

## Naive Bayes

Suppose categories indexed by  $c$ , and features represented in a vector  $\mathbf{x}$ .

Assume features in  $\mathbf{x}$  are independent given the category.


(Feature independence is the “naive” part).

$$p(\mathbf{x}|c) = \prod_i p(x_i|c)$$

## Naive Bayes

$$p(\mathbf{x}, c) = p(c) \prod_i p(x_i|c)$$

Graphical model is:



## Naive Bayes

$$p(\mathbf{x}, c) = p(c) \prod_i p(x_i|c)$$

Note that:

The forms of  $p(x_i|c)$  need not all be the same (but usually are)

If  $p(\mathbf{x}|c)$  is a Gaussian, then it has diagonal covariance matrix.  
(This simplifying assumption is nearly always needed with Gaussians if the dimension,  $D$ , is large).

## Naive Bayes

Typically,  $p(x_i|c)$  come from training data linked to known (labeled) classes (supervised learning).

Example (1) fit a univariate Gaussian to each variable,  $x_i$ , for each class,  $c$ .

Example (2), record a histogram for each variable,  $x_i$ , for each class,  $c$ .

## Inference using Naive Bayes

$$p(\mathbf{x}|c) = \prod_i p(x_i|c) \quad (\text{forward model})$$

$$p(c|\mathbf{x}) \propto p(\mathbf{x}|c)p(c) \quad (\text{the Bayes part})$$

This leaves us with simple, and often very effective model and associated inference. We combine the likelihood  $p(\mathbf{x}|c)$  with the prior  $p(c)$  over categories.

## Naive Bayes for face identification

- Example features
  - Location, color, texture, of left eye
  - Location, color, texture, of right eye
  - Location, color, texture, of mouth
  - Location, color, texture, of nose
- We can imagine training these with different facial expressions, lighting conditions, etc.
- Notice that these are not independent.
- This sort of thing often works pretty well anyway.
- Possible explanation is that, while the model allows for the eyes to be different, this rarely occurs in training or testing data.

## Clustering

Clustering is the canonical case of “unsupervised” learning.

Given the data, what are the categories (clusters),  $c$ ?

(Given a cluster, the features might be independent like Naive Bayes, or they might not be).

We will focus on clustering based on statistical models, but first review clustering in general.

## Why is clustering hard?

### Main reason

- The number of possible clusterings is exponential in the number of data points

### Other important issues

- The number of clusters (and a good way to check) is usually **not** known
- A good distance function between points may not be known
- A good model explaining the existence of clusters may not be known.
- High dimensionality

## Clustering based on distance measure


- Most common data representation is an N dimensional “feature” vector.
- Most common distance is Euclidian distance.
- Be careful with scaling and units!
- Probabilistic models can finesse scaling and multiple modalities
- Problems with correlated variables can be mitigated using transformations and data reduction methods such as PCA, ICA.

## Clustering approaches

- Agglomerative clustering
  - initialize: every item is a cluster
  - attach item that is “closest” to a cluster to that cluster
  - repeat
- Divisive clustering
  - split cluster along best boundary
  - repeat
- Probabilistic clustering
  - Define a probabilistic grouping model

## Simple agglomerative approaches

- Point-Cluster or Cluster-Cluster distance
  - single-link clustering (minimum distance from point to points in clusters or among pairs of points, one from each cluster)
  - complete-link clustering (maximum)
  - group-average clustering (average)
  - (terms are not important, but concepts are worth thinking about)
- Dendrograms
  - classic picture of output as clustering process continues



## K-Means

- Choose a fixed number of clusters (“K”)
- Choose cluster centers (**means**) and point-cluster allocations (membership) to minimize the error

$$\sum_{i \in \text{clusters}} \left\{ \sum_{j \in \text{elements of } i\text{'th cluster}} \|x_j - \mu_i\|^2 \right\}$$

