TU München
Fakultät für Informatik
PD Dr. Rudolph Triebel
John Chiotellis, Maximilian Denninger

Machine Learning for Computer Vision

January 23, 2019
Topic: Clustering

Exercise 1: K-Means Compression (Programming)

a) K-Means finds the parameters (cluster means) that minimize the assignment cost:

$$\boldsymbol{\theta}^* = \arg\min_{\boldsymbol{\theta}} \sum_{i=1}^{N} \sum_{j=1}^{K} A_{ij} ||\boldsymbol{x}_i - \boldsymbol{\mu}_j||^2 \,, \tag{1}$$

where $\boldsymbol{\theta} = \{\boldsymbol{\mu}_k : 1 \leq k \leq K\}$ and A_{ij} is 1 if x_i is assigned to cluster j and 0 otherwise.

b) See code.

Exercise 2: Expectation-Maximization for GMM (Programming)

a) EM finds the parameters (means, covariances and mixture coefficients) that maximize the conditional data log-likelihood:

$$\boldsymbol{\theta}^* = \arg\max_{\boldsymbol{\theta}} \log p(X|\boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\pi}) \tag{2}$$

where $\boldsymbol{\theta} = \{\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k, \boldsymbol{\pi}_k : 1 \leq k \leq K\}$

b) See code.

Exercise 3: Expectation-Maximization for GMM

In the standard EM algorithm, we first define the responsibilities γ as

$$\gamma_{nk} = p(z_{nk} = 1|x_n) = \frac{\pi_k \mathcal{N}(x_n|\mu_k, \Sigma_k)}{\sum_{j=1}^{K} \pi_j \mathcal{N}(x_n|\mu_j, \Sigma_j)} \quad , z_{nk} \in \{0, 1\}, \sum_{k=1}^{K} z_{nk} = 1$$

a) Find the optimal means, covariances and mixing coefficients that maximize the data likelihood. How can we interpret the results?

We want to maximize the data likelihood, so as usual we minimize the negative log-likelihood:

$$-\mathcal{LL} = -\log p(X|\boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\pi}) = -\log \prod_n \sum_k \pi_k \mathcal{N}(x_n|\mu_k, \Sigma_k) \tag{3}$$

This time we minimize 3 times independently with respect to the means, the covariances and the mixture coefficients:

$$\mu_k^* = \arg\min_{\mu_k} -\mathcal{LL} \tag{4}$$
$$\Sigma_k^* = \arg\min_{\Sigma_k} -\mathcal{LL} \tag{5}$$
$$\pi_k^* = \arg\min_{\pi_k} -\mathcal{LL} \tag{6}$$

In the following, to avoid confusion of sums and covariances, we denote covariance Σ_k as C_k. To simplify some expressions, let us agree on the following notation:

$$\mathcal{N}_{nk} \equiv \mathcal{N}(x_n|\mu_k, C_k) \tag{7}$$
$$Z_k \equiv ((2\pi)^d |C_k|)^{1/2} \tag{8}$$
$$D_{nk} \equiv (x_n - \mu_k)^T C_k^{-1}(x_n - \mu_k) \tag{9}$$
$$\text{Therefore } \mathcal{N}_{nk} = Z_k^{-1} \exp\{-\frac{1}{2} D_{nk}\} \tag{10}$$

Thus, we have:

$$\begin{aligned} -\mathcal{LL} &= -\sum_n \log \sum_k \pi_k \mathcal{N}_{nk} \\ &= -\sum_n \log \sum_k \pi_k Z_k^{-1} \exp(-\frac{1}{2} D_{nk}) \end{aligned}$$

Solving for the means:

$$\frac{\partial \mathcal{LL}}{\partial \mu_k} = \sum_n \frac{1}{\sum_j \pi_j \mathcal{N}_{nj}} \frac{\partial \sum_k \pi_k Z_k^{-1} \exp(-\frac{1}{2} D_{nk})}{\partial \mu_k} \tag{11}$$
$$= \sum_n \frac{1}{\sum_j \pi_j \mathcal{N}_{nj}} \pi_k Z_k^{-1} \frac{\partial \exp(-\frac{1}{2} D_{nk})}{\partial \mu_k} \tag{12}$$
$$= \sum_n \frac{1}{\sum_j \pi_j \mathcal{N}_{nj}} \pi_k Z_k^{-1} \exp(-\frac{1}{2} D_{nk}) C_k^{-1}(x_n - \mu_k) \tag{13}$$
$$= \sum_n \frac{\pi_k \mathcal{N}_{nk}}{\sum_j \pi_j \mathcal{N}_{nj}} C_k^{-1}(x_n - \mu_k) \tag{14}$$
$$= \sum_n \gamma_{nk} C_k^{-1}(x_n - \mu_k) \tag{15}$$
$$\tag{16}$$

