

CS434a/541a: Pattern Recognition
Prof. Olga Veksler

Lecture 17

Today


- Parametric Unsupervised Learning
 - Expectation Maximization (EM)
 - one of the most useful statistical methods
 - oldest version in 1958 (*Hartley*)
 - seminal paper in 1977 (*Dempster et al.*)
 - can also be used when some samples are missing features

Unsupervised Learning

- In unsupervised learning, where we are only given samples $\mathbf{x}_1, \dots, \mathbf{x}_n$ without class labels
- Last 2 lectures: nonparametric approach (clustering)
- Today, parametric approach
 - assume parametric distribution of data
 - estimate parameters of this distribution
 - much “harder” than the supervised learning case

Parametric Unsupervised Learning

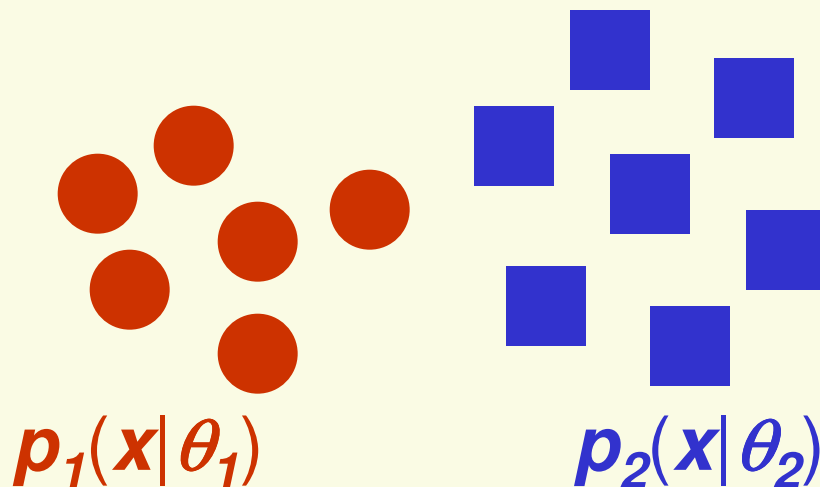
- Assume the data was generated by a model with known shape but unknown parameters



- Advantages of having a model
 - Gives a meaningful way to cluster data
 - adjust the parameters of the model to maximize the probability that the model produced the observed data
 - Can sensibly measure if a clustering is good
 - compute the likelihood of data induced by clustering
 - Can compare 2 clustering algorithms
 - which one gives the higher likelihood of the observed data?

Parametric Supervised Learning

- Let us recall supervised parametric learning
 - have m classes
 - have samples $\mathbf{x}_1, \dots, \mathbf{x}_n$ each of class $1, 2, \dots, m$
 - suppose D_i holds samples from class i
 - probability distribution for class i is $p_i(\mathbf{x}|\theta_i)$




Parametric Supervised Learning

- Use the ML method to estimate parameters θ_i
 - find θ_i which maximizes the likelihood function $F(\theta_i)$


$$p(\mathbf{D}_i | \theta_i) = \prod_{x \in \mathbf{D}_i} p(x | \theta_i) = F(\theta_i)$$

- or, equivalently, find θ_i which maximizes the log likelihood $l(\theta_i)$

$$l(\theta_i) = \ln p(\mathbf{D}_i | \theta_i) = \sum_{x \in \mathbf{D}_i} \ln p(x | \theta_i)$$




$$\hat{\theta}_1 = \operatorname{argmax}_{\theta_1} [\ln p(\mathbf{D}_1 | \theta_1)]$$



$$\hat{\theta}_2 = \operatorname{argmax}_{\theta_2} [\ln p(\mathbf{D}_2 | \theta_2)]$$

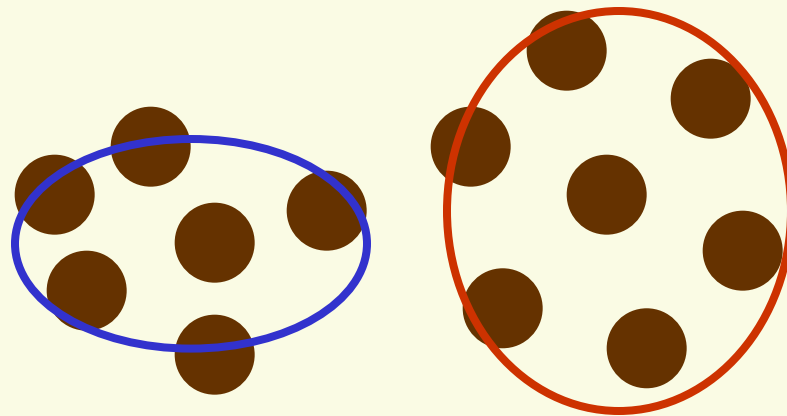
Parametric Supervised Learning

- now the distributions are fully specified
- can classify unknown sample using MAP classifier





