# CS434a/541a: Pattern Recognition Prof. Olga Veksler 

## Lecture 7

## Today

- Problems of high dimensional data, "the curse of dimensionality"
- running time
- overfitting
- number of samples required

Dimensionality Reduction Methods

- Principle Component Analysis (today)
- Fisher Linear Discriminant (next time)


## Dimensionality on the Course Road Map

1. Bayesian Decision theory (rare case)

- Know probability distribution of the categories
- Do not even need training data
- Can design optimal classifier

2. ML and Bayesian parameter estimation

- Need to estimate Parameters of probability dist.
- Need training data

3. Non-Parametric Methods

- No probability distribution, labeled data

4. Linear discriminant functions and Neural Nets

- The shape of discriminant functions is known
- Need to estimate parameters of discriminant functions

5. Unsupervised Learning and Clustering

- No probability distribution and unlabeled data


## Curse of Dimensionality: Complexity

- Complexity (running time) increases with dimension d
- A lot of methods have at least $\mathrm{O}\left(\boldsymbol{n} \boldsymbol{d}^{2}\right)$ complexity, where $\boldsymbol{n}$ is the number of samples
- For example if we need to estimate covariance matrix
- So as $\boldsymbol{d}$ becomes large, $\mathrm{O}\left(\boldsymbol{n d}^{2}\right)$ complexity may be too costly


## Curse of Dimensionality: Overfitting

- If $\boldsymbol{d}$ is large, $\boldsymbol{n}$, the number of samples, may be too small for accurate parameter estimation
- For example, covariance matrix has $\boldsymbol{d}^{2}$ parameters:

$$
\Sigma=\left[\begin{array}{ccc}
\sigma_{1}^{2} & \cdots & \sigma_{1 d} \\
\vdots & \ddots & 1 \\
\sigma_{d 1} & \cdots & \sigma_{d}^{2}
\end{array}\right]
$$

- For accurate estimation, $\boldsymbol{n}$ should be much bigger than $\boldsymbol{d}^{2}$
- Otherwise model is too complicated for the data, overfitting:



## Curse of Dimensionality: Overfitting

- Paradox: If $\boldsymbol{n}<\boldsymbol{d}^{2}$ we are better off assuming that features are uncorrelated, even if we know this assumption is wrong
- In this case, the covariance matrix has only d parameters:

$$
\Sigma=\left[\begin{array}{ccc}
\sigma_{1}^{2} & \cdots & 0 \\
\vdots & \ddots & \vdots \\
0 & \cdots & \sigma_{d}^{2}
\end{array}\right]
$$

- We are likely to avoid overfitting because we fit a model with less parameters:



## Curse of Dimensionality: Number of Samples

- Suppose we want to use the nearest neighbor approach with $k=1$ (1NM)
- Suppose we start with only one feature

- This feature is not discriminative, i.e. it does not separate the classes well
- We decide to use 2 features. For the 1 NN method to work well, need a lot of samples, i.e. samples have to be dense
- To maintain the same density as in 1D (9 samples per unit length), how many samples do we need?


## Curse of Dimensionality: Number of Samples

- We need $9^{2}$ samples to maintain the same density as in 1D



## Curse of Dimensionality: Number of Samples

- Of course, when we go from 1 feature to 2 , no one gives us more samples, we still have 9

- This is way too sparse for $\mathbf{1 N N}$ to work well


## Curse of Dimensionality: Number of Samples

- Things go from bad to worse if we decide to use 3 features:

- If 9 was dense enough in 1D, in 3D we need $9^{3}=729$ samples!