Setting $-\frac{\partial \mathcal{LL}}{\partial \mu_k} \stackrel{!}{=} 0$ gives us:

$$\sum_n \gamma_{nk} C_k^{-1} \mu_k = \sum_n \gamma_{nk} C_k^{-1} x_n \tag{17}$$

$$C_k^{-1} \mu_k \sum_n \gamma_{nk} = C_k^{-1} \sum_n \gamma_{nk} x_n \tag{18}$$

$$C_k^{-1} \mu_k \sum_n \gamma_{nk} = C_k^{-1} \sum_n \gamma_{nk} x_n \tag{19}$$

$$\mu_k \sum_n \gamma_{nk} = \sum_n \gamma_{nk} x_n \tag{20}$$

$$\mu_k = \frac{\sum_n \gamma_{nk} x_n}{\sum_n \gamma_{nk}} \tag{21}$$

Solving for the covariances:

$$\frac{\partial \mathcal{LL}}{\partial C_k} = \sum_n \frac{1}{\sum_j \pi_j \mathcal{N}_{nj}} \frac{\partial \sum_k \pi_k Z_k^{-1} \exp(-\frac{1}{2} D_{nk})}{\partial C_k} \tag{22}$$

$$= \sum_n \frac{1}{\sum_j \pi_j \mathcal{N}_{nj}} \pi_k \frac{\partial Z_k^{-1} \exp(-\frac{1}{2} D_{nk})}{\partial C_k} \tag{23}$$

$$= \sum_n \frac{1}{\sum_j \pi_j \mathcal{N}_{nj}} \pi_k \left(\frac{\partial Z_k^{-1}}{\partial C_k} \exp(-\frac{1}{2} D_{nk}) + Z_k^{-1} \frac{\partial \exp(-\frac{1}{2} D_{nk})}{\partial C_k} \right) \tag{24}$$

$$= \sum_n \frac{1}{\sum_j \pi_j \mathcal{N}_{nj}} \pi_k \left((-\frac{1}{2} Z_k^{-1} C_k^{-1}) \exp(-\frac{1}{2} D_{nk}) + \frac{1}{2} Z_k^{-1} \exp(-\frac{1}{2} D_{nk}) C_k^{-1} (x_n - \mu_k)(x_n - \mu_k)^T C_k^{-1} \right) \tag{25}$$

$$= (-\frac{1}{2}) \sum_n \frac{1}{\sum_j \pi_j \mathcal{N}_{nj}} \pi_k Z_k^{-1} \exp(-\frac{1}{2} D_{nk}) \left(C_k^{-1} - C_k^{-1} (x_n - \mu_k)(x_n - \mu_k)^T C_k^{-1} \right) \tag{26}$$

$$= (-\frac{1}{2}) \sum_n \gamma_{nk} \left(C_k^{-1} - C_k^{-1} (x_n - \mu_k)(x_n - \mu_k)^T C_k^{-1} \right) \tag{27}$$

$$\tag{28}$$

Here, we used the derivative of the determinant as follows:

$$\frac{\partial Z_k^{-1}}{\partial C_k} = \frac{\partial ((2\pi)^d |C_k|)^{-\frac{1}{2}}}{\partial C_k} = ((2\pi)^d)^{-\frac{1}{2}} \frac{\partial (|C_k|)^{-\frac{1}{2}}}{\partial C_k} \tag{29}$$

$$= ((2\pi)^d)^{-\frac{1}{2}} (-\frac{1}{2}) |C_k|^{-\frac{3}{2}} \frac{\partial (|C_k|)}{\partial C_k} = ((2\pi)^d)^{-\frac{1}{2}} (-\frac{1}{2}) |C_k|^{-\frac{3}{2}} |C_k| (C_k^{-1})^T \tag{30}$$

$$= (-\frac{1}{2}) ((2\pi)^d)^{-\frac{1}{2}} |C_k|^{-\frac{1}{2}} C_k^{-1} = -\frac{1}{2} Z_k^{-1} C_k^{-1} \tag{31}$$

and the derivative of the Mahalanobis distance as:

$$\frac{\partial x^T C^{-1} x}{\partial C} = -C^{-T} x x^T C^{-T} = -C^{-1} x x^T C^{-1} \tag{32}$$