- $x$ 's could be any set of features for which we can compute a distance (careful with scaling)

## K-Means

- Want to minimize

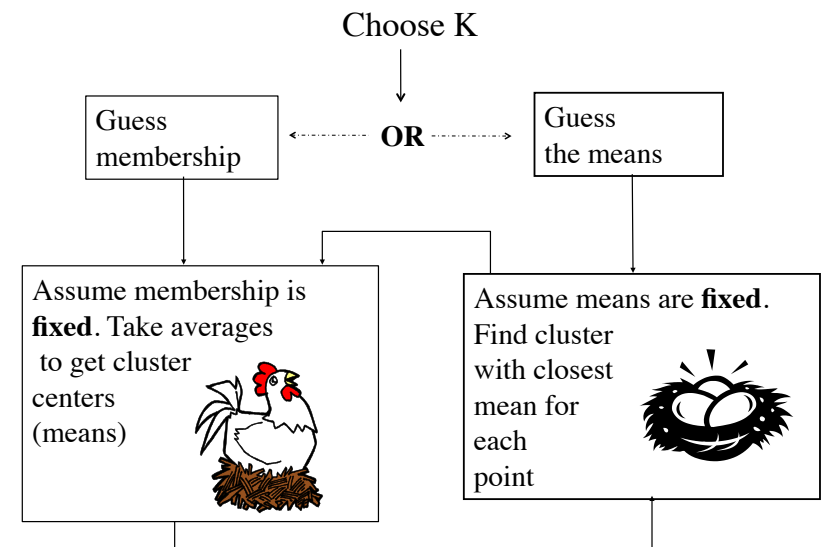
$$\sum_{i \in \text{clusters}} \left\{ \sum_{j \in \text{elements of } i\text{'th cluster}} \|x_j - \mu_i\|^2 \right\}$$

- Note that if we know the cluster allocation, we can compute the means
- **Cannot** do this optimization by search, because there are too many possible allocations.
- Standard difficulty which we handle with an iterative process (chicken and egg)

## K-Means algorithm (intuition)

- If we know the cluster centers, the best cluster for each point is easy to compute
  - Just compute the distance to each to find the closest
- If we know the best cluster for each point, the cluster centers are also easy to compute
  - Just average the points in each cluster
- Algorithm
  - 1) Guess one of the two.
  - 2) Alternatively re-compute the values for each

### K-means flow chart



## Notes on K-Means

- K-means is “hard” clustering. This means that each point is completely in exactly one cluster
- What you get is a function of starting “guess”  
you should be able to argue why this is true
- The error goes down with every iteration
  - This means you get a local minimum, but not necessarily a global one.
- Unfortunately, the dimension of the space is usually large, and high-dimensional space have lots of local maximum (standard problem!)
  - Dimensionality here is  $K \cdot \dim(\mathbf{x})$
- Finding the global minimum for a real problem is very optimistic!

## Clustering using a generative statistical model

Associate each cluster with the same model type, but with different parameters.

Example (Gaussian Mixture Model (GMM)),

$$p(\mathbf{x}|c) = N(\mathbf{u}_c, \Sigma_c)$$


or, assuming feature independence,

$$p(\mathbf{x}|c) = N(\mathbf{u}_c, \sigma_c^2)$$

$p(\mathbf{x}|c)$  could also be a product of independent multinomials, or, even a product of different distributions (roll your own!).

## Clustering using a generative statistical model

Graphical model (and for independent features)



## Inference given a clustering

Given a learned clustering model (either supervised or unsupervised), we can compute a posterior probability of which cluster an instance belongs to.

$$p(c|\mathbf{x}) \propto p(\mathbf{x}|c)p(c)$$

Easily normalized since the number of clusters is finite:

$$p(c|\mathbf{x}) = \frac{p(\mathbf{x}|c)p(c)}{\sum_c p(\mathbf{x}|c)p(c)}$$

## Clustering models representing data statistics

What is the distribution of data best described by clusters?  
(Example, data coming from a bimodal distribution?)

$$\begin{aligned} p(\mathbf{x}) &= \sum_c p(\mathbf{x}, c) \\ &= \sum_c p(c) p(\mathbf{x}|c) \end{aligned}$$

Generative story:


- 1) choose a cluster with probability,  $p(c)$ .
- 2) sample from  $p(\mathbf{x}|c)$ .
- 3) rinse and repeat.

