AI2613 随机过程 2020–2021 春季学期

Lecture 2 – Disceret Markov Chain (I)

2021 年 3 月 1 日

Lecturer: 张驰豪 *Scribe*: 杨宽

1 Discrete Markov Chain

Definition 1 (Discrete Markov chain). Suppose there is a sequence of random variables

$$X_0, X_1, \dots, X_t, X_{t+1}, \dots$$

where $\mathrm{Range}(X_i) \subseteq \mathcal{S}$ for some countable set $\mathcal{S}$. Then $\{X_n\}$ is a *discrete Markov chain* if $\forall\, t \ge 0$ and $\forall\, a_0, a_1, \dots, a_{t+1} \in \mathcal{S}$,

$$\mathbf{Pr}[X_{t+1} = a_{t+1} \mid X_t = a_t, X_{t-1} = a_{t-1}, \dots, X_0 = a_0] = \mathbf{Pr}[X_{t+1} = a_{t+1} \mid X_t = a_t]\,.$$

Remark. If for all $i, j \in \mathcal{S}$, there exists a constant $p_{i,j}$ such that

$$\forall\, t \ge 0, \qquad \mathbf{Pr}\left[X_{t+1} = j \mid X_t = i\right] = p_{i,j},$$

the Markov chain is called a *time-homogeneous* Markov chain.

Example 2 (Gambler's ruin). Consider a gambler who starts with an initial fortune of 1 and then on each successive gamble either wins 1 or loses 1 independent of the past with probabilities p and $q = 1 - p$ respectively. The gamble ends when the gambler reaches the total fortune of N (the gambler *wins*) or gets ruined (the gambler *loses*).

Let X_n be the total fortune after the n-th gamble. Then $X_0 = 1$ and for all $t \ge 0$,

$$\mathbf{Pr}\left[X_{t+1} = j \mid X_t = i\right] = \begin{cases} 1, & \text{if } i = j = 0; \\ 1, & \text{if } i = j = N; \\ p, & \text{if } 1 \le i \le N-1 \text{ and } j = i+1; \\ q, & \text{if } 1 \le i \le N-1 \text{ and } j = i-1. \end{cases}$$

We can use a state-transition graph or an automaton to describe the Markov chain:

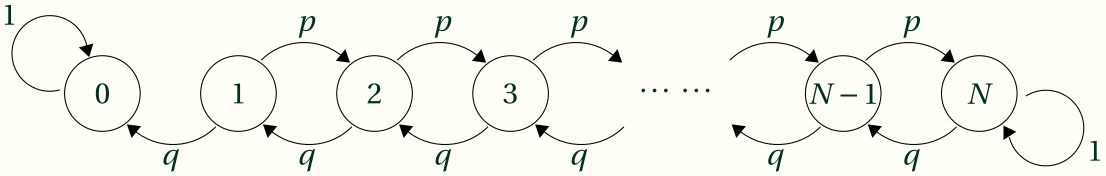

Example 3 (Random walk on $\mathbb{Z}$). The set of state $\mathcal{S}$ is $\mathbb{Z}$, and the transition probability is given by

$$p_{i,j} = \mathbf{Pr}\left[X_{t+1} = j \mid X_t = i\right] = \begin{cases} 1/2, & \text{if } \left|i - j\right| = 1; \\ 0, & \text{otherwise}. \end{cases}$$

Remark. If the state space $\mathcal{S}$ is a finite set, we can use an $|\mathcal{S}| \times |\mathcal{S}|$ matrix $\mathbf{P} \in [0,1]^{\mathcal{S}\times\mathcal{S}}$ to denote transition probabilities, where $\mathbf{P}(i, j) = p_{i,j}$.