Parametric Unsupervised Learning

- In unsupervised learning, no one tells us the true classes for samples. We still know
 - have m classes
 - have samples $\mathbf{x}_1, \dots, \mathbf{x}_n$ each of **unknown** class
 - probability distribution for class i is $p_i(\mathbf{x}|\theta_i)$
- Can we determine the classes and parameters simultaneously?



Example: MRI Brain Segmentation



Picture from M. Leventon

- In MRI brain image, different brain tissues have different intensities
- Know that brain has 6 major types of tissues
- Each type of tissue can be modeled by a Gaussian $\mathbf{N}(\mu_i, \sigma_i^2)$ reasonably well, parameters μ_i, σ_i^2 are unknown
- Segmenting (classifying) the brain image into different tissue classes is very useful
 - don't know which image pixel corresponds to which tissue (class)
 - don't know parameters for each $\mathbf{N}(\mu_i, \sigma_i^2)$


Mixture Density Model

- Model data with **mixture density**

component densities

$$p(\mathbf{x} | \theta) = \sum_{j=1}^m \underbrace{p(\mathbf{x} | \mathbf{c}_j, \theta_j)}_{\text{component densities}} \underbrace{P(\mathbf{c}_j)}_{\text{mixing parameters}}$$

- where $\theta = \{\theta_1, \dots, \theta_m\}$
- $P(\mathbf{c}_1) + P(\mathbf{c}_2) + \dots + P(\mathbf{c}_m) = 1$
- To generate a sample from distribution $p(\mathbf{x} | \theta)$
 - first select class j with probability $P(\mathbf{c}_j)$
 - then generate \mathbf{x} according to probability law $p(\mathbf{x} | \mathbf{c}_j, \theta_j)$




Example: Gaussian Mixture Density

- Mixture of 3 Gaussians

$$p_1(\mathbf{x}) \cong N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}\right)$$

$$p_2(\mathbf{x}) \cong N\left(\begin{bmatrix} 6 \\ 6 \end{bmatrix}, \begin{bmatrix} 4 & 0 \\ 0 & 4 \end{bmatrix}\right)$$

$$p_3(\mathbf{x}) \cong N\left(\begin{bmatrix} 7 \\ -7 \end{bmatrix}, \begin{bmatrix} 6 & 0 \\ 0 & 6 \end{bmatrix}\right)$$



$$p(\mathbf{x}) = 0.2p_1(\mathbf{x}) + 0.3p_2(\mathbf{x}) + 0.5p_3(\mathbf{x})$$

Mixture Density

$$p(\mathbf{x} | \theta) = \sum_{j=1}^m p(\mathbf{x} | \mathbf{c}_j, \theta_j) P(\mathbf{c}_j)$$

- $P(\mathbf{c}_1), \dots, P(\mathbf{c}_m)$ can be known or unknown
- Suppose we know how to estimate $\theta_1, \dots, \theta_m$ and $P(\mathbf{c}_1), \dots, P(\mathbf{c}_m)$
- Can “break apart” mixture $p(\mathbf{x} | \theta)$ for classification
- To classify sample \mathbf{x} , use MAP estimation, that is choose class i which maximizes

$$P(\mathbf{c}_i | \mathbf{x}, \theta_i) \propto \underbrace{p(\mathbf{x} | \mathbf{c}_i, \theta_i)}_{\text{probability of component } i \text{ to generate } \mathbf{x}} \underbrace{P(\mathbf{c}_i)}_{\text{probability of component } i}$$

*probability of component i
to generate \mathbf{x}*

*probability of
component i*

ML Estimation for Mixture Density

$$p(\mathbf{x} | \theta, \rho) = \sum_{j=1}^m p(\mathbf{x} | \mathbf{c}_j, \theta_j) P(\mathbf{c}_j) = \sum_{j=1}^m p(\mathbf{x} | \mathbf{c}_j, \theta_j) \rho_j$$

