

Lecture 2, 3: PCA, FA and EM

4F13: Machine Learning


Zoubin Ghahramani and Carl Edward Rasmussen

Department of Engineering, University of Cambridge

January 20th, 21st 2010

Latent Variable Models

Explain correlations in y by assuming a **generative model** with **latent** (hidden) variables x



$$\mathbf{x} \sim p(\mathbf{x}|\theta_x)$$

$$y|\mathbf{x} \sim p(y|\mathbf{x}, \theta_y)$$

$$p(\mathbf{x}, y|\theta_x, \theta_y) = p(y|\mathbf{x}, \theta_y)p(\mathbf{x}|\theta_x)$$

$$p(y|\theta_x, \theta_y) = \int p(y|\mathbf{x}, \theta_y)p(\mathbf{x}|\theta_x)d\mathbf{x}$$

Inference and Learning; latent variables and parameters

Some of the variables in our model are termed **latent variables** and some are called **parameters**, why?

- each example has separate latent values; so there are many
- we usually try to integrate out latent values

The process of finding the distribution over latent variables is called **inference**.
The process of finding parameters is called **learning**.

Ideally, we would want to integrate over all unknown quantities

$$p(y) = \iint p(y|x, \theta) p(x|\theta) p(\theta) dx d\theta,$$

but unfortunately, this is not trivial — **but we will show you how to do this approximately in a few weeks time.**

Today, we concentrate on the simpler task where parameters are treated as deterministic quantities

$$p(y) \simeq p(y|\hat{\theta}), \text{ where } \hat{\theta} = \operatorname{argmax} \int p(y|x, \theta) p(x|\theta) dx.$$

Linear Regression

1-D linear regression, with inputs x and outputs y .

Observed data:

$$\mathcal{D} = \{(x^{(1)}, y^{(1)}), \dots, (x^{(N)}, y^{(N)})\}$$

Model:

$$y^{(n)} = ax^{(n)} + b$$

Is this reasonable? No!

Let's try again...

Model:

$$y^{(n)} = ax^{(n)} + b + \epsilon_n$$

where ϵ_n is a model for the noise.

- What distribution should we assume for the noise?
- How should we fit a and b ?

Linear Regression

1-D linear regression, with inputs x and outputs y .

Observed data: $\mathcal{D} = \{(x^{(1)}, y^{(1)}), \dots, (x^{(N)}, y^{(N)})\}$

Model: $y^{(n)} = ax^{(n)} + b + \epsilon_n$ where ϵ_n is a model for the noise. Assume:

- that the noise is Gaussian with mean zero and variance σ^2 , and
- that the data was independently and identically distributed (iid) from this model.

Let $y = (y^{(1)}, \dots, y^{(N)})$ and $x = (x^{(1)}, \dots, x^{(N)})$. What is the probability of the observed outputs, y , given the inputs, x , and parameters $\theta = (a, b, \sigma^2)$?

$$p(y|x, \theta) = \prod_n p(y^{(n)}|x^{(n)}, a, b, \sigma^2),$$

where

$$p(y^{(n)}|x^{(n)}, a, b, \sigma^2) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{1}{2\sigma^2}(y^{(n)} - ax^{(n)} - b)^2\right).$$

$p(y|x, \theta)$ is known as the **likelihood** for the parameters θ .

Solving for Maximum Likelihood (ML) Parameters

The ML method is one way of finding a **point estimate** for the parameters:

$$\max_{\theta} p(y|x, \theta) = \max_{\theta} \log p(y|x, \theta) \stackrel{\text{def}}{=} \max_{\theta} \mathcal{L}(\theta)$$

Let $\tilde{\mathbf{x}}^{(n)} = \begin{bmatrix} x^{(n)} \\ 1 \end{bmatrix}$ and $\beta = \begin{bmatrix} a \\ b \end{bmatrix}$.

