Lecture Notes for EE306
Module II

Department of Electrical and Electronics Engineering
Middle East Technical University

Elif Uysal

April 5, 2022

Chapter 1

Discrete Stochastic Processes

1.1 Introduction

Consider observing the value of a measured quantity at times $n = 1, 2, 3, \ldots$ such that X_n is the value measured at time n.

Ex: Starting $t = 0$, a temperature sensor reports a measurement every 10 minutes. Let X_n be the measurement reported at $t = 10n$ minutes.

Ex: An autonomous vehicle sends the GPS measurement of its location every second. X_n is the value of the location at the n^{th} second.

Ex: You check the price of Gold each day. X_n is the price on day n.

We want to model X_n as a random variable. But we do not want to stop at just modeling a single value measured at a single time n: we want to consider the evolution of the sequence of values measured at a sequence of times $n = 1, 2, \ldots$. In other words, we want a model for the whole *process* $\{X_1, X_2, X_3, \ldots\}$, so that we can make some computations regarding it, and make decisions such as:

- Given the vehicle's location 3 seconds ago, what are the chances it is dangerously close to another vehicle at the moment?
- The price of an ounce of Gold is 1963 USD today. Based on its price history over the last six months, is it a good time to sell?
- Looking at the most recent blood test of a patient, a doctor decides whether or not to change the dosage of their medication.

Such decisions occur all the time in daily life as well as in engineering problems. Especially with the proliferation of the Internet of Things (IoT) technology, autonomous driving, robotics, remote monitoring, social networking, and the fact that each of us run applications on our devices that monitor distant phenomena and make decisions constantly, such questions are perhaps more relevant than ever.

Definition 1 *A **discrete stochastic process** is a sequence of random variables, $\{X_n\}$, indexed by n. Typically, $n = 1, 2, \ldots$.*

Ex: The Bernoulli process with rate p is a sequence of IID Bernoulli random variables with parameter p:

$$X_i = \begin{cases} 1 & \text{, with probability p} \\ 0 & \text{, with probability } 1 - \text{p} \end{cases}$$

The Bernoulli process is memoryless: X_n does not depend on the *history* of the process, i.e. $\{X_i, i < n\}$, or the *future*, $\{X_i, i > n\}$.

In the rest, we will explore a type of stochastic process that has memory.

1.2 Introduction to Markov Chains

The Bernoulli process considered above was the simplest nontrivial case of stochastic process, with no memory. In many real-life and engineering problems this would be a poor model. Consider, for example, the location of a vehicle, X_t, at times $t = 1$ and $t = 2$. If the vehicle has finite speed, knowing X_1 will give us some information about X_2. Similarly, the price of gold is not determined independently each day! It either rises or falls from its value the previous day. The peak temperature today exhibits dependence on not only yesterday's peak temperature, but perhaps on the temperature of the previous day, through the weather system that has been in effect for several days. As in these examples, many discrete stochastic processes have *memory*.

In some cases, the value of the stochastic process X_{n+1} at time $n + 1$ depends on the past history $X_n, X_{n-1}, X_{n-2}, \ldots, X_0$, only through the present value, X_n. We call this a *Markov Chain* (MC).

Markov Chains are among the most interesting models in probability, built on the notion of conditional probability. Here, we briefly study the topic through a sequence of simple examples.

Ex: Suppose that if it rains today, then it will rain tomorrow with probability α, and if it does not rain today, then it will rain tomorrow with probability β, independently of the weather in previous days. Define a state space, S and a Markov Chain to model this situation.

Solution: Let $S = \{0, 1\}$. Suppose we say that the process is in state 1 when it rains and state 0 when it does not rain. Let $X_n \in S$ be the state at time n. $P(X_n = 0|X_{n-1} = 0, X_{n-2}, \ldots, X_1\}$ =P(X$_n$ = $0|X_{n-1} = 0) = 1 - \beta$, and $P(X_n = 0|X_{n-1} = 1, X_{n-2}, \ldots, X_1\}$ =P(X$_n$ $= 0|X_{n-1} = 1) = 1 - \alpha$.

Draw a pictorial representation of this 2-state Markov Chain.

Specify values for α and β such that the resulting process is a Bernoulli process.

