

Section 8

Review

- **Law of Total Probability (partition based on value of a r.v.):** If X is a discrete random variable, then

$$\mathbb{P}(A) = \sum_{x \in \Omega_X} \mathbb{P}(A|X = x)p_X(x)$$

If X is a continuous random variable, then

$$\mathbb{P}(A) = \int_{-\infty}^{\infty} \mathbb{P}(A|X = x)f_X(x) dx$$

- **Conditional Expectation:** Let X and Y be random variables. Then, the conditional expectation of X given $Y = y$ is

$$\mathbb{E}[X|Y = y] = \sum_{x \in \Omega_X} x \cdot \mathbb{P}(X = x|Y = y) \quad X \text{ discrete}$$

and for any event A ,

$$\mathbb{E}[X|A] = \sum_{x \in \Omega_X} x \cdot \mathbb{P}(X = x|A) \quad X \text{ discrete}$$

Note that linearity of expectation still applies to conditional expectation: $\mathbb{E}[X + Y|A] = \mathbb{E}[X|A] + \mathbb{E}[Y|A]$

- **Law of Total Expectation (Event Version):** Let X be a random variable, and let events A_1, \dots, A_n partition the sample space. Then,

$$\mathbb{E}[X] = \sum_{i=1}^n \mathbb{E}[X|A_i]\mathbb{P}(A_i)$$

- **Law of Total Expectation (RV Version):** Suppose X and Y are random variables. Then,

$$\mathbb{E}[X] = \sum_y \mathbb{E}[X|Y = y]p_Y(y) \quad Y \text{ discrete r.v.}$$

$$\mathbb{E}[X] = \int_{-\infty}^{\infty} \mathbb{E}[X|Y = y]f_Y(y)dy \quad Y \text{ continuous r.v.}$$

Maximum Likelihood Estimation

- 1) **Realization/Sample:** A realization/sample x of a random variable X is the value that is actually observed.
- 2) **Likelihood:** Let x_1, \dots, x_n be iid realizations from probability mass function $p_X(x; \theta)$ (if X discrete) or density $f_X(x; \theta)$ (if X continuous), where θ is a parameter (or a vector of parameters). We define the likelihood function to be the probability of seeing the data.

If X is discrete:

$$L(x_1, \dots, x_n; \theta) = \prod_{i=1}^n p_X(x_i; \theta)$$

If X is continuous:

$$L(x_1, \dots, x_n; \theta) = \prod_{i=1}^n f_X(x_i; \theta)$$

- 3) **Maximum Likelihood Estimator (MLE):** We denote the MLE of θ as $\hat{\theta}_{\text{MLE}}$ or simply $\hat{\theta}$, the parameter (or vector of parameters) that maximizes the likelihood function (probability of seeing the data).

$$\hat{\theta}_{\text{MLE}} = \arg \max_{\theta} L(x_1, \dots, x_n; \theta) = \arg \max_{\theta} \ln L(x_1, \dots, x_n; \theta)$$

- 4) **Log-Likelihood:** We define the log-likelihood as the natural logarithm of the likelihood function. Since the logarithm is a strictly increasing function, the value of θ that maximizes the likelihood will be exactly the same as the value that maximizes the log-likelihood.

If X is discrete:

$$\ln L(x_1, \dots, x_n; \theta) = \sum_{i=1}^n \ln p_X(x_i; \theta)$$

If X is continuous:

$$\ln L(x_1, \dots, x_n; \theta) = \sum_{i=1}^n \ln f_X(x_i; \theta)$$

- 5) **Steps to find the maximum likelihood estimator, $\hat{\theta}$:**

- Find the likelihood and log-likelihood of the data.
- Take the derivative of the log-likelihood and set it to 0 to find a candidate for the MLE, $\hat{\theta}$.
- Take the second derivative and show that $\hat{\theta}$ indeed is a maximizer, that $\frac{\partial^2 L}{\partial \theta^2} < 0$ at $\hat{\theta}$. Also ensure that it is the global maximizer: check points of non-differentiability and boundary values.
- If we are finding the MLE for a set of parameters, then we set up the system of equations obtained by taking the partial derivative of the log-likelihood function with respect to each of the parameters and setting it equal to 0. We then solve this system to get the MLEs. (And again, second order conditions need to be checked.)

