

CSE 312: Foundations of Computing II

Section 9: Maximum Likelihood, Markov Chains Solutions

1. Review of Main Concepts

- (a) **Realization/Sample:** A realization/sample x of a random variable X is the value that is actually observed.
- (b) **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:

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

- (c) **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).

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

- (d) **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:

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

- (e) **Bias:** The bias of an estimator $\hat{\theta}$ for a true parameter θ is defined as $\text{Bias}(\hat{\theta}, \theta) = \mathbb{E}[\hat{\theta}] - \theta$. An estimator $\hat{\theta}$ of θ is unbiased iff $\text{Bias}(\hat{\theta}, \theta) = 0$, or equivalently $\mathbb{E}[\hat{\theta}] = \theta$.

- (f) **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.

- (g) A **discrete-time stochastic process (DTSP)** is a sequence of random variables X_0, X_1, X_2, \dots , where X_t is the value at time t . For example, the temperature in Seattle or stock price of TESLA each day, or which node you are at after each time step on a random walk on a graph.

- (h) A **Markov Chain** is a DTSP, with the additional following three properties:

- I. ...has a finite (or countably infinite) **state space** $\mathcal{S} = \{s_1, \dots, s_n\}$ which it bounces between, so each $X_t \in \mathcal{S}$.
- II. ...satisfies the **Markov property**. A DTSP satisfies the Markov property if the future is (conditionally) independent of the past given the present. Mathematically, it means, $P(X_{t+1} = x_{t+1} | X_0 = x_0, X_1 = x_1, \dots, X_{t-1} = x_{t-1}, X_t = x_t) = P(X_{t+1} = x_{t+1} | X_t = x_t)$.
- III. ...has **stationary transition probabilities**. Meaning, if we are at some state s_i , we transition to another state s_j with probability *independent* of the current time. Due to this property and the previous, the transitions are governed by n^2 probabilities: the probability of transitioning from one of n current states to one of n next states. These are stored in a square $n \times n$ **transition probability matrix (TPM)** P , where $P_{ij} = P(X_{t+1} = s_j | X_t = s_i)$ is the probability of transitioning from state s_i to state s_j for any/every value of t .

2. 312 Grades

Suppose Professor Karlin loses everyone's grades for 312 and decides to make it up by assigning grades randomly according to the following probability distribution, and hoping the n students won't notice: give an A with probability 0.5, a B with probability θ , a C with probability 2θ , and an F with probability $0.5 - 3\theta$. Each student is assigned a grade independently. Let x_A be the number of people who received an A, x_B the number of people who received a B, etc, where $x_A + x_B + x_C + x_F = n$. Find the MLE for θ , $\hat{\theta}$.

Solution:

The data tells us, for each student in the class, what their grade was. We begin by computing the likelihood of seeing the given data given our parameter θ . Because each student is assigned a grade independently, the likelihood is equal to the product over students of the chance they got the particular grade they got, which gives us:

$$L(x|\theta) = 0.5^{x_A} \theta^{x_B} (2\theta)^{x_C} (0.5 - 3\theta)^{x_F}$$

From there, we just use the MLE process to get the log-likelihood, take the first derivative, set it equal to 0, and solve for $\hat{\theta}$.

$$\ln L(x|\theta) = x_A \ln(0.5) + x_B \ln(\theta) + x_C \ln(2\theta) + x_F \ln(0.5 - 3\theta)$$

$$\frac{\partial}{\partial \theta} \ln L(x|\theta) = \frac{x_B}{\theta} + \frac{x_C}{\theta} - \frac{3x_F}{0.5 - 3\theta} = 0$$

Solving yields $\hat{\theta} = \frac{x_B + x_C}{6(x_B + x_C + x_F)}$.

3. A Red Poisson

Suppose that x_1, \dots, x_n are i.i.d. samples from a $\text{Poisson}(\theta)$ random variable, where θ is unknown. Find the MLE of θ .