- Can use Maximum Likelihood estimation for a mixture density; need to estimate
 - $\theta_1, \dots, \theta_m$
 - $\rho_1 = P(\mathbf{c}_1), \dots, \rho_m = P(\mathbf{c}_m)$, and $\rho = \{\rho_1, \dots, \rho_m\}$
- As in the supervised case, form the logarithm likelihood function

$$l(\theta, \rho) = \ln p(\mathbf{D} | \theta, \rho) = \sum_{k=1}^n \ln p(\mathbf{x}_k | \theta, \rho) = \sum_{k=1}^n \ln \left[\sum_{j=1}^m p(\mathbf{x} | \mathbf{c}_j, \theta_j) \rho_j \right]$$

ML Estimation for Mixture Density

$$l(\theta, \rho) = \sum_{k=1}^n \ln \left[\sum_{j=1}^m p(\mathbf{x} | \mathbf{c}_j, \theta_j) \rho_j \right]$$

- need to maximize $l(\theta, \rho)$ with respect to θ and ρ
- As you may have guessed, $l(\theta, \rho)$ is not the easiest function to maximize
 - If we take partial derivatives with respect to θ, ρ and set them to $\mathbf{0}$, typically we have a “coupled” nonlinear system of equation
 - usually closed form solution cannot be found
- We could use the gradient ascent method
 - in general, it is not the greatest method to use, should only be used as last resort
- There is a better algorithm, called **EM**

Mixture Density

- Before EM, let's look at the mixture density again

$$p(\mathbf{x} | \theta, \rho) = \sum_{j=1}^m p(\mathbf{x} | \mathbf{c}_j, \theta_j) \rho_j$$

- Suppose we know how to estimate $\theta_1, \dots, \theta_m$ and ρ_1, \dots, ρ_m

- Estimating the class of \mathbf{x} is easy with MAP, maximize

$$p(\mathbf{x} | \mathbf{c}_i, \theta_i) P(\mathbf{c}_i) = p(\mathbf{x} | \mathbf{c}_i, \theta_i) \rho_i$$

- Suppose we know the class of samples $\mathbf{x}_1, \dots, \mathbf{x}_n$

- This is just the supervised learning case, so estimating $\theta_1, \dots, \theta_m$ and ρ_1, \dots, ρ_m is easy

$$\hat{\theta}_i = \operatorname{argmax}_{\theta_i} [\ln p(\mathbf{D}_i | \theta_i)] \quad \hat{\rho}_i = \frac{|\mathbf{D}_i|}{n}$$

- This is an example of chicken-and-egg problem

- *ME* algorithm approaches this problem by adding “hidden” variables

Expectation Maximization Algorithm

- **EM** is an algorithm for ML parameter estimation when the data has missing values. It is used when
 1. data is incomplete (has missing values)
 - some features are missing for some samples due to data corruption, partial survey responses, etc.
 - This scenario is very useful, covered in section 3.9
 2. Suppose data \mathbf{X} is complete, but $p(\mathbf{X}|\theta)$ is hard to optimize. Suppose further that introducing certain hidden variables \mathbf{U} whose values are missing, and suppose it is easier to optimize the “complete” likelihood function $p(\mathbf{X}, \mathbf{U}|\theta)$. Then EM is useful.
 - This scenario is useful for the mixture density estimation, and is subject of our lecture today
- Notice that after we introduce artificial (hidden) variables \mathbf{U} with missing values, case **2** is completely equivalent to case **1**


EM: Hidden Variables for Mixture Density

$$p(\mathbf{x} | \theta) = \sum_{j=1}^m p(\mathbf{x} | \mathbf{c}_j, \theta_j) \rho_j$$

- For simplicity, assume component densities are

$$p(\mathbf{x} | \mathbf{c}_j, \theta_j) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(\mathbf{x} - \mu_j)^2}{2\sigma^2}\right)$$

- assume for now that the variance is known
- need to estimate $\theta = \{\mu_1, \dots, \mu_m\}$



- If we knew which sample came from which component (that is the class label), the ML parameter estimation is easy
- Thus to get an easier problem, introduce hidden variables which indicate which component each sample belongs to

EM: Hidden Variables for Mixture Density

- For $1 \leq i \leq n$, $1 \leq k \leq m$, define hidden variables $\mathbf{z}_i^{(k)}$

$$\mathbf{z}_i^{(k)} = \begin{cases} 1 & \text{if sample } i \text{ was generated by component } k \\ 0 & \text{otherwise} \end{cases}$$

$$\mathbf{x}_i \rightarrow \{\mathbf{x}_i, \mathbf{z}_i^{(1)}, \dots, \mathbf{z}_i^{(m)}\}$$

