STA 431s23 Assignment Three[1]

For the Quiz on Friday Feb. 3d, please bring printouts of your full R input for Question 16. The other problems are not to be handed in. They are practice for the Quiz.

---

1. Let $cov(\mathbf{x}) = \mathbf{\Sigma} = \mathbf{CDC}^\top$. The random vector $cov(\mathbf{x})$ has four elements, so that the matrix of eigenvectors may be written
$$\mathbf{C} = \left(\, \mathbf{v}_1 \, \middle| \, \mathbf{v}_2 \, \middle| \, \mathbf{v}_3 \, \middle| \, \mathbf{v}_4 \, \right),$$
where the $\mathbf{v}_j$ are the eigenvectors. What is $\mathbf{v}_3^\top \mathbf{C}$?

2. Let $\mathbf{x} \sim N_p(\mathbf{0}, \mathbf{\Sigma})$, with $\mathbf{\Sigma} = \mathbf{CDC}^\top$. Let $\mathbf{v}_j$ be the eigenvector (a column of $\mathbf{C}$) corresponding to the eigenvalue λ_j of $\mathbf{\Sigma}$.
   (a) What is the distribution of the scalar random variable $y = \mathbf{v}_j^\top \mathbf{x}$?
   (b) What is the distribution of $\mathbf{y} = \mathbf{C}^\top \mathbf{x}$?
   (c) How do you know that the elements of $\mathbf{y}$ are independent?

   The elements of $\mathbf{y}$ are called the *principal components* of $\mathbf{x}$.

3. Let $\mathbf{x} = (x_1, x_2, x_3)^\top$ be multivariate normal with
$$\boldsymbol{\mu} = \left(\begin{array}{c} 1 \\ 0 \\ 6 \end{array} \right) \text{ and } \mathbf{\Sigma} = \left(\begin{array}{ccc} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 1 \end{array} \right).$$
   (a) Derive the joint distribution of x_1 and x_2 using matrices.
   (b) Let $y_1 = x_1 + x_2$ and $y_2 = x_2 + x_3$. Find the joint distribution of y_1 and y_2 using matrices.

4. Let x_1 be Normal(μ_1, σ_1^2), and x_2 be Normal(μ_2, σ_2^2), independent of x_1. What is the joint distribution of $y_1 = x_1 + x_2$ and $y_2 = x_1 - x_2$? What is required for y_1 and y_2 to be independent? Hint: Use matrices.

5. If $\mathbf{x} \sim N_p(\boldsymbol{\mu}, \mathbf{\Sigma})$, then you know the distribution of $\mathbf{Ax}$ from the formula sheet. Use this result to obtain the distribution of the sample mean under normal random sampling. That is, let $x_1, \ldots, x_n$ be a random sample from a $N(\mu, \sigma^2)$ distribution. Find the distribution of $\overline{x}$. You might want to use $\mathbf{1}$ to represent an $n \times 1$ column vector of ones.

---

6. This problem will guide you through the proof that if $\mathbf{w} \sim N_p(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ with $\boldsymbol{\Sigma}$ positive definite, $y = (\mathbf{w} - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{w} - \boldsymbol{\mu})$ has a chi-squared distribution with p degrees of freedom.

   (a) What is the distribution of $\mathbf{w} - \boldsymbol{\mu}$? Just write down the answer.

   (b) Because $\boldsymbol{\Sigma}$ is positive definite, we know that $\boldsymbol{\Sigma}^{-1/2}$ exists; there is no need to prove it. What is the distribution of $\mathbf{z} = \boldsymbol{\Sigma}^{-1/2}(\mathbf{w} - \boldsymbol{\mu})$? Show some work.

   (c) How do you know that $\mathbf{z}$ is made up of *independent* standard normals?

   (d) What is the distribution of $\mathbf{z}^\top \mathbf{z}$? Hint: What is the distribution of a squared standard normal? What is the distribution of a sum of independent chi-squares?

   (e) Calculate $\mathbf{z}^\top \mathbf{z} = (\mathbf{w} - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{w} - \boldsymbol{\mu})$. Proved.

---

7. Let $x_1, \ldots, x_n$ be a random sample from a Poisson distribution with expected value $\lambda > 0$.

   (a) What is the parameter of this model?

   (b) What is the parameter space? See the lecture slides for how to write it.

8. Let $x_1, \ldots, x_n$ be a random sample from a normal distribution with expected value μ and variance σ^2.

   (a) What is the parameter space for this model?

   (b) Obtain the Maximum Likelihood Estimator of the pair $\theta = (\mu, \sigma^2)$ by specializing a result on the formula sheet. You can just write down the answer.

   (c) Find a Method of Moments estimator of θ. Use the fact that $E(x_i) = \mu$ and $Var(x_i) = \sigma^2$. This is very quick. Don't waste time and effort doing unnecessary things.

   (d) In the following R output, data are in the vector x. Based on this, give $\widehat{\theta}$. Your answer is a pair of numbers. I needed a calculator because R's `var` function uses $n - 1$ in the denominator.

```
> c(length(x),mean(x),var(x))
[1]  20.0000  94.3800 155.1554
```

   (e) Give the maximum likelihood estimator of the standard deviation σ. The answer is a number. Do it the easy way. How do you know that this is okay?

9. Let $x_1, \ldots, x_n$ be a random sample from a continuous distribution with density

$$f(x;\theta) = \frac{1}{\theta^{1/2}\sqrt{2\pi}}\, e^{-\frac{x^2}{2\theta}},$$

where the parameter $\theta > 0$. Propose a reasonable estimator for the parameter θ, and use the Law of Large Numbers to show that your estimator is consistent.

