11

Conditional Expectation, Markov Chains, Continuous Random Variables and Distribution

Computer Science Mentors 70

Apr 26 - Apr 30
Spring 2021

Amogh Gupta, Sylvia Jin, Aekus Bhathal, Abinav Routhu, Debayan Bandyopadhyay
Roast us here: `https://tinyurl.com/csm70-feedback20`

1 Markov Chains

1. **One State at a Time**

(a) A sequence of random variables $X_0, X_1, X_2, X_3, \ldots$ is a Markov chain if ____________________ and ____________________. (Fill in the blanks with equations)

> **Solution:** $\Pr[X_{t+1}|X_t] = \Pr[X_{t+1}|X_t, X_{t-1}, \ldots, X_0]$ for all t and $\Pr[X_{t'+1}|X_{t'}] = \Pr[X_{t+1}|X_t]$.

(b) Any irreducible Markov chain where one state has a self-loop is aperiodic. (True/False)

> **Solution:** True. Since one can reach this state from any other state (and vice versa), a cycle of any length through this state can be constructed. This means that the period, which is the gcd of all possible cycle lengths through this state, is 1.

(c) Given a Markov Chain, let the random variables $X_1, X_2, X_3, \ldots$, where X_t is the state visited at time t in the Markov Chain. Then $E[X_t|X_{t1} = x] = E[X_t|X_{t-1} = x, _{t-2} = x']$.(True/False)

> **Solution:** True. This isn't the Markov property since the full history isn't present on the RHS, but this can be derived from the Markov property. The value of X_t conditioned on X_{t-1} is independent of *any* subset of the history, $X_{t-2}, \ldots, X_0$.

2. **Types of Markov Chains**

(a) Draw a Markov Chain that is reducible.

> **Solution:** Here is one example.
>
> 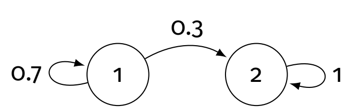
>
>
> The Markov chain can never leave State 2. Intuitively, reducible Markov chains have states of no return – the chain may leave some state and never be able to return.

(b) From here on, we will exclusively discuss (and assume) irreducible Markov chains. Draw a periodic Markov Chain. Does it have an invariant distribution? Does it always converge to this invariant distribution? What does the invariant distribution represent?

> **Solution:** For the following example, we see that this chain is periodic because in order to get back to a state, you must travel some multiple of 2 steps. While a unique invariant distribution exists since the chain is irreducible, these types of Markov chains **do not always converge** to the invariant distribution. Clearly, in this case the invariant distribution

is $\pi = [0.5 \;\; 0.5]$. However, if we start with the distribution $\pi_0 = [1 \;\; 0]$, this distribution will never converge to π. When the chain is finite and irreducible, one can think of the invariant distribution as average amount of time spent in each state.

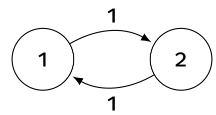

(c) Draw an aperiodic Markov Chain. Does it converge to an invariant distribution?

Solution: These types of Markov chains (aperiodic and irreducible) **always** converge to the invariant distribution. The Markov chain cannot get "stuck" oscillating between intermediate distributions.

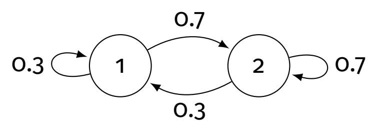

3. Sylvia is stuck in jail with her life savings of \$3 and needs \$8 for bail to get out so she can teach her students again. Professor Rao gives her a chance to redeem herself through a series of bets – if Sylvia bets \$$A$, she wins \$$A$ with probability 0.4 and loses \$$A$ with probability 0.6. If she would like to get to \$8 without losing all her money, draw a Markov chain representing the following situations:

(a) Choosing to play it safe, Sylvia bets \$1 at a time. (Hint: the states are the possible dollar amounts Sylvia has at any point)

Solution:

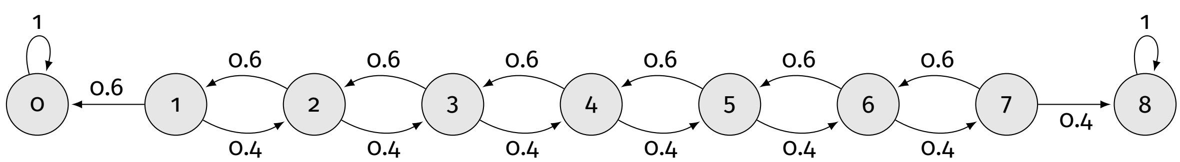

(b) Set up, but DO NOT SOLVE, the equations to determine the probability Sylvia reaches \$8 before she runs out of money.

