

## Silly Example

Let events be "grades in a class
$\mathrm{w}_{1}=$ Gets an $A$
$\mathrm{w}_{2}=$ Gets a B
$w_{3}=$ Gets a $\underline{C}$
$\mathrm{w}_{4}=$ Gets a D


Assume we want to estipratefrem data. In a given class there were


What's the maximum likelinood estimate of $\mu$ given $a, b, c, d$ ?

$$
\hat{\mu}_{\text {MLE }}
$$

## Trivial Statistics

$$
P(A)=1 / 2 \quad P(B)=\mu \quad P(C)=2 \mu \quad P(D)=1 / 2-3 \mu
$$

$$
P(a, b, c, d \mid \mu)=K\left(\frac{1}{2}\right)^{a}(\mu)^{b}(2 \mu)^{c}(1 / 2-3 \mu)^{d}
$$

$$
\log P(a, b, c, d \mid \mu)=\log K^{2}+a l \log 1 / 2+b \log \mu+c \log 2 \mu+d \log (1 / 2-3 \mu)
$$

$$
\text { FOR MAX LIKE } \mu \text {, SET } \frac{\partial \operatorname{LogP}}{\partial \mu}=0
$$

$$
\frac{\partial \log P}{\partial \mu}=\frac{b}{\mu}+\frac{2 c}{2 \mu}-\frac{3 d}{1 / 2-3 \mu}=0
$$

Gives max like $\mu=\frac{b+c}{6(b+c+d)}$
So if class got
Max lik $\mu=\frac{1}{10}$

| A | B | C | D |
| :---: | :---: | :---: | :---: |
| 14 | 6 | 9 | 10 |

## Same Problem with Hidden Information

Someone tells us that
Number of High grades (A's + B's) $=\underline{h}$
Number of C's $=c$
Number of D's $=d$
What is the max. like estimate of $\mu$ now?

$$
\begin{aligned}
& \text { REMEMBER } \\
& P(A)=1 / 2 \\
& P(B)=\mu \\
& P(C)=2 \mu \\
& P(D)=1 / 2-3 \mu \\
& \hline
\end{aligned}
$$

We can answer this question circularly:

## EXPECTATION

If we know the value of mpe could compute the expected value ffa and $b$

$$
\bar{a}=\frac{1 / 2}{1 / 2+\mu} h \quad \bar{b}=\frac{\mu}{1 / 2+\mu} h
$$

## MAXIMIZATION

If we know the expected values of $\bar{a}$ and $\bar{b}$ we could compute the maximum likelihood value of $\mu$

$$
\left\lvert\, \hat{\mu}=\frac{\vec{b}+c}{6(\vec{b}+c+d)}\right.
$$

## E.M. for our Trivial Problem

REMEMBER
$P(A)=1 / 2$
$P(B)=\mu$
$P(C)=2 \mu$
$P(D)=1 / 2-3 \mu$

We begin with a guess for $\mu$
We iterate between EXPECTATION and MAXIMALIZATION to improve our estimates of $\mu$ and $a$ and $b$.

Define $\mu^{(t)}$ the estimate of $\mu$ on the t'th iteration
$\underline{b}^{(t)}$ the estimate of $b$ on lath $^{\prime}$ iteration
 $\mu^{(0)}=$ initial guess

$$
\begin{aligned}
& \underline{\underline{\mu}}^{(0)}=\text { initial guess } \\
& b^{(t)}=\frac{\mu^{(t)} b}{1 / 2+\mu^{(t)}}=\mathrm{E}\left[b \mid \mu^{(t)}\right]
\end{aligned}
$$



max like est. of $\mu \mathrm{g}$ ven $b^{(t)}$

## Continue iterating until converged.

## Good news: Converging to local optimum is assured.

Bad news: I said "local" optimußbe-2007 Carlos Guestin

## E.M. Convergence

- Convergence proof based on fact that $\operatorname{Prob}($ data $\mid \mu$ ) must increase or remain same between each iteration [not obvious]
- But it can never exceed 1 [obvious]

So it must therefore converge [obvious]

In our example,
suppose we had
$h=20$
$c=10$
d $=10$
$\mu^{(0)}=\underline{0}$
Convergence is generally linear: error decreases by a constant factor each time step. $\qquad$