- $\mathbf{z}_i^{(k)}$ are indicator **random** variables, they indicate which Gaussian component generated sample \mathbf{x}_i
- Let $\mathbf{z}_i = \{\mathbf{z}_i^{(1)}, \dots, \mathbf{z}_i^{(m)}\}$, indicator r.v. corresponding to sample \mathbf{x}_i
- Conditioned on \mathbf{z}_i , distribution of \mathbf{x}_i is Gaussian

$$p(\mathbf{x}_i | \mathbf{z}_i, \theta) \sim N(\mu_k, \sigma^2)$$

- where k is s.t. $\mathbf{z}_i^{(k)} = 1$

EM: Joint Likelihood

- Let $\mathbf{z}_i = \{\mathbf{z}_i^{(1)}, \dots, \mathbf{z}_i^{(m)}\}$, and $\mathbf{Z} = \{\mathbf{z}_1, \dots, \mathbf{z}_n\}$
- The complete likelihood is

$$\begin{aligned} p(\mathbf{X}, \mathbf{Z} | \theta) &= p(\mathbf{x}_1, \dots, \mathbf{x}_n, \mathbf{z}_1, \dots, \mathbf{z}_n | \theta) = \prod_{i=1}^n p(\mathbf{x}_i, \mathbf{z}_i | \theta) \\ &= \prod_{i=1}^n \underbrace{p(\mathbf{x}_i | \mathbf{z}_i, \theta)}_{\text{gaussian}} \underbrace{p(\mathbf{z}_i)}_{\text{part of } \rho_c} \end{aligned}$$

- If we actually observed \mathbf{Z} , the log likelihood $\ln[p(\mathbf{X}, \mathbf{Z} | \theta)]$ would be trivial to maximize with respect to θ and ρ_i
- The problem, is, of course, is that the values of \mathbf{Z} are missing, since we made it up (that is \mathbf{Z} is hidden)

EM Derivation

- Instead of maximizing $\ln[p(\mathbf{X}, \mathbf{Z} | \theta)]$ the idea behind *EM* is to maximize some function of $\ln[p(\mathbf{X}, \mathbf{Z} | \theta)]$, usually its expected value

$$E_{\mathbf{Z}}[\ln p(\mathbf{X}, \mathbf{Z} | \theta)]$$

- If θ makes $\ln[p(\mathbf{X}, \mathbf{Z} | \theta)]$ large, then θ tends to make $E[\ln p(\mathbf{X}, \mathbf{Z} | \theta)]$ large
 - the expectation is with respect to the missing data \mathbf{Z}
 - that is with respect to density $p(\mathbf{Z} | \mathbf{X}, \theta)$
- however θ is our ultimate goal, we don't know θ !

EM Algorithm

- EM solution is to iterate

1. start with initial parameters $\theta^{(0)}$

iterate the following 2 step until convergence


E. compute the expectation of log likelihood with respect to current estimate $\theta^{(t)}$ and \mathbf{X} . Let's call it $\mathbf{Q}(\theta | \theta^{(t)})$

$$\mathbf{Q}(\theta | \theta^{(t)}) = \mathbf{E}_{\mathbf{Z}} [\ln p(\mathbf{X}, \mathbf{Z} | \theta) | \mathbf{X}, \theta^{(t)}]$$

M. maximize $\mathbf{Q}(\theta | \theta^{(t)})$


$$\theta^{(t+1)} = \underset{\theta}{\operatorname{argmax}} \mathbf{Q}(\theta | \theta^{(t)})$$

EM Algorithm: Picture




*optimal value for θ
we'd like to find it but
optimizing $p(\mathbf{X} | \theta)$ is
very difficult*

EM Algorithm: Picture



EM Algorithm: Picture



EM Algorithm

- It can be proven that EM algorithm converges to the local maximum of the log-likelihood

$$\ln p(\mathbf{X} | \theta)$$

- Why is it better than gradient ascent?
 - Convergence of EM is usually **significantly** faster, in the beginning, very large steps are made (that is likelihood function increases rapidly), as opposed to gradient ascent which usually takes tiny steps
 - gradient descent is not guaranteed to converge
 - recall all the difficulties of choosing the appropriate learning rate