Solution: At states 0 and 8, her chances are 0 and 1, respectively. For every other state, there's a 0.4 chance she goes up by \$1 and 0.6 chance she goes down by \$1.

$$\alpha(0) = 0$$
$$\alpha(8) = 1$$
$$\alpha(i) = 0.4\alpha(i+1) + 0.6\alpha(i-1), \quad \text{for } i = 1, \ldots, 7$$

(c) She realizes that with the strategy in the previous parts, she has a $< 10\%$ chance of getting out! She decides on a more risky strategy: on each bet, she bets the exact amount needed to reach \$8.

Solution: Notice that this time, she can only bet \$3 the first time, then if she wins, bet \$2, etc. In other words, with

$\$A$, she can only bet $\min\{A, 8-A\}$ each time. By this logic, we have the following chain:

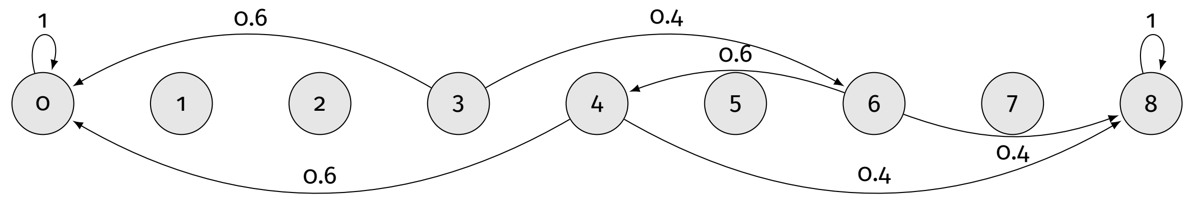

(d) Given that the first strategy gives a $\approx 9.64\%$ chance of escaping, which strategy gives her a better chance of getting out?

Solution: To compare the effectiveness to the second strategy, we need to compute the corresponding probability for the second strategy. The relevant equations are:

$$\begin{aligned}
\alpha(0) &= 0 \\
\alpha(3) &= 0.6\alpha(0) + 0.4\alpha(6) \\
\alpha(4) &= 0.6\alpha(0) + 0.4\alpha(8) \\
\alpha(6) &= 0.6\alpha(4) + 0.4\alpha(8) \\
\alpha(8) &= 1
\end{aligned}$$

After some simplification, we get

$$\begin{aligned}
&\alpha(0) = 0, \alpha(8) = 1 \\
&\alpha(4) = 0.6(0) + 0.4(1) = 0.4 \\
&\alpha(6) = 0.6(0.4) + 0.4(1) = 0.64 \\
&\alpha(3) = 0.6(0) + 0.4(0.64) = 0.256
\end{aligned}$$

and since $\alpha(3) = 0.256 > 0.0964$, the second “bolder” strategy gives a higher probability of escape!

4. **Life of Alex**

Alex is enjoying college life. She spends a day either studying, partying, or looking for housing for the next year. If she is studying, the chances of her studying the next day are 30%, the chances of her partying the next day are 50%, and the chances of her looking for housing the next day are 20%. If she is partying, the chances of her partying the next day are 10%, the chances of her studying the next day are 60%, and the chances of her looking for housing the next day are 30%. If she is looking for housing, the chances of her looking for housing the next day are 50%, the chances of her partying the next day are 30% and the chances of her studying the next day are 20%.

(a) Draw a Markov chain to visualize Alex's life.

