STA 410/2102 — Practice Questions for First Test

Please note that these questions do **not** cover all the topics that may be on the test.

Question 1: Calculate the result of each of the following arithmetic operations, assuming that the numbers are represented in *decimal* (base ten) floating-point, with three significant decimal digits in the mantissa. Assume that all results are properly rounded to the nearest possible number, and that overflow and underflow do not occur. Write your answers in the form $0.d_1d_2d_3 \times 10^E$, as in the question.

1. $(0.999 \times 10^{-5}) \ + \ (0.200 \times 10^{-3})$
2. $(0.100 \times 10^{20}) \ * \ (0.101 \times 10^{-30})$
3. $(0.788 \times 10^{12}) \ - \ (0.786 \times 10^{12})$
4. $(0.999 \times 10^{15}) \ / \ (0.300 \times 10^{-5})$

Question 2: For each of the following true statements about arithmetic on real numbers, give an example showing that it is not necessarily true for computer arithmetic on floating-point numbers in which the results are rounded to three decimal digits (ie, the numbers have the form $0.d_1d_2d_2 \times 10^E$ for some decimal digits d_1, d_2, and d_3 and integer exponent E. In your examples, assume that overflow and underflow do not occur. (In other words, these should be examples of the effects of round-off error, not of the effects of overflow or underflow.)

1. If $x \neq y$ then $x + z \neq y + z$.
2. $(1 + x) - x = 1$.
3. $x + (y + z) = (x + y) + z$.
4. If $x \neq 0$ then $x \times (1/x) = 1$.

Question 3: Suppose we have n binary (0/1) observations, $y_1, \ldots, y_n$, that we model as being i.i.d. samples from a Bernoulli distribution with the probability of 1 being p. Suppose that 1/10 of the observations are 1s and 9/10 are 0s, so the maximum likelihood estimate of p is $\hat{p} = 1/10$. Suppose we use floating-point numbers of the form $0.d_1d_2d_3 \times 10^E$, where d_1, d_2, d_3 are decimal digits and E is an integer exponent in the range -100 to $+100$. For what values of n can we compute the probability of $y_1, \ldots, y_n$ using the maximum likelihood estimate $\hat{p}$ without the result underflowing to zero?

An approximate answer is good enough. Recall that $(1 - \epsilon)^{1/\epsilon} \approx e^{-1} \approx 0.37$ for ϵ close to zero.

Question 4: Let the data points $x_1, \ldots, x_n$ be modeled as being independent, with each x_i coming from a Cauchy distribution with location parameter 0 and scale parameter e^ρ. The density function for this Cauchy distribution is

$$f(x) \;=\; \frac{1}{\pi e^\rho (1 + x^2 e^{-2\rho})}$$

Write an R function called `mle` to find the maximum likelihood estimate for ρ, using Newton-Raphson iteration. The arguments of `mle` should be the data vector, x, an initial value for ρ, and the number of Newton-Raphson iterations to do. The value returned by `mle` should be the final maximum likelihood estimate found. You need not check the arguments for validity.

Question 5: Suppose we try to solve the equation $x^4 - 81 = 0$ using Newton-Raphson iteration.

1. How will we find the next guess at the solution, $x^{(t+1)}$, from the current guess, $x^{(t)}$?
2. Suppose we start from an initial value of $x^{(0)} = 5$. Here are the values found in the first four iterations:

$$\begin{aligned} x^{(1)} &= 3.912 \\ x^{(2)} &= 3.2722427442656 \\ x^{(3)} &= 3.03212968848923 \\ x^{(4)} &= 3.00050709092522 \end{aligned}$$

   The exact answer is of course 3 (for the solution in this neighborhood).

   Estimate what value of $x^{(5)}$ will be the result of doing one more iteration, without actually doing this iteration. Try to get as accurate an answer as you can by considering in detail the rate at which the error ought to be going down.

Question 6: Suppose we have n i.i.d. data points, $x_1, \ldots, x_n$, that are positive real numbers, with each having the distribution with density function

$$f(x) = \frac{1}{\theta\,(1 + x/\theta)^2}$$

where θ is an unknown positive model parameter.

Derive the formulas needed to use Newton-Raphson iteration to find the maximum likelihood estimate for θ, and write an R program that takes as arguments a data vector x, an initial guess for θ, and the number of Newton-Raphson iterations to do, and which returns a list consisting of the maximum likelihood estimate for θ along with its standard error, found using the observed information.

Question 7: Non-negative integer counts $Y_1, \ldots, Y_n$ were generated independently from a Poisson distribution with some unknown mean λ, which we wish to estimate. Data was collected by writing each Y_i on a separate slip of paper, all of which were put in a box. Unfortunately, only now, when it is time to try to estimate λ, has it been realized that it's not possible to tell the difference between the number 6 and the number 9, because we don't know which way up each slip of paper is supposed to go. Therefore, all these numbers have been entered into the computer as 6, even though some of them may have actually been 9. Call the data as entered $X_1, \ldots, X_n$. The X_i are independent, but, due the confusion of 6 and 9, the probability distribution for X_i is not Poisson, but is instead:

$$P(X_i = k \,|\, \lambda) = \begin{cases} e^{-\lambda}\lambda^k/k! & \text{if } k \neq 6 \text{ and } k \neq 9 \\ e^{-\lambda}\lambda^6/6! + e^{-\lambda}\lambda^9/9! & \text{if } k = 6 \\ 0 & \text{if } k = 9 \end{cases}$$

Show how to derive the formulas for an EM algorithm to find the maximum likelihood estimate for λ from $X_1, \ldots, X_n$, with the unobserved data being the true counts associated with the values recorded as 6. Write an R function to implement this EM algorithm. It should take as arguments a vector of `x` values, and the number of iterations of EM to do. Use some reasonable initial value for λ to begin the EM algorithm.