## Curse of Dimensionality: Number of Samples

- In general, if $\boldsymbol{n}$ samples is dense enough in 1D
- Then in $\boldsymbol{d}$ dimensions we need $\boldsymbol{n}^{\boldsymbol{d}}$ samples!
- And $\boldsymbol{n}^{\boldsymbol{d}}$ grows really really fast as a function of $\boldsymbol{d}$
- Common pitfall:
- If we can't solve a problem with a few features, adding more features seems like a good idea
- However the number of samples usually stays the same
- The method with more features is likely to perform worse instead of expected better


## Curse of Dimensionality: Number of Samples

- For a fixed number of samples, as we add features, the graph of classification error:

- Thus for each fixed sample size $\boldsymbol{n}$, there is the optimal number of features to use


## The Curse of Dimensionality

- We should try to avoid creating lot of features
- Often no choice, problem starts with many features
- Example: Face Detection
- One sample point is $\boldsymbol{k}$ by $\boldsymbol{m}$ array of pixels
- Feature extraction is not trivial, usually every pixel is taken as a feature
- Typical dimension is 20 by $20=400$
- Suppose 10 samples are dense enough for 1 dimension. Need only $\mathbf{1 0 4 0 0}$ samples


## The Curse of Dimensionality

- Face Detection, dimension of one sample point is $\mathbf{k m}$
- The fact that we set up the problem with $\boldsymbol{k m}$ dimensions (features) does not mean it is really a $\mathbf{k m}$-dimensional problem
- Space of all $\boldsymbol{k}$ by $\boldsymbol{m}$ images has $\boldsymbol{k m}$ dimensions
- Space of all $\boldsymbol{k}$ by $\boldsymbol{m}$ faces must be much smaller, since faces form a tiny fraction of all possible images
- Most likely we are not setting the problem up with the right features
- If we used better features, we are likely need much less than $\boldsymbol{k m}$-dimensions


## Dimensionality Reduction

- High dimensionality is challenging and redundant
- It is natural to try to reduce dimensionality
- Reduce dimensionality by feature combination: combine old features $\boldsymbol{x}$ to create new features $\boldsymbol{y}$

$$
x=\left[\begin{array}{c}
x_{1} \\
x_{2} \\
\vdots \\
x_{d}
\end{array}\right] \rightarrow f\left(\left[\begin{array}{c}
x_{1} \\
x_{2} \\
\vdots \\
x_{d}
\end{array}\right]\right)=\left[\begin{array}{c}
y_{1} \\
\vdots \\
y_{k}
\end{array}\right]=y \quad \text { with } k<d
$$

- For example,

$$
x=\left[\begin{array}{l}
x_{1} \\
x_{2} \\
x_{3} \\
x_{4}
\end{array}\right] \rightarrow\left[\begin{array}{l}
x_{1}+x_{2} \\
x_{3}+x_{4}
\end{array}\right]=y
$$

- Ideally, the new vector $\boldsymbol{y}$ should retain from $\boldsymbol{x}$ all information important for classification


## Dimensionality Reduction

- The best $\boldsymbol{f}(\boldsymbol{x})$ is most likely a non-linear function
- Linear functions are easier to find though
- For now, assume that $\boldsymbol{f}(\boldsymbol{x})$ is a linear mapping
- Thus it can be represented by a matrix W:

$$
\left[\begin{array}{c}
x_{1} \\
x_{2} \\
\vdots \\
x_{d}
\end{array}\right] \Rightarrow W\left[\begin{array}{c}
x_{1} \\
x_{2} \\
\vdots \\
x_{d}
\end{array}\right]=\left[\begin{array}{clc}
w_{11} & \cdots & w_{1 d} \\
\vdots & & \vdots \\
w_{k 1} & \cdots & w_{k d}
\end{array}\right]\left[\begin{array}{c}
x_{1} \\
x_{2} \\
\vdots \\
x_{d}
\end{array}\right]=\left[\begin{array}{c}
y_{1} \\
\vdots \\
y_{k}
\end{array}\right] \quad \text { with } k<d
$$

## Feature Combination

- We will look at 2 methods for feature combination
- Principle Component Analysis (PCA)
- Fischer Linear Discriminant (next lecture)


