

Computational Methods in Stochastics

Lecture III

This time:

1. Conditional Probability and Expectation
2. Markov Processes and Stochastic Models

Conditional Probability and Conditional Expectation

Conditional Probability and Expectation

In what follows, the main properties of Markov processes needed for understanding advanced Monte Carlo methods are presented.

(Notation: $\Pr\{\cdot\} = P\{\cdot\} =$ Probability of event \cdot)

[Chapter 5](#) of the online book is worth skimming through.

Conditional Probability and Expectation

Discrete variables

The conditional probability of the event A given the event B :

$$\Pr\{A|B\} = \frac{\Pr\{A \cap B\}}{\Pr\{B\}} \quad \text{if } \Pr\{B\} > 0.$$

$\Pr\{A|B\}$ is not defined when $\Pr\{B\} = 0$.

Let X and Y be random variables that can attain *countably many* values. The *probability mass function* of X given $Y = y$:

$$p_{X|Y}(x|y) = \frac{\Pr\{X = x \text{ and } Y = y\}}{\Pr\{Y = y\}} \quad \text{if } \Pr\{Y = y\} > 0.$$

$p_{X|Y}(x|y)$ is not defined when $\Pr\{Y = y\} = 0$.

Conditional Probability and Expectation

Or, in terms of joint and marginal probability mass functions:

$$p_{X|Y}(x|y) = \frac{p_{XY}(x, y)}{p_Y(y)} \quad \text{if } p_Y(y) > 0; \quad x, y = 0, 1, \dots$$

Note: In ratios Pr's, PMFs, and PDFs all work equally.

Two variates in the conditional that can be summed over:

1. The “given” variate → probability of a value

The law of total probabilities

$$\Pr\{X = x\} = \sum_{y=0}^{\infty} p_{X|Y}(x|y)p_Y(y)$$

The law of total probabilities is used to obtain marginal probabilities (for examples, see [Ch5](#) of the online book.) →

Conditional Probability and Expectation

Example. Let X have a binomial distribution with parameters p and N , where N has a binomial distribution with parameters q and M . What is the *marginal distribution* of X ?

Conditional prob. mass function

$$p_{X|N}(k|n) = \binom{n}{k} p^k (1-p)^{n-k}, k = 0, 1, \dots, n.$$

Marginal distribution

$$p_N(n) = p_{XN}(k, n) = \binom{M}{n} q^n (1-q)^{M-n}, n = 0, 1, \dots, M.$$

The marginal distribution of X (the law of total probability)

$$\begin{aligned} \Pr\{X = k\} &= \sum_{n=0}^M p_{X|N}(k|n) p_N(n) = \dots = && \text{(see Introduction to} \\ &&& \text{Stochastic Modeling:} \\ &&& \text{Ch2)} \\ &= \frac{M!}{k! (M-k)!} (pq)^k (1-pq)^{M-k}, k = 0, 1, \dots, M. \end{aligned}$$

→ “ X has a binomial distribution with parameters M and pq .”

Conditional Probability and Expectation

Example 2. Let X have a binomial distribution with parameters p and N , where N has a Poisson distribution with mean λ . What is the *marginal distribution* of X ?

Conditional probability mass function

$$p_{X|N}(k|n) = \binom{n}{k} p^k (1-p)^{n-k}, k = 0, 1, \dots, n.$$

Marginal distribution

$$p_N(n) = \frac{\lambda^n e^{-\lambda}}{n!}, \quad n = 0, 1, \dots$$

The marginal distribution of X (the law of total probability)

$$\begin{aligned} \Pr\{X = k\} &= \sum_{n=0}^{\infty} p_{X|N}(k|n) p_N(n) = \\ &= \sum_{n=0}^{\infty} \frac{n!}{k! (n-k)!} p^k (1-p)^{n-k} \frac{\lambda^n e^{-\lambda}}{n!} = \dots = \frac{(\lambda p)^k e^{-\lambda p}}{k!} \quad \text{for } k = 0, 1, \dots \end{aligned}$$

→ “ X has a Poisson distribution with mean λp .”

