Name: ____

Problem	Points	Score
1(a)	20	
1(b)	10	
1(c)	5	
2	35	
3(a)	25	
3(b)	5	
Total	100	

Notes:

(1) The exam is closed books and notes except for one double-sided sheet of notes.

(2) Please indicate clearly your answer to the problem.

(3) If I can't read or follow your solution, it is wrong and no partial credit will be awarded.

Problem No. 1: Consider a two-class discrete distribution problem:

$$\omega_1: \{[0,0], [2,0], [2,2], [0,2]\}$$
$$\omega_2: \{[1,1], [2,1], [1,2], [3,3]\}$$

(20 pts) (a) Compute the minimum achievable error rate by a linear machine (hint: draw a picture of the data). Assume the classes are equiprobable.

(10 pts) (b) Assume the priors for each class are: $P(\omega_1) = \alpha$ and $P(\omega_2) = 1 - \alpha$. Sketch $P(E)$ as a function of α for a maximum likelihood classifier based on the assumption that each class is drawn from a multivariate Gaussian distribution. Compare and contrast your answer with your answer to (a). Be very specific in your sketch and label all critical points. Unlabeled plots will receive no partial credit.

(5 pts) (c) Assume you are not constrained to a linear machine. What is the minimum achievable error rate that can be achieved for this data? Is this value different than (a)? If so, why? How might you achieve such a solution? Compare and contrast this solution to (a).

Problem No. 2: Suppose we have a random sample $X_1, X_2, \ldots, X_n$ where:

- $X_i = 0$ if a randomly selected student does not own a laptop, and
- $X_i = 1$ if a randomly selected student does own a laptop.

(35 pts) (a) Assuming that the X_i are independent Bernoulli random variables with unknown parameter p:

$$p(x;p) = (p)^{x_i}(1-p)^{1-x_i}$$

where $x_i = 0 \; or \; 1$ and $0 < p < 1$. Find the maximum likelihood estimator of p, the proportion of students who own a laptop.

Problem No. 3: Let's assume you have a 2D Gaussian source which generates random vectors of the form $[x_1, x_2]$. You observe the following data: [1,1], [2,2], [3,3]. You were told the mean of this source was 0 and the standard deviation was 1.

(25 pts) (a) Using Bayesian estimation techniques, what is your best estimate of the mean based on these observations?

(5 pts) (b) Now, suppose you observe a 4th value: [0,0]. How does this impact your estimate of the mean? Explain, being as specific as possible. Support your explanation with calculations and equations.