$$\begin{aligned} \mathcal{L}(\theta) &= \sum_n \log p(y^{(n)} | x^{(n)}, a, b, \sigma^2) \\ &= -\frac{1}{2\sigma^2} \sum_n (y^{(n)} - \beta^\top \tilde{\mathbf{x}}^{(n)})^2 - \frac{N}{2} \log(2\pi\sigma^2) \\ &= -\frac{1}{2\sigma^2} \left[\sum_n y^{(n)2} - 2 \left(\sum_n y^{(n)} \tilde{\mathbf{x}}^{(n)\top} \right) \beta + \beta^\top \left(\sum_n \tilde{\mathbf{x}}^{(n)} \tilde{\mathbf{x}}^{(n)\top} \right) \beta \right] \\ &\quad - \frac{N}{2} \log(2\pi\sigma^2) \end{aligned}$$

Solving for Maximum Likelihood (ML) Parameters

Solve by taking derivatives and setting to zero...

$$\begin{aligned}\frac{\partial \mathcal{L}(\theta)}{\partial \beta} &= \frac{1}{\sigma^2} \left[\left(\sum_n y^{(n)} \tilde{\mathbf{x}}^{(n)\top} \right) - \beta^\top \left(\sum_n \tilde{\mathbf{x}}^{(n)} \tilde{\mathbf{x}}^{(n)\top} \right) \right] = 0 \\ \hat{\beta} &= \left(\sum_n \tilde{\mathbf{x}}^{(n)} \tilde{\mathbf{x}}^{(n)\top} \right)^{-1} \left(\sum_n y^{(n)} \tilde{\mathbf{x}}^{(n)} \right) \quad \text{“Normal Equations”} \\ \frac{\partial \mathcal{L}(\theta)}{\partial \sigma} &= \sigma^{-3} \left[\sum_n \dots \right] - N\sigma^{-1} = 0 \\ \hat{\sigma}^2 &= \left[\sum_n \dots \right] / N\end{aligned}$$

Linear *in the parameters* Regression

Inputs $\mathbf{x} \in \mathbb{R}^D$

Model:


$$\begin{aligned} y^{(n)} &= \beta_0 + \beta_1 x_1^{(n)} + \dots + \beta_D x_D^{(n)} + \epsilon_n \\ y^{(n)} &= \beta^\top \tilde{\mathbf{x}}^{(n)} + \epsilon_n \end{aligned}$$

where ϵ_n is Gaussian noise, and $\tilde{\mathbf{x}}^{(n)} = \begin{bmatrix} \mathbf{x}^{(n)} \\ 1 \end{bmatrix}$.

- Vector Valued Inputs: **Easy!** We've solved this already.
- Polynomial (Non-linear) Regression: **Easy!** We've solved this already.
- Basis Function (Non-linear) Regression: **Easy!** We've solved this already.

Factor Analysis

Latent variable models are useful even when both latent and observed variables are Gaussian.



Linear generative model: $y_d = \sum_{k=1}^K \Lambda_{dk} x_k + \epsilon_d$

- x_k are independent $\mathcal{N}(0, 1)$ Gaussian **factors**
- ϵ_d are independent $\mathcal{N}(0, \Psi_{dd})$ Gaussian **noise**
- the parameters are $\theta = \{\Lambda, \Psi\}$
- $K < D$

So, $y|\theta$ is Gaussian with: $p(y) = \int p(x)p(y|x, \Lambda, \Psi)dx = \mathcal{N}(0, \Lambda\Lambda^\top + \Psi)$
where Λ is a $D \times K$ matrix, and Ψ is diagonal.

Dimensionality Reduction: Finds a low-dimensional projection of high dimensional data that captures the **correlation structure** of the data.

Inference in the Factor Analysis Model

Given an observation, \mathbf{y} , what is the *distribution* of the latent cause?