Solution:

Because each Poisson RV is i.i.d., the likelihood of seeing that data is just the PMF of the Poisson distribution multiplied together for every x_i . From there, take the log-likelihood, then the first derivative, set it equal to 0

and solve for for $\hat{\theta}$.

$$\begin{aligned}
 L(x_1, \dots, x_n \mid \theta) &= \prod_{i=1}^n e^{-\theta} \frac{\theta^{x_i}}{x_i!} \\
 \ln L(x_1, \dots, x_n \mid \theta) &= \sum_{i=1}^n [-\theta - \ln(x_i!) + x_i \ln(\theta)] \\
 \frac{\partial}{\partial \theta} \ln L(x_1, \dots, x_n \mid \theta) &= \sum_{i=1}^n \left[-1 + \frac{x_i}{\theta}\right] = 0 \\
 -n + \frac{\sum_{i=1}^n x_i}{\hat{\theta}} &= 0 \\
 \hat{\theta} &= \frac{\sum_{i=1}^n x_i}{n}
 \end{aligned}$$

4. Independent Shreds, You Say?

You are given 100 independent samples x_1, x_2, \dots, x_{100} from Bernoulli(θ), where θ is unknown. (Each sample is either a 0 or a 1). These 100 samples sum to 30. You would like to estimate the distribution's parameter θ . Give all answers to 3 significant digits.

(a) What is the maximum likelihood estimator $\hat{\theta}$ of θ ?

Solution:

Note that $\sum_{i \in [n]} x_i = 30$, as given in the problem spec. Therefore, there are 30 **1**s and 70 **0**s. (Note that they come in some specific order.) Therefore, we can setup L as follows, because there is a θ chance of getting a 1, and a $(1 - \theta)$ chance of getting a 0 and they are each i.i.d. From there, take the log-likelihood, then the first derivative, set it equal to 0 and solve for for $\hat{\theta}$.

$$\begin{aligned}
 L(x_1, \dots, x_n \mid \theta) &= (1 - \theta)^{70} \theta^{30} \\
 \ln L(x_1, \dots, x_n \mid \theta) &= 70 \ln(1 - \theta) + 30 \ln \theta \\
 \frac{\partial}{\partial \theta} \ln L(x_1, \dots, x_n \mid \theta) &= -\frac{70}{1 - \theta} + \frac{30}{\theta} = 0 \\
 \frac{30}{\hat{\theta}} &= \frac{70}{1 - \hat{\theta}} \\
 30 - 30\hat{\theta} &= 70\hat{\theta} \\
 \hat{\theta} &= \frac{30}{100}
 \end{aligned}$$

(b) Is $\hat{\theta}$ an unbiased estimator of θ ?

Solution:

An estimator is unbiased if the expectation of the estimator is equal to the original parameter, i.e.: $E[\hat{\theta}] = \theta$. Setting up the expectation of our estimator and plugging it in for the generic case, we get the following, which we can then reduce with linearity of expectation:

$$\begin{aligned}
 \mathbb{E}[\hat{\theta}] &= \mathbb{E}\left[\frac{1}{100} \sum_{i=1}^{100} X_i\right] \\
 &= \frac{1}{100} \sum_{i=1}^{100} \mathbb{E}[X_i] \\
 &= \frac{1}{100} \cdot 100\theta = \theta.
 \end{aligned}$$

so it is unbiased.

5. 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 .

Solution:

Since the samples are i.i.d., the likelihood of seeing n samples of them is just their PDFs multiplied together. From there, take the log-likelihood, then the first derivative, set it equal to 0 and solve for for $\hat{\theta}$.

$$\begin{aligned} L(y_1, \dots, y_n | \theta) &= \prod_{i=1}^n \frac{1}{2\theta} \exp\left(-\frac{|y_i|}{\theta}\right) \\ \ln L(y_1, \dots, y_n | \theta) &= \sum_{i=1}^n \left[-\ln 2 - \ln \theta - \frac{|y_i|}{\theta} \right] \\ \frac{\partial}{\partial \theta} \ln L(y_1, \dots, y_n | \theta) &= \sum_{i=1}^n \left[-\frac{1}{\theta} + \frac{|y_i|}{\theta^2} \right] = 0 \\ -\frac{n}{\theta} + \frac{\sum_{i=1}^n |y_i|}{\theta^2} &= 0 \\ \hat{\theta} &= \frac{\sum_{i=1}^n |y_i|}{n} \end{aligned}$$

6. 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?