$\left(\right.$| t | $\mu^{(\mathrm{t})}$ |  |
| :---: | :--- | :--- |
| 0 | 0 |  |
| 1 | $\underline{0.0833}$ | 1 |
| 2 | $\underline{0.0937}$ | 1 |
| 3 | $\underline{0.0947}$ | 1 |
| 4 | 0.0948 | $\prime$ |
| 5 | 0.0948 | $\vdots$ |
| 6 | 0.0948 | $i$ |


| $\frac{2.857}{3.158}$ |
| :--- |
| 3.185 |
| 3.187 |
| 3.187 |
| 3.187 |



EM for simple case of GMMs: The
E-step expectide value of hidenvars

- If we know $\mu_{1}, \ldots, \mu_{\mathrm{k}} \rightarrow$ easily compute prob.
point $x_{j}$ belongs to class $y=i$
Bayes Rube

$$
\underbrace{\mathrm{p}\left(y=\mid x_{j}, \mu_{1} \ldots \mu_{k}\right) x \exp \left(-\frac{1}{2 \sigma^{\sigma}} \| x_{j}-\left.\mu_{i}\right|^{2}\right) \mathrm{P}(y=i)})
$$

for each point $j$

$$
\begin{aligned}
& P\left(y=1 \mid x_{j}, \mu_{1} \ldots \mu_{k}\right)=0.7 \\
& P\left(y=2 \mid x_{j}, \mu_{1} \ldots \mu_{k}\right)=0.2 \\
& P\left(y=3 \mid x_{j}, \mu_{1} \ldots \mu_{k}\right)=0.1
\end{aligned}
$$

## EM for simple case of GMMs: The M-step

- If we know prob. point $x_{j}$ belongs to class $y=i$
$\rightarrow$ MLE for $\mu_{\mathrm{i}}$ is weighted average
$\square$ imagine $k$ copies of each $x_{j}$, each with weight $P\left(y=i \mid x_{j}\right)$ :



## E.M. for GMMs

E-step
each
Compute "expected" classes of datapoint for each class

$$
\mathrm{p}\left(y=i \mid x_{j}, \mu_{1} \ldots \mu_{k}\right) \propto \exp \left(-\frac{1}{2 \sigma^{2}}\left\|x_{j}-\mu_{i}\right\|^{2}\right) \mathrm{P}(y=i)
$$

Just evaluate a Gaussian at $x_{j}$

M-step
Compute Max. like $\boldsymbol{\mu}$ given our data's class membership distributions
$\mu_{i}=\frac{\sum_{j=1}^{m} P\left(y=i \mid x_{j}\right) x_{j}}{\sum_{j=1}^{m} P\left(y=i \mid x_{j}\right)}$


- This algorithm is REALLY USED. And in high dimensional state spaces, too. E.G. Vector Quantization for Speech Data


## E.M. for axis-aligned GMN

Iterate. On the $t$ 'th iteration let our estimates be


Compute Max. like $\mu$ given our data's class membership distributions $-\frac{m-i}{m}$


## E.M. for General GMMs

Iterate. On the $t$ 'th iteration let our estimates be

$$
\lambda_{t}=\left\{\mu_{1}^{(t)}, \mu_{2}^{(t)} \ldots \mu_{k}^{(t)}, \sum_{1}^{(t)}, \sum_{2}^{(t)} \ldots \Sigma_{k}^{(t)}, p_{1}^{(t)}, p_{2}{ }^{(t)} \ldots p_{k}^{(t)}\right\}
$$

E-step
Compute "expected" classes of all datapoints for each class


Compute Max. like $\mu$ given our data's class membership distributions

$$
\mu_{i}^{(t+1)}=\frac{\sum_{j} \mathrm{P}\left(y=i \mid x_{j}, \lambda_{t}\right) x_{j}}{\sum_{j} \mathrm{P}\left(y=i \mid x_{j}, \lambda_{t}\right)} \quad \Sigma_{i}^{(t+1)}=\frac{\left.\sum_{j} \mathrm{P}\left(y=i \mid x_{j}, \lambda_{t}\right)\left[x_{j}-\mu_{i}^{(t+1)}\right] x_{j}-\mu_{i}^{(t+1)}\right]^{\mathrm{T}}}{\sum_{j} \mathrm{P}\left(y=i \mid x_{j}, \lambda_{t}\right)}
$$

$$
p_{i}^{(t+1)}=\frac{\sum_{j} \mathrm{P}\left(y=i \mid x_{j}, \lambda_{t}\right)}{m}=m=\text { \#records }
$$