Use Bayes' rule:

$$p(\mathbf{x}|\mathbf{y}) = \frac{p(\mathbf{y}|\mathbf{x})p(\mathbf{x})}{p(\mathbf{y})} \propto p(\mathbf{y}|\mathbf{x})p(\mathbf{x}).$$

In detail:

$$\begin{aligned} p(\mathbf{x}|\mathbf{y}) &\propto p(\mathbf{x}) p(\mathbf{y}|\mathbf{x}) \\ &= (2\pi)^{-K/2} \exp\left[-\frac{1}{2}\mathbf{x}^\top \mathbf{x}\right] |2\pi\Psi|^{-1/2} \exp\left[-\frac{1}{2}(\mathbf{y} - \Lambda\mathbf{x})^\top \Psi^{-1}(\mathbf{y} - \Lambda\mathbf{x})\right] \\ &\propto \exp\left(-\frac{1}{2}[\mathbf{x}^\top \mathbf{x} + (\mathbf{y} - \Lambda\mathbf{x})^\top \Psi^{-1}(\mathbf{y} - \Lambda\mathbf{x})]\right) \\ &\propto \exp\left(-\frac{1}{2}[\mathbf{x}^\top (\mathbf{I} + \Lambda^\top \Psi^{-1} \Lambda) \mathbf{x} - 2\mathbf{x}^\top \Lambda^\top \Psi^{-1} \mathbf{y}]\right) \\ &\propto \exp\left(-\frac{1}{2}[\mathbf{x}^\top \Sigma^{-1} \mathbf{x} - 2\mathbf{x}^\top \Sigma^{-1} \mu + \mu^\top \Sigma^{-1} \mu]\right) \propto \mathcal{N}(\mu, \Sigma), \end{aligned}$$

where $\Sigma = (\mathbf{I} + \Lambda^\top \Psi^{-1} \Lambda)^{-1} = \mathbf{I} - \beta \Lambda$, $\mu = \Sigma \Lambda^\top \Psi^{-1} \mathbf{y} = \beta \mathbf{y}$, and $\beta = \Sigma \Lambda^\top \Psi^{-1}$.

Note that μ is a linear function of \mathbf{y} and Σ does not depend on \mathbf{y} .

Learning in the Factor Analysis Model

Integrating over the latent variables, we obtained:

$$p(\mathbf{y}|\Lambda, \Psi) = \int p(\mathbf{x})p(\mathbf{y}|\mathbf{x})d\mathbf{x} = \mathcal{N}(0, \Lambda\Lambda^\top + \Psi).$$

Note that Factor Analysis is a restricted form of Gaussian model.

The unrestricted (zero mean) Gaussian model has a closed form solution

$$p(\mathbf{y}) \sim \mathcal{N}(0, \frac{1}{n} \sum_c \mathbf{y}^{(c)} \mathbf{y}^{(c)\top}),$$

but there is no closed form solution for the Factor Analysis model.

We could try to use the gradients to do maximum likelihood wrt the parameters Λ and Ψ .

It turns out that there is a better algorithm for learning (the EM algorithm).

Eigenvalues and Eigenvectors

λ is an **eigenvalue** and \mathbf{x} is an **eigenvector** of A if:

$$A\mathbf{x} = \lambda\mathbf{x}$$

and \mathbf{x} is a unit vector ($\mathbf{x}^\top \mathbf{x} = 1$).

Interpretation: the operation of A in direction \mathbf{x} is a scaling by λ .


The K Principal Components are the K eigenvectors with the largest eigenvalues of the **data covariance matrix** (i.e. K directions with the largest variance).

Note: Σ can be decomposed:

$$\Sigma = \mathbf{U}\mathbf{S}\mathbf{U}^\top,$$


where \mathbf{S} is $\text{diag}(\sigma_1^2, \dots, \sigma_D^2)$ and \mathbf{U} is an orthonormal matrix.