## Clustering models representing data statistics

Distribution of data described by clusters.

$$p(\mathbf{x}) = \sum_c p(c) p(\mathbf{x}|c)$$

Distribution modeled by  
3 multivariate Gaussians.



Even if we know the exact model, we cannot be sure from the data which point comes from which cluster. We only have the distribution for this.

## Learning the parameters from data

For concreteness, assume GMM

Assume K clusters

The goal is to learn mixing coefficients,  $p(c)$ , and cluster parameters for  $p(\mathbf{x}|c)$  for all K clusters indexed by  $c$ .

## Learning the parameters from data

The goal is to learn mixing coefficients,  $p(c)$ , and cluster parameters for  $p(\mathbf{x}|c)$  for all K clusters indexed by  $c$ .

From previous arguments, given  $p(\mathbf{x}|c)$ , we know the distribution over clusters for each data point.

We simultaneously cluster points and learn the cluster model.

## Learning the parameters from data

$$p(\mathbf{x}_i|\theta) = \sum_c p(c) p(\mathbf{x}_i|c, \theta_c)$$

Probability of all observed data will be the objective function

$$p(\{\mathbf{x}_i\}|\theta) = \prod_i \left( \sum_c p(c) p(\mathbf{x}_i|c, \theta_c) \right) \quad (\text{want this to be large})$$

or

$$\sum_i \log \left( \sum_c p(c) p(\mathbf{x}_i|c, \theta_c) \right) \quad (\text{should be large})$$

## Expectation Maximization (EM)

Operationally this is similar to K-means.

Observe that:

If we knew the cluster assignments,  
we could estimate the parameters for  $p(\mathbf{x}|c)$ .

If we knew  $p(\mathbf{x}|c)$ , we could make cluster  
assignments by computing the distribution  $p(c|\mathbf{x})$

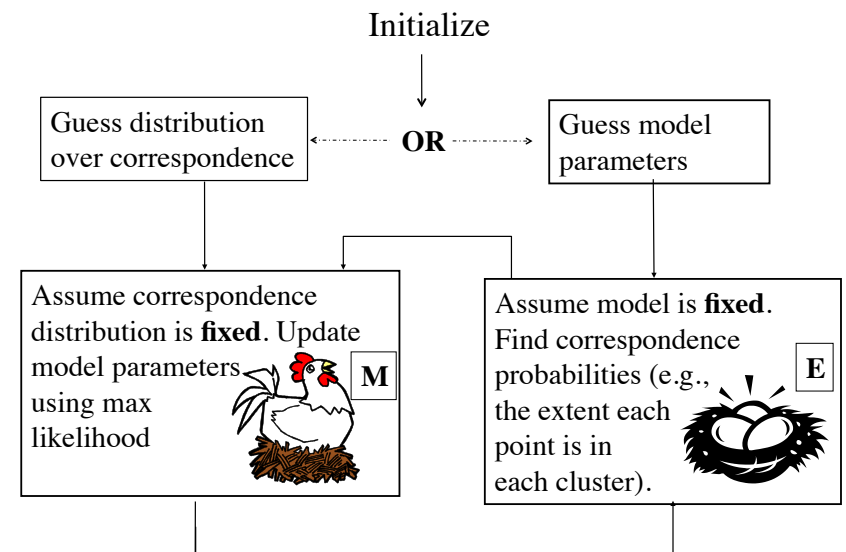
## Expectation Maximization (EM)

Difference with K-means.