- 6) An estimator $\hat{\theta}$ for a parameter θ of a probability distribution is **unbiased** iff $\mathbb{E}[\hat{\theta}(X_1, \dots, X_n)] = \theta$

Task 1 – Content Review

- True or False:** The Log-Likelihood gives a slightly different estimate, but because it is close enough and easier to compute we use it for our estimate of θ .
- True or False:** $\hat{\theta}$ is the true parameter and θ is our estimate.
- True or False:** An estimator is unbiased if $\text{Bias}(\hat{\theta}, \theta) = \mathbb{E}[\hat{\theta}] - \theta = 0$ or equivalently $\mathbb{E}[\hat{\theta}] = \theta$
- You flip a coin 10 times and observe HHHHTHHTHH (8 heads, 2 tails). What is the MLE of θ , where θ is the true probability of seeing tails?
 - $\hat{\theta} = .2$
 - $\hat{\theta} = .25$
 - $\hat{\theta} = .8$
 - $\hat{\theta} = .3$

Task 2 – Mystery Dish!

A fancy new restaurant has opened up that features only 4 dishes. The unique feature of dining here is that they will serve you any of the four dishes randomly according to the following probability distribution: give dish A with probability 0.5, dish B with probability θ , dish C with probability 2θ , and dish D with probability $0.5 - 3\theta$. Each diner is served a dish independently. Let x_A be the number of people who received dish A, x_B the number of people who received dish B, etc, where $x_A + x_B + x_C + x_D = n$. Find the MLE $\hat{\theta}$ for θ .

Task 3 – A Red Poisson

Suppose that x_1, \dots, x_n are i.i.d. samples from a Poisson(θ) random variable, where θ is unknown. In other words, they follow the distributions $\mathbb{P}(k; \theta) = \theta^k e^{-\theta} / k!$, where $k \in \mathbb{N}$ and $\theta > 0$ is a positive real number.

Find the MLE of θ .

Task 4 – A biased estimator

In class, we showed that the maximum likelihood estimate of the variance θ_2 of a normal distribution (when both the true mean μ and true variance σ^2 are unknown) is what's called the *population variance*. That is

$$\hat{\theta}_2 = \left(\frac{1}{n} \sum_{i=1}^n (x_i - \hat{\theta}_1)^2 \right)$$

where $\hat{\theta}_1 = \frac{1}{n} \sum_{i=1}^n x_i$ is the MLE of the mean. Is $\hat{\theta}_2$ unbiased?

Task 5 – Weather Forecast

A weather forecaster predicts sun with probability θ_1 , clouds with probability $\theta_2 - \theta_1$, rain with probability $\frac{1}{2}$ and snow with probability $\frac{1}{2} - \theta_2$. This year, there have been 55 sunny days, 100 cloudy days, 160 rainy days and 50 snowy days. What is the maximum likelihood estimator for θ_1 and θ_2 ?

Task 6 – Elections

Individuals in a certain country are voting in an election between 3 candidates: A , B and C . Suppose that each person makes their choice independent of others and votes for candidate A with probability θ_1 , for candidate B with probability θ_2 and for candidate C with probability $1 - \theta_1 - \theta_2$. (Thus, $0 \leq \theta_1 + \theta_2 \leq 1$.) The parameters θ_1, θ_2 are unknown.

Let n_A, n_B , and n_C be the number of votes for candidate A, B , and C , respectively. What are the maximum likelihood estimates for θ_1 and θ_2 in terms of n_A, n_B , and n_C ?

(You don't need to check second order conditions.)

Task 7 – Continuous Law of Total Probability

Suppose that the time until server 1 crashes is $X \sim \text{Exp}(\lambda)$ and the time until server 2 crashes is independent, with $Y \sim \text{Exp}(\mu)$.

What is the probability that server 1 crashes before server 2?

Task 8 – Y Me?

Let y_1, y_2, \dots, y_n be i.i.d. samples of a random variable with density function

$$f_Y(y; \theta) = \frac{1}{2\theta} \exp\left(-\frac{|y|}{\theta}\right).$$

Find the MLE for θ in terms of $|y_i|$ and n .

Task 9 – Elevator rides

[This is the problem we did in class.] The number X of people who enter an elevator on the ground floor is a Poisson random variable with mean 10. If there are N floors above the ground floor, and if each person is equally likely to get off at any one of the N floors, independently of where others get off, compute the expected number of stops the elevator will make before discharging all the passengers.