Ex: Burak is determined to take Calculus I until he passes the course. There is no eviction from the undergraduate program, which means he is offered an unlimited number of chances to take the course. However, Burak does not study at all, and does not learn the material. So, every semester that he takes Calculus, the probability that he passes is p, irrespective of how many times he took the course before. Let X_n be a random variable indicating whether Burak passed the course at the end of semester n or not. $X_n = 0$ if Burak fails at the end of the n^{th} semester, and $X_n = 1$ otherwise.

(a) Consider possible sample paths of the process and contrast with that of a Bernoulli process with rate p.

(b) Argue why this is NOT a Bernoulli process.

(c) Show that this process is Markov.

(d) Note that, the Bernoulli process is also a special case of a Markov Chain, but not all binary valued Markov Chains are Bernoulli processes.

Discussion: If Burak passes on the k^{th} attempt, we can set $X_{k+1} = X_{k+2} = ... = 1$. Imagine a *sequence* of X_ns that continue forever, even after Burak passes the course. Sample paths will look like

0000111111111111111111111111...

That is, an infinite number of ones, possibly preceded by a number of zeroes.
The following follow from the description of the problem:

$$\begin{aligned} P(X_{n+1} = 0|X_n = i, X_{n-1}, X_{n-2}, \ldots, X_1) &= P(X_{n+1} = 0|X_n = i) \\ &= \begin{cases} 0, & \text{if } i = 1 \\ 1-p, & \text{if } i = 0 \end{cases} \end{aligned}$$

and

$$\begin{aligned} P(X_{n+1} = 1|X_n = i, X_{n-1}, X_{n-2}, \ldots, X_1) &= P(X_{n+1} = 1|X_n = i) \\ &= \begin{cases} 1, & \text{if } i = 1 \\ p, & \text{if } i = 0 \end{cases} \end{aligned}$$

From the above, given the present value of the process, X_n, the future of the process, X_{n+1} depends only on X_n. That is, conditioned on the present, the future is independent of the past. In other words, the process is *Markov*.

1.2.1 Transition Probabilities

In the above example, the probability of going from *State 0* (*Did not pass the course yet*) to *State 1* (Done with Calculus) is

$$P(X_{n+1} = 1|X_n = 0) = p$$

The probability of going from *State 0* back to the same state (in other words, staying in that state) is:

$$P(X_{n+1} = 0|X_n = 0) = 1 - p$$

These are called *transition probabilities*. When in state 1, the probability of going to state 0 is zero, hence the only possibility is a *self-transition*, that is, staying in state 1.

The diagram in Fig.1.2.1, summarizes the MC visually. Go ahead and mark the transition probabilities on the diagram. Note that the transition probabilities out of every state add to 1.

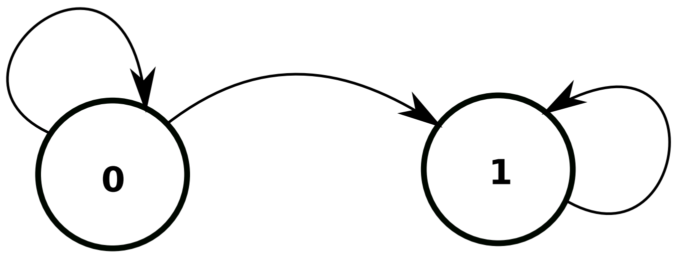

Figure 1.1: Markov Chain with one transient and one absorbing state.

In general, we define the transition probabilities[1] as:

$$P_{ij} = P(X_{n+1} = j | X_n = i)$$

For $i, j \in S$. Here $P_{ij} \geq 0$ is the probability that, when in state i, the process will next go to state j (j can be the same as i, in that case the process would simply be staying in i). As the process must make a transition into some state, we have, for all states i,:

$$\sum_{j \in S} P_{ij} = 1$$

The **transition probability matrix P** is defined as the matrix of values $\{P_{ij}\}$ for all $i, j \in S$. Due to the above, each row of **P** adds to 1.

$$\mathbf{P} = \begin{bmatrix} P_{00} & P_{01} & P_{02} & \dots \\ P_{10} & P_{11} & P_{12} & \dots \\ . & . & . & . \\ . & . & . & . \\ P_{i0} & P_{i1} & P_{i2} & \dots . \\ . & . & . & \\ . & . & . & . \end{bmatrix}$$

Ex: For the rain/no rain example above,

$$\mathbf{P} = \begin{bmatrix} \alpha & 1-\alpha \\ \beta & 1-\beta \end{bmatrix}$$

[1]Our scope is limited to the case of a time-homogeneous MC.

