Name: ____

Problem	Points	Score
1(a)	30	
1(b)	10	
1(c)	10	
2	25	
3	25	
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 two probability distributions representing a 2-class problem:

$$p(x|\omega_1) = \begin{cases} 1/\alpha & 0 \le x < \alpha \\ 0 & elsewhere \end{cases} \qquad p(x|\omega_1) = \begin{cases} 1 & (\beta - 1/2) \le x < (\beta + 1/2) \\ 0 & elsewhere \end{cases}$$

(a) Sketch the probability of error for a maximum likelihood classifier (assume equal priors) as a function of α and β. Think carefully how these parameters influence the result and show a set of plots that are representative of the behavior. Under what condition is the error minimum? maximum? What are the minimum and maximum error rates that can be achieved?

(b) Pick a value of α and β for which the error rate is approximately 25%. Now assume you are doing maximum a posteriori classification. Plot the probability of error as the prior for class 1 increases and the prior for class 2 decreases (these must sum to one, so when one increases the other decreases).

(c) Suppose now that you are going to model each class as a Gaussian random variable. How would the result for (a) change? Be precise.

Problem No. 2: A coin is flipped 100 times. Given that there were 55 heads, find the maximum likelihood estimate for the probability p of heads on a single toss.

Problem No. 3: Maximum likelihood methods apply to estimates of prior probabilities as well. Let samples be drawn by successive, independent selections of a state of nature ω_i with unknown probability $P(\omega_i)$. Let $z_{ik} = 1$ if the state of nature for the kth sample is ω_i and $z_{ik} = 0$ otherwise. Derive the maximum likelihood estimate of the prior probability for class i, $P(\omega_i)$, and discuss why this makes sense. (Hint: write an expression for $P(z_{i1},...,z_{in}|P(\omega_i)))$. How would this estimate change if Bayesian methods were employed?