## Gaussian Mixture Example: Start



## After first iteration

## $\square-\square \square$



## After 2nd iteration



## After 3rd iteration

 ■-

## After 4th iteration




## After 6th iteration




## Some Bio Assay data



## GMM clustering of the assay data






## The general learning problem with missing data

- Marginal likelihood $-\mathbf{x}$ is observed, $\mathbf{z}$ is missing:

$$
\begin{aligned}
& \text { wait locex } \\
& \text { to max } \stackrel{i i d}{=} \log \prod_{j=1}^{m} P\left(\mathbf{x}_{j} \mid \theta\right) \\
&=\frac{\sum_{j=1}^{m} \log P\left(\mathbf{x}_{j} \mid \theta\right)}{} \\
&=\sum_{j=1}^{m} \log \sum_{\mathbf{z}} P\left(\mathbf{x}_{j}, \mathbf{z} \mid \theta\right)
\end{aligned}
$$



## E-step

- $\mathbf{x}$ is observed, $\mathbf{z}$ is missing
- Compute probability of missing data given current choice of $\theta$
$\square \underline{Q}\left(\underline{z} \mid \mathbf{x}_{\mathbf{j}}\right)$ for each $\mathbf{x}_{\mathrm{j}}$
- e.g., probability computed during classification step
- corresponds to "classification step" in K-means
$\underline{Q^{(t+1)}\left(\mathbf{z} \mid \mathbf{x}_{j}\right)=P\left(\mathbf{z} \mid \mathbf{x}_{j}, \theta^{(t)}\right)}$


## Jensen's inequality

$$
\underline{\ell(\theta: \mathcal{D}})=\sum_{j=1}^{m} \log \sum_{\underline{z}} P\left(\mathbf{z} \mid \mathbf{x}_{j}\right) P\left(\mathbf{x}_{j} \mid \theta\right)
$$

- Theorem: $\log \sum_{z} P(z) f(z) \geq \sum_{z} P(z) \log f(z)$



## Applying Jensen's inequality $\log \frac{4}{b}$ $=\log a-\log b$

- Use: $\log \sum_{\mathbf{z}} P(\mathbf{z}) f(\mathbf{z}) \geq \sum_{\mathbf{z}} P(\mathbf{z}) \log f(\mathbf{z})$ $\ell\left(\theta^{(t)}: \mathcal{D}\right)=\sum_{j=1}^{m} \log \sum_{\mathbf{z}} Q^{(t+1)}\left(\mathbf{z} \mid \mathbf{x}_{j}\right) \frac{P\left(\mathbf{z}, \mathbf{x}_{j} \mid \theta^{(t)}\right)}{Q^{(t+1)}\left(\mathbf{z} \mid \mathbf{x}_{j}\right)}$
$\geq \sum_{j=1}^{m} \sum_{z} Q^{(f+1)}\left(z \mid x_{j}\right) \log \frac{P\left(z, x_{j}| |^{(t)}\right)}{Q^{(++1)}\left(z \mid x_{j}\right)}$
$=\sum_{j=1}^{m} \sum Q^{(t+1)}\left(z-x_{j}\right) \log P\left(z, x_{j} \mid \theta^{(t)}\right)$
$-\sum_{j=1}^{m} \underbrace{\sum_{z} Q^{(t+1)}\left(z \mid x_{j}\right) \log Q^{(t+1)}\left(z \mid x_{j}\right)}_{-H\left(Q^{(t+1)} \mid x_{j}\right)}$


## The M-step maximizes lower bound on weighted data <br> - Lower bound from Jensen's: <br> wat it to max $\ell\left(\theta^{(t)}: \mathcal{D}\right)$$\sum_{\text {I don't Know }}^{m} \sum_{\text {, }}^{\sum_{\mathbf{z}} Q^{(t+1)}\left(\mathbf{z} \mid \mathbf{x}_{j}\right)} \underbrace{\log P\left(\mathbf{z}, \mathbf{x}_{j}\right.}$ but I introduce $|z|$ data points with eight $Q^{(t+1)}\left(z \mid x_{j}\right)$

- Corresponds to weighted dataset:

```
<x},\mathbf{z=1>}\mathrm{ with weight Q }\mp@subsup{Q}{}{(t+1)}(z=1|\mp@subsup{x}{1}{}
<x},\mp@code{z=2> with weight Q Q }\mp@subsup{}{(t+1)}{(z=2|\mp@subsup{x}{1}{})
<x},\mathbf{z=3>}\mathrm{ with weight Q Q (t+1)}(\mathbf{z}=3|\mp@subsup{\mathbf{x}}{1}{}
<x
<x},\mathbf{z}=2> with weight Q Q (t+1)(z=2|x ( )
<x},\mp@code{z=3> with weight Q }\mp@subsup{Q}{}{(t+1)}(z=3|\mp@subsup{x}{2}{}
```

$\square$.