Conditional Probability and Expectation

2. Summing over the “outcome” variate → expectation

Let $g(x)$ be a function for which the expectation of $g(X)$ is finite.

Then the *conditional expected value* of $g(X)$ given $Y = y$ is

$$E[g(X)|Y = y] = \sum_x g(x)p_{X|Y}(x|y) \quad \text{if } p_Y(y) > 0$$

(This conditional mean is not defined at y for which $p_Y(y) = 0$.)

The law of total probability →

$$E[g(X)] = E\{E[g(X)|Y]\}$$

(Summing over both variates.)

Conditional Probability and Expectation

Discrete and continuous variables

Let X and N be jointly distributed random variables, where X is continuous and N has the discrete set of values $n = 0, 1, 2, \dots$

The conditional distribution function $F_{X|N}(x|n)$ of the random variable X , given that $N = n$:

$$F_{X|N}(x|n) = \frac{\Pr\{X \leq x \text{ and } N = n\}}{\Pr\{N = n\}} \quad \text{if } \Pr\{N = n\} > 0.$$

$F_{X|N}(x|n)$ is not defined at values of n for which $\Pr\{N = n\} = 0$.

The *conditional probability density function*:

$$f_{X|N}(x|n) = \frac{d}{dx} F_{X|N}(x|n) \quad \text{if } \Pr\{N = n\} > 0.$$

Conditional Probability and Expectation

$$\Rightarrow \Pr\{a \leq X < b, N = n\} = \int_a^b f_{X|N}(x|n) p_N(n) dx$$

for $a < b$ and where $p_N(n) = \Pr\{N = n\}$.

The law of total probability $\rightarrow f_X(x) = \sum_{n=0}^{\infty} f_{X|N}(x|n) p_N(n)$

For the function g such that $E[g(X)] < \infty$, the *conditional expectation* of $g(X)$ given that $N = n$:

$$E[g(X)|N = n] = \int g(x) f_{X|N}(x|n) dx$$

The law of total probability:

$$E[g(X)] = \sum_{n=0}^{\infty} E[g(X)|N = n] p_N(n) = E\{E[g(X)|N]\}$$

Conditional Probability and Expectation

The previous derivations show that the same rules apply for discrete and continuous variables, one only needs to use $\sum \dots$ or $\int \dots$, respectively. This is the basic interpretation of the Lebesgue-Stieltjes integral $I = \int f(x)dg(x)$, where $g(x)$ is of bounded variation within the integration interval (see Lecture I: Random Variables). Remember, in computation you draw samples from g to compute I ; dg defines your (probability) measure. (Lecture II, p. 3.)

Likelihood

Suppose X_1, \dots, X_n are i.i.d. observations from the joint density $f(\mathbf{x}|\theta_1, \dots, \theta_k)$. The *likelihood function* is

$$L(\theta|\mathbf{x}) = L(\theta_1, \dots, \theta_k | x_1, \dots, x_n) = \prod_{i=1}^n f(x_i | \theta_1, \dots, \theta_k).$$

More generally, when the X_i 's are not i.i.d., the likelihood is defined as the joint density $f(x_1, \dots, x_n | \theta)$. The value of $\theta = \hat{\theta}$ at which $L(\theta|\mathbf{x})$ attains its maximum with \mathbf{x} held fixed is the *maximum likelihood estimator* (MLE).

Example. Gamma MLE. i.i.d. observations X_1, \dots, X_n from the *gamma density*

$$f(x|\alpha, \beta) = \frac{1}{\Gamma(\alpha)\beta^\alpha} x^{\alpha-1} e^{-x/\beta} \quad , \text{ where } \alpha \text{ is known.}$$

Likelihood

The *log likelihood*:

$$\begin{aligned}\log L(\alpha, \beta | x_1, \dots, x_n) &= \log \prod_{i=1}^n f(x_i | \alpha, \beta) = \\ &= -n \log \Gamma(\alpha) - n\alpha \log \beta + (\alpha - 1) \sum_{i=1}^n \log x_i - \sum_{i=1}^n x_i / \beta.\end{aligned}$$

Solving for $\frac{\partial}{\partial \beta} \log L(\alpha, \beta | x_1, \dots, x_n) = 0$ yields the MLE of β ,

$$\hat{\beta} = \sum_{i=1}^n \frac{x_i}{n\alpha}.$$

Random Sums

One often encounters *conditional distributions of discrete* (i.e. the random number of summands) *and continuous variables* in connection with *random sums*.