## Principle Component Analysis (PCA)

- Main idea: seek most accurate data representation in a lower dimensional space
- Example in 2-D
- Project data to 1-D subspace (a line) which minimize the projection error

large projection errors, bad line to project to

small projection errors, good line to project to
- Notice that the the good line to use for projection lies in the direction of largest variance
- After the data is projected on the best line, need to transform the coordinate system to get 1D representation for vector $\boldsymbol{y}$

- Note that new data $\boldsymbol{y}$ has the same variance as old data $\boldsymbol{x}$ in the direction of the green line
- PCA preserves largest variances in the data. We will prove this statement, for now it is just an intuition of what PCA will do


## PCA: Approximation of Elliptical Cloud in 3D




- What is the direction of largest variance in data?
- Recall that if $\boldsymbol{x}$ has multivariate distribution $\mathrm{N}(\mu, \Sigma)$, direction of largest variance is given by eigenvector corresponding to the largest eigenvalue of $\Sigma$

- This is a hint that we should be looking at the covariance matrix of the data (note that PCA can be applied to distributions other than Gaussian)


## PCA: Linear Algebra for Derivation

- Let $\boldsymbol{V}$ be a $\boldsymbol{d}$ dimensional linear space, and $\boldsymbol{W}$ be a $\boldsymbol{k}$ dimensional linear subspace of $V$
- We can always find a set of dimensional vectors $\left\{\boldsymbol{e}_{1}, \boldsymbol{e}_{2}, \ldots, \boldsymbol{e}_{k}\right\}$ which forms an orthonormal basis for $\boldsymbol{W}$ - $\left\langle\boldsymbol{e}_{i}, \boldsymbol{e}_{j}\right\rangle=0$ if $\boldsymbol{i}$ is not equal to $\boldsymbol{j}$ and $\left\langle\boldsymbol{e}_{i}, \boldsymbol{e}_{\boldsymbol{i}}\right\rangle=1$
- Thus any vector in $W$ can be written as

$$
\alpha_{1} \boldsymbol{e}_{1}+\alpha_{2} \boldsymbol{e}_{2}+\ldots+\alpha_{k} \boldsymbol{e}_{k}=\sum_{i=1}^{k} \alpha_{i} \boldsymbol{e}_{i} \quad \text { for scalars } \alpha_{1}, \ldots, \alpha_{k}
$$



Let $\boldsymbol{V}=\boldsymbol{R}^{\mathbf{2}}$ and $\boldsymbol{W}$ be the line
$x-2 y=0$. Then the orthonormal basis for W is

$$
\left\{\left[\begin{array}{l}
1 / \sqrt{5} \\
-2 / \sqrt{5}
\end{array}\right]\right\}
$$

## PCA: Linear Algebra for Derivation

- Recall that subspace $W$ contains the zero vector, i.e. it goes through the origin

- For derivation, it will be convenient to project to subspace $W$ : thus we need to shift everything



## PCA Derivation: Shift by the Mean Vector

- Before PCA, subtract sample mean from the data

$$
x-\frac{1}{n} \sum_{i=1}^{n} x_{i}=x-\hat{\mu}
$$

- The new data has zero mean: $E(\boldsymbol{X}-\mathrm{E}(\boldsymbol{X}))=\mathrm{E}(\boldsymbol{X})-\mathrm{E}(\boldsymbol{X})=0$
- All we did is change the coordinate system

- Another way to look at it:
- first step of getting $\boldsymbol{y}$ is to subtract the mean of $\boldsymbol{x}$

$$
x \rightarrow y=f(x)=g(x-\hat{\mu})
$$

## PCA: Derivation

- We want to find the most accurate representation of data $\boldsymbol{D}=\left\{\boldsymbol{x}_{1}, \boldsymbol{x}_{2}, \ldots, \boldsymbol{x}_{\boldsymbol{n}}\right\}$ in some subspace $\boldsymbol{W}$ which has dimension $\boldsymbol{k}<\boldsymbol{d}$
- Let $\left\{\boldsymbol{e}_{1}, \boldsymbol{e}_{2}, \ldots, \boldsymbol{e}_{k}\right\}$ be the orthonormal basis for $\boldsymbol{W}$. Any vector in $\boldsymbol{W}$ can be written as $\sum_{i=1}^{k} \alpha_{i} \boldsymbol{e}_{i}$
- Thus $\boldsymbol{x}_{\boldsymbol{1}}$ will be represented by some vector in $\boldsymbol{W}$

