CSE 632 (Fall 2019): Analysis of Algorithms II : Randomized Algorithms

Lecture 17 (10/23/2019): Markov Chains

Lecturer: Shi Li *Scriber: Sen Lin*

1 Markov Chains

Markov chains can be used to model dynamic systems such as Brownian motion, market trends (which can be treated as random processes) and languages (which is more "deterministic"). A state diagram is a directed graph where each nodes correspond to a possible state of the system and arcs give the possible transitions among those states. Then intuitively, a (discrete-time finite space) Markov chain can be viewed as a state diagram with probabilities on the edges : a weight $p_{(u,v)}$ on the edge (u, v) indicates the probability that we transit to state v, given that we are currently at state u. See Figure 1 for two examples. From the definition of weights on edges, we have

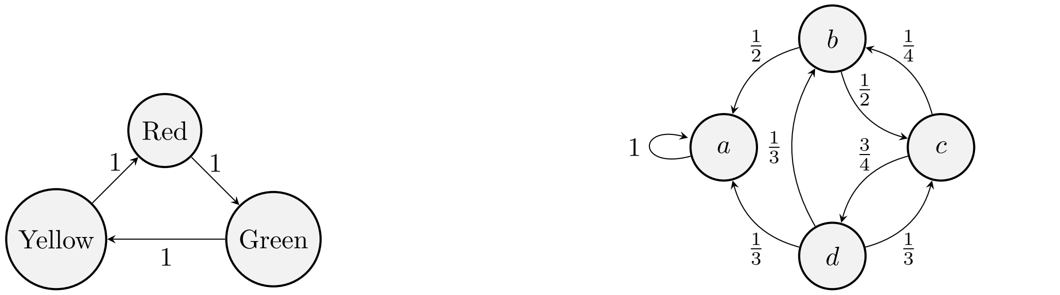

(a) Traffic light system. The traffic light can change from red to green, green to yellow and yellow to red. Each transition edge has weight 1, indicating that the transitions among states are deterministic.

(b) An abstract state diagram. The transitions among states are random. For example, if the system is currently at state b, then with probability 1/2, the state will be changed to a, and with probability 1/2 it will be changed to c.

Figure 1: Two Examples of Markov Chains. Weights on the edges denote the transition probabilities.

Observation 1. *In a state diagram for a Markov Chian, all weights are non-negative and the sum of weights for all outgoing edges of a state is 1.*

1.1 Formal Definition

To formally define a Markov chain, we need to define a *stochastic process* first.

Definition 2 (Stochastic Processes)**.** *Let T be a set denoting the times. Then a stochastic process is a family of random variables $\{X_t : t \in T\}$, one for each time $t \in T$. If T is discrete, we say the process has a discrete time. If all X_t's can take values from a discrete space, then we say the process has a discrete space.*

Definition 3 (Markov Chains)**.** *A discrete time stochastic process $X = (X_0, X_1, \dots)$ is called a (discrete time) Markov chain if for every $t \geq 1$ and $a_0, a_1, \dots, a_t$, we have*

$$\Pr[X_t = a_t | X_{t-1} = a_{t-1}, X_{t-2} = a_{t-2}, \dots, X_0 = a_0] = \Pr[X_t = a_t | X_{t-1} = a_{t-1}]$$

I.e., for every a_{t-1}, conditioned on $X_{t-1} = a_{t-1}$, X_t is independent of $X_{t-2}, X_{t-3}, \dots, X_0$. But does that mean X_t and X_{t-2} are independent? No. If we do not condition on a specific value of X_{t-1}, X_t and X_{t-2} may be dependent.

Compare this with the definition of Martingales. They both have a sequence of random variables; also, conditioned on $X_1, X_2, \cdots, X_{t-1}$, some property of the random variable X_t only depends on X_{t-1}:

- For a Markov Chain, the distribution for X_t only depends on X_{t-1}.
- For a Martingale, the distribution for X_t may depend on $X_1, X_2, \cdots, X_{t-2}$, even if X_{t-1} is fixed. However, we always have that expected value of X_t is X_{t-1}.