Solution:

Let $\bar{X} = \frac{1}{n} \sum_{i=1}^n X_i$. Then

$$E(\hat{\theta}_2) = E\left(\frac{1}{n} \sum_{i=1}^n (X_i - \bar{X})^2\right) = E\left(\frac{1}{n} \sum_{i=1}^n (X_i^2 - 2X_i\bar{X} + \bar{X}^2)\right)$$

which by linearity of expectation (and distributing the sum) is

$$\begin{aligned} &= \frac{1}{n} \sum_{i=1}^n E(X_i^2) - E\left(\frac{2}{n} \bar{X} \sum_{i=1}^n X_i\right) + E(\bar{X}^2) \\ &= \frac{1}{n} \sum_{i=1}^n E(X_i^2) - 2E(\bar{X}^2) + E(\bar{X}^2) \\ &= \frac{1}{n} \sum_{i=1}^n E(X_i^2) - E(\bar{X}^2). \quad (**) \end{aligned}$$

We know that for any random variable Y , since $Var(Y) = E(Y^2) - (E(Y))^2$ it holds that

$$E(Y^2) = Var(Y) + (E(Y))^2.$$

Also, we have $E(X_i) = \mu$, $Var(X_i) = \sigma^2 \forall i$ and $E(\bar{X}) = \mu$, $Var(\bar{X}) = \frac{\sigma^2}{n}$. Combining these facts, we get

$$E(X_i^2) = \sigma^2 + \mu^2 \quad \forall i \quad \text{and} \quad E(\bar{X}^2) = \frac{\sigma^2}{n} + \mu^2.$$

Substituting these equations into (**) we get

$$\begin{aligned} E\left(\frac{1}{n} \sum_{i=1}^n (X_i - \bar{X})^2\right) &= \frac{1}{n} \sum_{i=1}^n E(X_i^2) - E(\bar{X}^2) = \sigma^2 + \mu^2 - \left(\frac{\sigma^2}{n} + \mu^2\right) \\ &= \left(1 - \frac{1}{n}\right) \sigma^2. \end{aligned}$$

Thus $\hat{\theta}_2$ is not unbiased.

7. Faulty Machines

You are trying to use a machine that only works on some days. If on a given day, the machine is working it will break down the next day with probability $0 < b < 1$, and works on the next day with probability $1 - b$. If it is not working on a given day, it will work on the next day with probability $0 < r < 1$ and not work the next day with probability $1 - r$.

- (a) In this problem we will formulate this process as a Markov chain. First, let X_t be a random variable that denotes the state of the machine at time t . Then, define a state space \mathcal{S} that includes all the possible states that the machine can be in. Lastly, for all $A, B \in \mathcal{S}$ find $\mathbb{P}(X_{t+1} = A \mid X_t = B)$ (A and B can be the same state).

Solution:

Formally, a Markov chain is defined by a state space \mathcal{S} and a transition probability matrix. The two possible states of the machine are “working” and “broken”. So, $\mathcal{S} = \{W, B\}$. Let X_t be the state of the process at time t . Then we can define the following transition probabilities:

$$\mathbb{P}(X_{t+1} = W \mid X_t = W) = 1 - b \quad \mathbb{P}(X_{t+1} = B \mid X_t = W) = b$$

$$\mathbb{P}(X_{t+1} = W \mid X_t = B) = r \quad \mathbb{P}(X_{t+1} = B \mid X_t = B) = 1 - r$$

We can also represent the TPM with the following matrix:

$$P = \begin{bmatrix} 1 - b & b \\ r & 1 - r \end{bmatrix}$$

where the ij th entry is probability that the machine is in the j th state at time $t + 1$ given it was in state i at time t . (Here state 1 is working and state 2 is broken.)

- (b) Suppose that on day 1, the machine is working. What is the probability that it is working on day 3?

