Note to other teachers and users of these slides. Andrew would be delighted if you found this source material useful in giving your own lectures. Feel free to use these slides verbatim, or to modify them to fit your own needs. PowerPoint originals are available. If you make use of a significant portion of these slides in your own lecture, please include this message, or the following link to the source repository of Andrew's tutorials: http://www.cs.cmu.edu/~awm/tutorials . Comments and corrections gratefully received.

Clustering with Gaussian Mixtures

Andrew W. Moore
Associate Professor
School of Computer Science
Carnegie Mellon University
www.cs.cmu.edu/~awm
awm@cs.cmu.edu
412-268-7599

 Nov 10th, 2001

Unsupervised Learning

- You walk into a bar.
  A stranger approaches and tells you:
  "I've got data from k classes. Each class produces observations with a normal distribution and variance σ^2I . Standard simple multivariate gaussian assumptions. I can tell you all the P(w_i)'s ."
- So far, looks straightforward.
  "I need a maximum likelihood estimate of the μ_i's ."
- No problem:
  "There's just one thing. None of the data are labeled. I have datapoints, but I don't know what class they're from (any of them!)
- Uh oh!!

Some data from a GMM

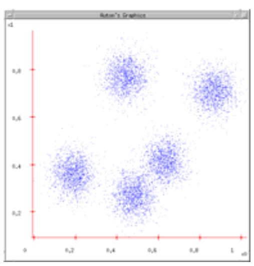

The GMM assumption

- There are k components. The i'th component is called ω_i
- Component ω_i has an associated mean vector μ_i

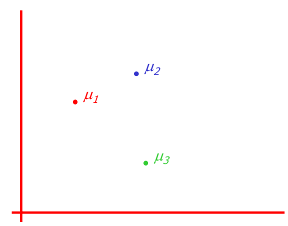

The GMM assumption

- There are k components. The i'th component is called ω_i
- Component ω_i has an associated mean vector μ_i
- Each component generates data from a Gaussian with mean μ_i and covariance matrix $\sigma^2 \mathbf{I}$

Assume that each datapoint is generated according to the following recipe:

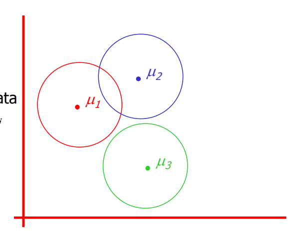

The GMM assumption

- There are k components. The i'th component is called ω_i
- Component ω_i has an associated mean vector μ_i
- Each component generates data from a Gaussian with mean μ_i and covariance matrix $\sigma^2 \mathbf{I}$

Assume that each datapoint is generated according to the following recipe:

1. Pick a component at random. Choose component i with probability $P(\omega_i)$.

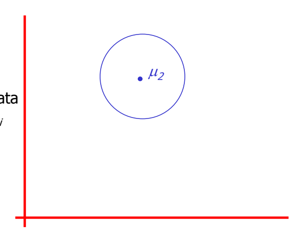

The GMM assumption

- There are k components. The i'th component is called ω_i
- Component ω_i has an associated mean vector μ_i
- Each component generates data from a Gaussian with mean μ_i and covariance matrix $\sigma^2 \mathbf{I}$

Assume that each datapoint is generated according to the following recipe:

1. Pick a component at random. Choose component i with probability $P(\omega_i)$.
2. Datapoint ~ $N(\mu_i, \sigma^2 \mathbf{I})$

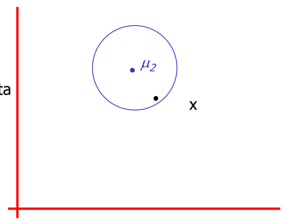

The General GMM assumption

- There are k components. The i'th component is called ω_i
- Component ω_i has an associated mean vector μ_i
- Each component generates data from a Gaussian with mean μ_i and covariance matrix Σ_i

Assume that each datapoint is generated according to the following recipe:

1. Pick a component at random. Choose component i with probability $P(\omega_i)$.
2. Datapoint ~ $N(\mu_i, \Sigma_i)$

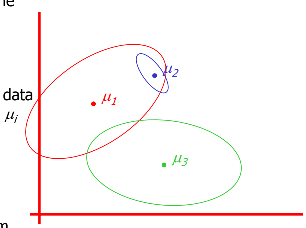

Unsupervised Learning: not as hard as it looks

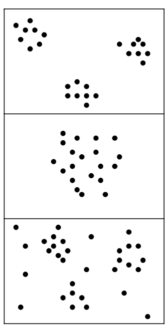

Sometimes easy

Sometimes impossible

and sometimes in between

*IN CASE YOU'RE WONDERING WHAT THESE DIAGRAMS ARE, THEY SHOW 2-d UNLABELED DATA (**X** VECTORS) DISTRIBUTED IN 2-d SPACE. THE TOP ONE HAS THREE VERY CLEAR GAUSSIAN CENTERS*

Computing likelihoods in unsupervised case

We have $\mathbf{x}_1, \mathbf{x}_2, \ldots \mathbf{x}_N$

We know $P(w_1)$ $P(w_2)$ .. $P(w_k)$

We know σ

$P(\mathbf{x}|w_i, \boldsymbol{\mu}_i, \ldots \boldsymbol{\mu}_k)$ = Prob that an observation from class w_i would have value $\mathbf{x}$ given class means $\boldsymbol{\mu}_1 \ldots \boldsymbol{\mu}_x$

Can we write an expression for that?

likelihoods in unsupervised case

We have $\mathbf{x}_1\ \mathbf{x}_2 \ldots \mathbf{x}_n$

We have $P(w_1)$ .. $P(w_k)$. We have σ.

We can define, for any $\mathbf{x}$, $P(\mathbf{x}|w_i, \boldsymbol{\mu}_1, \boldsymbol{\mu}_2 .. \boldsymbol{\mu}_k)$

Can we define $P(\mathbf{x} \mid \boldsymbol{\mu}_1, \boldsymbol{\mu}_2 .. \boldsymbol{\mu}_k)$?

Can we define $P(\mathbf{x}_1, \mathbf{x}_2, .. \mathbf{x}_n \mid \boldsymbol{\mu}_1, \boldsymbol{\mu}_2 .. \boldsymbol{\mu}_k)$?

[YES, IF WE ASSUME THE X_i'S WERE DRAWN INDEPENDENTLY]

Unsupervised Learning: Mediumly Good News

We now have a procedure s.t. if you give me a guess at $\boldsymbol{\mu}_1, \boldsymbol{\mu}_2 .. \boldsymbol{\mu}_k$,

I can tell you the prob of the unlabeled data given those $\boldsymbol{\mu}$'s.

Suppose x's are 1-dimensional.

(From Duda and Hart)

There are two classes; w_1 and w_2

$P(w_1) = 1/3$ $P(w_2) = 2/3$ $\sigma = 1$.

There are 25 unlabeled datapoints

$x_1 = 0.608$
$x_2 = -1.590$
$x_3 = 0.235$
$x_4 = 3.949$
:
$x_{25} = -0.712$

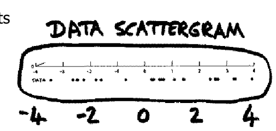

Duda & Hart's Example

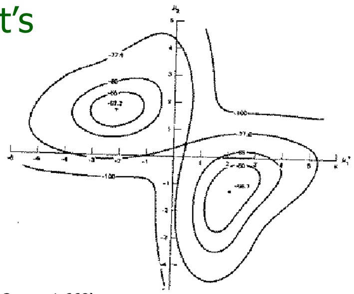

Graph of
log P(x_1, x_2 .. x_{25} | μ_1, μ_2)
against μ_1(→) and μ_2(↑)

Max likelihood = (μ_1 =-2.13, μ_2 =1.668)

Local minimum, but very close to global at (μ_1 =2.085, μ_2 =-1.257)*

* corresponds to switching w_1 + w_2.

Duda & Hart's Example

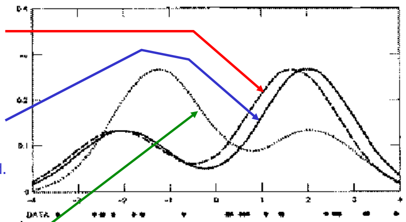

- They are close. Good.
- The 2nd solution tries to put the "2/3" hump where the "1/3" hump should go, and vice versa.
- In this example unsupervised is almost as good as supervised. If the x_1 .. x_{25} are given the class which was used to learn them, then the results are (μ_1=-2.176, μ_2=1.684). Unsupervised got (μ_1=-2.13, μ_2=1.668).