$$
\sum_{i=1}^{k} \alpha_{1 i} e_{i}
$$

- Error this representation:

$$
\text { error }=\left\|x_{1}-\sum_{i=1}^{k} \alpha_{1 i} e_{i}\right\|^{2}
$$



## PCA: Derivation

- To find the total error, we need to sum over all $\boldsymbol{x}_{\boldsymbol{j}}$ 's
- Any $\boldsymbol{x}_{\boldsymbol{j}}$ can be written as $\sum_{i=1}^{k} \alpha_{j i} \boldsymbol{e}_{i}$
- Thus the total error for representation of all data $\boldsymbol{D}$ is:
sum over all data points
$J(\underbrace{\boldsymbol{e}_{1}, \ldots, \boldsymbol{e}_{k}, \alpha_{11}, \ldots \alpha_{n k}}_{\text {unknowns }})=\sum_{j=1}^{n}\left\|x_{i}-\sum_{i=1}^{k} \alpha_{j i} \boldsymbol{e}_{i}\right\|^{2}$


## PCA: Derivation

- To minimize $J$, need to take partial derivatives and also enforce constraint that $\left\{\boldsymbol{e}_{1}, \boldsymbol{e}_{2}, \ldots, \boldsymbol{e}_{k}\right\}$ are orthogonal

$$
J\left(e_{1}, \ldots, e_{k}, \alpha_{11}, \ldots \alpha_{n k}\right)=\sum_{j=1}^{n}\left\|x_{j}-\sum_{i=1}^{k} \alpha_{j i} \boldsymbol{e}_{i}\right\|^{2}
$$

- Let us simplify $\boldsymbol{J}$ first

$$
\begin{aligned}
J\left(\boldsymbol{e}_{1}, \ldots, \boldsymbol{e}_{k}, \alpha_{11}, \ldots \alpha_{n k}\right) & =\sum_{j=1}^{n}\left\|\boldsymbol{x}_{j}\right\|^{2}-2 \sum_{j=1}^{n} \boldsymbol{x}_{j}^{t}\left(\sum_{i=1}^{k} \alpha_{j i} \boldsymbol{e}_{i}\right)+\sum_{j=1}^{n} \sum_{i=1}^{k} \alpha_{j i}^{2} \\
& =\sum_{j=1}^{n}\left\|\boldsymbol{x}_{j}\right\|^{2}-2 \sum_{j=1}^{n} \sum_{i=1}^{k} \alpha_{j i} \boldsymbol{x}_{j}^{t} \boldsymbol{e}_{i}+\sum_{j=1}^{n} \sum_{i=1}^{k} \alpha_{j i}^{2}
\end{aligned}
$$

## PCA: Derivation

$J\left(\boldsymbol{e}_{1}, \ldots, \boldsymbol{e}_{k}, \alpha_{11}, \ldots \alpha_{n k}\right)=\sum_{j=1}^{n}\left|\boldsymbol{x}_{j}\right|^{2}-2 \sum_{j=1}^{n} \sum_{i=1}^{k} \alpha_{j i} \boldsymbol{x}_{j} \boldsymbol{e}_{i}+\sum_{j=1}^{n} \sum_{i=1}^{k} \alpha_{j i}^{2}$

- First take partial derivatives with respect to $\alpha_{m l}$

$$
\frac{\partial}{\partial \alpha_{m l}} J\left(e_{1}, \ldots, e_{k}, \alpha_{11}, \ldots \alpha_{n k}\right)=-2 x_{m}^{t} e_{l}+2 \alpha_{m l}
$$