We have **distributions** over the assignments,  $p(c|\mathbf{x})$ .

This leads us to work with expectations.

### EM flow chart



## EM for GMM

$$p(\mathbf{x}) = \sum_c p(c) p(\mathbf{x} | c) \quad \text{where} \quad p(\mathbf{x} | c) = \mathcal{N}(\boldsymbol{\mu}_c, \boldsymbol{\Sigma}_c)$$

$\searrow$   
 $\Theta = \{\Theta_c\}$

And, for multiple points

$$p(\{\mathbf{x}_i\} | \theta) = \prod_i \left( \sum_c p(c) p(\mathbf{x}_i | c) \right)$$

This is our objective function.

## EM for GMM


Assume we have estimates for the probability distribution over clusters for each point (the “egg”). Specifically we have:

$$p(c | \mathbf{x}_i, \Theta^{(s)}) \quad (\text{s indexes iteration (step)}).$$

These are called the *responsibilities*.

This is the extent to which each cluster explains the point. (Every point is in every cluster to some degree).

## Responsibilities illustrated



Points colored according to whether they were generated by the red, green, or blue clusters (normally not known).

Observed points without cluster information.

Points colored according to the degree that they are explained by the red, green, or blue clusters.

## EM for GMM

- We estimate the mean for each cluster naturally by:

Iteration (step)  $\searrow$

$$\mu_c^{(s+1)} = \frac{\sum_{i=1}^n \mathbf{x}_i \cdot p(c | \mathbf{x}_i, \Theta_c^{(s)})}{\sum_{i=1}^n p(c | \mathbf{x}_i, \Theta_c^{(s)})} \quad (\text{weighted average})$$

- Variances/covariances work similarly



## EM for GMM

- Also, intuitively,

$$p(c) = \frac{\sum_i p(c | \mathbf{x}_i, \Theta^{(s)})}{\sum_c \sum_i p(c | \mathbf{x}_i, \Theta^{(s)})} = \frac{\sum_i p(c | \mathbf{x}_i, \Theta^{(s)})}{N}$$

We can sort out the chicken!



## EM for GMM

Given the parameters (the chicken), the probability that a given point is associated with each cluster is computed by:

$$p(c | \mathbf{x}_i, \Theta^{(s)}) = \frac{\pi_c^{(s)} \cdot p(\mathbf{x}_i | \Theta_c^{(s)})}{\sum_{c'} \pi_{c'}^{(s)} \cdot p(\mathbf{x}_i | \Theta_{c'}^{(s)})} \quad (\text{Note that we select } \Theta_c^{(s)} \text{ from } \Theta^{(s)}).$$


where  $\pi_c^{(s)} = p(c | \Theta^{(s)})$  i.e.,  $\pi_c^{(s)}$  is part of  $\Theta_c^{(s)}$ .

This is the cluster membership discussed before,  
with less formal notation:  $p(c|x) \propto p(c)p(x|c)$

We can do the egg!



## EM illustrated



## EM (more formally)

Semi-optional technical material alert!

The formal treatment helps us use EM correctly in more complex situations. However, EM algorithms for many problems can “guessed at” using intuition.

The more formal treatment is not needed for the homework.

## EM (more formally)

- Assume K clusters. Index over clusters by  $k$ , over points by  $n$ .
- New notation for cluster membership:

For each point,  $n$ ,  $z_n$  is a vector of K values where exactly one  $z_{n,k} = 1$ , and all others are 0. Note that  $\sum_k z_{n,k} = 1$ .

