CS 688 Pattern Recognition

Midterm Exam (Fall 2010)

Instructor: Carlotta Domeniconi

October 19, 2010

Student's name:

This test is governed by the GMU Honor Code. The paper you turn in must be your sole work. Help may be obtained from the instructor to understand the description of the problem, but the solution must be the student's own work. Any deviation from this is considered a Honor Code violation.

.

1. [25 points]
(a) Write the mathematical expression of the sample covariance matrix for N points in a two dimensional space. Write each component of the matrix, and define any quantity involved.

(b) Compute the sample covariance matrix for the following four two-dimensional points:
$\mathbf{a} = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$, $\mathbf{b} = \begin{pmatrix} 3 \\ 1 \end{pmatrix}$, $\mathbf{c} = \begin{pmatrix} 1 \\ 5 \end{pmatrix}$, $\mathbf{d} = \begin{pmatrix} 3 \\ 5 \end{pmatrix}$.

2. [25 points]
Consider a two-class classification problem in two dimensions with
$p(\mathbf{x}|c_1) \sim N(\mathbf{x}|\mathbf{0}, \mathbf{I})$, $p(\mathbf{x}|c_2) \sim N(\mathbf{x}|\mathbf{1}, \mathbf{I})$, $P(c_1) = P(c_2) = 1/2$,
where:
$\mathbf{0} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$, $\mathbf{1} = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$, $\mathbf{I} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$, and N is the Gaussian density function.

(a) Calculate the Bayes decision boundary. Show your calculations.

Note: The Gaussian distribution defined over a D-dimensional vector $\mathbf{x}$ of continuous variables is given by:

$$N(\mathbf{x}|\boldsymbol{\mu}, \Sigma) = \frac{1}{(2\pi)^{D/2}} \frac{1}{|\Sigma|^{1/2}} e^{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^T \Sigma^{-1} (\mathbf{x}-\boldsymbol{\mu})}$$

.

3. [20 points]
Describe (in English) the criterion optimized by *linear discriminant analysis* (LDA). Be precise and concise. What are the fundamental characteristics of LDA?

4. [30 points]
Consider the estimation $\hat{f}$ of an unknown function f. Suppose $\hat{f}$ is computed using a given training data set. In general, the mean squared error corresponding to the estimated value $\hat{f}$ can be written as follows:

$$E[(\hat{f} - f)^2] = E[(\hat{f} - E(\hat{f}))^2] + (E(\hat{f}) - f)^2 \tag{1}$$

that is, the mean squared error can be written as the sum of the *variance* and the squared *bias*.

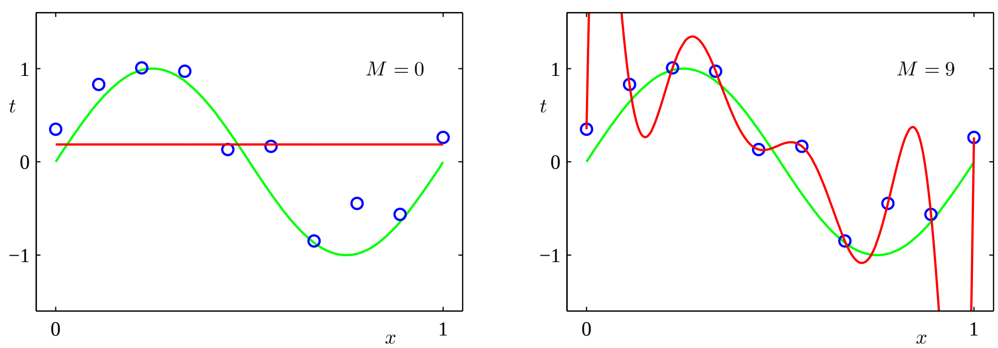

(a) (b)

(1) Derive equation (1) above.
(2) Explain (in English) what the variance and bias measure.
(3) Comment on the variance and bias components of the error corresponding to the fitted function $\hat{f}$ in Figures (a) and (b) above (f corresponds to the sinusoidal function, and M is the order of the polynomial fitted to the data). Motivate your answer.