On the other hand, in a Markov Chain, the expectation of random variables may not be defined. Consider the traffic light example, where a random variable has its value in {"red", "green", "yellow"}. Even if the expectation is defined, we may not have $\mathbb{E}[X_t|X_1, X_2, \cdots, X_{t-1}] = X_{t-1}$.

In the course, we only focus on discrete time finite space Markov chains that are *time-homogeneous*, defined as follows.

Definition 4 (Time Homogeneous Markov Chains)**.** *A Markov chain is time-homogeneous if for every $t \geq 0$, x, y, we have*

$$\Pr[X_t = x|X_{t-1} = y] = \Pr[X_{t+1} = x|X_t = y].$$

1.2 Transition Matrix

Definition 5. *Given a (discrete time finite space time-homogeneous) Markov chain, we define its transition matrix P to be a matrix with rows and columns indexed by the values (or states) in the space of the random variables such that*

$$P_{i,j} = \Pr[X_t = j|X_{t-1} = i] \qquad \forall t, i, j.$$

(1) is the transition matrix for the Markov Chain in Figure (1b).

$$\begin{array}{c c} & \begin{array}{cccc} a & b & c & d \end{array} \\ \begin{array}{c} a \\ b \\ c \\ d \end{array} & \begin{pmatrix} 1 & 0 & 0 & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ 0 & \frac{1}{4} & 0 & \frac{3}{4} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 \end{pmatrix} \end{array} \tag{1}$$

Notice that a transition matrix is a "row-stochastic matrix", which is a matrix satisfying the following properties:

- every entry is non-negative, and
- every row sum is 1.

We know $\Pr[X_1 = b|X_0 = c]$ is 1/4. From the matrix, we can also compute a two-step transition probability. For example,

$$\begin{aligned} \Pr[X_2 = d|X_0 = b] &= \sum_i \Pr[X_1 = i|X_0 = b] \Pr[X_2 = d|X_1 = i, X_0 = b] \\ &= \sum_i \Pr[X_1 = i|X_0 = b] \Pr[X_2 = d|X_1 = i] && \text{by defintion of Markov Chains} \\ &= \frac{1}{2} \times \frac{3}{4} = \frac{3}{8}. \end{aligned}$$

The probabilities can be represented more concisely as follows. Let $\pi^{(t)}$ be the row vector for the probabilities of distribution X_t. That is $\pi_a^{(t)} = \Pr[X_t = a]$ for every a. Then we have

Lemma 6. *Let P the transition matrix for a Markov Chain, and $\pi^{(t)}$ be the row vector for the probabilities of distribution X_t. Then for every $t \geq 1$, we have $\pi^{(t)} = \pi^{(t-1)} P$.*

Proof.

$$\begin{aligned}\pi_j^{(t)} &= \Pr[X_t = j] \\ &= \sum_i \Pr[X_{t-1} = i] \Pr[X_t = j | X_{t-1} = i] \\ &= \sum_i \pi_i^{(t-1)} P_{i,j} \\ &= \pi^{(t-1)} \underbrace{P_{\star,j}}_{j\text{-th column}}\end{aligned}$$

So, $\pi^{(t)} = \pi^{(t-1)} P$. □

Applying the lemma repeatedly we have

Lemma 7. *For every $t \geq 0$ and $m \geq 1$, we have $\pi^{(t+m)} = \pi^{(t)} P^m$.*

2 Example: Another Gambler's Game

Consider the following game that a gambler plays in a casino. Initially, he has a dollars. He pays 1 dollar for each round. Then he will get 2 dollars back with probability 1/2 (with the remaining 1/2 probability, he will not get anything back). The gambler leaves the casino when either he losses all his money, or he has b dollars in his pocket, where $b > a$ is some integer. The gambler loses the whole game in the former case, and wins the game in the latter. We are interested in the probability that the gambler wins the game. We prove the following lemma

Lemma 8. $\Pr[$*"Win"*$] = \frac{a}{b}$ *and* $\Pr[$*"Lose"*$] = 1 - \frac{a}{b}$.

