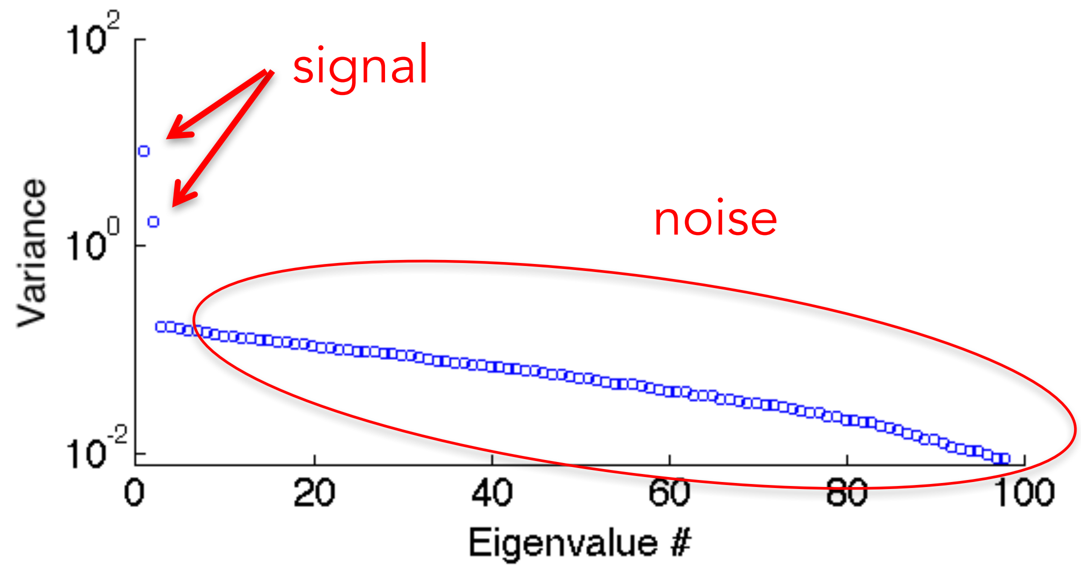

- The first two eigen-values are much larger than all the rest

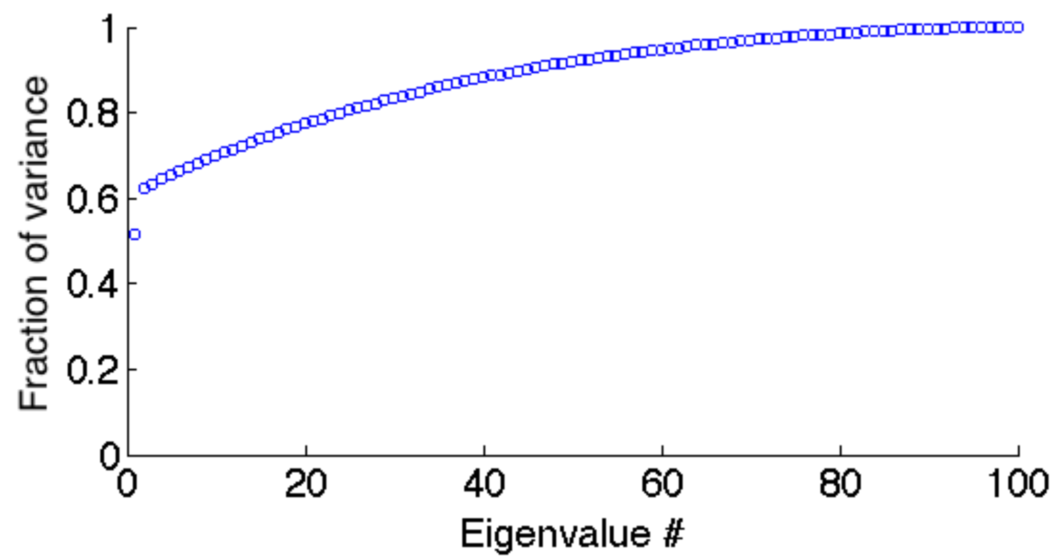

- The first two eigen-values explain over 60% of the total variance.

