


Exercise 3



Instruction: These are for your own practice. There is no need to submit your answer.

Question 1

Observations $X_1, \dots, X_n$ are independent identically distributed, following the PDF $f_{X_i}(x_i) = \theta x_i^{\theta-1}$, and that $0 < X_i < 1$ for all i . The parameter θ is an unknown positive number. Find the ML estimator of θ .

Question 2

Observations $X_1, \dots, X_n$ are independent identically distributed, being Rayleigh with $f_X(x) = \frac{x}{\sigma^2} e^{-x^2/(2\sigma^2)}$, where σ^2 is unknown. Find the ML estimator of σ^2 .

Question 3

In a certain production line, it is known that the lifetime (measured in hours) of the products follows a Poisson distribution with parameter λ . Five such products were sampled and tested, and they lasted, respectively, 5.2, 5.3, 5.5, 5.7, 5.3 hours. Find the ML estimator of λ .

Question 4

Observations $X_1, \dots, X_n$ are independent identically distributed, following the Gamma PDF given by

$$f_{X_i}(x_i) = \frac{\beta^\alpha x_i^{\alpha-1} e^{-\beta x_i}}{\Gamma(\alpha)},$$

where α and β are parameters of the distribution, and $\Gamma(\cdot)$ is the Gamma function (see Lecture 1 if you don't remember what this function is), with derivative denoted by $\Gamma'(\cdot)$. We wish to obtain the ML estimators of α and β , respectively denoted by $\hat{\alpha}$ and $\hat{\beta}$, based on the observations.


1. Show that, by maximizing the log-likelihood, we have the relationship

$$\hat{\beta} = \frac{\hat{\alpha}}{\frac{1}{n} \sum x_i}$$

where the summation ranges from $i = 1$ to $i = n$.

2. Show that, by maximizing the log-likelihood and using the relationship above, we have

$$\log \hat{\alpha} - \frac{\Gamma'(\hat{\alpha})}{\Gamma(\hat{\alpha})} = \log \left(\frac{1}{n} \sum x_i \right) - \frac{1}{n} \sum \log x_i$$



Question 5

Observations $X_1, X_2, \dots, X_n$, where $n > 2$, are independent identically distributed and drawn from a normal distribution $\mathcal{N} \sim (0, \sigma^2)$ with $\sigma > 0$.

1. Show that the estimator $\tilde{\sigma}_n^2 = \frac{1}{n} \sum_{i=1}^n X_i^2$ is an unbiased estimator of σ^2 .
2. Let $Y \sim \mathcal{N}(\mu, \sigma^2)$. Show that $\mathbb{E}((Y - \mu)^4) = 3\sigma^4$.
Hint: Let $Z = (Y - \mu)/\sigma$. What is $\mathbb{E}(Z^4)$ and how can this help with solving this part? You may need to evaluate some integrals directly.
3. Compute $\text{var}((Y - \mu)^2)$.

Question 6

A researcher is conducting an experiment. The probability that he will succeed is a continuous random variable denoted by X , where

$$f_X(x) = \begin{cases} 2x & 0 \leq x \leq 1 \\ 0 & \text{otherwise.} \end{cases}$$

During a research project, he will continue this experiment for Y times until he succeeds for the first time. Now we know that he has succeeded in the third attempt. Find the MAP estimate of X .

Question 7

Observations $X = \{X_1, \dots, X_n\}$ are independent identically distributed, being Gaussian with $\mathcal{N}(\mu, \sigma^2)$, where μ is unknown while σ^2 is known. Suppose that μ itself (i.e. the prior distribution) is a Gaussian random variable with mean ν and variance β^2 . We would like to find and analyze the MAP estimate of μ .

1. Let $C = f_X(x)$. Using Bayes' rule, show that the posterior probability can be written as

$$f_{\mu|X}(\mu|x) = \frac{1}{C} \left(\prod_{i=1}^N \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x_i - \mu)^2}{2\sigma^2}} \right) \left(\frac{1}{\sqrt{2\pi}\beta} e^{-\frac{(\mu - \nu)^2}{2\beta^2}} \right)$$

2. Show that

$$\sum_{i=1}^N \frac{x_i - \mu}{\sigma^2} = \frac{\mu - \nu}{\beta^2}$$

by maximizing the log posterior.