Setting $-\frac{\partial \mathcal{LL}}{\partial C_k} \stackrel{!}{=} 0$ gives us:

$$\sum_n \gamma_{nk} C_k^{-1} = \sum_n \gamma_{nk} C_k^{-1}(x_n - \mu_k)(x_n - \mu_k)^T C_k^{-1} \tag{33}$$

$$C_k^{-1} \sum_n \gamma_{nk} = C_k^{-1} \sum_n \gamma_{nk}(x_n - \mu_k)(x_n - \mu_k)^T C_k^{-1} \tag{34}$$

$$\sum_n \gamma_{nk} = \sum_n \gamma_{nk}(x_n - \mu_k)(x_n - \mu_k)^T C_k^{-1} \tag{35}$$

$$C_k = \frac{\sum_n \gamma_{nk}(x_n - \mu_k)(x_n - \mu_k)^T}{\sum_n \gamma_{nk}} \tag{36}$$

Solving for the mixture coefficients: Here we must take into account that $\sum_k \pi_k = 1$. We enforce this constraint with a Lagrange multiplier. Our objective then becomes:

$$\mathcal{LL}' = \mathcal{LL} + \lambda(\sum_k \pi_k - 1) \tag{37}$$

where $\lambda < 0$.

Deriving w.r.t. π_k, we get

$$\frac{\partial \mathcal{LL}'}{\partial \pi_k} = \sum_n \frac{1}{\sum_j \pi_j \mathcal{N}_{nj}} \frac{\partial \sum_k \pi_k \mathcal{N}_{nk}}{\partial \pi_k} + \lambda \tag{38}$$

$$= \sum_n \frac{1}{\sum_j \pi_j \mathcal{N}_{nj}} \mathcal{N}_{nk} + \lambda \tag{39}$$

$$= \sum_n \frac{\gamma_{nk}}{\pi_k} + \lambda \tag{40}$$

Setting equal to zero and solving for λ, we get

$$\lambda = -\sum_n \frac{\gamma_{nk}}{\pi_k} \tag{41}$$

$$\lambda \pi_k = -\sum_n \gamma_{nk} \tag{42}$$

$$\sum_k \lambda \pi_k = -\sum_k \sum_n \gamma_{nk} \tag{43}$$

$$\lambda = -N \tag{44}$$

Now we can plug this back to the objective and actually solve for π_k:

$$\frac{\partial \mathcal{LL}'}{\partial \pi_k} = \sum_n \frac{\gamma_{nk}}{\pi_k} - N \stackrel{!}{=} 0 \tag{45}$$

$$\frac{1}{\pi_k} \sum_n \gamma_{nk} = N \tag{46}$$

$$\pi_k = \frac{\sum_n \gamma_{nk}}{N} = \frac{N_k}{N} \tag{47}$$

We can interpret these results as weighted averages of means and covariances, the weights corresponding to the responsibilities γ_{nk}. The mixture coefficients π_k are simply the ratio of data points explained by each component.

b) Define the complete-data-log-likelihood. What is the difference to the standard log-likelihood?

Assuming we observe not only the data but also the binary latent variables Z we define the complete data likelihood as:

$$p(X, Z|\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{C}) = \prod_n p(z_n|\pi)p(x_n|z_n, \boldsymbol{\mu}, \boldsymbol{C}) \tag{48}$$

where $p(z_n|\boldsymbol{\pi}) = \prod_k \pi_k^{z_{nk}}$ and $p(x_n|z_n, \boldsymbol{\mu}, \boldsymbol{C}) = \prod_k \mathcal{N}(x_n|\mu_k, C_k)^{z_{nk}}$.

Remember that $\sum_k z_{nk} = 1$.

Since now we only have products, we can more easily compute the logarithm:

$$\log p(X, Z|\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{C}) = \sum_n \sum_k z_{nk}(\log \pi_k + \log \mathcal{N}(x_n|\mu_k, C_k)) \tag{49}$$

Of course in practice, the latent variables are not known, so we maximize the *expectation*:

$$\mathbb{E}[\log p(X, Z|\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{C})] = \sum_n \sum_k \mathbb{E}[z_{nk}](\log \pi_k + \log \mathcal{N}(x_n|\mu_k, C_k)) \tag{50}$$

where we know that $\mathbb{E}[z_{nk}] = \gamma_{nk}$.

The theory says that the log-marginal is also maximized implicitly!