- Thus the optimal value for $\alpha_{m l}$ is

$$
-2 x_{m}^{t} e_{l}+2 \alpha_{m l}=0 \Rightarrow \alpha_{m l}=x_{m}^{t} e_{l}
$$

## PCA: Derivation

$\boldsymbol{J}\left(\boldsymbol{e}_{1}, \ldots, \boldsymbol{e}_{k}, \alpha_{11}, \ldots \alpha_{n k}\right)=\sum_{j=1}^{n}\left\|\boldsymbol{x}_{j}\right\|^{2}-2 \sum_{j=1}^{n} \sum_{i=1}^{k} \alpha_{j i} \boldsymbol{x}_{j}^{t} \boldsymbol{e}_{i}+\sum_{j=1}^{n} \sum_{i=1}^{k} \alpha_{j i}^{2}$

- Plug the optimal value for $\alpha_{m l}=\boldsymbol{x}_{\boldsymbol{m}}^{\boldsymbol{t}} \boldsymbol{e}_{\boldsymbol{l}}$ back into $\boldsymbol{J}$

$$
J\left(e_{1}, \ldots, e_{k}\right)=\sum_{j=1}^{n}\left\|x_{j}\right\|^{2}-2 \sum_{j=1}^{n} \sum_{i=1}^{k}\left(x_{j}^{t} e_{i}\right) x_{j}^{t} e_{i}+\sum_{j=1}^{n} \sum_{i=1}^{k}\left(x_{j}^{t} e_{i}\right)^{2}
$$

- Can simplify J

$$
J\left(\boldsymbol{e}_{1}, \ldots, \boldsymbol{e}_{k}\right)=\sum_{j=1}^{n}\left\|\boldsymbol{x}_{j}\right\|^{2}-\sum_{j=1}^{n} \sum_{i=1}^{k}\left(\boldsymbol{x}_{j}^{t} \boldsymbol{e}_{i}\right)^{2}
$$

## PCA: Derivation

$$
J\left(e_{1}, \ldots, e_{k}\right)=\sum_{j=1}^{n}\left\|x_{j}\right\|^{2}-\sum_{j=1}^{n} \sum_{i=1}^{k}\left(x_{j}^{t} e_{i}\right)^{2}
$$

- Rewrite $\boldsymbol{J}$ using $\left(\boldsymbol{a}^{t} \boldsymbol{b}\right)^{2}=\left(\boldsymbol{a}^{t} \boldsymbol{b}\right)\left(\boldsymbol{a}^{t} \boldsymbol{b}\right)=\left(\boldsymbol{b}^{t} \boldsymbol{a}\right)\left(\boldsymbol{a}^{t} \boldsymbol{b}\right)=\boldsymbol{b}^{\boldsymbol{t}}\left(\boldsymbol{a} \boldsymbol{a}^{t}\right) \boldsymbol{b}$

$$
\begin{aligned}
J\left(e_{1}, \ldots, e_{k}\right) & =\sum_{j=1}^{n}\left\|x_{j}\right\|^{2}-\sum_{i=1}^{k} \boldsymbol{e}_{i}^{t}\left(\sum_{j=1}^{n}\left(x_{j} x_{j}^{t}\right)\right) \boldsymbol{e}_{i} \\
& =\sum_{j=1}^{n}\left\|\boldsymbol{x}_{j}\right\|^{2}-\sum_{i=1}^{k} \boldsymbol{e}_{i}^{t} \boldsymbol{S} \boldsymbol{e}_{i}
\end{aligned}
$$

- Where $\boldsymbol{S}=\sum_{j=1}^{n} \boldsymbol{x}_{j} \boldsymbol{x}_{j}^{t}$
- $S$ is called the scatter matrix, it is just n-1 times the sample covariance matrix we have seen before

$$
\hat{\Sigma}=\frac{1}{n-1} \sum_{j=1}^{n}\left(x_{j}-\hat{\mu}\right)\left(x_{j}-\hat{\mu}\right)^{t}
$$