3. By rearranging the terms, show that

$$\hat{\mu} = \frac{\beta^2 \sum_{i=1}^N x_i + \sigma^2 \nu}{\beta^2 N + \sigma^2}$$

4. What happens when N grows large? Note that we have

$$\mu_{ML} = \frac{1}{N} \sum_{i=1}^N x_i$$

which is the maximum likelihood estimate of μ .



Question 8

The photon arrival at a particular photodetector is modeled as a Poisson distribution, with $X \sim \text{Poisson}(\theta)$, i.e.,

$$f_X(x) = \frac{e^{-\theta} \theta^x}{x!},$$

where θ is an unknown to be estimated from some observation data. We assume that θ follows a Gamma distribution denoted as $\Theta \sim \text{Gamma}(\alpha, \beta)$, given by

$$f_{\Theta}(\theta; \alpha, \beta) = \frac{\beta^{\alpha} \theta^{\alpha-1} e^{-\beta\theta}}{\Gamma(\alpha)},$$

where α and β are parameters of the distribution, and $\Gamma(\cdot)$ is the Gamma function. It is known that the mode of the Gamma distribution (with $\alpha \geq 1$) is given by $(\alpha - 1)/\beta$.


Assume we make a series of observations, X_i for $1 \leq i \leq n$, which are independent identically distributed, following the same Poisson distribution.

1. Show that the posterior is given by

$$f_{\theta|X_1, \dots, X_n}(\theta|x_1, \dots, x_n) = c \cdot e^{-(\beta+n)\theta} \cdot \theta^{(\alpha+s)-1},$$

where $s = \sum_{i=1}^n x_i$, and c is a certain constant.

2. Express the posterior in terms of a Gamma distribution.
3. Find the MAP estimate of θ .
4. What happens to the MAP estimate when n is large?



Question 9

We have developed an optical imaging system for single-cell inspection. Each cell can be either normal or cancerous. Through extensive studies, we find that the two types of cells have different physiological characteristics. In particular, it is possible to differentiate them in terms of the cell diameter.

For a normal cell, the mean diameter is $\mu_1 = 3 \mu\text{m}$, with standard deviation $\sigma_1 = 0.4 \mu\text{m}$.

For a cancerous cell, the mean diameter is $\mu_2 = 4.5 \mu\text{m}$, with standard deviation $\sigma_2 = 0.6 \mu\text{m}$.

In an experiment, we have a mixture containing 80% of normal and 20% of cancerous cells. We now image a specific cell and find that it has a diameter of $3.7 \mu\text{m}$.

1. We assume that both normal and cancerous cells follow Gaussian distributions, with mean and standard deviation as stated above.
 - (a) What is the posterior probability that the cell is normal?
 - (b) What is the posterior probability that the cell is cancerous?
 - (c) Using MAP, would you conclude the cell is normal or cancerous?
2. On the basis of some new theory, we change our models, and now both cell distributions are Laplacian centered around the mean value stated above. Again, their mean and standard deviation remain unchanged.

- (a) The PDF of normal cell is

$$f_X(x) = \frac{\lambda_1}{2} e^{-\lambda_1 |x - m_1|}.$$


Compute λ_1 and m_1 .

- (b) The PDF of cancerous cell is

$$f_X(x) = \frac{\lambda_2}{2} e^{-\lambda_2 |x - m_2|}.$$

Compute λ_2 and m_2 .

- (c) What is the posterior probability that the cell is normal?
- (d) What is the posterior probability that the cell is cancerous?
- (e) Using MAP, would you conclude the cell is normal or cancerous?




Question 10

Imagine that you are a data scientist working for an e-commerce platform. The company has recently implemented a new recommendation system, and you are tasked with evaluating its performance. Specifically, you are interested in the probability, $0 \leq \theta \leq 1$, that a user will click on a recommended product. This probability is unknown and you wish to estimate it using data collected from user interactions.