10. Let $x_1, \ldots, x_n$ be a random sample from a Gamma distribution with $\alpha = \beta = \theta > 0$. That is, the density is

$$f(x;\theta) = \frac{1}{\theta^\theta \Gamma(\theta)} e^{-x/\theta} x^{\theta-1},$$

for $x > 0$. Let $\widehat{\theta} = \overline{x}_n$. Is $\widehat{\theta}$ consistent for θ? Answer Yes or No and prove your answer. Hint: The expected value of a Gamma random variable is $\alpha\beta$.

11. Let $x_1, \ldots, x_n$ be a random sample from a distribution with mean μ_x and variance σ^2_x. The formula sheet has a formula for the sample variance $\widehat{\sigma}^2$. Show that $\widehat{\sigma}^2_x$ is a consistent estimator of σ^2_x.

12. Let $(x_1, y_1), \ldots, (x_n, y_n)$ be a random sample from a bivariate distribution with $E(x_i) = \mu_x$, $E(y_i) = \mu_y$, $Var(x_i) = \sigma^2_x$, $Var(y_i) = \sigma^2_y$, and $Cov(x_i, y_i) = \sigma_{xy}$. The formula sheet has a formula for the sample covariance $\widehat{\sigma}_{xy}$. Show that $\widehat{\sigma}_{xy}$ is a consistent estimator of σ_{xy}.

13. Let $x_1, \ldots, x_n$ be a random sample from a distribution with expected value μ and variance σ^2_x. Independently of $x_1, \ldots, x_n$, let $y_1, \ldots, y_n$ be a random sample from a distribution with the same expected value μ, and a variance σ^2_y that might be different from σ^2_x. Let $t_n = \alpha\overline{x}_n + (1-\alpha)\overline{y}_n$, where $0 \leq \alpha \leq 1$. Is t_n always a consistent estimator of μ? Answer Yes or No and show your work.

14. Independently for $i = 1, \ldots, n$, let $y_i = \beta x_i + \epsilon_i$, where $x_i \sim N(\mu_x, \sigma^2_x)$, $\epsilon_i \sim N(0, \sigma^2_\epsilon)$, and x_i and ϵ_i are independent.

    (a) What are the parameters of this model?

    (b) What is the parameter space?

    (c) Write the joint distribution of x_i and ϵ_i in matrix form.

    (d) Obtain the joint distribution of x_i and y_i by writing

    $$\left(\begin{array}{c} x_i \\ y_i \end{array} \right) = \mathbf{A} \left(\begin{array}{c} x_i \\ \epsilon_i \end{array} \right).$$

    Give the matrix $\mathbf{A}$, and calculate the vector of expected values and the covariance matrix.

    (e) To check your work, obtain $E(y_i)$, $Var(y_i)$ and $Cov(x_i, y_i)$ with scalar (non-matrix) calculations.

(f) Based on your work so far, you have *two* potential method of moments estimators for β, one based on the vector of expected values, and the other based on the covariance matrix. Let $\widehat{\beta}_1$ denote the estimator based on the expected values, and let $\widehat{\beta}_2$ denote the estimator based on the covariance matrix.

   i. Give the formula for $\widehat{\beta}_2$, and show it is consistent.

   ii. Show that $\widehat{\beta}_1 \stackrel{p}{\rightarrow} \beta$ in most of the parameter space.

   iii. However, consistency means that the estimator converges to the parameter in probability *everywhere* in the parameter space. Where in the parameter space does $\widehat{\beta}_1$ fail?

   iv. *This last item is optional, and will not be on the quiz or on the final exam.* To see exactly how $\widehat{\beta}_1$ fails, use the fact that the ratio of two independent standard normal random variables is a standard Cauchy. Start by simplifying $\overline{y}_n$. You can take it for granted that functions of independent random variables are still independent. If you have followed this path without getting lost, you will conclude that if $\mu_x = 0$, the distribution of $\widehat{\beta}_1$ is Cauchy, but multiplied by a constant and centered on β_1. Notably, the distribution of $\widehat{\beta}_1$ is the same for all n. As $n \rightarrow \infty$, it stays exactly the same, never changing at all. It certainly does not shrink down to any constant, including β.

15. The formula sheet has a useful expression for the multivariate normal likelihood.

   (a) Show that you understand the notation by giving the univariate version, in which $x_1, \ldots, x_n \stackrel{i.i.d}{\sim} N(\mu, \sigma^2)$. Your answer will have no matrix notation for the trace, transpose or inverse.

   (b) Now starting with the univariate normal density $f(x;\mu,\sigma^2) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left\{-\frac{1}{2}\frac{(x-\mu)^2}{\sigma^2}\right\}$, show that the univariate normal likelihood is the same as your answer to the previous question. Hint: Add and subtract $\overline{x}$.

   (c) How does this expression allow you to see *without differentiating* that the MLE of μ is $\overline{x}$?

16. Let $x_1, \ldots, x_n$ be a random sample from a distribution with density

$$f(x) = \frac{\theta e^{\theta(x-\mu)}}{(1+e^{\theta(x-\mu)})^2}$$

for x real, where $-\infty < \mu < \infty$ and $\theta > 0$. Numerical data are available at `http://www.utstat.toronto.edu/brunner/openSEM/data/mystery2.data.txt`.

   (a) Find the maximum likelihood estimates of μ and θ.

   (b) Obtain an approximate 95% confidence interval for θ.

   (c) Test $H_0 : \mu = 2.1$ at the $\alpha = 0.05$ significance level with a large-sample z-test.

Please bring a printout of your full R input and output to the quiz.