Math4310/Biol6317, Fall 2011
Problem Set 4, due Thursday, Sep 22

Problem 1. This part is to get started simulating in R. Simulate the outcomes for $1,000$ random, uniformly distributed variables with density

$$f(x) = \begin{cases} 0.1 & \text{, if } 0 \leq x \leq 10 \\ 0 & \text{, else} \end{cases}$$

and do the following:

a. Calculate their sample mean, sample variance and sample standard error of the mean.

b. Compare with the mean and variance of the random variable with density f, and compare the sample standard error of the mean with the standard error of the mean obtained from the density.

c. Plot a histogram of the outcomes.

Hint: Use the functions `runif`, `mean`, `var` and `hist`.

Problem 2. The Chinese Mini-Mental State Test (CMMS) is a test used to identify people with Alzheimer's disease (AD) and dementia among people in China. An extensive clinical trial was performed of this test, whereby participants were evaluated with the test and a definitive (clinical) diagnosis of AD was made. The table below show the counts obtained on the subgroup of people with some basic education requirement. Suppose a cutoff value of ≤ 20 on the test is used to identify people with AD.

	Confirmed as AD	
CMMS score	No	Yes
0-5	0	2
6-10	0	1
11-15	3	4
16-20	9	5
21-25	16	3
26-30	18	1

(a) What is the sensitivity and specificity of the CMMS test using the 20 cutoff?

(b) Create a plot of the sensitivity versus (1 - specificity), which is the true positive rate versus the false positive rate for all of the possible cut-offs between 0 and 30. This is called an Receiver Operating Characteristic (ROC) curve.

(c) No information of prevalence is given. Based on the sensitivity and specificity of the CMMS test in (a) (with cut-off value 20), plot the positive predictive value as a function of the prevalence of AD. Do the same for the negative predictive value.

Hint: Use the R command `plot` with suitable arguments.

Problem 3. The clown Flip has an oddly deformed coin and tries the following experiment. Flip flips his coin 10 times, 7 of which are heads. You think maybe Flip's coin is biased towards having a greater probability of yielding a head.

a. What is the maximum likelihood estimate of p, the true probability of heads associated with this coin, given the observed number of heads?

b. Plot the likelihood associated with this experiment. Renormalize the likelihood so that its maximum is one. Does the likelihood suggest that the coin is fair?

c. What is the probability of seeing 7 or more heads out of ten coin flips if the coin was fair? Does this probability suggest that the coin is fair? Note this number is called a P-value.

Hint: For plotting, use the R functions `dbinom` and `plot`.