Example of PCA: Eigenfaces



from <http://vismod.media.mit.edu/vismod/demos/facerec/basic.html>

Principal Components Analysis (PCA)



Noise variable becomes infinitesimal compared to the scale of the data:

$$\Psi = \lim_{\sigma^2 \rightarrow 0} \sigma^2 I$$


Equivalently: reconstruction cost becomes infinite compared to the cost of coding the hidden units under the prior.

$$p(\mathbf{x}|\mathbf{y}) = \mathcal{N}(\beta\mathbf{y}, \mathbf{I} - \beta\Lambda),$$

where

$$\beta = \lim_{\sigma^2 \rightarrow 0} \Lambda^\top (\Lambda\Lambda^\top + \sigma^2 \mathbf{I})^{-1} = (\Lambda^\top \Lambda)^{-1} \Lambda^\top.$$

From Supervised Learning to PCA: linear autoencoder




A linear autoencoder neural network trained to minimise squared error learns to perform PCA (Baldi & Hornik, 1989).

FA vs PCA

- PCA is rotationally invariant; FA is not
- FA is measurement scale invariant; PCA is not
- FA defines a probabilistic model; PCA does not
- PCA can be computed in closed form; FA can not

Jensen's Inequality

For any concave function, such as $\log(x)$



For $\alpha_i \geq 0$, $\sum \alpha_i = 1$ and any $\{x_i > 0\}$

$$\log \left(\sum_i \alpha_i x_i \right) \geq \sum_i \alpha_i \log(x_i)$$

Equality if and only if $\alpha_i = 1$ for some i (and therefore all others are 0).

The Expectation Maximization (EM) algorithm

Given a set of observed (visible) variables V , a set of unobserved (hidden / latent / missing) variables H , and model parameters θ , optimize the log likelihood:

$$\mathcal{L}(\theta) = \log p(V|\theta) = \log \int p(H, V|\theta) dH, \quad (1)$$

where we have written the marginal for the visibles in terms of an integral over the joint distribution for hidden and visible variables.

Using *Jensen's inequality* for **any** distribution of hidden states $q(H)$ we have:

$$\mathcal{L} = \log \int q(H) \frac{p(H, V|\theta)}{q(H)} dH \geq \int q(H) \log \frac{p(H, V|\theta)}{q(H)} dH = \mathcal{F}(q, \theta), \quad (2)$$

defining the $\mathcal{F}(q, \theta)$ functional, which is a **lower bound** on the log likelihood.

In the EM algorithm, we alternately optimize $\mathcal{F}(q, \theta)$ wrt q and θ , and we can prove that this will never decrease \mathcal{L} .

The E and M steps of EM

The lower bound on the log likelihood:

$$\mathcal{F}(q, \theta) = \int q(H) \log \frac{p(H, V|\theta)}{q(H)} dH = \int q(H) \log p(H, V|\theta) dH + \mathcal{H}(q), \quad (3)$$

where $\mathcal{H}(q) = - \int q(H) \log q(H) dH$ is the **entropy** of q . We iteratively alternate:

E step: optimize $\mathcal{F}(q, \theta)$ wrt the distribution over hidden variables given the parameters:


$$q^{(k)}(H) := \operatorname{argmax}_{q(H)} \mathcal{F}(q(H), \theta^{(k-1)}). \quad (4)$$

M step: maximize $\mathcal{F}(q, \theta)$ wrt the parameters given the hidden distribution:

$$\theta^{(k)} := \operatorname{argmax}_{\theta} \mathcal{F}(q^{(k)}(H), \theta) = \operatorname{argmax}_{\theta} \int q^{(k)}(H) \log p(H, V|\theta) dH, \quad (5)$$

which is equivalent to optimizing the expected complete-data likelihood $p(H, V|\theta)$, since the **entropy of $q(H)$** does not depend on θ .