## EM (more formally)

- Denote cluster priors by:

$$\pi_k \equiv p(z_k = 1)$$

- Denote the responsibilities that each cluster has for each point by:

$$\gamma(z_{n,k}) \equiv p(z_{n,k} = 1 | x_n, \theta^{(s)}) = \frac{\pi_k p(x_n | z_{n,k} = 1, \theta_k^{(s)})}{\sum_{k'} \pi_{k'} p(x_n | z_{n,k'} = 1, \theta_{k'}^{(s)})}$$

## EM (more formally)

- Responsibilities for GMM using this notation:

$$\gamma(z_{n,k}) = \frac{\pi_k N(x_n | u_k^{(s)}, \Sigma_k^{(s)})}{\sum_{k'} \pi_{k'} N(x_n | u_{k'}^{(s)}, \Sigma_{k'}^{(s)})}$$

## EM (more formally)

Represent the entire data set of N points,  $\mathbf{x}_n$ , as a matrix X with rows  $\mathbf{x}_n^T$ .

Represent the latent variable assignments with a matrix Z. (For the true assignment, each row is zero except for a single element that is 1.)

We call  $\{Z, X\}$  the *complete* data set (everything is known). The observed data,  $\{X\}$ , is called the *incomplete* data set.

## EM (more formally)

We assume that computing the MLE of parameters,

$$\arg \max_{\theta} \left\{ \log \left\{ p(Z, X | \theta) \right\} \right\}$$

with complete data is relatively easy.

Recall our intuitive treatment of EM for GMM. If we knew the cluster membership, we would know how to compute the means.

Since we did not know the cluster membership we did a weighted computation, which happens to be an expectation of the complete log likelihood, over the assignment (responsibility) distribution.

## EM (more formally)

For the E step, we compute the responsibilities which is straightforward.

$$\text{Define } Q(\theta^{(s+1)}, \theta^{(s)}) = \sum_Z p(Z | X, \theta^{(s)}) \log \left( p(X, Z | \theta^{(s+1)}) \right)$$

(Expectation of  $\log \left( p(X, Z | \theta^{(s+1)}) \right)$  over  $p(Z | X, \theta^{(s)})$ ).

$$\text{The M step then computes } \theta^{(s+1)} = \arg \max_{\theta} \left\{ Q(\theta^{(s+1)}, \theta^{(s)}) \right\}$$

Maximizing Q is generally feasible and corresponds to the intuitive development.

## EM (more formally)

Notice the complexity of the incomplete log likelihood:

$$\log(p(X | \theta)) = \sum_n \log \left( \underbrace{\sum_k \pi_k p(x_n | \theta)}_{\text{nasty sum in log}} \right)$$

By contrast, for complete log likelihood we can incorporate the assignment by:

$$p(X, Z | \theta) = \prod_n \prod_k \pi_k^{z_{n,k}} \left\{ p(x_n | \theta) \right\}^{z_{n,k}}$$

So

$$\log(p(X, Z | \theta)) = \sum_n \sum_k \left\{ z_{n,k} \left( \log(\pi_k) + \log(p(x_n | \theta)) \right) \right\}$$

(No nasty sum in log; well suited for the expectation calculation).

## General EM algorithm

1. Choose initial values for  $\theta^{(s=1)}$

(can also do assignments, but then jump to M step).

2. E step: Evaluate  $p(Z | X, \theta^{(s)})$

3. M step: Evaluate  $\theta^{(s+1)} = \arg \max_{\theta} \left\{ Q(\theta^{(s+1)}, \theta^{(s)}) \right\}$

$$\text{where } Q(\theta^{(s+1)}, \theta^{(s)}) = \sum_Z p(Z | X, \theta^{(s)}) \log \left( p(X, Z | \theta^{(s+1)}) \right)$$

4. Check for convergence; If not done, goto 2.

★ At each step, our objective function increases unless it is at a local maximum. It is important to check this is happening for debugging!

## General EM algorithm

- ★ At each step, our objective function (conditioned on the current model) increases unless it is at a local maximum. It is important to check this is happening for debugging!