EM for Mixture of Gaussians: E step

- Let's come back to our example $p(\mathbf{x} | \theta) = \sum_{j=1}^m p(\mathbf{x} | \mathbf{c}_j, \theta_j) \rho_j$

$$p(\mathbf{x} | \mathbf{c}_j, \theta_j) = \frac{1}{\sigma \sqrt{2\pi}} \exp\left(-\frac{(\mathbf{x} - \mu_j)^2}{2\sigma^2}\right)$$

- need to estimate $\theta = \{\mu_1, \dots, \mu_m\}$ and ρ_1, \dots, ρ_m
- for $1 \leq i \leq n$, $1 \leq k \leq m$, define $\mathbf{z}_i^{(k)}$
$$\mathbf{z}_i^{(k)} = \begin{cases} 1 & \text{if sample } i \text{ was generated by component } k \\ 0 & \text{otherwise} \end{cases}$$
- as before, $\mathbf{z}_i = \{\mathbf{z}_i^{(1)}, \dots, \mathbf{z}_i^{(m)}\}$, and $\mathbf{Z} = \{\mathbf{z}_1, \dots, \mathbf{z}_n\}$
- We need log-likelihood of observed \mathbf{X} and hidden \mathbf{Z}

$$\ln p(\mathbf{X}, \mathbf{Z} | \theta) = \ln \prod_{i=1}^n p(\mathbf{x}_i, \mathbf{z}_i | \theta) = \sum_{i=1}^n \ln p(\mathbf{x}_i | \mathbf{z}_i, \theta) P(\mathbf{z}_i)$$

EM for Mixture of Gaussians: E step

- We need log-likelihood of observed \mathbf{X} and hidden \mathbf{Z}

$$\ln p(\mathbf{X}, \mathbf{Z} | \theta) = \ln \prod_{i=1}^n p(\mathbf{x}_i, \mathbf{z}_i | \theta) = \sum_{i=1}^n \ln p(\mathbf{x}_i | \mathbf{z}_i, \theta) P(\mathbf{z}_i)$$

- First let's rewrite $p(\mathbf{x}_i | \mathbf{z}_i, \theta) P(\mathbf{z}_i)$

$$p(\mathbf{x}_i | \mathbf{z}_i, \theta) P(\mathbf{z}_i) = \begin{cases} p(\mathbf{x}_i | \mathbf{z}_i^{(1)} = \mathbf{1}, \theta) P(\mathbf{z}_i^{(1)} = \mathbf{1}) & \text{if } \mathbf{z}_i^{(1)} = \mathbf{1} \\ \vdots \\ p(\mathbf{x}_i | \mathbf{z}_i^{(m)} = \mathbf{1}, \theta) P(\mathbf{z}_i^{(m)} = \mathbf{1}) & \text{if } \mathbf{z}_i^{(m)} = \mathbf{1} \end{cases}$$

$$= \prod_{k=1}^m [p(\mathbf{x}_i | \mathbf{z}_i^{(k)} = \mathbf{1}, \theta) P(\mathbf{z}_i^{(k)} = \mathbf{1})]^{z_i^{(k)}}$$

$$= \prod_{k=1}^m \left[\frac{1}{\sigma \sqrt{2\pi}} \exp\left(-\frac{(\mathbf{x}_i - \mu_k)^2}{2\sigma^2}\right) P(\mathbf{z}_i^{(k)} = \mathbf{1}) \right]^{z_i^{(k)}}$$

EM for Mixture of Gaussians: E step

- log-likelihood of observed \mathbf{X} and hidden \mathbf{Z} is

$$\ln p(\mathbf{X}, \mathbf{Z} | \theta) = \sum_{i=1}^n \ln p(\mathbf{x}_i | \mathbf{z}_i, \theta) P(\mathbf{z}_i)$$

$$= \sum_{i=1}^n \ln \prod_{k=1}^m \left[\frac{1}{\sigma \sqrt{2\pi}} \exp\left(-\frac{(\mathbf{x}_i - \mu_k)^2}{2\sigma^2}\right) P(\mathbf{z}_i^{(k)} = 1) \right]^{z_i^{(k)}}$$

$$= \sum_{i=1}^n \sum_{k=1}^m \ln \left[\frac{1}{\sigma \sqrt{2\pi}} \exp\left(-\frac{(\mathbf{x}_i - \mu_k)^2}{2\sigma^2}\right) P(\mathbf{z}_i^{(k)} = 1) \right]^{z_i^{(k)}}$$

$$= \sum_{i=1}^n \sum_{k=1}^m z_i^{(k)} \left[\ln \frac{1}{\sigma \sqrt{2\pi}} - \frac{(\mathbf{x}_i - \mu_k)^2}{2\sigma^2} + \ln P(\mathbf{z}_i^{(k)} = 1) \right]$$