Let μ_t be the distribution of X_t, i.e. $X_t \sim \mu_t$. Then $\forall\, i \in \mathcal{S}$,

$$\begin{aligned} \mu_1(i) = \mathbf{Pr}\left[X_1 = j\right] &= \sum_{i\in\mathcal{S}} \mathbf{Pr}\left[X_1 = j \wedge X_0 = i\right] \\ &= \sum_{i\in\mathcal{S}} \mathbf{Pr}[X_0 = i] \cdot \mathbf{Pr}\left[X_1 = j \mid X_0 = i\right] \\ &= \sum_{i\in\mathcal{S}} \mu_0(i) \cdot p_{i,j} \end{aligned}$$

Suppose that $\mathcal{S} = \{0, 1, \ldots, N\}$, then we can use a column vector to denote μ_t, where

$$\mu_t = (\mu_t(0), \mu_t(1), \ldots, \mu_t(N))^\top.$$

Thus ${\mu_1}^\top = {\mu_0}^\top \cdot \mathbf{P}$, and by induction, it follows directly that

$$\forall\, t \geq 0, \qquad {\mu_t}^\top = {\mu_{t-1}}^\top \cdot \mathbf{P} = {\mu_{t-2}}^\top \cdot \mathbf{P}^2 = \cdots = {\mu_0}^\top \cdot \mathbf{P}^t.$$

Definition 4 (Stochastic matrix). A *stochastic matrix* $\mathbf{P}$ is a square matrix whose columns are probability vectors, i.e. $\mathbf{P} \in [0,1]^{\mathcal{S}\times\mathcal{S}}$ and $\forall\, i$, $\sum_{j\in S} \mathbf{P}(i, j) = 1$.

Question. Does $\mathbf{P}^t$ converge as $t \to \infty$?

2 Stationary Distribution

Definition 5 (Stationary distribution). Let $\{X_n\}$ be a Markov chain with transition matrix $\mathbf{P}$. Suppose π is a distribution such that

$$\pi \cdot \mathbf{P} = \pi.$$

Then π is called a *stationary distribution* of Markov chain $\{X_n\}$.

Remark. Note that the definition of the stationary distribution does not depends on the convergence of the Markov chain. We will see the connection between them later.

Example 6 (Random walk on undirected graph). Suppose $G = (V, E)$ is an undirected graph with n vertices and m edges. Let $d_i = \deg(i)$ denote the degree of vertex i. We define a Markov chain on G:

$$\mathbf{P}(i, j) = \begin{cases} 1/d_i, & \text{if } i \sim j; \\ 0, & \text{if } i \nsim j. \end{cases}$$

Then what is the stationary distribution π?
Suppose G is a regular graph, i.e., d_i is a constant d for all i. Then $\pi = (1/n, 1/n, \ldots, 1/n)^\top$. If G is not a regular graph, we claim that the stationary distribution is

$$\pi = \left(\frac{d_1}{\sum d_k}, \frac{d_2}{\sum d_k}, \ldots, \frac{d_n}{\sum d_k}\right)^\top.$$

Note that $\sum d_k = 2m$. So $\pi = \left(d_1/2m, d_2/2m, \ldots, d_n/2m\right)^\top$. It is easy to verify that π is indeed a stationary distribution. We will leave it as an exercise here.

Question. Why do we need to study stationary distribution?

One of the motivations comes from the *Markov chain Monte Carlo* method. The key to the method is to design a Markov chain whose stationary distribution is the desired one.

Example 7 (Cards shuffling). Let's consider a naïve "top-to-random" card shuffle:

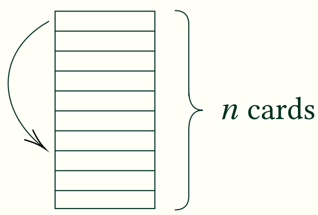

Suppose we have n cards, everytime we take the top card of the deck and insert it into the deck at one of the n distinct possible places uniformly at random.
Let i, j be two permutations on $[n]$. W.l.o.g. assume that $i = (1, 2, \ldots, n)$. Then $\mathbf{P}(i, j) > 0$ iff there exists k s.t. $j = (2, 3, \ldots, k, 1, k+1, \ldots, n)$.

Performing the shuffle repeatedly is a Markov chain. It is not difficult to verify that the uniform distribution $(1/n!, 1/n!, \ldots, 1/n!)^\top$ over all permutations is a stationary distribution. We leave it as an exercise again.

Question. For stationary distributions, we have the following three questions: under which condition can we prove

1. the existence of a stationary distribution π?
2. the uniqueness of π?
3. the convergence of the Markov chain?

