AI2613 随机过程 2020–2021 春季学期

Lecture 4 – Examples and Applications of Markov Chains (I)

2021 年 3 月 15 日

Lecturer: 张驰豪 *Scribe*: 杨宽

1 Galton–Watson Process

The model was formulated by F. Galton in the study of the survival and extinction of family names. In the nineteenth century, there was concern amongst the Victorians that aristocratic surnames were becoming extinct. In 1873, Galton originally posed the question regarding the probability of such an event, and later H. W. Watson replied with a solution.

Using more modern terms, the process can be defined formally as follows:

Definition 1 (Galton–Watson Process)**.** Suppose that all the individuals reproduce independently of each other and have the same offspring distribution. More precisely, let G_t denote the number of individuals of t-th generation.

- We start from the zero generation. For convenience, let $G_0 = 1$.
- Each individual of generation t gives birth to a random number of children of generation $t+1$: $\forall\, t \ge 0$ and $i \in [G_t]$, let $X_{t,i}$ denote the number of children of the i-th individual. Then $\{X_{t,i}\}$ is an array of i.i.d. random variables with $\mathbf{Pr}\left[X_{t,i} = k\right] = p_k$ for all $t \ge 0$ and $i \in [G_t]$.
- All individuals of generation $t+1$ are children of individuals of generation t:

$$G_{t+1} = \sum_{i=1}^{G_t} X_{t,i}\,.$$

The sequence $\{G_t\}_{t\ge0}$ is called the *Galton–Watson Process* with offspring distribution p.

Remark. It is clear that the process $\{G_t\}_{t\ge0}$ is a Markov chain.

Denote by ρ the probability of extinction, namely,

$$\rho \triangleq \mathbf{Pr}[\text{extinction}] = \mathbf{Pr}[\cup_{t\ge1}\{G_t = 0\}]\,.$$

Then the question is to determine the value of ρ.

Here are two simple examples:

- $p_0 = 0 \implies \rho = 0$;
- $p_0 > 0 \wedge p_0 + p_1 = 1 \implies \rho = 1$.

So from now on, we assume that $p_0 > 0$ and $p_0 + p_1 < 1$. Using the Markov property, we can calculate ρ as follows:

$$\begin{aligned}\rho = \mathbf{Pr}[\text{extinction}] &= \sum_{k=0}^{\infty} \mathbf{Pr}[\text{extinction} \wedge G_1 = k] \\ &= \sum_{k=0}^{\infty} \mathbf{Pr}[\text{extinction} \mid G_1 = k] \cdot p_k \\ &= \sum_{k=0}^{\infty} \rho^k \cdot p_k ,\end{aligned}$$

where in the last equation we use the independence of $X_{t,i}$. So ρ should satisfy that $\rho = \sum_k \rho^k \cdot p_k$. To analyze this equation, now we introduce a powerful tool: probability generating functions.

Definition 2 (Probability Generating Function). Let X be a discrete random variable defined on a probability space with probability measure $\mathbf{Pr}[\cdot]$. Assume that X has non-negative integer values. The *probability generating function* of X is given by

$$G_X(z) \triangleq \mathbb{E}\left[z^X\right] = \sum_{k=0}^{\infty} \mathbf{Pr}[X = k] \cdot z^k .$$

Let $\psi(z)$ be the probability generating function of the number of children, i.e.,

$$\psi(z) = \sum_{k=0}^{\infty} p_k \cdot z^k ,$$

then ρ satisfies that $\rho = \psi(\rho)$, namely, $z = \rho$ is a fixed point of $\psi(z)$.

Note that $\psi(z)$ has the following properties:

Proposition 1. *Properties of $\psi(z)$ on $[0, 1]$:*

- $\psi'(z) = \sum_{k=1}^{\infty} k \cdot p_k \cdot z^{k-1} \geq 0$. *So $\psi(z)$ is an increasing function.*
- $\psi(1) = 1$ *and* $\psi(0) = p_0 > 0$.
- $\psi''(z) = \sum_{k=2}^{\infty} k(k-1) \cdot p_k \cdot z^{k-2} \geq 0$. *So $\psi(z)$ is a convex function.*

Remark. In general, the generating function is a *formal power series* (形式幂级数). Usually generating functions allow us to use them without worrying about convergence, but still be careful when you are evaluating them at some point.

In particular, the series of probability generating functions converge for all $|z| \leq 1$.

How many fixed points can a convex increasing function have? The answer is at most 2. However, there are still two kinds of possible functions on $[0,1]$ (See Figure 1). If $\psi(z)$ is ψ_1-type, then $z=1$ is the only fixed point for $\psi(z)$. If $\psi(z)$ is ψ_2-type, then there is another fixed point $z=r\in[0,1]$. So there are two questions: which type of function $\psi(z)$ is and which fixed point ρ is.