In random sums the number of summands varies randomly.

$$X = \xi_1 + \cdots + \xi_N,$$

where N is random and ξ_1, ξ_2, \dots is a sequence of independent and identically distributed (i.i.d.) random variables. X is taken as 0 when $N = 0$.

Random sums arise frequently within stochastic processes (e.g. in queuing-type processes).

Martingales

An important class of stochastic processes is determined by them having the *Martingale property*.

Definition.

A stochastic process $\{X_n; n = 0, 1, \dots\}$ is a *martingale* if for $n = 0, 1, \dots$

$$(a) E[|X_n|] < \infty$$

and

$$(b) E[X_{n+1}|X_0, \dots, X_n] = X_n.$$

$$\Rightarrow E\{E[X_{n+1}|X_0, \dots, X_n]\} = E\{X_n\}.$$

Using the total probability in the form

$$E\{E[X_{n+1}|X_0, \dots, X_n]\} = E\{X_{n+1}\},$$

gives $E[X_{n+1}] = E[X_n]$.

Martingales

A martingale has constant mean

$$E[X_0] = E[X_k] = E[X_n], \quad \text{for } m \geq n.$$

The martingale equality (b) extends to future times:

$$E[X_m | X_0, \dots, X_n] = X_n \quad \text{for } m \geq n$$

Many important stochastic properties have the Martingale property, e.g. unbiased diffusion or random walk and gambler's fortune (a concept of importance in the stock market). In "a game" (like finance) the property is related to "fairness" – unpredictability; chances are the same regardless of the past.

Markov Processes and Stochastic Models

Markov Chains

Online book, [Ch 11.2](#).

Some definitions

Stochastic process

A stochastic process is a random variable, which evolves through time.

Markov process

A Markov process $\{X_t\}$ is a stochastic process with the property that, given the value of X_t , the values of X_s for $s > t$ are not influenced by the values of X_μ for $\mu < t$.

In words, the probability of any particular future behaviour of the process, when its current state is known exactly, is not altered by additional knowledge concerning its past behaviour.

Markov Chains

To make the transition matrix notation as clear as possible, we first adapt the notation in 'An Introduction to Stochastic Modeling'. Later, the notation varies slightly depending on the system, but we'll make a due note of that.

A discrete-time Markov chain: a Markov process whose state space is a finite or countable set and whose time index set is $T = (0, 1, 2, \dots)$.

Markov property:

$$\Pr\{X_{n+1} = j | X_0 = i_0, \dots, X_{n-1} = i_{n-1}, X_n = i\} = \Pr\{X_{n+1} = j | X_n = i\}$$

for all time points n and all states $i_0, \dots, i_{n-1}, i, j$.

Common convention: label the Markov chain by the non-negative integers $\{0, 1, 2, \dots\}$. " X_n is in state i ": $X_n = i$.


One-step transition probability = the probability of X_{n+1} to be in state j given that X_n is in state i :

$$P_{ij}^{n,n+1} = \Pr\{X_{n+1} = j | X_n = i\}.$$

Markov Chains

In general, any transition probability may vary as a function of the time of the transition. To keep things clear, in what follows we deal with *stationary transition probabilities* (independent of the time variable n): $P_{ij}^{n,n+1} = P_{ij}$.

The Markov matrix or *transition probability matrix* of the process:

$$\mathbf{P} = \begin{bmatrix} P_{00} & P_{01} & P_{02} & \dots \\ P_{10} & P_{11} & P_{12} & \dots \\ P_{20} & P_{21} & P_{22} & \dots \\ \vdots & \vdots & \vdots & \vdots \\ P_{i0} & P_{i1} & P_{i2} & \dots \\ \vdots & \vdots & \vdots & \vdots \end{bmatrix} = [P_{ij}]$$


The i th row is the probability distribution of the values of X_{n+1} under the condition that $X_n = i$.