Finding the max likelihood $\mu_1,\mu_2..\mu_k$

We can compute P(data | $\mu_1,\mu_2..\mu_k$)
How do we find the μ_i's which give max. likelihood?

- The normal max likelihood trick:
  Set $\frac{\partial}{\partial \mu_i}$ log Prob (....) = 0
  and solve for μ_i's.
  # Here you get non-linear non-analytically-solvable equations
- Use gradient descent
  Slow but doable
- Use a much faster, cuter, and recently very popular method…

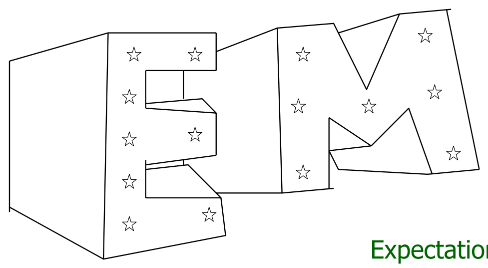

Expectation
Maximalization

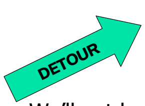

The E.M. Algorithm

- We'll get back to unsupervised learning soon.
- But now we'll look at an even simpler case with hidden information.
- The EM algorithm
  - Can do trivial things, such as the contents of the next few slides.
  - An excellent way of doing our unsupervised learning problem, as we'll see.
  - Many, many other uses, including inference of Hidden Markov Models (future lecture).

Silly Example

Let events be "grades in a class"

w_1 = Gets an A P(A) = ½
w_2 = Gets a B P(B) = μ
w_3 = Gets a C P(C) = 2μ
w_4 = Gets a D P(D) = ½-3μ
(Note 0 ≤ μ ≤1/6)

Assume we want to estimate μ from data. In a given class there were

a A's
b B's
c C's
d D's

What's the maximum likelihood estimate of μ given a,b,c,d ?

Computing

P(A) = ½ P(B) = μ P(C) = 2μ P(D) = ½-3μ

P(*a,b,c,d* | μ) = K(½)^a(μ)^b(2μ)^c(½-3μ)^d

log P(*a,b,c,d* | μ) = log K + *a*log ½ + *b*log μ + *c*log 2μ + *d*log (½-3μ)

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

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

$$\text{Gives max like } \mu = \frac{b+c}{6(b+c+d)}$$

So if class got

A	B	C	D
14	6	9	10

$$\text{Max like } \mu = \frac{1}{10}$$

 Clustering with Gaussian Mixtures: Slide 19

Same Problem with Hidden Information

Someone tells us that
Number of High grades (A's + B's) = *h*
Number of C's = *c*
Number of D's = *d*

What is the max. like estimate of μ now?

REMEMBER
P(A) = ½
P(B) = μ
P(C) = 2μ
P(D) = ½-3μ

 Clustering with Gaussian Mixtures: Slide 20

Same Problem with Hidden Information

Someone tells us that
Number of High grades (A's + B's) = *h*
Number of C's = *c*
Number of D's = *d*

REMEMBER
P(A) = ½
P(B) = μ
P(C) = 2μ
P(D) = ½-3μ

What is the max. like estimate of μ now?

We can answer this question circularly:

EXPECTATION

If we know the value of μ we could compute the expected value of *a* and *b*

Since the ratio a:b should be the same as the ratio ½ : μ

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

MAXIMIZATION

If we know the expected values of *a* and *b* we could compute the maximum likelihood value of μ

$$\mu = \frac{b+c}{6(b+c+d)}$$

 Clustering with Gaussian Mixtures: Slide 21

E.M. for our Trivial Problem

REMEMBER
P(A) = ½
P(B) = μ
P(C) = 2μ
P(D) = ½-3μ

We begin with a guess for μ
We iterate between EXPECTATION and MAXIMALIZATION to improve our estimates of μ and *a* and *b*.

Define μ(t) the estimate of μ on the t'th iteration
b(t) the estimate of *b* on t'th iteration

$$\mu(0) = \text{initial guess}$$

$$b(t) = \frac{\mu(t)h}{1/2 + \mu(t)} = \text{E}[b \mid \mu(t)]$$

E-step

$$\mu(t+1) = \frac{b(t)+c}{6(b(t)+c+d)}$$

$$= \text{max like est of } \mu \text{ given } b(t)$$

M-step

Continue iterating until converged.
Good news: Converging to local optimum is assured.
Bad news: I said "local" optimum.