Maximization step:

$$
\stackrel{\theta^{(t+1)}}{=} \leftarrow \underline{\theta} \arg \max _{\theta} \sum_{j=1}^{m} \sum_{\mathbf{z}} Q^{(t+1)}\left(\mathbf{z} \mid \mathbf{x}_{j}\right) \log P\left(\mathbf{z}, \mathbf{x}_{j} \mid \theta\right)
$$

IMLE w. weighted cate

- Use expected counts instead of counts:

If learning requires Count ( $\mathbf{x}, \mathbf{z}$ ) Use $\mathrm{E}_{\mathrm{Q}(\mathrm{t}+1)}[$ Count $(\mathbf{x}, \mathbf{z})]$

## Convergence of EM

- Define potential function $F(\theta, Q)$ :
- EM corresponds to coordinate ascent on F

Thus, maximizes lower bound on marginal log likelihood

## M-step is easy

$$
\theta^{(t+1)} \leftarrow \arg \max _{\theta} \sum_{j=1}^{m} \sum_{\mathbf{z}} Q^{(t+1)}\left(\mathbf{z} \mid \mathbf{x}_{j}\right) \log P\left(\mathbf{z}, \mathbf{x}_{j} \mid \theta\right)
$$

- Using potential function constant


## 

 potential function $\left.1 \quad Q^{(t+1)}(z \mid x)<P(z \mid x), \theta^{(t)}\right)$Fixing $\theta$ to $\theta^{(t)}$ :

$$
\begin{aligned}
& \ell\left(\theta^{(t)}: \mathcal{D}\right) \geq F\left(\theta^{(t)}, Q\right)=\sum_{j=1}^{m} \sum_{\mathbf{z}} Q\left(\mathbf{z} \mid \mathbf{x}_{j}\right) \log \frac{P\left(\mathbf{z}, \mathbf{x}_{j} \mid \theta^{(t)}\right)}{Q\left(\mathbf{z} \mid \mathbf{x}_{j}\right)} \\
& \stackrel{\text { chain vance }}{=} \sum_{j=1}^{n} \sum_{z} Q\left(z \mid x_{j}\right) \log \frac{P\left(z \mid x_{j} \theta^{(t)}\right) \cdot P\left(x_{j} \mid \theta^{(t)}\right)}{Q\left(z \mid x_{j}\right)=1}
\end{aligned}
$$

## KL-divergence

- Measures distance between distributions
$\underline{\underline{K L(Q \| P}})=\sum_{z} Q(z) \log \frac{Q(z)}{P(z)}$
- KL=zero if and only if $Q=P$


## E-step also doesn't decrease potential function 2

## -KL

Fixing $\theta$ to $\theta^{(t)}$ :

$$
\begin{aligned}
& \ell\left(\theta^{(t)}: \mathcal{D}\right) \geq F\left(\theta^{(t)}, Q\right)=\ell\left(\theta^{(t)}: \mathcal{D}\right)+\sum_{j=1}^{m} \sum_{\mathbf{z}} Q\left(\mathbf{z} \mid \mathbf{x}_{j}\right) \log \frac{P\left(\mathbf{z} \mid \mathbf{x}_{j}, \theta^{(t)}\right)}{Q\left(\mathbf{z} \mid \mathbf{x}_{j}\right)} \\
& \begin{array}{l}
=\ell\left(\theta^{(t)}: \mathcal{D}\right)-\underbrace{\sum_{j=1}^{m} K L\left(Q\left(\mathbf{z} \mid \mathbf{x}_{j}\right) \| P\left(\mathbf{z} \mid \mathbf{x}_{j}, \theta^{(t)}\right)\right)}_{\text {as small as possible }} \\
\text { the }
\end{array} \\
& \text { maximitl the } \\
& \text { right side, } \\
& \text { we know that } \\
& \rightarrow \text { by setting } \\
& K((Q \| P) \geqslant 0 \\
& Q\left(z \mid x_{j}\right)=P\left(z \mid x_{i j}, \theta^{(t)}\right) \\
& =0 \Leftrightarrow P=Q
\end{aligned}
$$