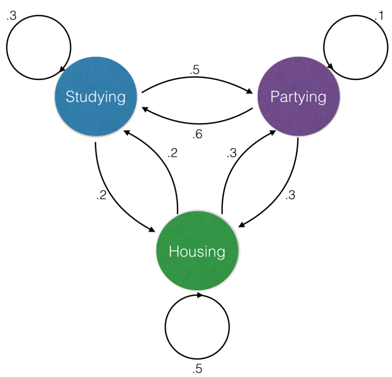

Solution:

(b) Write out a matrix to represent this Markov chain (In this solution (and in CS 70), the rows represent the source and the columns represent the destination).

Solution: In this solution (and in CS 70), the rows represent the source and the columns represent the destination.

$$\begin{bmatrix} .3 & .5 & .2 \\ .6 & .1 & .3 \\ .2 & .3 & .5 \end{bmatrix}$$

(c) If Alex studies on Monday, what is the chance that she is partying on Friday? (Don't do the math, just write out the expression that you would use to find it.)

Solution: She spends Monday studying so $\pi_0 = [1 \;\; 0 \;\; 0]$. Friday is 4 days later so we would like the distribution of her activities 4 steps later, π_4. If P is the matrix above, then $\pi_4 = \pi_0 P^4 \approx [0.36 \;\; 0.30 \;\; 0.33]$. The answer is the first entry of π_4: 0.36.

(d) What percentage of her time should Alex expect to spend looking for housing?

Solution: Solve the following system of equations: (first step equations)

$$\begin{aligned} S &= .3S + .6P + .2H \\ P &= .5S + .1P + .3H \\ H &= .2S + .3P + .5H \\ S + P + H &= 1 \end{aligned}$$

(e) If Alex parties on Monday, what is the chance of Alex partying again before studying?

Solution: Let $\alpha(i)$ denote the probability of partying before studying given Alice has just done activity i. Set up the following equations:

$$\alpha(S) = 0$$
$$\alpha(P) = .6\alpha(S) + .1\alpha(P) + .3\alpha(H)$$
$$\alpha(H) = .2\alpha(S) + .3\alpha(P) + .5\alpha(H)$$

Solving for $\alpha(P)$, we get 0.28.

5. **Riemann's Pontiac**

Your friend Riemann drives a 1994 Pontiac Firebird, an old car and that he hasn't taken care of it all too well. At every second in $\{t_1, \ldots, t_n\}$ with probability p he accelerates at 1 m/s^2 and with probability $1 - p$ his engine sputters and he stops immediately. Find the expected time it takes him to reach a velocity of n m/s.

Solution: To keep things simple (we're not in a physics class!), we're going to remove the units and look only at the numbers. Fortunately, everything is in meters and seconds so we won't have to worry about any conversions.

If we draw the markov chain with each state representing some velocity from 0 to n, then we see that the transition probability for any velocity v_k to v_{k+1} is p and the transition from v_k to v_0 occurs with probability $1 - p$.

Hence, we get the hitting time equation:

$$\beta(v_n) = 0 \quad (1)$$
$$\beta(v_k) = 1 + p\beta(v_{k+1}) + (1-p)\beta(v_0) \; \forall 0 \le k \le n \quad (2)$$

If we recurse backwards from v_n we get a rather ugly set of equations. For the first two steps,

$$\beta(v_{n-1}) = 1 + (1-p)\beta(v_0) \quad (3)$$
$$\beta(v_{n-2}) = 1 + p(1 + (1-p)\beta(v_0)) + (1-p)\beta(v_0) \quad (4)$$

These equations are quite nasty to deal with, so while hitting time would give us the correct answer, it would be too cumbersome to work out. Instead, let's use the random variable W_k to denote the the number of timesteps before reaching a velocity of k. Also denote the event where we accelerate on the next timestep with A. Then,

$$\mathbb{E}(W_k|W_{k-1}) = \mathbb{E}(W_k|W_{k-1}, A)\mathbb{P}(A) + \mathbb{E}(W_k|W_{k-1}, \bar{A})\mathbb{P}(\bar{A}) \quad (5)$$
$$= p(1 + \mathbb{E}(W_{k-1})) + (1-p)(1 + \mathbb{E}(W_{k-1}) + \mathbb{E}(W_k)) \quad (6)$$
$$= 1 + \mathbb{E}(W_{k-1}) + (1-p)\mathbb{E}(W_k) \quad (7)$$