 Clustering with Gaussian Mixtures: Slide 22

E.M. Convergence

- Convergence proof based on fact that Prob(data | μ) 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
μ(0) = 0

t	μ(t)	b(t)
0	0	0
1	0.0833	2.857
2	0.0937	3.158
3	0.0947	3.185
4	0.0948	3.187
5	0.0948	3.187
6	0.0948	3.187

Convergence is generally <u>linear</u>: error decreases by a constant factor each time step.

 Clustering with Gaussian Mixtures: Slide 23

Back to Unsupervised Learning of GMMs

Remember:
We have unlabeled data x_1 x_2 … x_R
We know there are k classes
We know P(w_1) P(w_2) P(w_3) … P(w_k)
We <u>don't</u> know $\boldsymbol{\mu}_1$ $\boldsymbol{\mu}_2$ .. $\boldsymbol{\mu}_k$

We can write P(data | $\boldsymbol{\mu}_1$…. $\boldsymbol{\mu}_k$)

$$= \text{p}(x_1 \ldots x_R \mid \mu_1 \ldots \mu_k)$$

$$= \prod_{i=1}^{R} \text{p}(x_i \mid \mu_1 \ldots \mu_k)$$

$$= \prod_{i=1}^{R} \sum_{j=1}^{k} \text{p}(x_i \mid w_j, \mu_1 \ldots \mu_k) \text{P}(w_j)$$

$$= \prod_{i=1}^{R} \sum_{j=1}^{k} \text{K} \exp\left(-\frac{1}{2\sigma^2}(x_i - \mu_j)^2\right) \text{P}(w_j)$$

 Clustering with Gaussian Mixtures: Slide 24

E.M. for GMMs

For Max likelihood we know $\frac{\partial}{\partial \mu_i} \log \Pr ob(\text{data}|\mu_1 ... \mu_k) = 0$

Some wild'n'crazy algebra turns this into : "For Max likelihood, for each j,

$$\mu_j = \frac{\sum_{i=1}^{R} P(w_j | x_i, \mu_1 ... \mu_k) x_i}{\sum_{i=1}^{R} P(w_j | x_i, \mu_1 ... \mu_k)}$$

This is n nonlinear equations in μ_j's."

If, for each x_i we knew that for each w_j the prob that μ_j was in class w_j is $P(w_j|x_i,\mu_1...\mu_k)$ Then… we would easily compute μ_j.

If we knew each μ_j then we could easily compute $P(w_j|x_i,\mu_1...\mu_j)$ for each w_j and x_i.

…I feel an EM experience coming on!!

E.M. for GMMs

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

$\lambda_t = \{ \mu_1(t), \mu_2(t) ... \mu_c(t) \}$

E-step

Compute "expected" classes of all datapoints for each class

Just evaluate a Gaussian at x_k

$$P(w_i|x_k,\lambda_t) = \frac{p(x_k|w_i,\lambda_t)P(w_i|\lambda_t)}{p(x_k|\lambda_t)} = \frac{p(x_k|w_i,\mu_i(t),\sigma^2\mathbf{I})p_i(t)}{\sum_{j=1}^{c} p(x_k|w_j,\mu_j(t),\sigma^2\mathbf{I})p_j(t)}$$

M-step.

Compute Max. like μ given our data's class membership distributions

$$\mu_i(t+1) = \frac{\sum_k P(w_i|x_k,\lambda_t)x_k}{\sum_k P(w_i|x_k,\lambda_t)}$$

E.M. Convergence

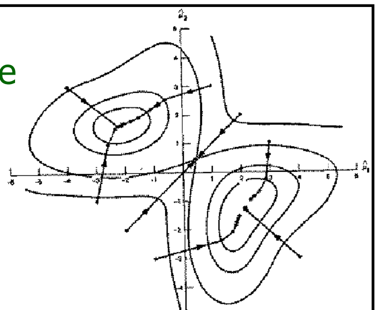

- As with all EM procedures, convergence to a local optimum guaranteed.
- This algorithm is REALLY USED. And in high dimensional state spaces, too. E.G. Vector Quantization for Speech Data

E.M. for General GMMs

$p_i(t)$ is shorthand for estimate of $P(\omega_i)$ on t'th iteration

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

$\lambda_t = \{ \mu_1(t), \mu_2(t) ... \mu_c(t), \Sigma_1(t), \Sigma_2(t) ... \Sigma_c(t), p_1(t), p_2(t) ... p_c(t) \}$

E-step

Compute "expected" classes of all datapoints for each class

Just evaluate a Gaussian at x_k

$$P(w_i|x_k,\lambda_t) = \frac{p(x_k|w_i,\lambda_t)P(w_i|\lambda_t)}{p(x_k|\lambda_t)} = \frac{p(x_k|w_i,\mu_i(t),\Sigma_i(t))p_i(t)}{\sum_{j=1}^{c} p(x_k|w_j,\mu_j(t),\Sigma_j(t))p_j(t)}$$

M-step.