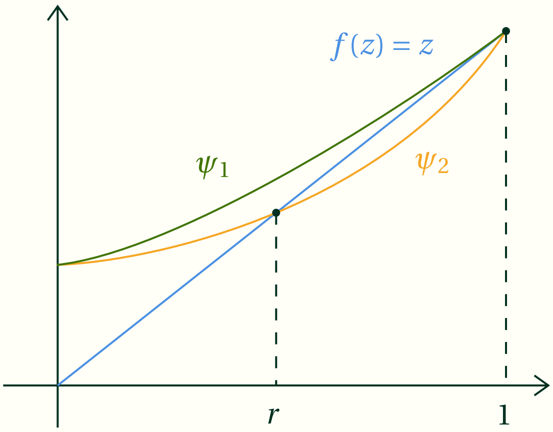

Figure 1: Which fixed point is ρ?

Here are the two cases:

1. ψ_1-type: $\psi'(1)\le 1 \implies \sum_{k\ge1} k\cdot p_k = \mathbb{E}\left[X_{t,i}\right]\le 1$.
2. ψ_2-type: $\psi'(1)>1 \implies \mathbb{E}\left[X_{t,i}\right]>1$. We claim that ρ is the smaller fixed point r.

Proof. Now we prove our claim. Let $q_t \triangleq \mathbf{Pr}[G_t=0]$. By definition it is clear that $q_t\le q_{t+1}<1$ and $\rho=\lim_{t\to\infty} q_t$. Denote by $q_t\uparrow\rho$ that $\{q_t\}$ is an increasing sequence and converges to ρ. Then we prove by induction that $q_t\le r$ for all $t\ge 0$.
The base case is $t=0$, where $q_t=0\le r$.
Then, assume that $t\ge0$ and $q_t\le r$. Note that

$$\begin{aligned} q_{t+1} &= \sum_{k=0}^{\infty} p_k\cdot q_t^k \\ &= \psi(q_t)\le\psi(r)=r\,. \end{aligned}$$

Thus by induction it yields that $q_t\le r$ for all $t\ge0$. Combining with $q_t\uparrow\rho$ and $\rho\in\{r,1\}$, we conclude that $\rho=r$. □

In summary, $\rho=\mathbf{Pr}[\text{extinction}]<1$ iff $\mathbb{E}\left[X_{t,i}\right]>1$.

2 Gambler's Ruin

Review the definition of Gambler's ruin, which we introduced in Lecture 2.

Definition 3 (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 $t \geq 0$,

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

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

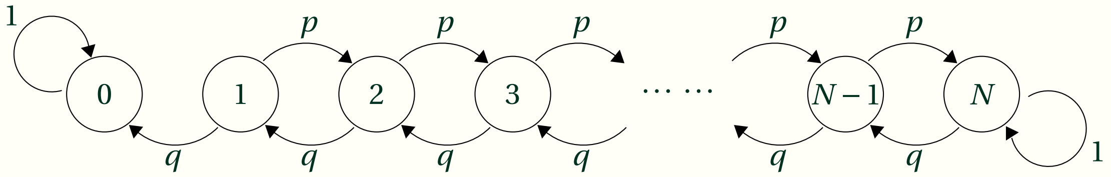

Figure 2: Gambler's ruin

Question. Suppose that start at $X_0 = i$. What is the probability of ending at N?

Let $Z_i = X_{i+1} - X_i$. Then $X_i = X_0 + \sum_{j=1}^{i-1} Z_j$. Let P_i be the probability that eventually win at N when starting from i, i.e. $P_i = \mathbf{Pr}[\text{win} \mid X_0 = i]$. Here are two simple cases: $P_0 = 0$ and $P_N = 1$. For $1 \leq i \leq n-1$, applying the total probability theorem, we obtain that

$$\begin{aligned} P_i &= \mathbf{Pr}[\text{win} \mid X_0 = i] \\ &= \mathbf{Pr}[\text{win} \wedge X_1 = i+1 \mid X_0 = i] + \mathbf{Pr}[\text{win} \wedge X_1 = i-1 \mid X_0 = i] \\ &= \mathbf{Pr}[\text{win} \mid X_0 = i \wedge Z_0 = 1] \cdot \mathbf{Pr}[Z_0 = 1 \mid X_0 = i] + \mathbf{Pr}[\text{win} \mid X_0 = i \wedge Z_0 = -1] \cdot \mathbf{Pr}[Z_0 = -1 \mid X_0 = i] \\ &= P_{i+1} \cdot p + P_{i-1} \cdot q, \end{aligned}$$

where we use the Markov property for the last equation. Thus,

$$P_{i+1} = \frac{1}{p} \cdot P_i - \frac{1-p}{p} \cdot P_{i-1},$$

which is a *second order linear recurrence.* Now we write the recurrence as multiplication of matrices:

$$\begin{pmatrix} P_{i+1} \\ P_i \end{pmatrix} = \begin{pmatrix} \frac{1}{p} & \frac{p-1}{p} \\ 1 & 0 \end{pmatrix} \begin{pmatrix} P_i \\ P_{i-1} \end{pmatrix} = \begin{pmatrix} \frac{1}{p} & \frac{p-1}{p} \\ 1 & 0 \end{pmatrix}^i \begin{pmatrix} P_1 \\ P_0 \end{pmatrix}.$$