## PCA: Derivation

$$
J\left(e_{1}, \ldots, e_{k}\right)=\sum_{\substack{j=1 \\ \text { constant }}}^{n}\left\|x_{i}\right\|^{2}-\sum_{i=1}^{k} e_{i}^{t} S e_{i}
$$

- Minimizing $\boldsymbol{J}$ is equivalent to maximizing $\sum_{i=1}^{k} \boldsymbol{e}_{i}^{t} \boldsymbol{S} \boldsymbol{e}_{i}$
- We should also enforce constraints $\boldsymbol{e}_{i}^{\boldsymbol{t}} \boldsymbol{e}_{\boldsymbol{i}}=1$ for all $\boldsymbol{i}$
- Use the method of Lagrange multipliers, incorporate the constraints with undetermined $\lambda_{1}, \ldots, \lambda_{k}$
- Need to maximize new function $\boldsymbol{u}$

$$
u\left(e_{1}, \ldots, e_{k}\right)=\sum_{i=1}^{k} e_{i}^{t} S e_{i}-\sum_{j=1}^{k} \lambda_{j}\left(e_{j}^{t} e_{j}-1\right)
$$

## PCA: Derivation

- If $\boldsymbol{x}$ is a vector and $\boldsymbol{f}(\boldsymbol{x})=\boldsymbol{f}\left(\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{\boldsymbol{d}}\right)$ is a function, to simplify notation, define

$$
\frac{\boldsymbol{d}}{\boldsymbol{d} \boldsymbol{x}} \boldsymbol{f}(\boldsymbol{x})=\left[\begin{array}{c}
\frac{\partial \boldsymbol{f}}{\partial x_{1}} \\
\vdots \\
\frac{\partial f}{\partial x_{d}}
\end{array}\right]
$$

- It can be shown that $\frac{\boldsymbol{d}}{d \boldsymbol{x}}\left(\boldsymbol{x}^{t} \boldsymbol{x}\right)=\mathbf{2 x}$
- If $\boldsymbol{A}$ is a symmetric matrix, it can be shown that

$$
\frac{d}{d x}\left(x^{t} A x\right)=2 A x
$$

$$
u\left(e_{1}, \ldots, e_{k}\right)=\sum_{i=1}^{k} e_{i}^{t} S e_{i}-\sum_{j=1}^{k} \lambda_{j}\left(e_{j}^{t} e_{j}-1\right)
$$

- Compute the partial derivatives with respect to $\boldsymbol{e}_{\boldsymbol{m}}$

$$
\frac{\partial}{\partial e_{m}} u\left(e_{1}, \ldots, e_{k}\right)=2 S e_{m}-2 \lambda_{m} e_{m}=0
$$

Note: $\boldsymbol{e}_{m}$ is a vector, what we are really doing here is taking partial derivatives with respect to each element of $\boldsymbol{e}_{\boldsymbol{m}}$ and then arranging them up in a linear equation

- Thus $\lambda_{m}$ and $\boldsymbol{e}_{m}$ are eigenvalues and eigenvectors of scatter matrix $\boldsymbol{S}$

$$
\boldsymbol{S} \boldsymbol{e}_{\boldsymbol{m}}=\lambda_{\boldsymbol{m}} \boldsymbol{e}_{\boldsymbol{m}}
$$

## PCA: Derivation

$$
J\left(e_{1}, \ldots, e_{k}\right)=\sum_{j=1}^{n}\left\|x_{j}\right\|^{2}-\sum_{i=1}^{k} e_{i}^{t} S e_{i}
$$

- Let's plug $\boldsymbol{e}_{\boldsymbol{m}}$ back into $\boldsymbol{J}$ and use $\boldsymbol{S e}_{\boldsymbol{m}}=\lambda_{m} \boldsymbol{e}_{m}$

$$
J\left(\boldsymbol{e}_{1}, \ldots, \boldsymbol{e}_{k}\right)=\sum_{j=1}^{n}\left\|\boldsymbol{x}_{j}\right\|^{2}-\sum_{i=1}^{k} \lambda_{i}\left\|\boldsymbol{e}_{i}\right\|^{2}=\sum_{\substack{j=1 \\ \text { constant }}}^{n}\left\|\boldsymbol{x}_{j}\right\|^{2}-\sum_{i=1}^{k} \lambda_{i}
$$

