Lahore University Of Management Sciences
BSc (Honours) Programme
MS Computer Science Programme

Roll # _____

Course Title	Pattern Recognition	**Quarter**	Spring
Course Code	CS533/CmpE533	**Academic Year**	2004-2005
Instructor	Sohaib A. Khan	**Date**	14-04-2005
Exam	Midterm	**Time Allowed**	75 minutes
		Total Marks	75 (25% of grade)

DO NOT OPEN THIS EXAM UNTIL TOLD TO DO SO.

The instructions below must be followed strictly. Failure to do so can result in serious grade loss.

⇒ *You may not*
- *talk to anyone once the exam begins.*
- *leave the examination room and then return.*

⇒ *Keep your eyes on your own paper.*

⇒ *Read all questions very carefully before answering them.*

Specific instructions:

1. Closed book / closed notes / no help sheet
2. Calculator usage: Calculator allowed
3. Any other instruction(s): **JUSTIFY each answer. Points will be deducted for ambiguous reasoning. If the solution of a problem is based on the solution of a previous problem that you were unable to attempt correctly, you may assume a non-trivial answer for the previous problem to get partial credit.**

1. [10 points] The following three matrices are given to you:

$$\sum_1 = \begin{bmatrix} 4 & 3 \\ 3 & 9 \end{bmatrix}, \sum_2 = \begin{bmatrix} 4 & -6 \\ -6 & 9 \end{bmatrix}, \sum_3 = \begin{bmatrix} 4 & 10 \\ 10 & 9 \end{bmatrix}$$

   Are each of these valid covariance matrices? If these denote the covariance of normally distributed data, what can you infer about the shape of the distribution, (without computing the eigenvalues and eigenvectors of the matrices)?

2. [15 points] Find the minimax decision boundary for a two category problem where $p(x|\omega_i)$ is $N(\mu_i, \sigma^2)$, i.e. both class conditional densities are 1-D normal densities with equal variance but different means. Assume a zero-one loss function in your solution.

3. [10 points] Consider the following experiment: I draw 100 samples each from two distributions, $N(\mu_1, \sigma_1^2)$ and $N(\mu_2, \sigma_2^2)$. Since the distributions are known to me, I compute the Chernoff Bound on error, using the values of μ_1, σ_1^2, μ_2 and σ_2^2. Next, I give these samples to my friend, without telling him the parameters. He estimates the parameters using MLE, and constructs a Bayes classifier based on his estimates. Is it possible that the misclassification rate that he reports on data is more than the Chernoff Bound that I have computed? Provide justification for your answer.

For the next three questions, use the following information.
$X_1, X_2, \ldots X_n$ are i.i.d. discrete random samples drawn from Poisson(λ) population.

4. [10 points] Compute the MLE of λ

5. [15 points] If the prior distribution of λ is known to be Gamma(α, β), find the Bayes Point Estimator of λ, and write is as a ratio of prior mean and the MLE.

6. [15 points] Since the variance of a Poisson random variable is also λ, the sample variance $S_n = \frac{1}{n-1}\sum_{i=1}^{n}(X_i - \overline{X})^2$ may also be used as an estimator of λ. Compute the bias of this estimator.