Massachusetts Institute of Technology

Department of Electrical Engineering & Computer Science

6.041/6.431: Probabilistic Systems Analysis

(Fall 2008)

Recitation 20

November 20, 2008

1. Let Θ be a Bernoulli random variable that indicates which of the two hypotheses is true, $\mathbf{P}(\Theta = 1) = p$. Under the null hypothesis ($\Theta = 0$), random variable X is uniformly distributed over the interval $[0, 1]$. Under the alternative hypothesis ($\Theta = 1$), the PDF of X is given by

$$f_{X|\Theta}(x \mid 1) = \begin{cases} 2x & \text{if } 0 \leq x \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

   (a) We are interested in figuring out whether X is uniformly distributed or if the PDF of X is a linear ramp as described above. Find the MAP rule for deciding which of the two hypotheses is true. Give the decision rule for $p = 1/2$, $p = 2/3$, $p = 1/3$.

   (b) For this question only, assume that $p = 2/3$. Find the probability of error for this decision rule given that the null hypothesis is true.

   (c) Find the probability of error for this decision rule as a function of p.

2. We wish to estimate the probability of heads of a biased coin, which we denote θ. We model θ as a value of a random variable Θ, which is uniformly distributed over the interval $[0, 1]$. Let X be the number of heads observed in n independent tosses of the coin.

   In your computations, you might find the following equality useful:

$$\int_0^1 t^a (1-t)^b \, dt = \frac{a!b!}{(a+b+1)!},$$

   for $a, b = 0, 1, 2, \ldots$.

   (a) Find the MAP estimator for the value of θ as a function of the number of observed heads, $\hat{\theta}_{\text{MAP}}(x) = \arg\max_\theta p_{\Theta|X}(\theta|x)$.

   (b) Find the conditional expectation estimator (i.e., the minimum mean squared error estimator) for the value of θ as a function of the number of observed heads, $\hat{\theta}_{\text{CE}}(x) = \mathbf{E}[\Theta \mid X = x]$.

   (c) Compare the two estimators.

   (d) Generalize your solution for the case when the prior PDF of Θ belongs to the Beta family:

$$f_\Theta(\theta) = \begin{cases} \frac{\theta^\alpha (1-\theta)^\beta}{B(\alpha,\beta)} & \text{if } 0 < \theta < 1 \\ 0 & \text{otherwise} \end{cases},$$

   for $\alpha, \beta = 0, 1, 2, \ldots$. As a reminder, $B(\alpha, \beta)$ is a normalizing constant, known as the Beta function.