$P(\text{sample } \mathbf{x}_i \text{ from class } k) = P(\mathbf{c}_k) = \rho_k$

$$= \sum_{i=1}^n \sum_{k=1}^m z_i^{(k)} \left[\ln \frac{1}{\sigma \sqrt{2\pi}} - \frac{(\mathbf{x}_i - \mu_k)^2}{2\sigma^2} + \ln \rho_k \right]$$

EM for Mixture of Gaussians: E step

- log-likelihood of observed \mathbf{X} and hidden \mathbf{Z} is

$$\ln p(\mathbf{X}, \mathbf{Z} | \theta) = \sum_{i=1}^n \sum_{k=1}^m z_i^{(k)} \left[\ln \frac{1}{\sigma \sqrt{2\pi}} - \frac{(\mathbf{x}_i - \mu_k)^2}{2\sigma^2} + \ln \rho_k \right]$$

- For the E step, we must compute

$$Q(\theta | \theta^{(t)}) = Q(\theta | \mu_1^{(t)}, \dots, \mu_m^{(t)}, \rho_1^{(t)}, \dots, \rho_m^{(t)}) = \mathbf{E}_{\mathbf{Z}} [\ln p(\mathbf{X}, \mathbf{Z} | \theta) | \mathbf{X}, \theta^{(t)}]$$

$$= \mathbf{E}_{\mathbf{Z}} \left(\sum_{i=1}^n \sum_{k=1}^m z_i^{(k)} \left[\ln \frac{1}{\sigma \sqrt{2\pi}} - \frac{(\mathbf{x}_i - \mu_k)^2}{2\sigma^2} + \ln \rho_k^{(t)} \right] \right)$$



$$\mathbf{E}_X \left[\sum_i a_i x_i + b \right] = \sum_i a_i \mathbf{E}_X [x_i] + b$$

$$= \sum_{i=1}^n \sum_{k=1}^m \mathbf{E}_{\mathbf{Z}} [z_i^{(k)}] \left(\ln \frac{1}{\sigma \sqrt{2\pi}} - \frac{(\mathbf{x}_i - \mu_k)^2}{2\sigma^2} + \ln \rho_k \right)$$

EM for Mixture of Gaussians: E step

$$Q(\theta | \theta^{(t)}) = \sum_{i=1}^n \sum_{k=1}^m \mathbf{E}_Z[\mathbf{z}_i^{(k)}] \left(\ln \frac{1}{\sigma \sqrt{2\pi}} - \frac{(\mathbf{x}_i - \mu_k)^2}{2\sigma^2} + \ln \rho_k \right)$$

- need to compute $\mathbf{E}_Z[\mathbf{z}_i^{(k)}]$ in the above expression

$$\mathbf{E}_Z[\mathbf{z}_i^{(k)}] = \mathbf{0} * P(\mathbf{z}_i^{(k)} = \mathbf{0} | \theta^{(t)}, \mathbf{x}_i) + \mathbf{1} * P(\mathbf{z}_i^{(k)} = \mathbf{1} | \theta^{(t)}, \mathbf{x}_i)$$

$$= P(\mathbf{z}_i^{(k)} = \mathbf{1} | \theta^{(t)}, \mathbf{x}_i) = \frac{p(\mathbf{x}_i | \theta^{(t)}, \mathbf{z}_i^{(k)} = \mathbf{1}) P(\mathbf{z}_i^{(k)} = \mathbf{1} | \theta^{(t)})}{p(\mathbf{x}_i | \theta^{(t)})}$$

$$= \frac{\rho_k^{(t)} \exp\left(-\frac{1}{2}(\mathbf{x}_i - \mu_k^{(t)})^2\right)}{\sum_{j=1}^m P(\mathbf{x}_i | \theta^{(t)}, \mathbf{z}_i^{(j)} = \mathbf{1}) P(\mathbf{z}_i^{(j)} = \mathbf{1} | \theta^{(t)})} = \frac{\rho_k^{(t)} \exp\left(-\frac{1}{2\sigma^2}(\mathbf{x}_i - \mu_k^{(t)})^2\right)}{\sum_{j=1}^m \rho_j^{(t)} \exp\left(-\frac{1}{2\sigma^2}(\mathbf{x}_i - \mu_j^{(t)})^2\right)}$$