## E-step also doesn't decrease potential function 3

$\left.\begin{array}{l}\ell\left(\theta^{(t)}: \mathcal{D}\right) \geq F\left(\theta^{(t)}, Q\right)=\ell\left(\theta^{(t)}: \mathcal{D}\right) \text { - 缶 } \sum_{j=1}^{m} K L\left(Q\left(\mathbf{z} \mid \mathbf{x}_{j}\right)\right. \\ \text { - Fixing } \theta \text { to } \theta^{(t)} \\ \text { - Maximizing } \mathrm{F}\left(\theta^{(t)}, \mathrm{Q}\right)\end{array}\right)$ over $\mathrm{Q} \rightarrow$ set Q to posterior probability:

- Note that

$$
Q^{(t+1)}\left(\mathbf{z} \mid \mathbf{x}_{j}\right) \leftarrow P\left(\mathbf{z} \mid \mathbf{x}_{j}, \theta^{(t)}\right)
$$

## EM is coordinate ascent

$\ell(\theta: \mathcal{D}) \geq F(\theta, Q)=\sum_{j=1}^{m} \sum_{\mathbf{z}} Q\left(\mathbf{z} \mid \mathbf{x}_{j}\right) \log \frac{P\left(\mathbf{z}, \mathbf{x}_{j} \mid \theta\right)}{Q\left(\mathbf{z} \mid \mathbf{x}_{j}\right)}$

- M-step: Fix Q , maximize F over $\theta$ (a lower bound on $\ell(\theta: \mathcal{D})$ ):
$\ell(\theta: \mathcal{D}) \geq F\left(\theta, Q^{(t)}\right)=\sum_{j=1}^{m} \sum_{\mathbf{z}} Q^{(t)}\left(\mathbf{z} \mid \mathbf{x}_{j}\right) \log P\left(\mathbf{z}, \mathbf{x}_{j} \mid \theta\right)+m \cdot H\left(Q^{(t)}\right)$
- E-step: Fix $\theta$, maximize F over Q:
$\ell\left(\theta^{(t)}: \mathcal{D}\right) \geq F\left(\theta^{(t)}, Q\right)=\ell\left(\theta^{(t)}: \mathcal{D}\right)-\psi \sum_{j=1}^{m} K L\left(Q\left(\mathbf{z} \mid \mathbf{x}_{j}\right) \| P\left(\mathbf{z} \mid \mathbf{x}_{j}, \theta^{(t)}\right)\right)$
"Realigns" F with likelihood:

$$
F\left(\theta^{(t)}, Q^{(t+1)}\right)=\ell\left(\theta^{(t)}: \mathcal{D}\right)
$$

## What you should know

K-means for clustering:
$\square$ algorithm
$\square$ converges because it's coordinate ascent

- EM for mixture of Gaussians:

How to "learn" maximum likelihood parameters (locally max. like.) in the case of unlabeled data

- Be happy with this kind of probabilistic analysis

Remember, E.M. can get stuck in local minima, and empirically it DOES

- EM is coordinate ascent
- General case for EM


## Acknowledgements

- K-means \& Gaussian mixture models presentation contains material from excellent tutorial by Andrew Moore:
$\square$ http://www.autonlab.org/tutorials/
- K-means Applet:
$\square$ http://www.elet.polimi.it/upload/matteucc/Clustering/tu torial html/AppletKM.html
- Gaussian mixture models Applet:
$\square$ http://www.neurosci.aist.go.jp/\~akaho/MixtureEM. html


## Dimensionality reduction

- Input data may have thousands or millions of dimensions!
e.g., text data has
- Dimensionality reduction: represent data with fewer dimensions
easier learning - fewer parametersvisualization - hard to visualize more than 3D or 4D
discover "intrinsic dimensionality" of data
- high dimensional data that is truly lower dimensional


## Feature selection

- Want to learn $\mathrm{f}: \mathrm{X} \rightarrow \mathrm{Y}$
$\square X=<X_{1}, \ldots, X_{n} \geqslant 40 \mathrm{~K}$
$\square$ but some features are more important than others
- Approach: select subset of features to be used by learning algorithm
$\square$ Score each feature (or sets of features)
Select set of features with best score


## Simple greedy forward feature selection algorithm

- Pick a dictionary of features
e.g., polynomials for linear regression

Greedy heuristic:
$\square$ Start from empty (or simple) set of features $F_{0}=\varnothing$
Run learning algorithm for current set of features $F_{t}$