Compute Max. like μ given our data's class membership distributions

$$\mu_i(t+1) = \frac{\sum_k P(w_i|x_k,\lambda_t)x_k}{\sum_k P(w_i|x_k,\lambda_t)} \qquad \Sigma_i(t+1) = \frac{\sum_k P(w_i|x_k,\lambda_t)[x_k - \mu_i(t+1)][x_k - \mu_i(t+1)]^T}{\sum_k P(w_i|x_k,\lambda_t)}$$

$$p_i(t+1) = \frac{\sum_k P(w_i|x_k,\lambda_t)}{R}$$

R = #records

Gaussian Mixture Example: Start

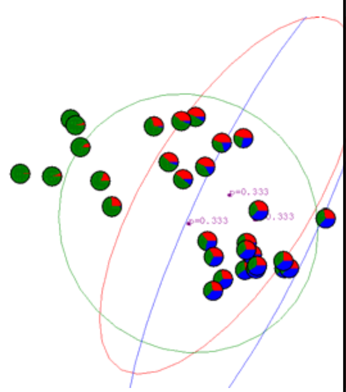

Advance apologies: in Black and White this example will be incomprehensible

After first iteration

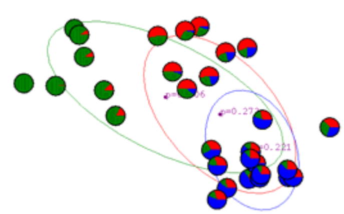

After 2nd iteration

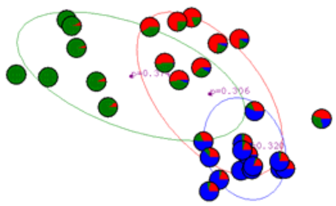

After 3rd iteration

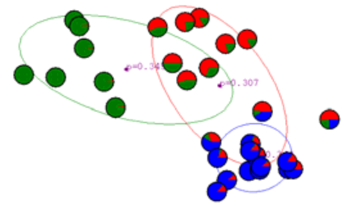

After 4th iteration

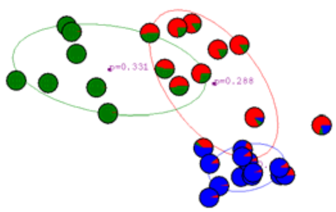

After 5th iteration

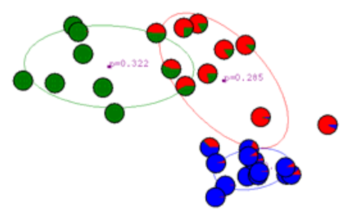

After 6th iteration

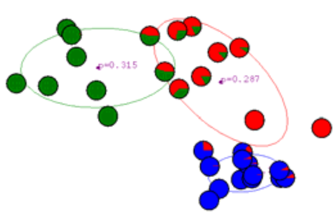

After 20th iteration

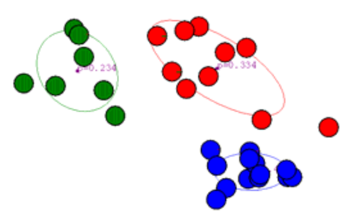

Some Bio Assay data

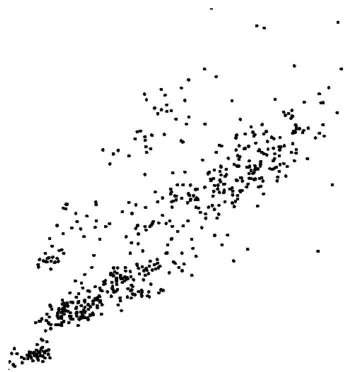

GMM clustering of the assay data

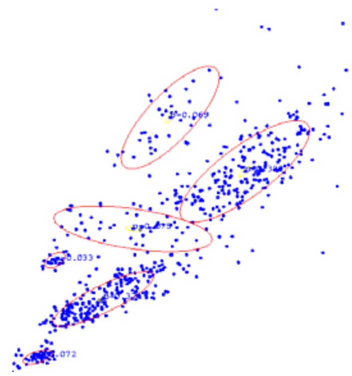

Resulting Density Estimator

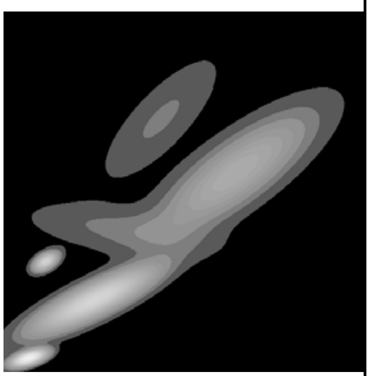

Three classes of assay

(each learned with it's own mixture model)

(Sorry, this will again be semi-useless in black and white)

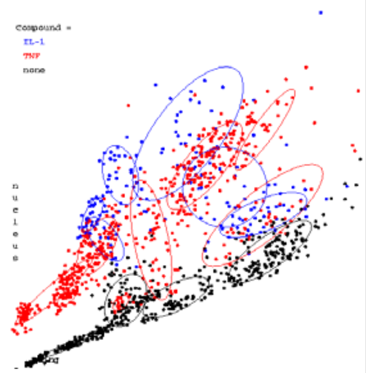

Resulting Bayes Classifier

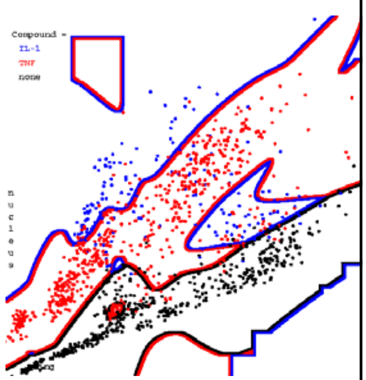

Resulting Bayes Classifier, using posterior probabilities to alert about ambiguity and anomalousness