EM as Coordinate Ascent in \mathcal{F}



The EM algorithm never decreases the log likelihood

The difference between the cost functions:


$$\begin{aligned}\mathcal{L}(\theta) - \mathcal{F}(q, \theta) &= \log p(V|\theta) - \int q(H) \log \frac{p(H, V|\theta)}{q(H)} dH \\ &= \log p(V|\theta) - \int q(H) \log \frac{p(H|V, \theta)p(V|\theta)}{q(H)} dH \\ &= - \int q(H) \log \frac{p(H|V, \theta)}{q(H)} dH = \mathcal{KL}(q(H), p(H|V, \theta)),\end{aligned}$$

is called the Kullback-Liebler divergence; it is non-negative and only zero if and only if $q(H) = p(H|V, \theta)$ (thus this is the E step). Although we are working with the **wrong cost function**, the likelihood is still increased in every iteration:

$$\mathcal{L}(\theta^{(k-1)}) \underset{\text{E step}}{=} \mathcal{F}(q^{(k)}, \theta^{(k-1)}) \underset{\text{M step}}{\leq} \mathcal{F}(q^{(k)}, \theta^{(k)}) \underset{\text{Jensen}}{\leq} \mathcal{L}(\theta^{(k)}),$$

where the first equality holds because of the E step, and the first inequality comes from the M step and the final inequality from Jensen. Usually EM converges to a local optimum of \mathcal{L} (although there are exceptions).

EM for Factor Analysis



The model for \mathbf{y} :

$$p(\mathbf{y}|\theta) = \int p(\mathbf{x}|\theta)p(\mathbf{y}|\mathbf{x}, \theta)d\mathbf{x} = \mathcal{N}(0, \Lambda\Lambda^\top + \Psi)$$

Model parameters: $\theta = \{\Lambda, \Psi\}$.

E step: For each data point \mathbf{y}_c , compute the posterior distribution of hidden factors given the observed data: $q_c(\mathbf{x}_c) = p(\mathbf{x}_c|\mathbf{y}_c, \theta^{(t)})$.

M step: Find the $\theta^{(t+1)}$ that maximises $\mathcal{F}(q, \theta)$:

$$\begin{aligned}\mathcal{F}(q, \theta) &= \sum_c \int q_c(\mathbf{x}_c) \left[\log p(\mathbf{x}_c|\theta) + \log p(\mathbf{y}_c|\mathbf{x}_c, \theta) - \log q_c(\mathbf{x}_c) \right] d\mathbf{x}_c \\ &= \sum_c \int q_c(\mathbf{x}) \left[\log p(\mathbf{x}_c|\theta) + \log p(\mathbf{y}_c|\mathbf{x}_c, \theta) \right] d\mathbf{x} + \text{const.}\end{aligned}$$

The M step for Factor Analysis

M step: Find $\theta^{(t+1)}$ maximising $\mathcal{F} = \sum_c \int q_c(\mathbf{x}) [\log p(\mathbf{x}|\theta) + \log p(\mathbf{y}_c|\mathbf{x}, \theta)] d\mathbf{x}$.

$$\log p(\mathbf{x}|\theta) + \log p(\mathbf{y}_c|\mathbf{x}, \theta)$$

$$\begin{aligned} &= \text{const} - \frac{1}{2} \mathbf{x}_c^\top \mathbf{x}_c - \frac{1}{2} \log |\Psi| - \frac{1}{2} (\mathbf{y}_c - \Lambda \mathbf{x}_c)^\top \Psi^{-1} (\mathbf{y}_c - \Lambda \mathbf{x}_c) \\ &= \text{const}' - \frac{1}{2} \log |\Psi| - \frac{1}{2} [\mathbf{y}_c^\top \Psi^{-1} \mathbf{y}_c - 2 \mathbf{y}_c^\top \Psi^{-1} \Lambda \mathbf{x}_c + \mathbf{x}_c^\top \Lambda^\top \Psi^{-1} \Lambda \mathbf{x}_c] \\ &= \text{const}' - \frac{1}{2} \log |\Psi| - \frac{1}{2} [\mathbf{y}_c^\top \Psi^{-1} \mathbf{y}_c - 2 \mathbf{y}_c^\top \Psi^{-1} \Lambda \mathbf{x}_c + \text{trace} \Lambda^\top \Psi^{-1} \Lambda \mathbf{x}_c \mathbf{x}_c^\top] \end{aligned}$$