Recall our objective function:

$$p(X) = \prod_n \sum_k p(k) p(x_n | k)$$

or

$$\log(p(X)) = \sum_n \log \left( \sum_k p(k) p(x_n | k) \right)$$

## Deriving the GMM M-step

$$\text{Evaluate } \theta^{(s+1)} = \arg \max_{\theta} \{Q(\theta^{(s+1)}, \theta^{(s)})\}$$

$$\text{where } Q(\theta^{(s+1)}, \theta^{(s)}) = \sum_Z p(Z | X, \theta^{(s)}) \log(p(X, Z | \theta^{(s)}))$$

$$\text{Recall that } \log(p(X, Z | \theta)) = \sum_n \sum_k \{z_{n,k} (\log(\pi_k) + \log(p(x_n | \theta_k)))\}$$

$$\text{So } Q(\theta^{(s+1)}, \theta^{(s)}) = \sum_Z p(Z | X, \theta^{(s)}) \sum_n \sum_k \{z_{n,k} (\log(\pi_k) + \log(p(x_n | \theta_k)))\}$$

## Deriving the GMM M-step

$$\begin{aligned} Q(\theta^{(s+1)}, \theta^{(s)}) &= \sum_Z p(Z | X, \theta^{(s)}) \sum_n \sum_k \{z_{n,k} (\log(\pi_k) + \log(p(x_n | \theta_k)))\} \\ &= \sum_n \sum_k \sum_Z p(Z | X, \theta^{(s)}) \{z_{n,k} (\log(\pi_k) + \log(p(x_n | \theta_k)))\} \\ &= \sum_n \sum_k \{\gamma(z_{n,k}) (\log(\pi_k) + \log(p(x_n | \theta_k)))\} \end{aligned}$$

(Intuitive?)

## Deriving the GMM M-step

$$\begin{aligned} &\sum_Z p(Z | X, \theta^{(s)}) f(n, k) \\ &= \sum_Z \prod_{n'} p(z_{n'} | X, \theta^{(s)}) f(n, k) \quad (\text{independence!}) \\ &= \sum_{z_n} p(z_n | X, \theta^{(s)}) f(n, k) \underbrace{\sum_{z_1} \dots \sum_{z_{n-1}} \sum_{z_{n+1}} \dots \sum_N \prod_{n' \neq n} p(z_{n'} | X, \theta^{(s)})}_{\text{all possibilities without point n.}} \\ &= \sum_{z_n} p(z_n | X, \theta^{(s)}) f(n, k) \end{aligned}$$

## Deriving the GMM M-step

$$\begin{aligned}
 Q(\theta^{(s+1)}, \theta^{(s)}) &= \sum_Z p(Z|X, \theta^{(s)}) \sum_n \sum_k \left\{ z_{n,k} \left( \log(\pi_k) + \log(p(x_n | \theta_k)) \right) \right\} \\
 &= \sum_n \sum_k \sum_{z_n} p(z_n | X, \theta^{(s)}) \left\{ z_{n,k} \left( \log(\pi_k) + \log(p(x_n | \theta_k)) \right) \right\} \\
 &= \sum_n \sum_k \sum_{z_n} p(z_n | X, \theta^{(s)}) z_{n,k} \left( \log(\pi_k) + \log(p(x_n | \theta_k)) \right) \\
 &= \sum_n \sum_k \gamma(n, k) \left( \log(\pi_k) + \log(p(x_n | \theta_k)) \right)
 \end{aligned}$$

Here we use the special nature of the indicator variables. They pick the value in the sum corresponding to the index.

