Hw 1, Math 472, Spring 2018

Last name, first name, ID number : Score:

1. (5 points) Let Y_1 and Y_2 be two independent random variables with respective probability distributions Poisson(λ_1) and Poisson(λ_2). Let $Y = Y_1 + Y_2$. Use moment generating functions to show that Y has probability distribution Poisson($\lambda_1 + \lambda_2$).
   Solution: Denote $m_1(t)$, $m_2(t)$ and $m(t)$ the moment generating functions of Y_1, Y_2 and Y, respectively. Since Y_1 and Y_2 are independent, we have

$$m(t) = m_1(t)m_2(t) = e^{(\lambda_1+\lambda_2)(e^t-1)}$$

   which is the moment generating function for a random variable that has a Poisson($\lambda_1 + \lambda_2$) distribution. Therefore, Y has a Poisson($\lambda_1 + \lambda_2$) distribution.

2. (5 points) An anthropologist wishes to estimate the average height of men for a certain race of people. If the population standard deviation is assumed to be 2.5 inches and if she randomly samples 100 men, find the probability that the difference between the sample mean and the true mean will not exceed 0.5 inch.
   Solution: Let $n = 100$ and $(X_i)_{1\le i\le n}$, be the height of the n men. By assumption, the random variables $(X_i)_{1\le i\le n}$, have the same true mean μ, which is unknown, and the same standard deviation $\sigma = 2.5$. Denote $\bar{X}_n$ the sample mean of the $(X_i)_{1\le i\le n}$. The number of observations n being large, we can apply the central limit theorem that gives

$$\begin{aligned} P(-0.5 \le \bar{X}_n - \mu \le 0.5) &= P(-\frac{0.5\sqrt{n}}{\sigma} \le \frac{\sqrt{n}(\bar{X}_n - \mu)}{\sigma} \le \frac{0.5\sqrt{n}}{\sigma}) \\ &= P(-2 \le \frac{\sqrt{n}(\bar{X}_n - \mu)}{\sigma} \le 2) \\ &\approx P(-2 \le Z \le 2) \quad \text{where } Z \sim \mathcal{N}(0,1) \\ &\approx 0.9544 \text{ (from table or online calculator.)} \end{aligned}$$

3. (5 points) Suppose that a single observation X is to be taken from the uniform distribution on the interval $[\theta - \frac{1}{2}, \theta + \frac{1}{2}]$, the value of θ is unknown and the prior distribution of θ is the uniform distribution on the interval $[10, 20]$. If the observed value of X is 12, what is the posterior distribution of θ?

   Solution: By assumption, the p.d.f. of X given θ is $f(x|\theta) = \mathbb{1}_{[\theta-\frac{1}{2},\theta+\frac{1}{2}]}(x)$. Notice that

$$\theta - \frac{1}{2} < x < \theta + \frac{1}{2} \iff |x - \theta| < \frac{1}{2} \iff |\theta - x| < \frac{1}{2} \iff x - \frac{1}{2} < \theta < x + \frac{1}{2}$$

   therefore $f(x|\theta) = \mathbb{1}_{[x-\frac{1}{2},x+\frac{1}{2}]}(\theta)$. Moreover, the p.d.f. of the prior distribution of θ is $\xi(\theta) = \frac{1}{10}\mathbb{1}_{[10,20]}(\theta)$. Therefore, the p.d.f. of the posterior distribution of θ given the observation $X = 12$ satisfies

$$\begin{aligned} \xi(\theta|X = 12) &\propto \mathbb{1}_{[12-\frac{1}{2},12+\frac{1}{2}]}(\theta)\mathbb{1}_{[10,20]}(\theta) \\ &\propto \mathbb{1}_{[11.5,12.5]}(\theta). \end{aligned}$$

   Thus, the posterior distribution of θ given $X = 12$ is the uniform distribution on the interval $[11.5, 12.5]$.

4. (5 points) Suppose that the heights of the individuals in a population have a normal distribution for which the mean θ is unknown and the standard deviation is 2 inches. Suppose that the prior distribution of θ is a normal distribution for which the the mean is 68 inches and the standard deviation is 1 inch. Suppose that 10 people are selected at random and their average height is found to be 69.5 inches.