- Obtain $h_{\mathrm{t}}$

Select next best feature $X_{i}$

- e.g., $X_{j}$ that results in lowest crossvalidation error learner when learning with $F_{t} \cup\left\{\mathrm{X}_{\mathrm{j}}\right\}$
$\square F_{t+1} \leftarrow F_{t} \cup\left\{X_{i}\right\}$Recurse


## Simple greedy backward feature selection algorithm

Pick a dictionary of features
$\square$ e.g., polynomials for linear regression

- Greedy heuristic:
$\square$ Start from all features $F_{0}=F$
$\square$ Run learning algorithm for current set
of features $F_{t}$
- Obtain $h_{t}$
$\square$ Select next worst feature $X_{i}$
- e.g., $\overline{X_{j}}$ that results in lowest crossvalidation error learner when learning with $F_{t}-\left\{\mathrm{X}_{\mathrm{j}}\right\}$
$\square F_{t+1} \leftarrow F_{t}-\left\{X_{i}\right\}$
$\square$ Recurse
Impact of feature selection on classification of fMRI data [Pereria etal: 05]

| Accuracy classifying category of word read by subject |  |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| \#voxels | $\underset{\text { meant }}{\downarrow}$ | subjects $233 B$ | 329B | 332B | 424B | 474B | 496B | 77B | 86B |
| 50 | 0.735 | 0.783 | 0.817 | 0.55 | 0.783 | 0.75 | 0.8 | 0.65 | 0.75 |
| 100 | 0.742 | 0.767 | 0.8 | 0.533 | 0.817 | 0.85 | 0.783 | 0.6 | 0.783 |
| 200 | 0.737 | 0.783 | 0.783 | 0.517 | 0.817 | 0.883 | 0.75 | 0.583 | 0.783 |
| 300 | 0.75 | 0.8 | 0.817 | 0.567 | 0.833 | 0.888 | 0.75 | 0.583 | 0.767 |
| 400 | 0.742 | 0.8 | 0.783 | 0.583 | 0.85 | 0.833 | 0.75 | 0.583 | 0.75 |
| 800 | 0.735 | 0.833 | 0.817 | 0.567 | 0.833 | 0.833 | 0.7 | 0.55 | 0.75 |
| 1600 | 0.698 | 0.8 | 0.817 | 0.45 | 0.783 | 0.833 | 0.633 | 0.5 | 0.75 |
| all ( $\sim 2500$ ) | -0.638 | 0.767 | 0.767 | 0.25 | 0.75 | 0.833 | 0.567 | 0.433 | 0.733 |

Table 1: Average accuracy across all pairs of categories, restricting the procedure to use a certain number of voxels for each subject. The highlighted liue correxponds to the best mean accuracy, olstained using 300 vorels.

Voxels scored by p-value of regression to predict voxel value from the task

## Lower dimensional projections

Rather than picking a subset of the features, we can new features that are combinations of existing features

$$
\begin{aligned}
& \text { e.9., feature } \\
& \text { selection: } \\
& \text { use } x_{1}, x_{7}, x_{11}
\end{aligned}
$$

$$
\begin{aligned}
& \text { low. dim. proj. } \\
& \begin{aligned}
\tilde{x}= & =0.1 x_{1}+0.7 x_{2} \\
& -0.35 x_{3} \ldots
\end{aligned}
\end{aligned}
$$

- Let's see this in the unsupervised setting just X, but no Y



## Principal component analysis basic idea

- Project n -dimensional data into k -dimensional space while preserving information:
$\square$ e.g., project space of 10000 words into 3-dimensionse.g., project 3-d into 2-d

Choose projection with minimum reconstruction error

- Project a point into a (lower dimensional) space:
point: $\underline{x}=\left(x_{1}, \ldots, x_{n}\right)$
select a basis - set of basis vectors - ( $\left.\mathbf{u}_{1}, \ldots, \mathbf{u}_{k}\right)$
- we consider orthonormal basis:

$$
u_{i} \cdot u_{i}=1 \text { and } u_{i} u_{j}=0 \text { for } i \neq j
$$

select a center - $\overline{\mathbf{x}}$, defines offset of space
best coordinates in lower dimensional space defined by dot-products: $\left(z_{1}, \ldots, z_{k}\right), z_{i}=(x-\bar{x}) \bullet u_{i}$

- minimum squared error

©2005-2007 Carlos Guestrin