Applying iterated expectation,

$$\mathbb{E}(W_k) = 1 + \mathbb{E}(\mathbb{E}(W_{k-1})) + (1-p)\mathbb{E}(\mathbb{E}(W_k)) \quad (8)$$
$$\mathbb{E}(W_k) = 1 + \mathbb{E}(W_{k-1}) + (1-p)\mathbb{E}(W_k) \quad (9)$$

Hence,

$$\mathbb{E}(W_k) = \frac{1}{p} + \frac{1}{p}\mathbb{E}(W_{k-1}) \quad (10)$$

From here we can make a clever realization. Let $g(k) = \mathbb{E}(W_k) + \frac{1}{1-p}$. Then,

$$\begin{aligned} g(k) &= \frac{1}{p} + \frac{1}{p}\mathbb{E}(W_{k-1}) + \frac{1}{1-p} && (11)\\ &= \frac{1}{p}(\mathbb{E}(W_{k-1}) + \frac{1}{1-p}) && (12)\\ &= \frac{1}{p}g(k-1) && (13)\end{aligned}$$

Now that we have transformed this problem a geometric sequence, we have a closed form solution to the reccurence:

$$\begin{aligned} g(k) &= \frac{1}{p^k}g(0) && (14)\\ &= \frac{1}{p^k(1-p)} && (15)\end{aligned}$$

Now, we must transform back to $\mathbb{E}(W_k)$.

$$\begin{aligned} \mathbb{E}(W_{k-1}) &= g(k) - \frac{1}{1-p} && (16)\\ &= \frac{1}{p^k(1-p)} - \frac{1}{1-p} && (17)\end{aligned}$$

If you prefer, you can also try expanding a few terms, guessing a formula and using induction to prove it holds (the nice thing about recurrence relations are that they are pretty easy to prove with induction).

Thus, the expected time it takes to reach a velocity of nm/s is $\frac{1}{p^n(1-p)} - \frac{1}{1-p}$ seconds.

6. **Tired of Balls and Urns Yet?**
An urn contains six balls, of which three are red and three are green. In each step, two balls are selected at random. If one of them is red, and the other is green, then we discard them and replace them by two blue balls, and if both of the balls are blue, then we replace those blue balls with an equal amount of red and green balls. Otherwise, we do not do anything. Find the probability that if we start with an equal number of balls of every color, what is the probability that we reach 6 blue balls before 0 blue balls in the bag?

Solution: We notice that this is a Markov chain, and we have a total of 4 states, which we can uniquely represent by the number of red balls, since we know we have the same number of green balls at all times and the remaining balls are blue. The states are state 0 (0R 0G 6B), state 1 (1R 1G 4B), state 2 (2R 2G 4B), and state 3 (3R 3G 0B). If we have i each of red and green balls, then

$$P(1R1G) = \frac{i \cdot i}{\binom{6}{2}}$$

which is the transition from state i to $i-1$. Furthermore, if we have j blue balls, then

$$P(2B) = \frac{\binom{j}{2}}{\binom{6}{2}}$$

which is the transition from state i to $i+1$. The remaining probability are self loops. The transition matrix is then:

$$\begin{bmatrix} 0 & 1 & 0 & 0\\ \frac{1}{15} & \frac{8}{15} & \frac{6}{15} & 0\\ 0 & \frac{4}{15} & \frac{10}{15} & \frac{1}{15}\\ 0 & 0 & \frac{9}{15} & \frac{6}{15}\end{bmatrix}$$

The probability desired can be modeled recursively by considering the hitting probability of state 0 before state 3, given you are at state i, which we denote as $\alpha(i)$. Naturally, $\alpha(0) = 1$ and $\alpha(3) = 0$. Then our equations are:

$$\alpha(1) = \frac{1}{15}\alpha(0) + \frac{8}{15}\alpha(1) + \frac{6}{15}\alpha(2)$$
$$\alpha(2) = \frac{4}{15}\alpha(1) + \frac{10}{15}\alpha(2) + \frac{1}{15}\alpha(3)$$

Solving the equations above we get $\alpha(2) = \frac{4}{11}$

2 Continuous RV and Distribution

Discrete vs Continuous Probability

Here is a table illustrating the parallels between discrete and continuous probability.