After monitoring the new recommendation system, you gather a small sample of user interactions: $x = (0, 1, 0, 1, 0)$, where a "1" indicates a click and a "0" indicates no click.

1. You want to estimate the probability of a user clicking on a recommendation using the MLE method. What is the maximum likelihood (ML) estimate for θ based on this data?
2. The marketing team believes that the click probability could only be one of three values based on previous marketing studies: 0.3, 0.5, 0.6. With this new restriction, how does the ML for θ change?
3. To refine the estimate, the team decides to incorporate prior marketing knowledge into the estimation process. They assign different levels of belief (prior probabilities) to each of the three possible values of θ : $\mathbb{P}(\theta = 0.3) = 0.05$, $\mathbb{P}(\theta = 0.5) = 0.01$, and $\mathbb{P}(\theta = 0.6) = 0.94$. How do you compute the MAP for θ under these new assumptions?
4. You then decide to conduct an experiment: you want to demonstrate to the team that by altering the prior probabilities, you can control the outcome of the MAP estimation. How would you set the priors for θ if you want the MAP estimate to be 0.3?
5. As you collect more and more samples, what would be the relationship between using ML and MAP for the estimation? why?




Question 11


Speckle Generation: When a laser beam reflects off a rough surface, the scattered waves interfere and form a random pattern called ‘speckle’. The intensity distribution of the speckle pattern can be modeled probabilistically.


1. The electric field S of the scattered wave at a given point can be represented as a sum of N random ‘phasors’: $S = \sum_{n=1}^N A_n e^{i\phi_n}$, where A_n and ϕ_n are random amplitude and random phase for the n -th phasor. Assume that A_n and ϕ_n are independent of each other. Moreover, $\phi_n \sim U(-\pi, \pi)$ and A_n are i.i.d. for $n = 1, \dots, N$.


Let $S = X + iY$, where X and Y are the real and imaginary parts of the complex sum.

- (a) Find the expressions for X and Y in terms of the random variables A_n and ϕ_n .
 - (b) Show that for a very large N , both X and Y become approximately Gaussian distributed with $X, Y \sim \mathcal{N}(0, \sigma^2)$, where $\sigma^2 = \frac{N}{2} \mathbb{E}(A_n^2)$.
2. Let the intensity of the speckle be the magnitude of the complex sum: $R = |S| = \sqrt{X^2 + Y^2}$.
 - (a) Use the joint distribution of X and Y in polar coordinates to derive the PDF of the speckle intensity R .
 - (b) Derive the CDF, mean, median, and mode for the intensity of the speckle pattern.
 3. Suppose that you are given a set of M independent intensity measurements $R_1, R_2, \dots, R_M$ for the speckle pattern.
 - (a) Determine the Maximum Likelihood estimator to estimate σ and σ^2 .
 - (b) Determine whether the estimators are unbiased.









Question 12

In a game, the players need to control a game character to defeat as many monsters as possible. The players will only encounter one monster each round. A player's chance of defeating a monster in each round is p . In addition, the game has the following characteristics:

- The game character has unlimited number of opportunities to defeat the monsters, although in each round, there is still only one monster.
- Once a player has defeated r number of monsters, he or she is considered to have “won the game.”
- Let K denote the number of failures before winning the game. This is a discrete random variable.

Our goal for this problem is to estimate p , based on the observed values of K for many, many players. We denote these observations as $k_1, k_2, \dots, k_n$.


1. What is the probability mass function of K ?
2. What is the maximum likelihood estimate of p ?
3. Assume that p follows a beta distribution (i.e., the prior), with

$$f_p(p) \sim \text{Beta}(p; \alpha, \beta) = \begin{cases} \frac{1}{B(\alpha, \beta)} p^{\alpha-1} (1-p)^{\beta-1} & 0 < p < 1 \\ 0 & \text{otherwise,} \end{cases}$$

where $B(\alpha, \beta)$ is a normalizing constant, known as the beta function, given by

$$B(\alpha, \beta) = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha + \beta)} = \frac{(\alpha - 1)!(\beta - 1)!}{(\alpha + \beta - 1)!}$$

What is the MAP estimate of p ?