Eigenvectors

- Since there were only two large eigenvalues, we look at the eigenvectors associated with these eigenvalues

- These are just the first two columns of the F matrix

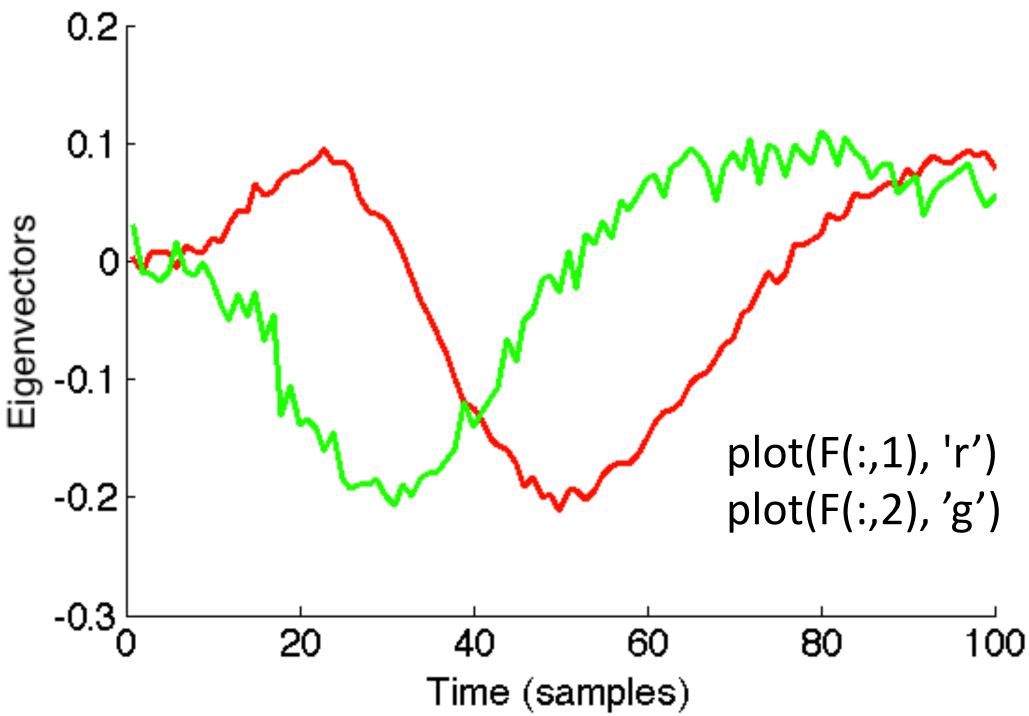

plot(F(:,3), 'r')
plot(F(:,4), 'g')

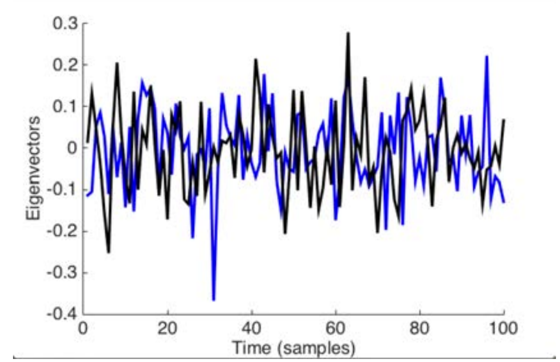

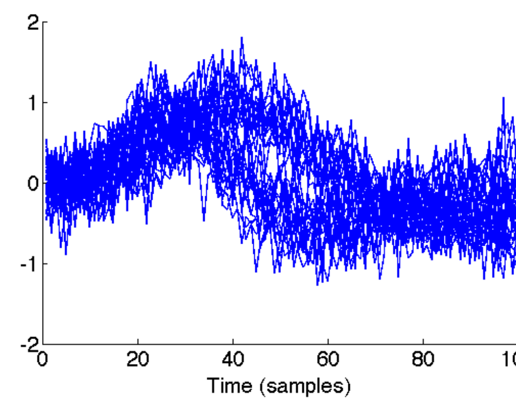

Principal components

- Principal components are just the projections of each of the original data vectors onto the two principal eigenvectors.

- Remember, this is just a change of basis using the matrix F

$$\vec{z}_f = F^T \vec{z}$$

```
Zf=F'*Z;
```

```
plot(Zf(1,:),Zf(2,:), 'o')
```