## Deriving the GMM M-step

$$\begin{aligned}
 Q(\theta^{(s+1)}, \theta^{(s)}) &= \sum_Z p(Z|X, \theta^{(s)}) \sum_n \sum_k \left\{ z_{n,k} \left( \log(\pi_k) + \log(p(x_n | \theta_k)) \right) \right\} \\
 &= \sum_n \sum_k \left\{ \gamma(z_{n,k}) \left( \log(\pi_k) + \log(p(x_n | \theta_k)) \right) \right\}
 \end{aligned}$$

We need to maximize this with respect to the parameters for each cluster,  $k$ . Notice that:

$$\frac{\delta}{\delta \theta_{k^*}} Q(\theta^{(s+1)}, \theta^{(s)}) = \sum_n \left\{ \gamma(z_{n,k^*}) \frac{\delta}{\delta \theta_{k^*}} \left( \log(\pi_{k^*}) + \log(p(x_n | \theta_{k^*})) \right) \right\}$$

## Deriving the GMM M-step

$$\begin{aligned}
 \frac{\delta}{\delta \mu_k} Q(\theta^{(s+1)}, \theta^{(s)}) &= \sum_n \left\{ \gamma(z_{n,k}) \frac{\delta}{\delta \mu_k} \left( \log(\pi_k) + \log(p(x_n | \theta_k)) \right) \right\} \\
 &= \sum_n \left\{ \gamma(z_{n,k}) \frac{\delta}{\delta \mu_k} \left( \log(p(x_n | \theta_k)) \right) \right\} \\
 &= \sum_n \left\{ \gamma(z_{n,k}) \frac{\delta}{\delta \mu_k} \left( \log(N(x_n | \mu_k, \Sigma_k)) \right) \right\}
 \end{aligned}$$

## Deriving the GMM M-step

$$\begin{aligned}
 N(\mathbf{x}_n | \mu_k, \Sigma_k) &= \frac{1}{(2\pi)^{D/2} |\Sigma_k|^{1/2}} \exp \left( -\frac{1}{2} (\mathbf{x}_n - \mu_k)^T \Sigma_k^{-1} (\mathbf{x}_n - \mu_k) \right) \\
 \log(N(\mathbf{x}_n | \mu_k, \Sigma_k)) &= \log \left( \frac{1}{(2\pi)^{D/2} |\Sigma_k|^{1/2}} \right) - \frac{1}{2} (\mathbf{x}_n - \mu_k)^T \Sigma_k^{-1} (\mathbf{x}_n - \mu_k)
 \end{aligned}$$

$$\frac{\delta}{\delta \mu_k} \log(N(\mathbf{x}_n | \mu_k, \Sigma_k)) = -\frac{1}{2} \Sigma_k^{-1} (\mathbf{x}_n - \mu_k)$$

## Deriving the GMM M-step

$$\frac{\partial}{\partial \mu_k} Q(\theta^{(s+1)}, \theta^{(s)}) = \sum_n \left\{ \gamma(z_{n,k}) \frac{\partial}{\partial \mu_k} \left( \log(N(x_n | \mu_k, \Sigma_k)) \right) \right\}$$

$$\frac{\partial}{\partial \mu_k} Q(\theta^{(s+1)}, \theta^{(s)}) = 0 \text{ means that}$$

$$-\frac{1}{2} \Sigma_k^{-1} \sum_n \left\{ \gamma(z_{n,k}) (\mathbf{x}_n - \mu_k) \right\} = 0$$

$$\text{So, } \sum_n \left\{ \gamma(z_{n,k}) (\mathbf{x}_n - \mu_k) \right\} = 0$$

## Deriving the GMM M-step

$$\text{So, } \sum_n \left\{ \gamma(z_{n,k}) (\mathbf{x}_n - \mu_k) \right\} = 0$$

$$\text{and } \mu_k \sum_n \left\{ \gamma(z_{n,k}) \right\} = \sum_n \left\{ \gamma(z_{n,k}) (\mathbf{x}_n) \right\}$$

$$\text{and } \mu_k = \frac{\sum_n \left\{ \gamma(z_{n,k}) (\mathbf{x}_n) \right\}}{\sum_n \left\{ \gamma(z_{n,k}) \right\}} \quad (\text{same as before})$$

## Deriving the GMM M-step

Finding variances/covariances is similar.

Finding the mixing coefficients is also similar, except we also need to enforce that they sum to one.