(a) If the square error loss function is used, what is the Bayes estimate of θ?
Solution: According to theorem 7.3.3, the posterior distribution of θ given the 10 observations is the normal distribution with mean $\mu_1 = \frac{4.68+10.69.5}{4+10} = \frac{967}{14}$ and variance $\sigma_1^2 = \frac{4}{4+10} = \frac{4}{14}$. Therefore, the Bayes estimatate of θ, which is is the expected value of the posterior distribution given the observations when the square error loss function is used, is $\delta^*(x) = \frac{967}{14}$.

(b) If the absolute error loss function is used, what is the Bayes estimate of θ?
Solution: When the absolute error loss function is used, the Bayes estimatate of θ is the median value of the posterior distribution given the observations.Since the median and mean values of a normal distribution are equal to each other, we find, again, that $\delta^*(x) = \frac{967}{14}$.

5. (5 points) Show that the family of the beta distributions is a conjugate family of prior distributions for samples from a negative binomial distribution with known parameter r and unknown parameter p $(0 < p < 1)$.
Solution: Let $X_1, \ldots, X_n$ be a random sample from the negative binomial distribution. For $\leq i \leq n$, the probability function for the observation $X_i = x_i$ is

$$f(x_i|p) = \binom{r + x_i - 1}{x_i} p^{x_i}(1-p)^r.$$

Thus the likelihood function is given by

$$f(x_1, \ldots, x_n|\theta) = \Pi_{i=1}^n \binom{r + x_i - 1}{x_i} p^{x_i}(1-p)^r \propto p^{nr}(1-p)^{\sum_{i=1}^n x_i}.$$

Moreover, the p.d.f of the prior distribution is

$$\xi(p) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} p^{\alpha-1}(1-p)^{\beta-1}.$$

Therefore, the p.d.f $\xi(p|(x_1, \ldots, x_n)$ satisfies

$$\xi(p|(x_1, \ldots, x_n) \propto f(x_1, \ldots, x_n|\theta)\xi(p) \propto p^{\alpha+nr-1}(1-p)^{\beta+\sum_{i=1}^n x_i - 1}$$

which implies that the posterior distribution follows a beta distribution with parameters $(\alpha + nr, beta + \sum_{i=1}^n x_i)$. Therefore the family of the beta distributions is a conjugate family of prior distributions for samples from a negative binomial distribution with known parameter r and unknown parameter p $(0 < p < 1)$.

6. Let θ be a real number and consider the function $f(x|\theta) = \frac{1}{2}e^{-|x-\theta|}$, for $x \in \mathbb{R}$.

(a) Show that $f(x|\theta)$ is a probability density function.
Solution Using the substitution $y = x - \theta$, we find

$$\int_{-\infty}^{+\infty} f(x|\theta)dx = \int_{\theta}^{+\infty} e^{-(x-\theta)}dx = \int_0^{+\infty} e^{-y}dy = -e^{-y}\Big|_0^{+\infty} = 1.$$

(b) Suppose that $X_1, \ldots, X_n$ form a random sample from a distribution for which the p.d.f. is $f(x|\theta)$ and suppose that the value of θ is unknown. Find the maximum likelihood estimator of θ.

Solution: For all observed values $x_1, \ldots, x_n$, denote $x = (x_1, \ldots, x_n)$. The likelihood function is given by $f_n(x|\theta) = \prod_{i+1}^n f(x_i|\theta) = \frac{1}{2^n}e^{-\sum_{i+1}^n |x_i - \theta|}$. The maximum likelihood estimate of θ, denoted $\hat{\theta}$, is the solution of the maximization problem

$$\max_{\theta} \frac{1}{2^n} e^{-\sum_{i+1}^n |x_i - \theta|}$$

whose solution is the same as the solution of the minimization problem

$$\min_{\theta} \sum_{i+1}^{n} |x_i - \theta|.$$

Now, denote Y the discrete random variable uniformly distributed on the set $\{x_1, \ldots, x_n\}$. We notice that $E(|Y - \theta|) = \frac{1}{n} \sum_{i+1}^{n} |x_i - \theta|$. Therefore, $\hat{\theta}$ is the solution of

$$\min_{\theta} E(|Y - \theta|)$$

which, according to theorem 4.5.3, is the median of the distribution of Y, that is to say the median value of the observations $x_1, \ldots, x_n$.