Ex: (Random Walk) Consider a MC whose state space is the set of integers, such that, for some $a > 0$, the transitions are either to the right with probability a, or to the "left" with probability $1 - a$.

$$P_{i,i+1} = a, \quad P_{i,i-1} = 1 - a$$

for $i = 0, \pm 1, \pm 2, \ldots$ Note that the above is a MC with a countably infinite state-space.

Ex: Burak finally passed Calculus I, and he is now taking Calculus II. There is a 20 percent probability that he will pass Calculus II this semester. If not, he will either take it again next semester, or change his major, those decisions being equally likely. Every time Burak takes Calculus, the probability of passing is 20 percent.

(a) Draw the diagram of a MC for this problem, where the states are $\{1, 2, 3\}$, corresponding to *taking Calculus*, *passed Calculus*, *quit.* Mark the transition probabilities

(b) Let X_n be the state in Semester n, where $X_1 = 1$, $n = 1, 2, \ldots$. Write down the transition probability matrix $\mathbf{P}$ with entries $P_{ij} = P(X_{n+1} = j | X_n = i)$, for all $i, j \in \{1, 2, 3\}$.

Ex: (Ross, Example 4.4) Now, assume that whether or not it rains today depends on previous weather conditions through the last two days. Specifically, suppose that if it has rained for the past two days, then it will rain tomorrow with probability 0.7; if it rained today but not yesterday, then it will rain tomorrow with probability 0.5; if it rained yesterday but not today, then it will rain tomorrow with probability 0.4; if it has not rained in the past two days, then it will rain tomorrow with probability 0.2. If we let the states be "rain, no rain" as before, is this a MC? If not, find a MC to represent this process.

Solution: We can form a MC with the following 4 states:

- state 0, if it rained both today and yesterday,
- state 1, if it rained today but not yesterday,
- state 2, if it rained yesterday but not today,
- state 3, if it did not rain either yesterday or today.

Please write down the transition probability matrix accordingly.

Ex: The 2 Umbrella Problem: I own two umbrellas. At time $t = 0$, one is at home, and the other is at the office. In the morning, if it is raining, I will take an umbrella with me as I walk from home to work. In the evening, as I walk back, I will again take an umbrella with me if it is raining. I never carry an umbrella if it is not raining. With this policy, I am interested in the steady-state probability of being caught in the rain with no umbrella. Construct a MC for this process, and specify $\mathbf{P}$.

1.2.2 n-step Transition Probabilities

We denote the probability of going from i to j in n steps as P_{ij}^n:

Definition 2

$$P_{ij}^n = P(X_{n+k} = j | X_k = i), \quad n \geq 0, i, j \in S$$

Note that,

$$\begin{aligned} P_{ij}^{n+m} &= \sum_k P(X_{n+m+k} = j | X_{m+k} = l, X_k = i) P(X_{m+k} = l | X_k = i), \\ & n, m \geq 0, i, j, l \in S \end{aligned}$$

which implies that:

$$P_{ij}^{(n+m)} = \sum_k P_{il}^m P_{lj}^n \quad n, m \geq 0, i, j, l \in S$$

In short,

$$P^{(n+m)} = P^{(n)} P^{(m)}$$

Ex: Burak is now taking EE202. His algorithm did not change, and is expressed by the 3-state MC represented below.

(a) Show that $P(X_{n+2} = j | X_n = i)$ is given by $\sum_{k=1}^{3} p_{ik} p_{kj}$. Note that this is the product of the i^{th} row and the j^{th} column of $\mathbf{P}$. In other words, the probability of going from state i to state j in two steps is given by the ij^{th} entry of the $\mathbf{P}^2$ matrix. Compute and write down the matrix $\mathbf{P}^2$.

(b) Argue similarly that the n step transition probabilities are given by entries of the matrix $\mathbf{P}^n$.

