CS 281A/Stat 241A Homework Assignment 3 (due October 9)

1. **Exponential family.**

   A probability distribution in the exponential family takes the following general form:

   $$p(x \mid \eta) = h(x) \exp\{\eta^T T(x) - A(\eta)\}$$

   for a parameter vector η, often referred to as the *natural parameter*, and for given functions T, A, and h.

   (a) Show that the following distributions are in the exponential family, exhibiting the T, A, and h functions in each case.

   i. Po(λ)—Poisson with parameter λ.

   ii. N(μ, I)—multivariate Gaussian with mean vector μ and identity covariance matrix.

   iii. Mu(θ)—multinomial with parameter vector $\theta = (\theta_1, \theta_2, \ldots, \theta_K)$. Use the fact that $\theta_K = 1 - \sum_{k=1}^{K-1} \theta_k$ and express the distribution using a $(K-1)$-dimensional parameter η.

   (b) The function $A(\eta)$ turns out to have moment-generating properties. In particular, show the following:

   $$\nabla_\eta A = \mathrm{E}[T(X)].$$

   (c) Demonstrate that the relationship in (b) holds for the three examples in part (a).

2. **Dirichlet expectations.**

   Recall the Dirichlet distribution:

   $$p(\theta \mid \alpha) = \frac{\Gamma(\sum_{i=1}^K \alpha_i)}{\prod_{i=1}^K \Gamma(\alpha_i)} \theta_1^{\alpha_1 - 1} \theta_2^{\alpha_2 - 1} \cdots \theta_K^{\alpha_K - 1},$$

   (a) Compute $\mathrm{E}[\theta_k]$. [Hint: do it directly.]

   (b) Compute $\mathrm{Cov}[\theta_j \theta_k]$. [Hint: do it directly.]

   (c) Compute $\mathrm{E}[\log \theta_k]$. [Hint: Show that the Dirichlet distribution is in the exponential family and use the results of problem (1).]

3. **Dirichlet-multinomial prediction.**

   Let $\theta \sim \mathrm{Di}(\alpha)$. Consider multinomial random variables $(X_1, X_2, \ldots, X_N)$, where $X_n \sim \mathrm{Mu}(\theta)$ for each n, and where the X_n are assumed conditionally independent given θ. Now consider a random variable $X_{\mathrm{new}} \sim \mathrm{Mu}(\theta)$ that is assumed conditionally independent of $(X_1, X_2, \ldots, X_N)$ given θ. Compute:

   $$p(x_{\mathrm{new}} \mid x_1, x_2, \ldots, x_N, \alpha)$$

   by integrating over θ. [Hint: Your result should take the form of a ratio of gamma functions.]

4. **The LMS algorithm.** The course homepage has a data set named "lms.dat" that contains twenty rows of three columns of numbers. The first two columns are the components of an input vector x and the last column is an output y value. (We will not use a constant term for this problem; thus the input vector and the parameter vector are both two dimensional.)

(a) Solve the normal equations for these data to find the optimal value of the parameter vector. (I recommend using MATLAB or SPlus.)

(b) Find the eigenvectors and eigenvalues of the covariance matrix of the input vectors and plot contours of the cost function J in the parameter space. These contours should of course be centered around the optimal value from part (a).

(c) Initializing the LMS algorithm at $\theta = 0$ plot the path taken in the parameter space by the algorithm for three different values of the step size ρ. In particular let ρ equal the inverse of the maximum eigenvalue of the covariance matrix, one-half of that value, and one-quarter of that value.