Taking expectations over $q_c(\mathbf{x}_c)$...

$$= \text{const}' - \frac{1}{2} \log |\Psi| - \frac{1}{2} [\mathbf{y}_c^\top \Psi^{-1} \mathbf{y}_c - 2 \mathbf{y}_c^\top \Psi^{-1} \Lambda \mu_c + \text{trace} \Lambda^\top \Psi^{-1} \Lambda (\mu_c \mu_c^\top + \Sigma)].$$

Note that we don't need to know everything about q , just the expectations of \mathbf{x} and $\mathbf{x}\mathbf{x}^\top$ under q (i.e. the expected sufficient statistics).

The M step for Factor Analysis (cont.)

$$\mathcal{F} =$$

$$\text{const}' - \frac{N}{2} \log |\Psi| - \frac{1}{2} \sum_c [y_c^\top \Psi^{-1} y_c - 2 y_c^\top \Psi^{-1} \Lambda \mu_c + \text{trace}(\Lambda^\top \Psi^{-1} \Lambda (\mu_c \mu_c^\top + \Sigma))]$$

Taking derivatives w.r.t. Λ and Ψ^{-1} , using $\frac{\partial \text{trace} AB}{\partial B} = A^\top$ and $\frac{\partial \log |A|}{\partial A} = A^{-\top}$:

$$\frac{\partial \mathcal{F}}{\partial \Lambda} = \Psi^{-1} \sum_c y_c \mu_c^\top - \Psi^{-1} \Lambda \left(N \Sigma + \sum_c \mu_c \mu_c^\top \right) = 0$$

$$\Lambda = \left(N \Sigma + \sum_c \mu_c \mu_c^\top \right)^{-1} \sum_c y_c \mu_c^\top$$

$$\frac{\partial \mathcal{F}}{\partial \Psi^{-1}} = \frac{N}{2} \Psi - \frac{1}{2} \sum_c [y_c y_c^\top - \Lambda \mu_c y_c^\top - y_c \mu_c^\top \Lambda^\top + \Lambda (\mu_c \mu_c^\top + \Sigma) \Lambda^\top]$$

$$\Psi = \frac{1}{N} \sum_c [y_c y_c^\top - \Lambda \mu_c y_c^\top - y_c \mu_c^\top \Lambda^\top + \Lambda (\mu_c \mu_c^\top + \Sigma) \Lambda^\top]$$


$$= \Lambda \Sigma \Lambda^\top + \frac{1}{N} \sum_c (y_c - \Lambda \mu_c)(y_c - \Lambda \mu_c)^\top \quad (\text{squared residuals})$$

Note: we should actually only take derivatives w.r.t. Ψ_{dd} since Ψ is diagonal.
When $\Sigma \rightarrow 0$ these become the equations for linear regression!


Limitations and Mixtures

So far, we have assumed that the data is reasonably well approximated by a Gaussian.


Uniform



Triangle



Heavy tails



A mixture distribution has a single discrete latent variable:

$$s_i \sim \text{Discrete}[\pi] \quad y_i|s_i \sim p(y_i|\theta_{s_i})$$


Mixtures arise naturally when observations from different sources have been collated.

They can also be used to *approximate* arbitrary distributions.

Clustering with MoG



Clustering with MoG



Likelihood for the Mixture of Gaussians model

Likelihood:

$$\begin{aligned} p(y_1, \dots, y_n | \theta) &= \prod_{c=1}^n \sum_{i=1}^k \pi_i \mathcal{N}(y_c | \mu_i, \Sigma_i) \\ &= \prod_{c=1}^n \sum_{i=1}^k \pi_i |2\pi\Sigma_i|^{-1/2} \exp \left(- (y_c - \mu_i)^\top \Sigma_i^{-1} (y_c - \mu_i) \right). \end{aligned}$$

Here, π_i are the **mixing proportions**, (non-negative, sum to one) and the parameters are collected in $\theta = (\pi, \mu_1, \dots, \mu_k, \Sigma_1, \dots, \Sigma_k)$.