In the next part of this lecture, we are going to answer these questions.

2.1 Finite States

We first consider the case that the state space is a finite set. Then we claim that the answer to Question 1 is always "yes", i.e., any finite Markov chain has stationary distribution.

Definition 8 (Spectral radius). Let $\mathbf{A}$ be a $n \times n$ nonnegative matrix. Then the *spectral radius* of $\mathbf{A}$, denoted by $\rho(\mathbf{A})$, is the maximum norm of its eigenvalues. Namely,

$$\rho(\mathbf{A}) = \max\{|\lambda| : \det(\lambda \mathbf{I} - \mathbf{A}) = 0\}.$$

Fact 1. *Let $\mathbf{A} = (a_{i,j}) \in \mathbb{R}_{\geq 0}^{n \times n}$ be a nonnegative matrix, i.e., $a_{i,j} \geq 0$. Then*

$$\min_{1 \leq i \leq n} \sum_{j=1}^{n} a_{i,j} \leq \rho(\mathbf{A}) \leq \max_{1 \leq i \leq n} \sum_{j=1}^{n} a_{i,j}.$$

Theorem 2 (Perron-Frobenius Theorem). *Let $\mathbf{A} = (a_{i,j}) \in \mathbb{R}_{\geq 0}^{n \times n}$ be a nonnegative matrix with spectral radius $\rho(\mathbf{A}) = \alpha$. Then α is a eigenvalue of $\mathbf{A}$, and has both left and right nonnegative eigenvectors.*

Perron-Frobenius Theorem answers Question 1: Let $\mathbf{P}$ be a stochastic matrix. Then $\mathbf{P} \cdot \mathbf{1} = \mathbf{1}$. Thus Fact 1 implies that $\rho(\mathbf{P}) = 1$. So $\mathbf{P}^{\top}$ has eigenvalue 1 and there exists $\pi \geq \mathbf{0}$ s.t. $\mathbf{P}^{\top} \cdot \pi = \pi$.

For Question 2 and 3, let's consider the following Markov chain.

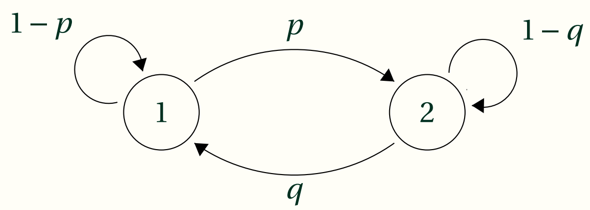

Then $\pi = \left(\frac{q}{p+q}, \frac{p}{p+q}\right)^{\top}$ is a stationary distribution. Now let $\Delta_t \triangleq \left|\mu_t(1) - \pi(1)\right|$. We have

$$\begin{aligned}
\Delta_t &= \left|\left({\mu_{t-1}}^{\top} \cdot \mathbf{P}\right)(1) - \pi(1)\right| \\
&= \left|(1-p) \cdot \mu_{t-1}(1) + q \cdot \left(1 - \mu_{t-1}(1)\right) - \frac{q}{p+q}\right| \\
&= \left|(1-p-q) \cdot \mu_{t-1}(1) + q \cdot \left(1 - \frac{1}{p+q}\right)\right| \\
&= \left|1-p-q\right| \cdot \Delta_{t-1}.
\end{aligned}$$

So $\Delta_t \to 0$ except for

1. $p = q = 0$;
2. $p = q = 1$.

Case 1: $p = q = 0$

In this case, we call the Markov chain *reducible* and the stationary distributions may not be unique.

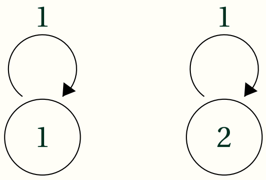

Definition 9 (Reducibility). We say that j is *accessible* from i iff $\exists\, t > 0$, s.t. $\mathbf{Pr}[X_t = j \mid X_0 = i] > 0$.
Moreover, i *communicates* with j if i is accessible from j and j is accessible from i.
We also define an equivalence relation $i \simeq j$: $i \simeq j$ iff i communicates with j.
Then a Markov chain is *irreducible* iff the number of equivalent classes is 1. In other words, the state graph is strongly connected. Otherwise, the Markov chain is called *reducible*.

