

Machine Learning for Computer Vision

June 15, 2018
Topic: Clustering

Exercise 1: K-Means Compression (Programming)

- a) Which objective function does k-means try to optimize?
- b) Download the *clustering.zip* found on the website and extract its contents. Implement K-Means and use it to compress the *mona-lisa.jpg* image, by using only $2, 3, \dots, 10$ colors.

Exercise 2: Expectation-Maximization for GMM (Programming)

- a) Which objective function does EM for GMM try to optimize?
- b) Implement the Expectation-Maximization algorithm (EM) for Gaussian Mixture Models (GMM) to cluster the *old faithful* dataset.

Exercise 3: Expectation-Maximization for GMM

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

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

- a) Find the optimal means, covariances and mixing coefficients that *maximize the data likelihood*. How can you interpret the results?
- b) Define the complete-data-log-likelihood. What is the difference to the standard log-likelihood?