

Section #3 Solutions

Problems by Chris

Reminder: Sign up for a PEP session. Sessions take place next week Mon April 27 - Wed April 29.

1 Better Evaluation of Eye Disease

When a patient has eye inflammation, eye doctors “grade” the inflammation. When “grading” inflammation they randomly look at a single 1 millimeter by 1 millimeter square in the patient’s eye and count how many “cells” they see.

There is uncertainty in these counts. If the true average number of cells for a given patient’s eye is 6, the doctor could get a different count (say 4, or 5, or 7) just by chance. As of 2021, modern eye medicine did not have a sense of uncertainty for their inflammation grades! In this problem we are going to change that. At the same time we are going to learn about poisson distributions over space.




Figure 1: A $1 \times 1 \text{mm}$ sample used for inflammation grading. Inflammation is graded by counting cells in a randomly chosen 1mm by 1mm square. This sample has 5 cells.

- Explain, as if teaching, why the number of cells observed in a 1×1 square is governed by a poisson process. Make sure to explain how a binomial distribution could approximate the count of cells. Explain what λ means in this context. Note: for a given person’s eye, the presence of a cell in a location is independent of the presence of a cell in another location.

We can approximate a distribution for the count by discretizing the square into a fixed number of equal sized buckets. Each bucket either has a cell or not. Therefore, the count of cells in the 1×1 square is a sum of Bernoulli random variables with equal p , and as such can be modeled as a Binomial random variable.

However, this is an approximation, because it doesn’t allow for two or more cells to be in the same bucket (the Binomial models binary events only). Just like with time, if we make the size of each bucket infinitely small, this limitation goes away and we converge on the true distribution of counts. The Binomial in the limit, i.e. a Binomial as $n \rightarrow \infty$, is truly represented by a Poisson random variable.

In this context, λ represents the average number of cells per 1×1 sample. See Figure 2 below for a visualization.




Figure 2: X is counts of events in discrete buckets. In the limit, as n (number of buckets) $\rightarrow \infty$, X becomes a Poisson.

- b. For a given patient the true average rate of cells is 5 cells per 1×1 sample. What is the probability that in a single 1×1 sample the doctor counts 4 cells?

Let X denote the number of cells in the 1×1 sample. We note that $X \sim \text{Poi}(\lambda = 5)$. We want to find $P(X = 4)$. We can plug into the Poisson PMF:

$$P(X = 4) = \frac{5^4 e^{-5}}{4!} \approx 0.175$$

2 Server Downloads

If this problem doesn't convince you that the Poisson and Exponential RVs are coupled, then I'm not sure what will!

A server set up for large file downloads only has enough bandwidth for one download at a time. We can assume there is an infinite queue of downloads requested by users, such that immediately after one download finishes, another download always begins. On average, downloads take 5 minutes to complete, and the time until a download completes is exponentially distributed.

- a. Using the random variable X , defined as the length of time a download takes to complete, what is the probability that a download takes longer than 10 mins?

The problem tells us that length of download time, X , is exponentially distributed. To define λ , we need to choose a unit time. If we choose 1 minute, $X \sim \text{Exp}(\lambda = \frac{1}{5})$:

$$P(X > 10) = 1 - F_X(10) = 1 - (1 - e^{-10\lambda}) = e^{-2} \approx 0.1353$$

We could also choose other units of time, such as 10 minutes. Then, $X \sim \text{Exp}(\lambda = 2)$:

$$P(X > 1) = 1 - F_X(1) = 1 - (1 - e^{-1\lambda}) = e^{-2} \approx 0.1353$$

- b. Using the random variable Y , defined as the number of downloads that finish over a 10-minute interval, what is the probability that a download takes longer than 10 mins?

Y is the number of downloads that finish in the next 10 minutes. If one download finishes every 5 minutes on average, then the average number of downloads finishing every 10 minutes is 2. $Y \sim \text{Poi}(\lambda = 2)$.

$$\begin{aligned} P(Y = 0) &= \frac{2^0 e^{-2}}{0!} \\ &= e^{-2} \approx 0.1353 \end{aligned}$$

In this case, we had only one choice for the unit time to define λ for, because only a unit time of 10 minutes allowed us to then use the PMF to ask for the probability of 0 downloads finishing in a 10 minute time window.

3 Popcorn

A bag of popcorn contains special kernels that, once placed in the microwave, pop after an amount of time that is normally distributed with mean $\mu = 100$ seconds and standard deviation $\sigma = 20$ seconds.

A particular kernel has not popped after 110 seconds. Given this information, what is the probability that the kernel will pop in the *next second*? (You can leave your answer in terms of Φ).

Let X be the amount of time it takes for a kernel to pop. We can model $X \sim N(\mu = 100, \sigma = 20)$. We know $X > 110$, and we want to find the probability that $X \leq 111$. This is the same

as calculating $P(X \leq 111|X > 110)$. Using conditional probability, we can express this as:

$$\begin{aligned} P(X \leq 111|X > 110) &= \frac{P(110 < X \leq 111)}{P(X > 110)} \\ &= \frac{P(110 < X \leq 111)}{1 - P(X \leq 110)} \\ &= \frac{\Phi\left(\frac{111-\mu}{\sigma}\right) - \Phi\left(\frac{110-\mu}{\sigma}\right)}{1 - \Phi\left(\frac{110-\mu}{\sigma}\right)} \\ &= \frac{\Phi(0.55) - \Phi(0.5)}{1 - \Phi(0.5)} \\ &\approx 0.056 \end{aligned}$$