Assignment (COSC6328 W08)

Due: in class on Feb. 11, 2008

You have to work individually. You are not allowed to view or exchange documents with your peers. Hand in a hardcopy of your answers before the deadline. No late submission will be accepted. No handwritting is accepted. Direct your queries to Hui Jiang (hj@cse.yorku.ca).

1. Assume we have a random vector $\mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}$ which follows a bivariate Gaussian distribution: $\mathcal{N}(\mathbf{x}|\mu, \Sigma)$, where $\mu = \begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}$ is the mean vector and $\Sigma = \begin{pmatrix} \sigma_1^2 & \rho\sigma_1\sigma 2 \\ \rho\sigma_1\sigma 2 & \sigma_2^2 \end{pmatrix}$ is the covariance matrix. Derive the formula to compute mutual information between x_1 and x_2, i.e., $I(x_1, x_2)$.

   Hints: Refer to the related sections in the reading assignment [B1]. Note there is a mistake in [B1]: equation (99) should be

   $$\Gamma(n+1) = n \cdot \Gamma(n).$$

2. In many pattern classification problems, one has the option either to assign the pattern to one of N classes, or to *reject* it as being unrecognizable. If the cost for rejection is not too high, rejection may be a desirable action. If we observe feature $\mathbf{x}$ of a pattern (assume its true class id is ω_i), let's define the loss function for all actions α_j as:

   $$\lambda(\alpha_j|\omega_i) = \begin{cases} 0 & : \quad j = i \ \text{(correct classification)} \\ \lambda_s & : \quad j \neq i \ \text{and} \ 1 \leq j \leq N \text{(wrong classification)} \\ \lambda_r & : \quad \text{rejection} \end{cases}$$

   where λ_s is the loss incurred for making any a wrong classification decision, and λ_r is the loss incurred for choosing the rejection action. Show the minimum risk is obtained by the following decision rule: we decide ω_i if $p(\omega_i|\mathbf{x}) \geq p(\omega_j|\mathbf{x})$ for all j and if $p(\omega_i|\mathbf{x}) \geq 1 - \lambda_r/\lambda_s$, and reject otherwise. What happens if $\lambda_r = 0$? What happens if $\lambda_r > \lambda_s$?

   (Hint: consider the average loss for each action.)

3. Assume we have c different classes, $\omega_1, \omega_2, \cdots, \omega_c$. Each class $\omega_i \ \ (i = 1, 2, \cdots, c)$ is modeled by a univariate Gaussian distribution with mean μ_i and variance σ, i.e., $p(x \mid \omega_i) = \mathcal{N}(x \mid \mu_i, \sigma^2)$, where σ is a common variance for all c classes. Suppose we have collected n data samples from these c classes, i.e., $\{x_1, x_2, \cdots, x_n\}$, and let $\{l_1, l_2, \cdots, l_n\}$ be their labels so that $l_k = i$ means the data sample x_k comes from the i-th class, ω_i.

   Based on the given data set, derive the maximum-likelihood estimates for all model parameters, i.e., all means $\mu_i \ (i = 1, 2, \cdots, c)$ and the common variance σ.