Markov Chains

Clearly,

$$P_{ij} \geq 0 \text{ and } \sum_{j=0}^{\infty} P_{ij} = 1 \quad \text{for } i = 0, 1, 2, \dots$$

These properties define the matrix as *stochastic*.

Note that $\sum_{j=0}^{\infty} P_{ij} = 1$ is simply stating that the sum of probabilities of all transitions starting from any state i equals 1. (This means that the rows of \mathbf{P} sum to 1.)

If also the sum of probabilities of all transitions ending in any state j equals 1, $\sum_{i=0}^{\infty} P_{ij} = 1$ (the columns of \mathbf{P} sum to 1), then the transition matrix \mathbf{P} is said to be *doubly stochastic*.

A **Markov process** is **completely defined** once its transition probability matrix and initial state X_0 (the probability distribution of X_0) are specified.

Markov Chains

Transition probability matrices of a Markov chain

The analysis of a Markov chain concerns mainly the calculation of the possible realisations of the process.

The n -step transition probability matrices: $\mathbf{P}^{(n)} = [P_{ij}^{(n)}]$.

Note: A product of stochastic matrices is a stochastic matrix.

$$P_{ij}^{(n)} = \Pr\{X_{m+n} = j | X_m = i\}$$

= the probability the process goes from state i to state j in n transitions.

Theorem. The n -step transition probabilities of a Markov chain

satisfy $P_{ij}^{(n)} = \sum_{k=0}^{\infty} P_{ik} P_{kj}^{(n-1)}$, where we define $P_{ij}^{(0)} = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{if } i \neq j. \end{cases}$

Markov Chains

This is matrix multiplication $\mathbf{P}^{(n)} = \mathbf{P} \times \mathbf{P}^{(n-1)}$.

By iteration, $\mathbf{P}^{(n)} = \mathbf{P} \times \mathbf{P} \times \dots \times \mathbf{P} = \mathbf{P}^n$.

So, the n -step transition probabilities $P_{ij}^{(n)}$ are the entries in the matrix \mathbf{P}^n , the n th power of \mathbf{P} .

Proof. The theorem can be proven via a *first-step analysis*, a breaking, or analysis, of the possible transitions on the first step, followed by an application of the Markov property:

$$\begin{aligned} P_{ij}^{(n)} &= \Pr\{X_n = j | X_0 = i\} = \sum_{k=0}^{\infty} \Pr\{X_n = j, X_1 = k | X_0 = i\} \\ &= \sum_{k=0}^{\infty} \Pr\{X_1 = k | X_0 = i\} \Pr\{X_n = j | X_0 = i, X_1 = k\} \\ &= \sum_{k=0}^{\infty} P_{ik} P_{kj}^{(n-1)}. \quad \blacksquare \end{aligned}$$

Markov Chains

Markov chains are visualised by **state transition diagrams**.

(Figure from the [online book](#).)

$$P = \begin{bmatrix} \frac{1}{4} & \frac{1}{2} & \frac{1}{4} \\ \frac{1}{3} & 0 & \frac{2}{3} \\ \frac{1}{2} & 0 & \frac{1}{2} \end{bmatrix}.$$

Figure 11.7 shows the state transition diagram for the above Markov chain. In this diagram, there are three possible states 1, 2, and 3, and the arrows from each state to other states show the transition probabilities p_{ij} . When there is no arrow from state i to state j , it means that $p_{ij} = 0$.




Figure 11.7 - A state transition diagram.

