Math 50 Stat Inf: Homework 6

due Wed Feb 15

I have created some better problems than ones in the book to teach you this week's stuff: properties of estimators and introduction to Bayesian parameter fitting (estimation). Note you can do question C, the only 'matlab' one, using material up to section 5.7 only.

5.4 : 18 (due to symmetry you do not even need any integrals)

A. Look back to question B.3 from HW2. You used the estimator $\hat{p} = \overline{X} = \sum_{i=1}^{N} X_i$ for the Poisson parameter, with $N = 1000$ data. Compute the standard deviation $\sqrt{\text{Var}\hat{p})}$ using the theory you've learned since then. What were the z-values of your sample(s) of $\hat{p}$?

B. Recall $\hat{\sigma^2} = \frac{1}{n-1}\sum_{i=1}^{n}(Y_i - \overline{Y})^2$ was an unbiased estimator for the variance of data given n samples from a normal model pdf, when the μ was not known. Assume instead μ is known; then is the estimator $\hat{\sigma^2} = \frac{1}{n-1}\sum_{i=1}^{n}(Y_i - \mu)^2$ biased? (Use the method of lecture to find the bias). If so, fix it.

5.5 : 2 (key),

3,

4,

5.

5.7 : 1 (review),

2 (using things you already know about combining normal variables).

C. So far you're getting the impression that the sample mean $\overline{Y}$ is a 'good' (unbiased, minimum-variance, consistent) estimator of the mean of a distribution (*e.g.* uniform, poisson, normal pdfs). Let's shake that up a bit with the 'lighthouse problem'. A lighthouse sits at location $(-1, \mu)$ in the plane and sends out pulses of light, at random angles uniformly distributed over $[-\pi/2, \pi/2]$, which travel in straight lines then are detected where they hit the y-axis (the 'shoreline') at locations y_i. (the 'shoreline'). Your job is to estimate the location μ using n samples $\{y_i\}$, $i = 1 \cdots n$.

   i) With $\mu = 0$ generate a list of $n = 10^4$ samples of y_i (you'll need to use elementary trigonometry). Histogram them with bins of width 0.1 over a suitable y range. They should have a *Cauchy* pdf $f_Y(y;\mu) = 1/[\pi(1 + (y-\mu)^2)]$.

   ii) Plot the sample mean $\bar{y}$ vs sample size n for a sequence of sizes $n = 10, 10^2, \ldots, 10^7$. [Hint use a log scale for n, via `semilogx` in matlab].

   iii) Is the estimate converging to the true $\mu = 0$? Comment. Why do you think Chebychev's law of large numbers breaks down?

   iv) Given $n = 10^2$ samples, plot the likelihood function $L(\mu)$ over the domain $-2 \leq \mu \leq 2$. Comment on the constistency of the ML estimate vs the sample mean.

D. A coin (Bernoulli pdf) of unknown bias α gives the independent samples $\{X_i\} = \{0, 0, 1, 1, 1, 1, 0, 1, 1, 1\}$. Find the likelihood function $L(\alpha|X)$.

   i) Assuming a uniform prior $p(\alpha)$ in $[0, 1]$, find the posterior $p(\alpha|X) = \text{const.}L(\alpha|X)p(\alpha)$. [Hint: normalizing the posterior is a simple way to get the const].

ii) Assume instead you already believe the coin has a low α so you use a beta pdf with $r = 2$, $s = 4$ (see Ex. 5.8.2). Now find the posterior.

5.8 : 1 (assume a single sample k is the measured data, so your posterior will depend on that k. You may want to consider only the θ and $(1 - \theta)$ factors, then normalize at the end),