

Weekly Exercise 4

Dr. Csaba Domokos and Lingni Ma
 Technische Universität München, Computer Vision Group
 April 26, 2016 (submission deadline: May 03, 2016)

The EM algorithm for mixtures of Gaussians (2 points)

Exercise 1 (M step for Σ , 2 points). Assuming n data samples $\{\mathbf{x}_n\}_{n=1}^N$, consider the log-likelihood function of a mixture of Gaussian model with K components

$$\mathcal{L}(\boldsymbol{\theta}) = \sum_{n=1}^N \sum_{k=1}^K \gamma_k(\mathbf{x}_n) (\ln \pi_k + \ln \mathcal{N}(\mathbf{x}_n | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)) .$$

Show that the optimal choice with respect to the covariance matrices $\boldsymbol{\Sigma}_k$ for all $k = 1, \dots, K$ is given as

$$\operatorname{argmax}_{\boldsymbol{\Sigma}_k} \mathcal{L}(\boldsymbol{\theta}) = \frac{\sum_{n=1}^N \gamma_k(\mathbf{x}_n) (\mathbf{x}_n - \boldsymbol{\mu}_k) (\mathbf{x}_n - \boldsymbol{\mu}_k)^T}{\sum_{m=1}^N \gamma_k(\mathbf{x}_m)} .$$

(Hint: for a symmetric $\mathbf{X} \in \mathbb{R}^{n \times n}$ matrix and vector $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$,


$$\frac{\partial}{\partial \mathbf{X}} \mathbf{a}^T \mathbf{X}^{-1} \mathbf{b} = -\mathbf{X}^{-T} \mathbf{a} \mathbf{b}^T \mathbf{X}^{-T} ,$$

and for a non-singular matrix $\mathbf{Y} \in \mathbb{R}^{n \times n}$,

$$\frac{\partial}{\partial \mathbf{Y}} |\mathbf{Y}| = |\mathbf{Y}| \mathbf{Y}^{-1} .$$

Graphical models (3 points)

Exercise 2 (Bayesian network, 1 point). Provide the factorization of $p(y_1, y_2, y_3, y_4, y_5)$ according to the following directed graphical model:



Exercise 3 (Proof of the Hammersley-Clifford theorem Σ , 2 points). Assume an undirected graphical model $G = (\mathcal{V}, \mathcal{E})$ satisfying the local Markov property. Consider two non-connected nodes $a, b \in \mathcal{V}$, i.e. $(a, b) \notin \mathcal{E}$, and a subset of nodes $w \subset \mathcal{V}$ such that $a, b \notin w$. Show that

$$q(y_a | \mathbf{y}_w) \stackrel{\Delta}{=} p(y_a | \mathbf{y}_w, \mathbf{y}_{\mathcal{V} \setminus (w \cup \{a\})}^*) = q(y_a | y_b, \mathbf{y}_w),$$

where $p(\mathbf{y}_z, \mathbf{y}_z^*)$ is a joint probability.

Programming (6 points)

Exercise 4 (Gaussian Mixture Model Estimation, 6 points). Train two Gaussian mixture models to model the foreground and background probabilities of the pixels based on their intensity. Apply the trained model to the input image and segment the foreground and background. The testing image is shown in Figure 1 (a) (you can download the images from supp_04.zip). The specific requirement is as follows.




Figure 1: (a) test image. (b) bounding box for foreground training. (c) foreground segmentation based the trained model $p_F(I)$.

1. Write a program to take in argument (x_1, y_1, x_2, y_2) , which specify the bounding box of the foreground as shown in in Figure 1 (b). Model the foreground probability $p_F(I), I \in \mathbb{R}^3$ using all pixels inside the bounding box, where $p_F(I)$ is a GMM model with 5 Gaussian components and I is the RGB values of the pixel.
2. Based on the same bounding box, model the background probability $p_B(I), I \in \mathbb{R}^3$ using all pixels outside the bounding box, where $p_B(I)$ is a GMM model with 5 Gaussian components.
3. Segment the input image into foreground and background use the two distributions you obtained, respectively. Compare the results.

Hints: the bounding box should mostly contain the banana. You can initialize the GMM with random Gaussian kernels. Note that the covariance matrix should not be singular, during any time of GMM training. If the covariance matrix does become singular, you can restart the estimation from a different initialization all over again. Alternatively, can you think of way to better initialize the GMM?