Markov Chains

It follows that if the probability of the process initially being in state j is p_j , i.e. the distribution law of X_0 is $\Pr\{X_0 = j\} = p_j$, then the probability of the process being in state k at time n is

$$p_k^{(n)} = \sum_{j=0}^{\infty} p_j P_{jk}^{(n)} = \Pr\{X_n = k\}.$$

Note: Matrix \mathbf{P} operates from the right, meaning that all vectors are row-vectors, not column-vectors like in almost any other field. This is the convention in stochastics.

$$\mathbf{P} = \begin{bmatrix} P_{00} & P_{01} & P_{02} & \dots \\ P_{10} & P_{11} & P_{12} & \dots \\ P_{20} & P_{21} & P_{22} & \dots \\ \vdots & \vdots & \vdots & \\ P_{i0} & P_{i1} & P_{i2} & \dots \\ \vdots & \vdots & \vdots & \end{bmatrix} = [P_{ij}] \quad (p_0 \ p_1 \ p_2 \ p_3 \ p_4 \ p_5) \begin{pmatrix} P_{00} & P_{01} & P_{02} & P_{03} & P_{04} & P_{05} \\ P_{10} & P_{11} & P_{12} & P_{13} & P_{14} & P_{15} \\ P_{20} & P_{21} & P_{22} & P_{23} & P_{24} & P_{25} \\ P_{30} & P_{31} & P_{32} & P_{33} & P_{34} & P_{35} \\ P_{40} & P_{41} & P_{42} & P_{43} & P_{44} & P_{45} \\ P_{50} & P_{51} & P_{52} & P_{53} & P_{54} & P_{55} \end{pmatrix}$$

Markov Chains

Note: Just to make sure you understand how to do calculations with Markov matrices you might want to do a few exercises from Chapter 3 of 'An Introduction to Stochastic Modeling'. (The answers are given in the end of the book.) The online book serves this purpose just as well.

Markov Chains

Markov Chain models

Markov chains are commonly used to model various stochastic processes in e.g. physics, biology, sociology and economics.

'An Introduction to Stochastic Modeling' gives a few examples that are useful to go through. We do this for one:

An inventory model

A commodity is stocked. The replenishment of stock takes place at the end of periods labelled $n = 0, 1, 2, \dots$. The total aggregate demand during period n is a random variable ξ_n whose distribution function is independent of n ,

$$\Pr\{\xi_n = k\} = a_k \quad \text{for } k = 0, 1, 2, \dots, \text{ where } a_k \geq 0 \text{ and } \sum_{k=0}^{\infty} a_k = 1.$$


Markov Chains

Replenishment policy: If the end-of period stock quantity is not greater than s , an amount sufficient to increase the quantity of stock up to the level S is procured. On the other hand, if the available stock is in excess of s , no replenishment is undertaken.

X_n = the quantity on hand at the end of period n .

The states of the process $\{X_n\}$: $S, S - 1, \dots, +1, 0, -1, -2, \dots$, where negative values are interpreted as unfilled demand that will be satisfied immediately after restocking.

The inventory process $\{X_n\}$:



Markov Chains

According to the inventory rules:

$$X_{n+1} = \begin{cases} X_n - \xi_{n+1} & \text{if } s < X_n \leq S, \\ S - \xi_{n+1} & \text{if } X_n \leq s, \end{cases}$$

where ξ_n is the quantity demanded in the n th period.

Assuming that the successive demands ξ_1, ξ_2, \dots are random variables, the stock values X_0, X_1, X_2, \dots constitute a Markov chain whose transition probability matrix can be determined from

$$P_{ij} = \Pr\{X_{n+1} = j | X_n = i\} = \begin{cases} \Pr\{\xi_{n+1} = i - j\} & \text{if } s < i \leq S, \\ \Pr\{\xi_{n+1} = S - j\} & \text{if } i \leq s. \end{cases}$$

Markov Chains

To illustrate, consider a spare parts inventory in which either 0, 1, or 2 repair parts are demanded and for which

$\Pr\{\xi_n = 0\} = 0.5$, $\Pr\{\xi_n = 1\} = 0.4$, $\Pr\{\xi_n = 2\} = 0.1$, and $s = 0$ and $S = 2$.

