Statistics 5444: Homework 2

For each homework assignment, turn in at the beginning of class on the indicated due date. Late assignments will only be accepted with special permission. Write each problem up *very* neatly (LaTeX is preferred). Show all of your work.

Problem 1

Consider the case where $x_i \sim N(\mu, \sigma^2)$. Denote $X = \{x_1, \ldots, x_N\}$. In class we derived that the posterior predictive distribution $\tilde{x}|X, \sigma^2 \sim N(\bar{x}, (1 + \frac{1}{N})\sigma^2)$, under the reference prior $p(\mu) \propto 1$.

Part 1

Under the reference prior $p(\phi) \propto 1/\phi$, find $p(\phi|X)$ (i.e. state the distribution and its underlying parameters).

Part 2

Show that $\tilde{x}|X \sim t_{\nu-1}(\bar{x}, (1 + \frac{1}{N})s^2)$.

Part 3

Using the iterative sampling scheme:

$$\begin{aligned} \text{sample } \phi|X &\leftarrow p(\phi|X), \\ \text{sample } \mu|X &\leftarrow p(\mu|\phi, X), \\ \text{sample } \tilde{x} &\leftarrow p(\tilde{x}|\mu, \phi), \end{aligned}$$

Sample $10,000$ samples. Plot a (normalized) histogram of your results, and 'overlay' the density function from Part 2.

Problem 2

This next result is very important and can be useful in many situations. We will expand on this in future exercises. Let $x \sim N(\mu, \sigma^2/\gamma)$. From this, we can write:

$$p(x|\mu, \sigma^2/\gamma) \propto (\gamma/\sigma^2)^{1/2} \exp(-\frac{1}{2}\gamma\frac{(x-\mu)^2}{\sigma^2}).$$

Part 1

Letting $\gamma \sim \text{Gamma}(\alpha, \beta)$ Find $p(x|\mu, \sigma^2) = \int p(x|\mu, \sigma^2/\gamma)p(\gamma)d\gamma$.

Part 2

State what α and β must be for $p(x|\mu, \sigma)$ to have a Cauchy distribution with *shift* μ and the *scale* σ (Note: a Cauchy is a "T-1" distribution).

Part 3

Recall that if $p(y) = \int p(y|\lambda)p(\lambda)d\lambda$, then we can generate p(y) via the following algorithm:

$$\begin{aligned} \text{sample } \tilde{\lambda} &\leftarrow p(\lambda) \\ \text{sample } \tilde{y} &\leftarrow p(y|\tilde{\lambda}). \end{aligned}$$

The resulting $\tilde{y}$ is a perfect sample from $p(y)$.
Use this method to simulate 1,000 Cauchy (t-1 distribution) random variables, using the result you obtained from part 1 and state the $(\min, \max)$ from your 1,000 sample draws. Repeat this exercise and comment on your findings.

Problem 3

A basic property of the MLE is that it is invariant to transformations. For example, let $\eta = \tau(\theta)$, and let $L(\hat{\theta}|x) = \max_\theta L(\theta|x)$. Denote the likelihood function $L^*(\eta|x)$ as the likelihood function under the transformation $\eta = \tau(\theta)$. Letting $L^*(\hat{\eta}|x) = \max_\eta L^*(\eta|x)$ we have that $\hat{\eta} = \tau(\hat{\theta})$. This result holds for *ALL* functions $\tau(\cdot)$.

The Map (Maximum A-Posteriori) estimator is defined to be $\hat{\theta}$ such that $p(\hat{\theta}|x) = \max_\theta p(\theta|x)$. $\hat{\theta}$ can also be referred to as the posterior mode. Is the MAP estimator invariant to transformations? That is, if we let $\eta = \tau(\theta)$, and denote $p(\hat{\eta}|x) = \max_\eta p(\eta|x)$, is $\hat{\eta} = \tau(\hat{\theta})$? If so, prove it. If not, disprove.

Problem 4

The iterated expectation and variance formulas follow as

$$\begin{aligned} E[X] &= E\left[E[X|Y]\right], \qquad \text{and} \\ V(X) &= V\left(E[X|Y]\right) + E\left[V(X|Y)\right] \end{aligned}$$

respectively. These are very useful in many cases. Recall, in class, we derived $p(\tilde{x}|X, \sigma^2) \sim N(\bar{x}, (1 + \frac{1}{n})\sigma^2)$. We will re-derive this result using the iterated formulas.

Part 1

Under $p(\mu) \propto 1$, explictly solve the predictive distribution $p(\tilde{x}|X, \sigma^2)$ to prove it is normally distributed. That is, solve:

$$\pi(\tilde{x}|X, \sigma^2) \propto \int_{-\infty}^{\infty} p(\tilde{x}|\mu, X, \sigma^2) * \pi(\mu|X, \sigma^2) d\mu$$

$$\pi(\tilde{x}|X, \sigma^2) \propto \int_{-\infty}^{\infty} e^{-\frac{1}{2}\left[\frac{(\tilde{x}-\mu)^2}{\sigma^2} + \frac{N(\mu-\bar{x})^2}{\sigma^2}\right]} d\mu.$$

Explicitly detail your integration steps.

Part 2

For using the iterated formulas above, choose and state the random variable that you are conditioning on. Derive the expectation and variance of $p(\tilde{x}|X, \sigma^2)$ using the iterated formulas.

Problem 5

Provided the Maximum Likelihood estimator, show that if $\pi(\theta)$ is proper, then $\pi(\theta|x)$ is also proper.