Discrete	Continuous
$P[X = k] = \sum_{\omega \in \Omega : X(\omega) = k} P(\omega)$	$P[k < X \leq k + \mathrm{d}x] = f_X(k)\mathrm{d}x$ (*)
$P[X \leq k] = \sum_{\omega \in \Omega : X(\omega) \leq k} P(\omega)$	$P[X \leq k] = F_X(k)$
$E[X] = \sum_{a \in A} a \cdot P[X = a]$	$E[X] = \int_{-\infty}^{\infty} x \cdot f_X(x)dx$
$E[\phi(X)] = \sum_{a \in A} \phi(a) \cdot P[X = a]$	$E[\phi(X)] = \int_{-\infty}^{\infty} \phi(x) \cdot f_X(x)dx$
$\sum_{\omega \in \Omega} P[\omega] = 1$	$\int_{-\infty}^{\infty} f_X(x)dx = 1$

(*) When solving problems with continuous distributions, you can think of $f_X(k)$ as being analogous to $P[X = k]$ in discrete distributions, but they are not equal.

1. **PDFs**

   Consider the following functions and determine whether or not they are valid probability density functions.

   (a) $f(x) = \sin(x)$

   Solution: This is not valid because $\sin(x)$ can be negative.

   (b) $f(x) = x$ for $0 \leq x \leq 1$, and $f(x) = 0$ everywhere else.

   Solution: This is not valid, since

$$\int_{-\infty}^{\infty} f(x)\mathrm{d}x = \int_{0}^{1} x\mathrm{d}x = \left[\frac{x^2}{2}\right]_0^1 = \frac{1}{2} \neq 1$$

(c) $f(x) = 1$ for $0 \leq x \leq 1$, and $f(x) = 0$ everywhere else.

Solution: This is valid, since $f(x) \geq 0$ for all x, and

$$\int_{-\infty}^{\infty} f(x)\mathrm{d}x = \int_{0}^{1} 1\mathrm{d}x = [x]_0^1 = 1 = 1.$$

(d) $f(x) = e^{-x}$ for $x \geq 0$, and $f(x) = 0$ everywhere else.

Solution: This is valid, since $f(x) \geq 0$ for all x, and

$$\int_{-\infty}^{\infty} e^{-x}\mathrm{d}x = \int_{0}^{\infty} e^{-x}\mathrm{d}x = [-e^{-x}]_0^{\infty} = 0 - (-1) = 1.$$

(This is the pdf of a Poisson(1) distribution.)

2. **Disk**

Define a continuous random variable R as follows: we pick a point uniformly at random on a disk of radius 1; the value of R is distance of this point from the center of the disk. We will find the probability density function of this random variable.

(a) Why is R not $U(0,1)$?

Solution: We can think of it somewhat in areas. There are more points that have larger radius than smaller radius, and since the likelihood of selecting any particular point is equal, it is more likely to get a larger radius than a smaller one.

(b) What is the probability that R is less than r, for any $0 \leq r \leq 1$? What is the CDF $F_R(r)$ of the random variable R?

Solution: r^2, because the area of the circle with distance between 0 and r is πr^2, and the area of the entire circle is π.
Thus, we have that $F_R(r) = r^2$ for $0 \leq r \leq 1$.

(c) What is the PDF $f_R(r)$ of the random variable R?

Solution: By definition,

$$f_R(r) = \frac{\mathrm{d}}{\mathrm{d}r}F_R(r) = \frac{\mathrm{d}}{\mathrm{d}r}r^2 = 2r$$

for $0 \leq r \leq 1$.

(d) Now say that $R \sim U(0,1)$. Are you more or less likely to hit closer to the center than before?