(Here the equations for the  $k$ 's are coupled).

So we use Lagrange Multipliers.

## Using Lagrange Multipliers


Now we find stationary points with respect to  $\{\pi_k, \lambda\}$  of

$$Q(\theta^{(s+1)}, \theta^{(s)}) + \lambda \left( \sum_k \pi_k - 1 \right)$$

Note that differentiating with respect to  $\lambda$ , and setting the result to zero puts the constraint into the equations.


But the real problem is doing the optimization under the constraint.

## Using Lagrange Multipliers



From Wikipedia

## Using Lagrange Multipliers



From Wikipedia

## Using Lagrange Multipliers

$$\nabla f \parallel \nabla g$$

$$\nabla f = \lambda \nabla g$$

$$\text{So, } \nabla(f - \lambda g) = 0$$

$$\text{or, } \nabla(f + \lambda g) = 0 \quad (\text{negate } \lambda)$$

## Using Lagrange Multipliers

Now we find stationary points with respect to  $\{\pi_k, \lambda\}$  of

$$Q(\theta^{(s+1)}, \theta^{(s)}) + \lambda \left( \sum_k \pi_k - 1 \right)$$

$$\begin{aligned} & \frac{\delta}{\delta \pi_k} \left\{ Q(\theta^{(s+1)}, \theta^{(s)}) + \lambda \left( \sum_k \pi_k - 1 \right) \right\} \\ &= \sum_n \left\{ \gamma(z_{n,k}) \frac{\delta}{\delta \pi_k} \left( \log(\pi_k) \log(N(x_n | \mu_k, \Sigma_k)) \right) \right\} + \lambda \\ &= \sum_n \left\{ \gamma(z_{n,k}) \frac{1}{\pi_k} \right\} + \lambda \end{aligned}$$

Setting the result to zero,  $\sum_n \left\{ \gamma(z_{n,k}) \frac{1}{\pi_k} \right\} + \lambda = 0$


So  $\pi_k = \frac{\sum \{ \gamma(z_{n,k}) \}}{-\lambda}$

Summing over  $k$  gives,  $1 = \frac{\sum_k \sum_n \{ \gamma(z_{n,k}) \}}{-\lambda} = \frac{N}{-\lambda}$

So,  $\lambda = -N$ , and  $\pi_k = \frac{\sum \{ \gamma(z_{n,k}) \}}{N}$  as before.

## EM in practice

- For GMM we need to consider clusters that have only one point:



- Easily fixed by adding a prior on the variance.

## EM in practice

- Tying parameters (using GMM as an example)
  - We can improve stability by assuming the variances (or covariances) for all clusters are the same.
  - Updates work as you expect. Instead of multiple weighted sums, you just use one big one.
  - But note that one advantage of GMM over K-means is that the scale is naturally taken care of, and the clusters **can** have different variances.

## EM in practice

- You must check that the log likelihood increases!
- A simple way to compute it during an iteration:

Recall our objective function:

$$p(X) = \prod_n \sum_k p(k) p(x_n | k)$$

Consider how we might compute the responsibilities

$$\gamma(n, k) \propto p(k) p(x_n | k)$$

(Then normalize once you have them all).

So, make a running sum of the unnormalized values



## EM in practice

- Precision problems --> must work with logs
- But we need to exponentiate to normalize --> rescaling tricks

Let  $P = \{p_i\}$ .

Suppose we want  $Q = \frac{1}{\sum_i p_i} \{p_i\}$

Where we need to use  $V = \{\log(p_i)\}$

and  $\exp(p_i)$  is too small, and the sum of them might be zero.

Let  $M = \max\{\log(p_i)\}$

Observe that working with  $V' = \{\log(p_i) - M\}$  does the trick.

## EM in practice

- Memory problems ---> we can compute means, etc., as running totals so that we do not need to store responsibilities for all points over all clusters.