Mixture of Gaussians, E-step

The likelihood (for a single data point) in the mixture of Gaussians model is

$$p(y|\theta) = \sum_{j=1}^k p(y, h=j|\theta) = \sum_{j=1}^k p(h=j|\theta)p(y|h=j, \theta),$$

where $\theta = (\pi_1, \dots, \pi_k, \mu_1, \dots, \mu_k, \Sigma_1, \dots, \Sigma_k)$ are the parameters, and h is the hidden, or latent, variable.

In the E-step we maximize the lower bound functional $\mathcal{F}(q(H), \theta)$ wrt $q(H)$ for fixed θ . We have seen that this is equivalent to setting $q(H)$ equal to the posterior:

$$q(H) = p(h|y, \theta) = \frac{p(y|h, \theta)p(h|\theta)}{p(y|\theta)} \propto p(y|h, \theta)p(h|\theta).$$

Thus, the *responsibilities* are:

$$r_j = p(h=j|y, \theta) \propto \pi_j |\Sigma_j|^{1/2} \exp \left(- (y - \mu_j)^\top \Sigma_j^{-1} (y - \mu_j) / 2 \right),$$

normalized such that $\sum_j r_j = 1$. For multiple data points, responsibilities are computed analogously: $r_{cj} = p(h_c=j|y_c, \theta)$.

Mixture of Gaussians, M-step

In the M-step, we maximize the lower bound functional $\mathcal{F}(q(H), \theta)$ wrt. θ . Recall, that this is equivalent to maximizing

$$E(\theta) = \int q(H) \log(p(H, V|\theta)) dH = \int q(H) \sum_c \log(p(H_c, V_c|\theta)) dH,$$

for a fixed $q(H)$. For mixtures of Gaussians:

$$E(\theta) = \sum_{c,j} r_{cj} \left[\log(\pi_j) - \frac{1}{2} \log |\Sigma_j| - (\mathbf{y}_c - \boldsymbol{\mu}_j)^\top \boldsymbol{\Sigma}_j^{-1} (\mathbf{y}_c - \boldsymbol{\mu}_j) / 2 \right].$$

Mixture of Gaussians, M-step, cont.

Optimizing wrt the θ parameters (and $S_j = \Sigma_j^{-1}$):

$$\begin{aligned}\frac{\partial E(\theta)}{\partial \mu_j} &= \sum_c r_{cj} \Sigma_j^{-1} (\mathbf{y}_c - \boldsymbol{\mu}) = \mathbf{0} \implies \boldsymbol{\mu}_j = \sum_c r_{cj} \mathbf{y}_c / (\sum_c r_{cj}) \\ \frac{\partial E(\theta)}{\partial S_j} &= \frac{1}{2} \sum_c r_{cj} \left[S_j^{-1} - (\mathbf{y}_c - \boldsymbol{\mu}_j)(\mathbf{y}_c - \boldsymbol{\mu}_j)^\top \right] = \mathbf{0} \\ &\implies \Sigma_j = \sum_{c=1}^n r_{cj} (\mathbf{y}_c - \boldsymbol{\mu}_j)(\mathbf{y}_c - \boldsymbol{\mu}_j)^\top / (\sum_c r_{cj}), \\ \frac{\partial E(\theta)}{\partial \pi_j} &= \sum_c \frac{r_{cj}}{\pi_j} - \lambda = 0 \implies \pi_j = \sum_c r_{cj} / (\sum_{c,j} r_{cj}),\end{aligned}$$

where λ is the Lagrange multiplier which ensures that the mixture normalizes $\sum_j \pi_j = 1$.

Mixture of Gaussian issues

The EM algorithm converges to a *local* maximum of the likelihood. There may be many local maxima.

There could be many bad local maxima, ie with low values of the likelihood.

In fact, we are not interested in the global optimum!

Another problem: “How many mixture components?”