Solution: More likely. Let's evaluate the probability that $R \leq c, c \in (0,1)$ in both cases. In the first case, $P(R \leq c) = c^2$. In the second case, $P(R \leq c) = c$. For $c \in (0,1), c \geq c^2$.

3. **Joint Density**

The joint density for the random variables X and Y is defined by $f(x,y) = \frac{2}{3}x + \frac{4}{3}y$ for $0 \leq x, y \leq c$ for some positive real number c, and $f(x,y) = 0$ for all other (x,y).

(a) For what value of c is this a valid joint density?

Solution: For the joint density to be valid, we need that the total integral equals 1, so

$$\begin{aligned}
1 &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f(x,y)\mathrm{d}x\mathrm{d}y \\
&= \frac{2}{3}\int_0^c\int_0^c x+2y\mathrm{d}x\mathrm{d}y \\
&= \frac{2}{3}\int_0^c \left[\frac{x^2}{2}+2xy\right]_0^c \mathrm{d}y \\
&= \frac{2}{3}\int_0^c \frac{c^2}{2}+2cy\mathrm{d}y \\
&= \frac{2}{3}\left[\frac{c^2}{2}y+cy^2\right]_0^c \\
&= c^3,
\end{aligned}$$

so we have that $c = 1$.

(b) Compute $P(X < Y)$.

Solution: To get $P(X < Y)$, we need to integrate the pdf over the region where $X < Y$, so

$$\begin{aligned}
P(X<Y) &= \int_0^c\int_0^y f(x,y)\mathrm{d}x\mathrm{d}y \\
&= \frac{2}{3}\int_0^1\int_0^y x+2y\mathrm{d}x\mathrm{d}y \\
&= \frac{2}{3}\int_0^1 \left[\frac{x^2}{2}+2xy\right]_0^y \mathrm{d}y \\
&= \frac{2}{3}\int_0^1 \frac{5y^2}{2}\mathrm{d}y \\
&= \frac{2}{3}\left[\frac{5y^3}{6}\right]_0^1 \\
&= \frac{5}{9}.
\end{aligned}$$

(c) Compute $E[X|Y=y]$ for $0 \le y \le c$.

Solution: We first need to compute $f_Y(y)$. We have that

$$\begin{aligned}
f_Y(y) &= \int_0^c f(x,y)\mathrm{d}x \\
&= \frac{2}{3}\int_0^1 x+2y\mathrm{d}x \\
&= \frac{2}{3}\left[\frac{x^2}{2}+2xy\right]_0^1 \\
&= \frac{4}{3}y+\frac{1}{3}.
\end{aligned}$$

Now, we have that

$$
\begin{aligned}
E[X|Y=y] &= \int_0^c xf(x|Y=y)\mathrm{d}x \\
&= \int_0^1 \frac{xf(x,y)}{f_Y(y)}\mathrm{d}x \\
&= \frac{2}{3}\int_0^1 \frac{x(x+2y)}{\frac{4}{3}y+\frac{1}{3}}\mathrm{d}x \\
&= \frac{2}{4y+1}\int_0^1 x^2+2xy\mathrm{d}x \\
&= \frac{2}{4y+1}\left[\frac{x^3}{3}+x^2y\right]_0^1 \\
&= \frac{6y+2}{12y+3}.
\end{aligned}
$$

(d) Compute $E[XY]$.

Solution: We have that

$$
\begin{aligned}
E[XY] &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} xyf(x,y)\mathrm{d}x\mathrm{d}y \\
&= \frac{2}{3}\int_0^c\int_0^c x^2y+2xy^2\mathrm{d}x\mathrm{d}y \\
&= \frac{2}{3}\int_0^1 \left[\frac{x^3y}{3}+x^2y^2\right]_0^1 \mathrm{d}y \\
&= \frac{2}{3}\int_0^1 \frac{y}{3}+y^2\mathrm{d}y \\
&= \frac{2}{3}\left[\frac{y^2}{6}+\frac{y^3}{3}\right]_0^1 \\
&= \frac{2}{3}\cdot\frac{1}{2} \\
&= \frac{1}{2}.
\end{aligned}
$$