Fact 3. *If a Markov chain is irreducible, then its stationary distribution is unique. Otherwise its stationary distributions may not be unique.*

Case 2: $p = q = 1$

In this case, $X_t = X_0$ if t is even and X_t is the other state if t is odd. Then the Markov chain is called a *periodic* chain.

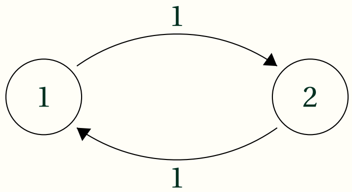

For all $i \in \mathcal{S}$, let $d_i \triangleq \gcd\{u : \mathbf{P}^u(i,i) > 0\}$. Namely, d_i is the greatest common divisor of the length of all loops starting from i and ending at i. Then we have the following lemma.

Lemma 4. *If i and j communicate with each other, then $d_i = d_j$.*

Proof. Suppose that $\mathbf{P}^{n_1}(i,j) > 0$, $\mathbf{P}^{n_2}(j,i) > 0$ and $\mathbf{P}^{n}(j,j) > 0$. Note that $d_i \mid (n_1 + n_2)$ and $d_i \mid (n_1 + n_2 + n)$. Thus $d_i \mid n$. It is easy to see that for all n that $\mathbf{P}^n(j,j) > 0$, $d_i \mid n$. So $d_i \mid d_j$, and vice versa. $\square$

Definition 10 (Periodicity). A Markov chain is *aperiodic* if $d_i = 1$ for all i, and is *periodic* otherwise.

Fact 5. *Periodic chains do not converge.*

In fact, these two cases are the only cases that the answer to Question 2 or 3 is "No".

Theorem 6 (Fundamental Theorem of Markov Chains). *If a finite Markov chain $\{X_n\}$ with transition matrix $\mathbf{P}$ is irreducible and aperiodic, then there exists a unique stationary distribution π, and*

$$\forall \mu, \qquad \lim_{t\to\infty} \mu^{\mathsf{T}} \cdot \mathbf{P}^t \to \pi^{\mathsf{T}}.$$

2.2 Countably Infinite States

Now we assume that the state space is a countably infinite set.
Note that if the state space is infinite, then we do not have Perron-Frobenius Theorem, so the answer to Question 1 is not always be "yes".
First we should introduce the concept of recurrence.

Definition 11 (Recurrence (常返)). Let T_i be the first hitting time of i, i.e. $T_i = \inf\{t > 0 : X_t = i\}$. Then

- i is *recurrent* if $P_i(T_i < \infty) = 1$, and
- i is *transient* if i is not recurrent.

Here P_x is defined as $P_x(A) = \mathbf{Pr}[A \mid X_0 = x]$.

Remark. What does "$\mathbf{Pr}[T < \infty] = 1$" mean?
$T < \infty$ is an event in the σ-algebra, specifically, $\{T < \infty\} = \cup_n \{T < n\}$. So

$$\mathbf{Pr}[T < \infty] = 1 \iff \lim_{n\to\infty} \mathbf{Pr}[T < n] = 1.$$

Let's see an example of Gambler's ruin:

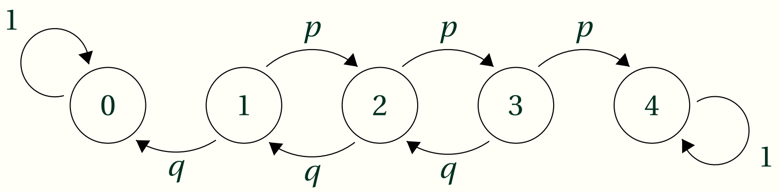

- recurrent: state 0, 4;
- transient: state 1, 2, 3.

For recurrent states, we have the following propositions.

Fact 7. *Let N_i be the total number of visits of the Markov chain to state i, that is, $N_i \triangleq \sum_{t=0}^{\infty} \mathbb{1}_{[X_t=i]}$. If i is recurrent, then $P_i[N_i = \infty] = 1$.*

(Why? The proof is left as an exercise.)

