STA 290 — Homework 7

Due Tuesday, November 27th

Please include all of your C code in an appendix.

1. We return to the `turns` data from homework 6. Read the data into C and find the MLE for the rate λ. Use the MLE and conduct a chi-squared test to see if it is reasonable that the data come from a Poisson distribution with rate λ. Recall that the chi-squared test statistic is of the form:

$$\chi^2 = \sum_{i=1}^{m} \frac{(obs_i - exp_i)^2}{exp_i}$$

   where obs_i is the observed count in bin i, exp_i is the expected count in bin i, and there are m total bins. For example, bin 1 could be the number of times no right-turns were observed in the three-minute interval, with $obs_1 = 14$, and exp_1 being the expected number of zero counts in 300 draws from a Poisson with rate λ. You need to choose the correct number of degrees of freedom, and then use IMSL to get the p-value for chi-squared test.

2. Write code to read in a (square) matrix, compute its inverse, and print both the matrix and its inverse, using code from LAPACK. Do this with dynamic or pseudo-dynamic memory allocation, i.e., do not hard code (either literally or with a `#define`) the size of the matrix (if desired, you may use a command line argument, or you may modify the data file to also specify the number of rows). Run your code on the matrix in the file `matrix.dat` from the course web page with all the rest of the C code (`http://www.stat.duke.edu/courses/Fall01/sta290/Slides/Ccode`).
   Possibly helpful notes:
   - With LAPACK, you will first need to use a routine to get the LU decomposition before calling the routine for matrix inversion.
   - Documentation on Compaq's implementation of BLAS and LAPACK (which should be the same as other implementations, except that they are optimized for the Alpha platform, but the syntax should all be the same) are on-line at: `http://www5.compaq.com/math/documentation/cxml/`

3. Recall the oxide pollution data (available on the datasets web page) from the regression examples (both in S-Plus and C).
   - Fit a Bayesian regression model on the full dataset using all covariates using MCMC in C (i.e., use the sample code from class). Use at least 10,000 iterations. Compare the posterior mean and posterior standard deviation of the coefficients to the least-squares coefficients and their standard errors.
   - Now fit a Bayesian model using only Wind Speed and Temperature as explanatory variables, again using MCMC in C. Compare the posterior mean and posterior standard deviations for this model and for the Bayesian model using all of the covariates.