Solution:

We are trying to find $\mathbb{P}(X_3 = W \mid X_1 = W)$. From the law of total probability, and then plugging in the values from our transition matrix:

$$\begin{aligned}
P(X_3 = W \mid X_1 = W) &= \sum_{i \in \mathcal{S}} \mathbb{P}(X_3 = W \mid X_1 = W, X_2 = i) \mathbb{P}(X_2 = i \mid X_1 = W) \\
&= \mathbb{P}(X_3 = W \mid X_2 = W) \mathbb{P}(X_2 = W \mid X_1 = W) + \mathbb{P}(X_3 = W \mid X_2 = B) \mathbb{P}(X_2 = B \mid X_1 = W) \\
&= \mathbb{P}(X_3 = W \mid X_2 = W) (1 - b) + \mathbb{P}(X_3 = W \mid X_2 = B) b \\
&= (1 - b)(1 - b) + rb \\
&= (1 - b)^2 + rb
\end{aligned}$$

- (c) As $n \rightarrow \infty$, what does the probability that the machine is working on day n converge to? To get the answer, solve for the *stationary distribution*.

Solution:

The stationary distribution is the row vector $\pi = [\pi_W \ \pi_B]$ such that $\pi P = \pi$. The entries in the vector π_W and π_B can be interpreted as the probabilities that the machine works or is broken converge to. As such, $\pi_W + \pi_B = 1$. Additionally, multiplying the stationary distribution by the TPM gives us the following two equations:

$$\pi_W = \pi_W(1 - b) + \pi_B r \quad \pi_B = \pi_W b + \pi_B(1 - r)$$

Solving each for π_W and π_B gives us the following solutions for the stationary distribution:

$$\pi_W = \frac{r}{b + r} \quad \pi_B = \frac{b}{b + r}$$

So, as $n \rightarrow \infty$ the probability that the machine works on day n is $\pi_W = \frac{r}{b+r}$

8. Three tails

You flip a fair coin until you see three tails in a row. Model this as a Markov chain with the following states:

- S : start state, which we are only in before flipping any coins.
- H : We see a heads, which means no streak of tails currently exists.
- T : We've seen exactly one tail in a row so far.
- TT : We've seen exactly two tails in a row so far.
- TTT : We've accomplished our goal of seeing three tails in a row and stop flipping.

- (a) Write down the transition probability matrix.

Solution:

$$P = \begin{bmatrix} 0 & 1/2 & 1/2 & 0 & 0 \\ 0 & 1/2 & 1/2 & 0 & 0 \\ 0 & 1/2 & 0 & 1/2 & 0 \\ 0 & 1/2 & 0 & 0 & 1/2 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

- (b) Write down the system of equations whose variables are $D(s)$ for each state $s \in \{S, H, T, TT, TTT\}$, where $D(s)$ is the expected number of steps until state TTT is reached starting from state s . Solve this system of equations to find $D(S)$.

Solution:

Using the law of total expectation and the TPM above we can set up and solve the following system of equations:

$$\begin{aligned}
D(TTT) &= 0 \\
D(TT) &= 1 + \frac{1}{2}D(H) + \frac{1}{2}D(TTT) = \frac{1}{2}D(H) + 1 \\
D(T) &= 1 + \frac{1}{2}D(H) + \frac{1}{2}D(TT) = \frac{3}{4}D(H) + \frac{3}{2} \\
D(H) &= 1 + \frac{1}{2}D(H) + \frac{1}{2}D(T) = \frac{7}{8}D(H) + \frac{7}{4} \\
D(S) &= 1 + \frac{1}{2}D(H) + \frac{1}{2}D(T) = \frac{7}{8}D(H) + \frac{7}{4}
\end{aligned}$$

Solving for $D(H)$ gives us that $D(H) = 14$, which allows us to solve for the rest of the expected number of steps, $D(TT) = 8$, $D(T) = 12$, $D(S) = 14$. So, we expect to flip 14 coins before we flip three tails in a row.

- (c) Write down the system of equations whose variables are $\gamma(s)$ for each state $s \in \{S, H, T, TT, TTT\}$, where $\gamma(s)$ is the expected number of heads seen before state TTT is reached. Solve this system to find $\gamma(S)$, the expected number of heads seen overall until getting three tails in a row.