Proposition 8. *If i is recurrent, and j is accessible from i, then*

1. *$P_i(T_j < \infty) = 1$;*
2. *$P_j(T_i < \infty) = 1$;*
3. *j is recurrent.*

Proof. Suppose $i \neq j$ since the result is trivial otherwise. Here we only give an informal proof but it is not difficult to rigorized it.

1. Let $q = P_i[\text{reach } j \text{ before return } i]$. It is easy to see that $q > 0$. Then

$$P_i[\text{reach } j \text{ before } n \text{ times return } i] = 1 - (1-q)^n.$$

By Fact 7, $P_i[N_i = \infty] = 1$, so

$$P_i[T_j < \infty] = \lim_{n\to\infty} 1 - (1-q)^n \to 1.$$

2. If $P_j[T_i < \infty] = 1 - q < 1$, then with probability q, it never comes to i from j. Thus with probability $p'q$, it never returns to i from i, where p' is the probability that access j from i.
3. $(1) + (2) \implies (3)$: starting from j the chain is certain to visit i eventually, and starting from i the chain is certain to visit j eventually, which implies that starting from j the chain is certain to visit i and will definitely get back to j after that. □

Corollary 9. *For a finite and irreducible Markov chain, every states is recurrent.*

Proposition 10. *i is recurrent $\iff \mathbb{E}_i[N_i] = \infty$.*

Proof. The "$\implies$" direction is clear. Since $P_i[N_i = \infty] = 1$ by Fact 7, we obtain that $\mathbb{E}_i[N_i] = \infty$. Now let's consider the "$\impliedby$" direction.
Suppose not, that is, i is transient. Then $P_i[T_i = \infty] = q > 0$. Namely, with probability q starting from i it will never return. So the distribution of N_i is the Geometric distribution and $\mathbb{E}[N_i] = 1/q$. □

Example 12 (Drunk person and drunk birds). Consider the d-dimensional random walk. Let $S_n = \sum_{i=1}^n X_i$ where $X_i \in (\pm 1, \pm 1, \dots, \pm 1) = \{+1, -1\}^d$. We want to understand whether this Markov chain is recurrent, or if a person standing at the origin can come back in a finite amount of time. By the theory developed so far, we only need to check whether the quantity $\mathbb{E}_0[N_0] = \sum_n \mathbf{P}^n(0,0)$ is finite or not. It turns out that this depends on how large d is.

First, we consider the case $d = 1$. Let $n = 2m$. Then

$$\mathbf{P}^n(0,0) = \binom{2m}{m} \cdot 2^{-2m}.$$

Using Stirling's approximation, where $n! \approx \sqrt{2\pi n}(n/\mathrm{e})^n$, it yields that

$$\begin{aligned}\binom{2m}{m} \cdot 2^{-2m} &= \frac{(2m)!}{(m!)^2} \cdot 2^{-2m} \\ &\approx \frac{\sqrt{4\pi m}\left(\frac{2m}{\mathrm{e}}\right)^{2m}}{2\pi m\left(\frac{m}{\mathrm{e}}\right)^{2m}} \cdot 2^{-2m} \\ &= \frac{1}{\sqrt{\pi m}}.\end{aligned}$$

So

$$\sum_n \mathbf{P}^n(0,0) = \sum_m \mathbf{P}^{2m}(0,0) \approx \sum_m \frac{1}{\sqrt{\pi m}} \to \infty.$$

For greater d, we have the following result since all dimension are mutually independent:

$$\mathbf{P}^{2m}(0,0) \approx \left(\frac{1}{\sqrt{\pi m}}\right)^d = \pi^{-d/2} \cdot m^{-d/2}.$$

Thus,

$$\sum_m \mathbf{P}^{2m}(0,0) = \begin{cases} \infty, & \text{if } d \leq 2; \\ \Theta(1), & \text{otherwise}. \end{cases}$$

Since a bird lives in a 3-dimensional space yet a person lives in a 2-dimensional space, our calculation justifies the following famous quote:

> "*A drunk person will always find their way home, while a drunk bird may get lost forever.*"