Yellow means anomalous

Cyan means ambiguous

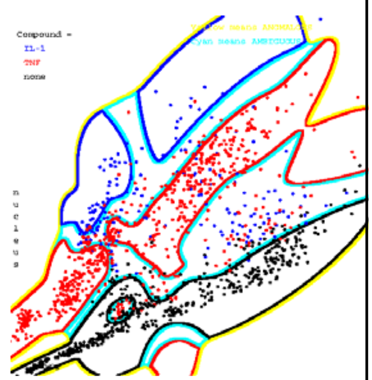

Unsupervised learning with symbolic attributes

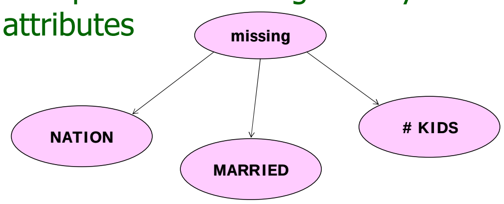

It's just a "learning Bayes net with known structure but hidden values" problem.

Can use Gradient Descent.

EASY, fun exercise to do an EM formulation for this case too.

Final Comments

- Remember, E.M. can get stuck in local minima, and empirically it DOES.
- Our unsupervised learning example assumed $P(w_i)$'s known, and variances fixed and known. Easy to relax this.
- It's possible to do Bayesian unsupervised learning instead of max. likelihood.
- There are other algorithms for unsupervised learning. We'll visit K-means soon. Hierarchical clustering is also interesting.
- Neural-net algorithms called "competitive learning" turn out to have interesting parallels with the EM method we saw.

What you should know

- How to "learn" maximum likelihood parameters (locally max. like.) in the case of unlabeled data.
- Be happy with this kind of probabilistic analysis.
- Understand the two examples of E.M. given in these notes.

For more info, see Duda + Hart. It's a great book. There's much more in the book than in your handout.

Other unsupervised learning methods

- K-means (see next lecture)
- Hierarchical clustering (e.g. Minimum spanning trees) (see next lecture)
- Principal Component Analysis
  - simple, useful tool
- Non-linear PCA
  - Neural Auto-Associators
  - Locally weighted PCA
  - Others…