Proof. We define a Markov chain as follows. Let X_t be the amount of money the gambler has after t rounds. Initially we have $X_0 = a$. For every $t \leftarrow 1, 2, 3, \ldots$, conditioned on $1 \leq X_{t-1} \leq b - 1$, we have

$$X_t = \begin{cases} X_{t-1} + 1 & \text{with probability } \frac{1}{2} \\ X_{t-1} - 1 & \text{with probability } \frac{1}{2} \end{cases}.$$

The game will be ended if at some time T we have $X_T = 0$ or $X_T = b$. The gambler wins the game if $X_T = b$ and loses the game otherwise. For simplicity, if a game ends at time T, for every $t' > T$ we assume $X_{t'} = X_T$. We can draw the state diagram (or the transition graph) as Figure 2 and have the transition matrix as (2).

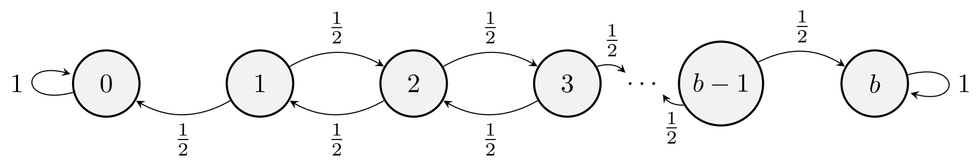

Figure 2: State Diagram (or Transition Graph) for Gambler's Game

$$\begin{bmatrix} 1 & & & & & & \\ \frac{1}{2} & 0 & \frac{1}{2} & & & & \\ & \frac{1}{2} & 0 & \frac{1}{2} & & & \\ & & & \ddots & & & \\ & & & \frac{1}{2} & 0 & \frac{1}{2} \\ & & & & & 1 \end{bmatrix} \tag{2}$$

We give two ways to prove the lemma. The first method is by solving a system of linear equations. Let us define P_i to be the probability of wining conditioned on that we are currently at state i

(meaning the gambler has i dollars), where i is an integer between 0 and b. Then we have

$$\begin{aligned} P_0 &= 0 \\ P_b &= 1 \\ P_i &= \frac{1}{2}P_{i-1} + \frac{1}{2}P_{i+1} && \forall i = 1, 2, 3, \cdots, b-1 \end{aligned}$$

To see the third equality, we know that the probability that the next state will be $i-1$ with probability $1/2$ and $i+1$ with probability $1/2$. Notice that the definition of P_j's is independent of the time. Thus, the equality holds.

The third equality is equivalent to $2P_i = P_{i-1} + P_{i+1}$, which is $P_1 - P_{i-1} = P_{i+1} - P_i$. Thus, we have that $P_0, P_1, P_2, \cdots, P_b$ form an arithmetic progression. Thus, $P_i = \frac{i}{b}$. That is, we have $\Pr[\text{"Win"}] = \frac{a}{b}$ and $\Pr[\text{"Lose"}] = 1 - \frac{a}{b}$.

Now we consider the second method. Without a formal proof, we make the following claim: the probability that the game will end in finite number of rounds is 1. Let $\pi_i^{(t)}$ be the probability that we are in state i at time t. Then mathematically, this means $\lim_{t\to\infty}\left(\pi_0^{(t)} + \pi_b^{(t)}\right) = 1$. It is easy to see that the sequence $(X_0, X_1, X_2, \cdots,)$ also forms a martingale, we have $\mathbb{E}[X_t] = X_0 = a$ for every t. Thus,

$$a = \lim_{t\to\infty} \mathbb{E}[X_t] = \lim_{t\to\infty} \sum_{i=0}^{b} \pi_i^{(t)} \cdot i = \sum_{i=0}^{b} i \cdot \lim_{t\to\infty} \pi_i^{(t)} = 0 \cdot \Pr[\text{lose}] + b \cdot \Pr[\text{win}].$$

Thus, we have $\Pr[\text{win}] = a/b$ and $\Pr[\text{lose}] = 1 - a/b$. Above we used that $\lim_{t\to\infty} \pi_i^{(t)} = 0$ for every $i \in \{1, 2, \cdots, b-1\}$, $\Pr[\text{win}] = \lim_{t\to\infty} \pi_b^{(t)}$ and $\Pr[\text{lose}] = \lim_{t\to\infty} \pi_0^{(t)}$. □