- Thus to minimize $\boldsymbol{J}$ take for the basis of $\boldsymbol{W}$ the $\boldsymbol{k}$ eigenvectors of $\boldsymbol{S}$ corresponding to the $\boldsymbol{k}$ largest eigenvalues
- The larger the eigenvalue of $\boldsymbol{S}$, the larger is the variance in the direction of corresponding eigenvector

- This result is exactly what we expected: project $\boldsymbol{x}$ into subspace of dimension $\boldsymbol{k}$ which has the largest variance
- This is very intuitive: restrict attention to directions where the scatter is the greatest
- Thus PCA can be thought of as finding new orthogonal basis by rotating the old axis until the directions of maximum variance are found



## PCA as Data Approximation

- Let $\left\{\boldsymbol{e}_{1}, \boldsymbol{e}_{2}, \ldots, \boldsymbol{e}_{d}\right\}$ be all $\boldsymbol{d}$ eigenvectors of the scatter matrix S, sorted in order of decreasing corresponding eigenvalue
- Without any approximation, for any sample $\boldsymbol{x}_{\boldsymbol{i}}$ :
error of approximation

$$
\boldsymbol{x}_{i}=\sum_{j=1}^{d} \alpha_{j} \boldsymbol{e}_{j}=\underbrace{\alpha_{1} \boldsymbol{e}_{1}+\ldots+\alpha_{k} \boldsymbol{e}_{k}}_{\text {approximation of } x_{i}}+\overbrace{\alpha_{k+1} \boldsymbol{e}_{k+1} \ldots+\alpha_{d} \boldsymbol{e}_{d}}
$$

- coefficients $\alpha_{m}=\boldsymbol{x}_{\boldsymbol{i}}^{\boldsymbol{t}} \boldsymbol{e}_{\boldsymbol{m}}$ are called principle components
- The larger $\boldsymbol{k}$, the better is the approximation
- Components are arranged in order of importance, more important components come first
- Thus PCA takes the first $\boldsymbol{k}$ most important components of $\boldsymbol{x}_{\boldsymbol{i}}$ as an approximation to $\boldsymbol{x}_{\boldsymbol{i}}$


## PCA: Last Step

- Now we know how to project the data
- Last step is to change the coordinates to get final $\boldsymbol{k}$-dimensional vector $\boldsymbol{y}$

- Let matrix $E=\left[\boldsymbol{e}_{1} \cdots \boldsymbol{e}_{k}\right]$
- Then the coordinate transformation is $\boldsymbol{y}=\boldsymbol{E}^{\boldsymbol{t}} \boldsymbol{x}$
- Under $E^{t}$, the eigenvectors become the standard basis:

$$
E^{t} e_{i}=\left[\begin{array}{c}
e_{1} \\
\vdots \\
e_{i} \\
\vdots \\
e_{k}
\end{array}\right] e_{i}=\left[\begin{array}{c}
0 \\
\vdots \\
\vdots \\
\vdots \\
0
\end{array}\right]
$$

## Recipe for Dimension Reduction with PCA

Data $D=\left\{\boldsymbol{x}_{1}, \boldsymbol{x}_{2}, \ldots, \boldsymbol{x}_{n}\right\}$. Each $\boldsymbol{x}_{\boldsymbol{i}}$ is a $\boldsymbol{d}$-dimensional vector. Wish to use PCA to reduce dimension to $\boldsymbol{k}$