- we are finally done with the **E** step
 - for implementation, just need to compute $\mathbf{E}_Z[\mathbf{z}_i^{(k)}]$'s don't need to compute **Q**

EM for Mixture of Gaussians: M step

$$Q(\theta | \theta^{(t)}) = \sum_{i=1}^n \sum_{k=1}^m E_Z[z_i^{(k)}] \left(\ln \frac{1}{\sigma \sqrt{2\pi}} - \frac{(x_i - \mu_k)^2}{2\sigma^2} + \ln \rho_k \right)$$

- Need to maximize Q with respect to all parameters
- First differentiate with respect to μ_k

$$\frac{\partial}{\partial \mu_k} Q(\theta | \theta^{(t)}) = \sum_{i=1}^n E_Z[z_i^{(k)}] \frac{(x_i - \mu_k)}{\sigma^2} = 0$$

$$\Rightarrow \text{new } \mu_k = \mu_k^{(t+1)} = \frac{1}{n} \sum_{i=1}^n E_Z[z_i^{(k)}] x_i$$



the mean for class k is weighted average of all samples, and this weight is proportional to the current estimate of probability that the sample belongs to class k

EM for Mixture of Gaussians: M step

$$Q(\theta | \theta^{(t)}) = \sum_{i=1}^n \sum_{k=1}^m \mathbf{E}_z[\mathbf{z}_i^{(k)}] \left(\ln \frac{1}{\sigma \sqrt{2\pi}} - \frac{(\mathbf{x}_i - \mu_k)^2}{2\sigma^2} + \ln \rho_k \right)$$

- For ρ_k we have to use Lagrange multipliers to preserve constraint

$$\sum_{j=1}^m \rho_j = 1$$

- Thus we need to differentiate $F(\lambda, \rho) = Q(\theta | \theta^{(t)}) - \lambda \left(\sum_{j=1}^m \rho_j - 1 \right)$

$$\frac{\partial}{\partial \rho_k} F(\lambda, \rho) = \sum_{i=1}^n \frac{1}{\rho_k} \mathbf{E}_z[\mathbf{z}_i^{(k)}] - \lambda = 0 \Rightarrow \sum_{i=1}^n \mathbf{E}_z[\mathbf{z}_i^{(k)}] - \lambda \rho_k = 0$$

- Summing up over all components: $\sum_{k=1}^m \sum_{i=1}^n \mathbf{E}_z[\mathbf{z}_i^{(k)}] = \sum_{k=1}^m \lambda \rho_k$

- Since $\sum_{k=1}^m \sum_{i=1}^n \mathbf{E}_z[\mathbf{z}_i^{(k)}] = n$ and $\sum_{k=1}^m \rho_k = 1$ we get $\lambda = n$

$$\rho_k^{(t+1)} = \frac{1}{n} \sum_{i=1}^n \mathbf{E}_z[\mathbf{z}_i^{(k)}]$$

EM Algorithm

The algorithm on this slide applies *ONLY* to univariate gaussian case with known variances

1. Randomly initialize $\mu_1, \dots, \mu_m, \rho_1, \dots, \rho_m$ (with constraint $\sum \rho_i = 1$)

iterate until no change in $\mu_1, \dots, \mu_m, \rho_1, \dots, \rho_m$

E. for all i, k , compute

$$\mathbf{E}_z[\mathbf{z}_i^{(k)}] = \frac{\rho_k \exp\left(-\frac{1}{2\sigma^2}(\mathbf{x}_i - \mu_k)^2\right)}{\sum_{j=1}^m \rho_j \exp\left(-\frac{1}{2\sigma^2}(\mathbf{x}_i - \mu_j)^2\right)}$$

M. for all k , do parameter update

$$\mu_k = \frac{1}{n} \sum_{i=1}^n \mathbf{E}_z[\mathbf{z}_i^{(k)}] \mathbf{x}_i \quad \rho_k = \frac{1}{n} \sum_{i=1}^n \mathbf{E}_z[\mathbf{z}_i^{(k)}]$$

EM Algorithm

- For the more general case of multivariate Gaussians with unknown means and variances