Solution:

Like in the previous part we can use the LoTE and the TPM to set up and solve the following system of equations:

$$\begin{aligned}
\gamma(TTT) &= 0 \\
\gamma(TT) &= 0.5\gamma(H) + 0.5\gamma(TTT) = 0.5\gamma(H) \\
\gamma(T) &= 0.5\gamma(H) + 0.5\gamma(TT) = 0.75\gamma(H) \\
\gamma(H) &= 1 + 0.5\gamma(H) + 0.5\gamma(T) = 0.875\gamma(H) + 1 \\
\gamma(S) &= 0.5\gamma(H) + 0.5\gamma(T) = 0.875\gamma(H)
\end{aligned}$$

Solving for $\gamma(H)$ gives us $\gamma(H) = 8$. This allows us to solve for the other expected values which are $\gamma(TT) = 4$, $\gamma(T) = 6$, $\gamma(S) = 7$. So, we expect to see 7 heads before we flip three tails in a row.

9. Another Markov chain

Suppose that the following is the transition probability matrix for a 4 state Markov chain (states 1,2,3,4).

$$P = \begin{bmatrix} 0 & 1/2 & 1/2 & 0 \\ 1/3 & 0 & 0 & 2/3 \\ 1/3 & 1/3 & 0 & 1/3 \\ 1/5 & 2/5 & 2/5 & 0 \end{bmatrix}$$

- (a) What is the probability that $X_2 = 4$ given that $X_0 = 4$?

Solution:

Let's denote the state space $\mathcal{S} = \{1, 2, 3, 4\}$. Using the law of total probability we can determine that

$$\begin{aligned}
\mathbb{P}(X_2 = 4 \mid X_0 = 4) &= \sum_{i \in \mathcal{S}} \mathbb{P}(X_2 = 4 \mid X_0 = 4, X_1 = i) \mathbb{P}(X_1 = i \mid X_0 = 4) \\
&= \sum_{i \in \mathcal{S}} \mathbb{P}(X_2 = 4 \mid X_1 = i) \mathbb{P}(X_1 = i \mid X_0 = 4) \\
&= 0 + \frac{2}{5} \cdot \frac{2}{3} + \frac{2}{5} \cdot \frac{1}{3} + 0 \\
&= \frac{2}{5}
\end{aligned}$$

(b) Write down the system of equations that the stationary distribution must satisfy and solve them.

Solution:

The stationary distribution is the row vector $\pi = [\pi_1 \ \pi_2 \ \pi_3 \ \pi_4]$ such that $\pi P = \pi$. We know that $\pi_1 + \pi_2 + \pi_3 + \pi_4 = 1$. Additionally, multiplying the stationary distribution by the TPM gives us the following equations:

$$\begin{aligned}
\pi_1 &= \frac{1}{3}\pi_2 + \frac{1}{3}\pi_3 + \frac{1}{5}\pi_4 \\
\pi_2 &= \frac{1}{2}\pi_1 + \frac{1}{3}\pi_3 + \frac{2}{5}\pi_4 \\
\pi_3 &= \frac{1}{2}\pi_1 + \frac{2}{5}\pi_4 \\
\pi_4 &= \frac{2}{3}\pi_2 + \frac{1}{3}\pi_3
\end{aligned}$$

Solving for each π_i gives us the following solutions for the stationary distribution:

$$\pi_1 = \frac{46}{206} \quad \pi_2 = \frac{60}{206} \quad \pi_3 = \frac{45}{206} \quad \pi_4 = \frac{55}{206}$$

10. Law of Total Probability Review

(a) (Discrete version) Suppose we flip a coin with probability U of heads, where U is equally likely to be one of $\Omega_U = \{0, \frac{1}{n}, \frac{2}{n}, \dots, 1\}$ (notice this set has size $n + 1$). Let H be the event that the coin comes up heads. What is $\mathbb{P}(H)$?