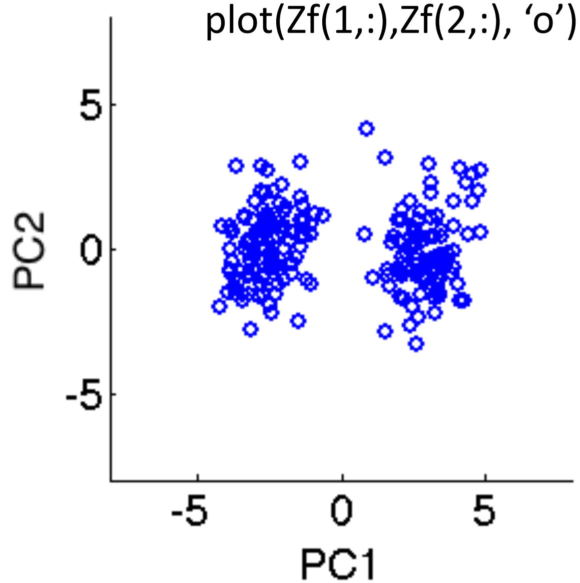


```
plot(Zf(2,:),Zf(3,:), 'o')
```

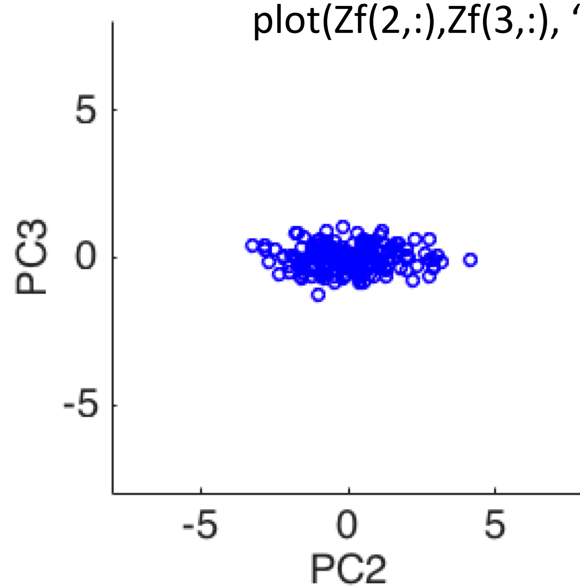

Filtering using PCA

- Only the first two entries in the column vectors Zf (in the rotated basis) have signal. So keep only the first two and set the rest to zero.

- Then rotate back to the original basis set

```
Zf=F'*Z;
Zffilt=Zf;
Zffilt(3:end,:)=0;
Zflt=F*Zffilt;
Xflt=Zflt+MU;
```

Before filtering

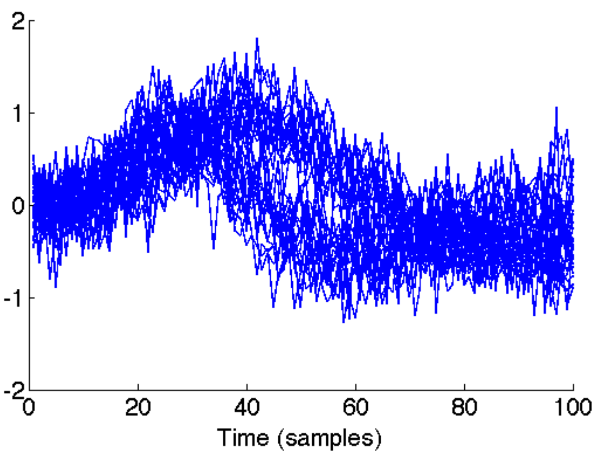

After filtering

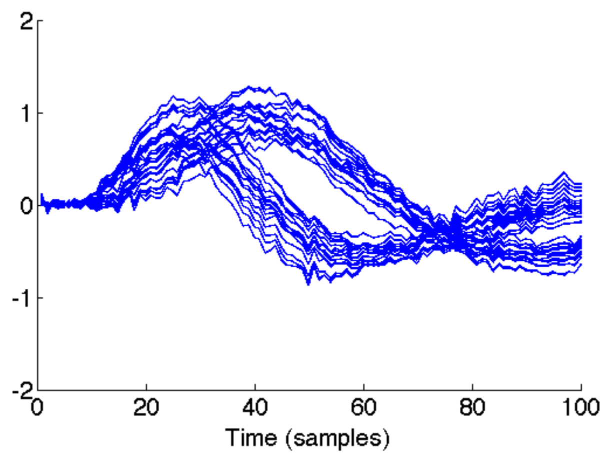

Learning Objectives for Lecture 17

- Eigenvectors and eigenvalues
- Variance and multivariate Gaussian distributions
- Computing a covariance matrix from data
- Principal Components Analysis (PCA)

MIT OpenCourseWare
https://ocw.mit.edu/

9.40 Introduction to Neural Computation
Spring 2018

For information about citing these materials or our Terms of Use, visit: https://ocw.mit.edu/terms.