Question 13


The number of photons X arriving at a detector within the detection duration t follows a Poisson distribution as

$$\mathbb{P}(X = k) = \frac{(\lambda t)^k}{k!} e^{-\lambda t} \quad k = 0, 1, \dots$$

where the parameter λ indicates the light intensity. In this question, we will explore two different ways of obtaining the light intensity λ via recoding and timing the arrival photons.

1. In the first way, the number of arrival photons within the detection duration t is encoded as a binary signal Y , with 0 indicating no photon arriving at the detector and 1 indicating there is at least one arrival photon.
 - (a) Show that Y follows a Bernoulli distribution, *i.e.*, $Y \sim \text{Ber}(\theta)$, where the parameter θ is determined by λ .
 - (b) Given a sequence of independent and identically distributed observations $y_1, \dots, y_n$ of the random variable Y , find the ML estimate of λ . (Hint: You can find the ML estimate of θ and determine λ using the relationship between θ and λ .)
2. In the second way, the detection duration is assumed to be long enough so that there are more than two arrival photons, and the time interval Z between the consecutive arrival photons is recorded.
 - (a) Show that the time interval Z follows an exponential distribution of which λ is the parameter, *i.e.*, $Z \sim \text{Exp}(\lambda)$.
 - (b) Given a set of independent and identically distributed observations $z_1, \dots, z_n$ of the random variable Z , find the ML estimate of λ .





Question 14

In cell biology, protein-protein interactions (PPIs) refer to contacts established between two protein molecules due to some biochemical events. They are foundational to carrying out various biological functions, such as DNA replication. In a collection of proteins, such interactions form a protein-protein interaction network (PPIN).

In this problem, we want to study a PPIN. Suppose certain drugs would facilitate the development of an interaction between two proteins. However, whether a PPI is formed between them is a random process, modeled by a Bernoulli random variable. Specifically, for protein P_i and protein P_j , where $i \neq j$ (a protein would not form a PPI with itself), their PPI is denoted as X_{ij} where


$$X_{ij} = \begin{cases} 1 & \text{if } (P_i, P_j) \text{ have interactions} \\ 0 & \text{otherwise.} \end{cases}$$

Assume, for simplicity, that they all follow the same distribution, i.e.,

$$X_{ij} \sim \text{Ber}(q) \quad \text{for } i \neq j.$$

The parameter q can be thought of as the *effectiveness* of the drug to form PPIs. Our goal is to estimate this q using maximum likelihood (ML).

1. What is the PMF of X_{ij} , denoted as $p_{X_{ij}}(x_{ij})$?
2. In a PPIN with N proteins, we observed the various x_{ij} . What is the log likelihood, denoted as $\log p_X(\mathbf{x}; q)$? (*note: X here would denote the list of all X_{ij} .*)
3. What is the ML estimate of q ?



Question 15

Suppose we collect n independent observations $Y_1, Y_2, \dots, Y_n$, each of which follows a normal distribution with mean equal to the product of a single unknown parameter θ and known constants $x_1, x_2, \dots, x_n$, and with a known variance σ^2 . In other words, $Y_i \sim \mathcal{N}(\theta x_i, \sigma^2)$, where x_i and σ^2 are fixed and known, while θ must be estimated.

1. Find the maximum likelihood estimator $\hat{\theta}$.
2. Before seeing any data, we assume a normal distribution prior on θ , with specified mean equal to θ_0 and variance equal to τ^2 .
 - (a) Derive the posterior distribution of θ . (Hint: It is also a normal distribution, but you need to derive expressions for the mean (θ_1) and variance (σ_1^2) of this normal distribution.)
 - (b) Show that we can express the mean of the posterior distribution as a weighted sum of the prior mean θ_0 and the maximum likelihood estimator $\hat{\theta}$.
 - (c) Find the limit of the posterior distribution as $\tau^2 \rightarrow \infty$, and compare the estimates given by ML and MAP.