Solution:

We can use the law of total probability, conditioning on $U = \frac{k}{n}$ for $k = 0, \dots, n$. Note that the probability of getting heads conditioning on a fixed U value is U , and that the probability of U taking on any value in its range is $\frac{1}{n+1}$ since it is discretely uniform.

$$\mathbb{P}(H) = \sum_{k=0}^n \mathbb{P}(H \mid U = \frac{k}{n}) \mathbb{P}(U = \frac{k}{n}) = \sum_{k=0}^n \frac{k}{n} \cdot \frac{1}{n+1} = \frac{1}{n(n+1)} \sum_{k=0}^n k = \frac{1}{n(n+1)} \frac{n(n+1)}{2} = \frac{1}{2}$$

(b) (Continuous version) Now suppose $U \sim \text{Uniform}(0,1)$ has the *continuous* uniform distribution over the interval $[0, 1]$. What is $\mathbb{P}(H)$?

Solution:

We do the same thing, this time using the continuous law of total probability. Note, this time, that we're conditioning on $U = u$ and taking the integral with respect to u , and that the density of U for any value in its range is 1 because it is uniformly random.

$$\mathbb{P}(H) = \int_{-\infty}^{\infty} \mathbb{P}(H|U = u)f_U(u)du$$

We can take the integral from 0 to 1 instead because outside of that range the density of U is 0.

$$= \int_0^1 \mathbb{P}(H|U = u)f_U(u)du = \int_0^1 u \cdot 1du = \frac{1}{2}[u^2]_0^1 = \frac{1}{2}$$

(c) Let's generalize the previous result we just used. Suppose E is an event, and X is a continuous random variable with density function $f_X(x)$. Write an expression for $\mathbb{P}(E)$, conditioning on X .

Solution:

We use the continuous law of total probability again, this time not deriving it any further and sticking with negative infinity to infinity because we don't know the range of the RV X .

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

11. Poisson CLT practice

Suppose X_1, \dots, X_n are iid $\text{Poisson}(\lambda)$ random variables, and let $\bar{X}_n = \frac{1}{n} \sum_{i=1}^n X_i$, the sample mean. How large should we choose n to be such that $\mathbb{P}(\frac{\lambda}{2} \leq \bar{X}_n \leq \frac{3\lambda}{2}) \geq 0.99$? Use the CLT and give an answer involving $\Phi^{-1}(\cdot)$. Then evaluate it exactly when $\lambda = 1/10$ using the Φ table on the last page.

Solution:

We know $\mathbb{E}[X_i] = \text{Var}(X_i) = \lambda$. By the CLT, $\bar{X}_n \approx \mathcal{N}(\lambda, \frac{\lambda}{n})$, so we can standardize this normal approximation.

$$\begin{aligned} \mathbb{P}(\frac{\lambda}{2} \leq \bar{X}_n \leq \frac{3\lambda}{2}) &\approx \mathbb{P}(\frac{-\lambda/2}{\sqrt{\lambda/n}} \leq Z \leq \frac{\lambda/2}{\sqrt{\lambda/n}}) = \Phi\left(\frac{\lambda/2}{\sqrt{\lambda/n}}\right) - \Phi\left(\frac{-\lambda/2}{\sqrt{\lambda/n}}\right) \\ &= \Phi\left(\frac{\lambda/2}{\sqrt{\lambda/n}}\right) - \left(1 - \Phi\left(\frac{\lambda/2}{\sqrt{\lambda/n}}\right)\right) = 2\Phi\left(\frac{\lambda/2}{\sqrt{\lambda/n}}\right) - 1 \geq 0.99 \rightarrow \Phi\left(\frac{\lambda/2}{\sqrt{\lambda/n}}\right) \geq 0.995 \\ &\rightarrow \frac{\sqrt{\lambda}}{2}\sqrt{n} \geq \Phi^{-1}(0.995) \rightarrow n \geq \frac{4}{\lambda} [\Phi^{-1}(0.995)]^2 \end{aligned}$$

We have $\lambda = \frac{1}{10}$ and from the table, $\Phi^{-1}(0.995) \approx 2.575$ so that $n \geq \frac{4}{1/10} \cdot 2.575^2 = 265.225$. So $n = 266$ is the smallest value that will satisfy the condition.