- E step:**
$$\mathbf{E}_Z[\mathbf{z}_i^{(k)}] = \frac{\rho_k \mathbf{p}(\mathbf{x} | \mu_k, \Sigma_k)}{\sum_{j=1}^m \rho_j \mathbf{p}(\mathbf{x} | \mu_j, \Sigma_j)}$$

where
$$\mathbf{p}(\mathbf{x} | \mu_k, \Sigma_k) = \frac{1}{(2\pi)^{d/2} |\Sigma_k^{-1}|^{1/2}} \exp\left[-\frac{1}{2}(\mathbf{x} - \mu_k)^t \Sigma_k^{-1}(\mathbf{x} - \mu_k)\right]$$

- M step:**

$$\rho_k = \frac{1}{n} \sum_{i=1}^n \mathbf{E}_Z[\mathbf{z}_i^{(k)}]$$

$$\mu_k = \frac{\sum_{i=1}^n \mathbf{E}_Z[\mathbf{z}_i^{(k)}] \mathbf{x}_i}{\sum_{i=1}^n \mathbf{E}_Z[\mathbf{z}_i^{(k)}]}$$

$$\Sigma_k = \frac{\sum_{i=1}^n \mathbf{E}_Z[\mathbf{z}_i^{(k)}] (\mathbf{x}_i - \mu_k)(\mathbf{x}_i - \mu_k)^T}{\sum_{i=1}^n \mathbf{E}_Z[\mathbf{z}_i^{(k)}]}$$

EM Algorithm and K-means

- *k*-means can be derived from EM algorithm
- Setting mixing parameters equal for all classes,

$$\mathbf{E}_z[\mathbf{z}_i^{(k)}] = \frac{\rho_k \exp\left(-\frac{1}{2\sigma^2}(\mathbf{x}_i - \mu_k)^2\right)}{\sum_{j=1}^m \rho_j \exp\left(-\frac{1}{2\sigma^2}(\mathbf{x}_i - \mu_j)^2\right)} = \frac{\exp\left(-\frac{1}{2\sigma^2}(\mathbf{x}_i - \mu_k)^2\right)}{\sum_{j=1}^m \exp\left(-\frac{1}{2\sigma^2}(\mathbf{x}_i - \mu_j)^2\right)}$$


- If we let $\sigma \rightarrow \infty$, then

$$\mathbf{E}_z[\mathbf{z}_i^{(k)}] = \begin{cases} 1 & \text{if } \forall j, \|\mathbf{x}_i - \mu_k\| > \|\mathbf{x}_i - \mu_j\| \\ 0 & \text{otherwise} \end{cases}$$

- so at the *E* step, for each current mean, we find all points closest to it and form new clusters
- at the *M* step, we compute the new means inside current clusters


$$\mu_k = \frac{1}{n} \sum_{i=1}^n \mathbf{E}_z[\mathbf{z}_i^{(k)}] \mathbf{x}_i$$

EM Gaussian Mixture Example




EM Gaussian Mixture Example

After first iteration




EM Gaussian Mixture Example

After second iteration




EM Gaussian Mixture Example

After third iteration



EM Gaussian Mixture Example

After 20th iteration




EM Example


- Example from R. Gutierrez-Osuna
- Training set of 900 examples forming an annulus
- Mixture model with $m = 30$ Gaussian components of unknown mean and variance is used
- Training:
 - Initialization:
 - means to 30 random examples
 - covariance matrices initialized to be diagonal, with large variances on the diagonal (compared to the training data variance)
 - During EM training, components with small mixing coefficients were trimmed
 - This is a trick to get in a more compact model, with fewer than 30 Gaussian components

EM Example


Iteration 0




Iteration 25




Iteration 50




Iteration 75



Iteration 275



Iteration 300



from R. Gutierrez-Osuna

EM Texture Segmentation Example




Figure from "Color and Texture Based Image Segmentation Using EM and Its Application to Content Based Image Retrieval", S.J. Belongie et al., ICCV 1998

EM Motion Segmentation Example

Three frames from the MPEG “flower garden” sequence




Figure from “Representing Images with layers,” by J. Wang and E.H. Adelson, IEEE Transactions on Image Processing, 1994, c 1994, IEEE

Summary

- Advantages
 - If the assumed data distribution is correct, the algorithm works well
- Disadvantages
 - If assumed data distribution is wrong, results can be quite bad.
 - In particular, bad results if use incorrect number of classes (i.e. the number of mixture components)