1. Find the sample mean $\hat{\mu}=\frac{1}{n} \sum_{i=1}^{n} x_{i}$
2. Subtract sample mean from the data $\boldsymbol{z}_{\boldsymbol{i}}=\boldsymbol{x}_{\boldsymbol{i}}-\hat{\mu}$
3. Compute the scatter matrix $S=\sum_{i=1}^{n} \boldsymbol{z}_{i} \boldsymbol{z}_{i}^{t}$
4. Compute eigenvectors $\boldsymbol{e}_{1}, \boldsymbol{e}_{2}, \ldots, \boldsymbol{e}_{\boldsymbol{k}}$ corresponding to the $\boldsymbol{k}$ largest eigenvalues of $\boldsymbol{S}$
5. Let $\boldsymbol{e}_{1}, \boldsymbol{e}_{2}, \ldots, \boldsymbol{e}_{\boldsymbol{k}}$ be the columns of matrix $\boldsymbol{E}=\left[\boldsymbol{e}_{1} \cdots \boldsymbol{e}_{k}\right]$
6. The desired $\boldsymbol{y}$ which is the closest approximation to $\boldsymbol{x}$ is $\boldsymbol{y}=\boldsymbol{E}^{\boldsymbol{t}} \boldsymbol{z}$

## PCA Example Using Matlab

- Let $\boldsymbol{D}=\{(1,2),(2,3),(3,2),(4,4),(5,4),(6,7),(7,6),(9,7)\}$
- Convenient to arrange data in array

$$
X=\left[\begin{array}{cc}
1 & 2 \\
\vdots & \vdots \\
9 & 7
\end{array}\right]=\left[\begin{array}{c}
x_{1} \\
\vdots \\
x_{8}
\end{array}\right]
$$

- Mean $\mu=\boldsymbol{m e a n}(X)=\left[\begin{array}{ll}4.6 & 4.4\end{array}\right]$

- Subtract mean from data to get new data array Z

$$
Z=X-\left[\begin{array}{c}
\mu \\
\vdots \\
\mu
\end{array}\right]=X-\operatorname{repmat}(\mu, 8,1)=\left[\begin{array}{cc}
-3.6 & -4.4 \\
\vdots & \vdots \\
4.4 & 2.6
\end{array}\right]
$$

- Compute the scatter matrix $S$
$S=7 * \operatorname{cov}(Z)=[-3.6-4.4]\left[\begin{array}{l}-3.6 \\ -4.4\end{array}\right]+\ldots+\left[\begin{array}{ll}4.4 & 2.6\end{array}\right]\left[\begin{array}{l}4.4 \\ 2.6\end{array}\right]=\left[\begin{array}{ll}57 & 40 \\ 40 & 34\end{array}\right]$
matlab uses unbiased estimate for covariance, so $S=(n-1)^{*} \operatorname{cov}(Z)$


## PCA Example Using Matlab

- Use $[\mathrm{V}, \mathrm{D}]=$ eig(S) to get eigenvalues and eigenvectors of $S$
$\lambda_{1}=87$ and $e_{1}=\left[\begin{array}{l}-0.8 \\ -0.6\end{array}\right]$
$\lambda_{2}=3.8$ and $e_{2}=\left[\begin{array}{l}0.6 \\ -0.8\end{array}\right]$

- Projection to 1D space in the direction of $\boldsymbol{e}_{1}$

$$
\begin{aligned}
Y=e_{1}^{t} Z^{t}=\left([-0.8-0.6]\left[\begin{array}{llll}
-3.6 & \cdots & 4.4 \\
-4.4 & \cdots & 2.6
\end{array}\right]\right) & =\left[\begin{array}{lll}
4.3 & \cdots & -5.1
\end{array}\right] \\
& =\left[\begin{array}{lll}
y_{1} & \cdots & y_{8}
\end{array}\right]
\end{aligned}
$$

## Drawbacks of PCA

- PCA was designed for accurate data representation, not for data classification
- Preserves as much variance in data as possible
- If directions of maximum variance is important for classification, will work
- However the directions of maximum variance may be useless for classification


- Next Lecture: Fisher Linear Discriminant
- preserve direction useful for discrimination