The possible values for X_n : $S = 2, 1, 0$, and -1 .

$$P_{10} = \Pr\{X_{n+1} = 0 | X_n = 1\}:$$

When $X_n = 1$, no replenishment $\Rightarrow X_{n+1} = 0$, when $\xi_{n+1} = 1$. This occurs with probability $P_{10} = 0.4$.

When $X_n = 0$, replenishment to $S = 2$ takes place $\Rightarrow X_{n+1} = 0$, when $\xi_{n+1} = 2$. This occurs with probability $P_{00} = 0.1$.

Markov Chains

Analysing each probability in this way, results in

$$\mathbf{P} = \begin{matrix} & \begin{matrix} -1 & 0 & +1 & +2 \end{matrix} \\ \begin{matrix} -1 \\ 0 \\ +1 \\ +2 \end{matrix} & \begin{bmatrix} 0 & 0.1 & 0.4 & 0.5 \\ 0 & 0.1 & 0.4 & 0.5 \\ 0.1 & 0.4 & 0.5 & 0 \\ 0 & 0.1 & 0.4 & 0.5 \end{bmatrix} \end{matrix} \begin{matrix} \leftarrow \text{Vector of states;} \\ \text{"state vector"} \\ \text{includes probabilities,} \\ \text{so strictly they are the} \\ \text{rows in the Markov} \\ \text{matrix} \end{matrix}$$

For the inventory model the important quantities are e.g. the long-term fraction of periods in which demand is not met ($X_n < 0$) and long-term average inventory level; using $p_j^{(n)} = \Pr\{X_n = j\}$, these quantities are given by $\lim_{n \rightarrow \infty} \sum_{j < 0} p_j^{(n)}$ and $\lim_{n \rightarrow \infty} \sum_{j > 0} p_j^{(n)}$, respectively.

Typically, determining conditions under which $p_j^{(n)}$ stabilise and the limiting probabilities $\pi_j = \lim_{n \rightarrow \infty} p_j^{(n)}$ is of importance.

Markov Chains

The simplest models can be solved analytically. To obtain limiting values for more complicated models one has to do stochastic simulation, that is, do many Monte Carlo runs starting from the given X_0 and compute the probabilities of different outcomes (distribution) of X_n as the number of times different values of S are obtained normalised by the total number of outcomes, e.g.

$$\Pr\{S_i\} = \frac{\#(S_i)}{\sum_{k=1}^N \#(S_k)}$$

Simple Stochastic Kinetics

As an example of a general process that is Markovian, let's look at a process where objects (people, cars, particles) arrive at a point randomly.

Here *event times* (arrival times) are governed by a stochastic process.

Poisson Process

Online book, [Ch 11.1](#).

In a Poisson process with rate λ defined on interval $[0, T]$ the inter-event times are $Exp(\lambda)$. To simulate: Initialise the process at time zero $X_0 = t_0 = 0$. Simulate $t_k \sim Exp(\lambda)$ and put $X_k = X_{k-1} + t_k$, where $k \in \{1, 2, 3, \dots\}$. Here t_k is the time from the $k - 1$ to the k th event. Stop when $X_k > T$ and keep the X_1, X_2, \dots as the realisation of the process.

Simple Stochastic Kinetics

The Poisson process is a *point process*. It is directly related to the corresponding *counting process*, where the process is defined via the number of events in a given time interval $(0, t]$ as $N_t \sim Po(\lambda t)$ (λ is the rate of events).

Stochastic simulations are often constructed by directly regarding them as counting processes rather than point processes.

Viewing Poisson as counting process, one needs to specify the smallest time interval Δt , or the resolution of time, at which the system is observed. Then the probability of an event during that time is, of course, $p = \lambda \Delta t$. Typically, one is interested in the mean and variance of inter-event times or times for a certain number of events to occur.

Simple Stochastic Kinetics

The stationary (and homogeneous) Poisson process is perhaps the most basic example of stochastic dynamics but describes surprisingly many real-world processes. In this form the state space and time are discrete. We will look at slightly more complicated and applicable variation, the *nonstationary* Poisson process in the next lecture. Here the process is defined in continuous time (but still simulated at discrete times.)