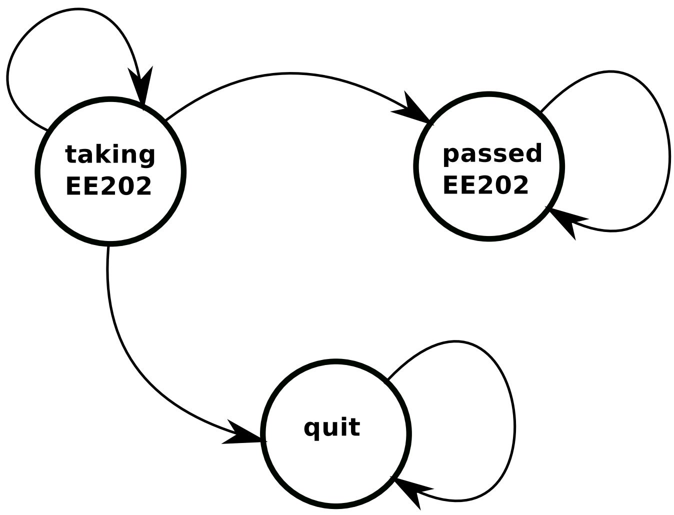

Figure 1.2: Markov Chain with two absorbing states.

(c) Notice that the second and third rows of P^n will always be fixed, while the first row changes with n. Show that the first entry of the first row, P_{11}^n is equal to 0.4^n. Similarly, let the second entry be $P_{12}^n = a_n$. Observe that $a_n = 0.4a_{n-1} + 0.2$. Let $P_{13}^n = b_n$. Observe that $b_n = 0.4b_{n-1} + 0.4$.

(d) Based on the observations just made, what is the matrix P^n converging to, as n grows? Interpret the results: what is the probability that Burak will eventually pass the course? What is the probability that he will change his major?

The above example is a Markov Chain with two *absorbing states*: states 2 and 3. If we start in State 1, we will eventually leave this state and get absorbed in either state 2 or 3.

Now, we will compute the probability that Burak eventually passes EE202, through an alternative approach:

(a) Let a be the probability that the eventual state is passing EE202, given that we start at time $n = 1$ by taking EE202: that is, $a = P(\lim_{n\to\infty} X_n = 2|X_1 = 1)$, and b be the probability that the eventual state is quitting: that is, $b = P(\lim_{n\to\infty} X_n = 3|X_1 = 1)$. Argue that $b = 1 - a$.
(Hint: We have three types of sequences possible: those that hit 2 and get stuck there, those that hit 3 and get stuck there, and those that always stay at 1. The probability of the last kind of sequence is zero. To see this, consider the first column of the matrix $\mathbf{P}^n$ computed above, and see that it converges to the all-zero vector.)

(b) Show that we can write the following equation for a:

$$a = 1 \times p_{12} + 0 \times p_{13} + a \times p_{11}$$

(Hint: Just like we found two step transition probabilities above, we can condition on the next move we make starting from state 1, and apply the law of total probability. $P(\lim_{n\to\infty} X_n = 2|X_1 = 1) = \sum_{k=1}^{3} P(\lim_{n\to\infty} X_n = 2|X_1 = 1, X_2 = k)P(X_2 = k|X_1 = 1)$)

(c) Solve the above equation, determine the values of a and b.

Ex: Recall the 2 Umbrella Problem, where I started on the morning of day 1 with one umbrella at each location. Suppose it independently rains each morning and evening with probability 0.5. Compute the probability that, on the way back home on day 2, I am left in the rain with no umbrella.

Ex: There is an insect walking (from what I gather, randomly) on my windowsill. Out of boredom, and in order to observe the movements of the insect better, I marked 4 equally spaced regions on the windowsill. At location 1 at the left hand corner, there is a spider web. If the fly gets to that location, it will be captured. Upon careful observation, I determined that the transition probabilities between locations are as in the matrix $\mathbf{P}$ given below:

$$\mathbf{P} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0.3 & 0.4 & 0.3 & 0 \\ 0 & 0 & 0.5 & 0.5 \\ 0 & 0 & 0.5 & 0.5 \end{bmatrix}$$

(a) Given that the fly starts at location 2, what is the probability, that after 2 steps, it is at location 4?

(b) Given that the fly starts at location 2, what is the probability, that after infinitely many steps (in practice, after a sufficiently long time), it is observed at location 4?

Ex: Random Walk with Two Thresholds: Consider a random walk on the set of integers: The probability of moving right or left are each equal to 1/3. We stay in place with probability 1/3. Given $X_0 = 0$, compute the probability of the walk reaching the value 2 before ever hitting -2.

Ex: Random Walk on a Finite Set: Now, truncate the random walk in the previous example at -2 and 2, such that when we reach these endpoints, the probability of staying in place is 2/3. Given $X_0 = 0$, compute the expected number of transitions until the walk hits 2.