Let $\mathbf{A} = \begin{pmatrix} \frac{1}{p} & \frac{p-1}{p} \\ 1 & 0 \end{pmatrix}$. But how can we calculate $\mathbf{A}^i$? A natural idea is to diagonalize $\mathbf{A}$. Then we should calculate $\mathbf{A}$'s eigenvalues first:

$$\begin{aligned} |\mathbf{A} - \lambda \mathbf{I}| &= \begin{vmatrix} \frac{1}{p} - \lambda & \frac{p-1}{p} \\ 1 & -\lambda \end{vmatrix} \\ &= \lambda^2 - \frac{1}{p} \cdot \lambda + \frac{1-p}{p}. \end{aligned}$$

Thus $|\mathbf{A} - \lambda \mathbf{I}| = 0 \iff p\lambda^2 - \lambda + (1-p) = 0 \iff (\lambda - 1)(p\lambda + p - 1) = 0$, which implies that $\mathbf{A}$ has two eigenvalues $\lambda_1 = 1$ and $\lambda_2 = (1-p)/p$.

Note that if $\mathbf{A}$ has 2 different eigenvalues then $\mathbf{A}$ has 2 linearly independent eigenvectors. Hence the following theorem tells us that $\mathbf{A}$ is diagonalizable.

Theorem 2 (Diagonalization Theorem). *An $n \times n$ matrix is diagonalizable if and only if it has n linearly independent eigenvectors.*

So there are 2 cases for diagonalizing $\mathbf{A}$: $p \neq 1/2$ or $p = 1/2$.

- Case 1: $p \neq 1/2$. Then $\mathbf{A}$ is diagonalizable, that is, $\exists$ an invertible 2×2 matrix Λ s.t.

  $$\mathbf{A} = \Lambda \begin{pmatrix} 1 & 0 \\ 0 & \frac{1-p}{p} \end{pmatrix} \Lambda^{-1},$$

  thus we obtain that

  $$\mathbf{A}^i = \Lambda \begin{pmatrix} 1 & 0 \\ 0 & \left(\frac{1-p}{p}\right)^i \end{pmatrix} \Lambda^{-1} \quad \text{and} \quad P_i = a + b \cdot \left(\frac{1-p}{p}\right)^i \text{ for some } a, b.$$

  Since $P_0 = 0$ and $P_N = 1$, we have

  $$\left.\begin{array}{l} a + b = 0 \\ a + b \cdot \left(\frac{q}{p}\right)^N = 1 \end{array}\right\} \implies \left\{\begin{array}{l} a = -b \\ b = \frac{1}{(q/p)^N - 1} \end{array}\right. .$$

  Therefore,

  $$P_i = \frac{1}{1 - \left(\frac{q}{p}\right)^N} \cdot \left(1 - \left(\frac{q}{p}\right)^i\right).$$

- Case 2: $p = 1/2$. In this case $P_{i+1} = 2P_i - P_{i-1}$. Then $P_{i+1} - P_i = P_i - P_{i-1}$, which implies that $\{P_i\}$ is an arithmetic progression. So $P_i = a + b \cdot i$ for some a, b. Substituting it by $P_0 = 0$ and $P_N = 1$ we obtain that $P_i = i/N$.

Overall, it turns out that

$$P_i = \begin{cases} \frac{1-(q/p)^i}{1-(q/p)^N}, & \text{if } p \neq 1/2; \\ \frac{i}{N}, & \text{if } p = 1/2. \end{cases}$$

Now we introduce another example and see how to apply this result.

Suppose that we have a kind of new drug drug$_1$ and a classical drug drug$_2$. Let P_i be the proportion of cured patients after using drug$_i$. Assume that P_2 is well-known while P_1 is unknown. Our goal is to determine whether $P_1 > P_2$.

Suppose $(X_1, Y_1), (X_2, Y_2), \ldots, (X_n, Y_n), \ldots$ is a sequence of pairs of patients. Let patient X_i take drug$_1$ and patient Y_i take drug$_2$ for all i. With abuse of notations here, we also let X_i, Y_i be the indicator random variables that indicates whether the patient recovers after taking the medicine. Now let $Z_i = X_i - Y_i$ and $S_n = \sum_{i=1}^n Z_i$. The test ends as long as $S_n = M$ or $S_n = -M$ for some threshold value M. If the test ends with $S_n = M$, then we believe $P_1 > P_2$, otherwise we believe $P_1 < P_2$. So the question is, what is the probability that we are wrong?

W.l.o.g. assume that $P_1 > P_2$. The test result is wrong if and only if the test ends with $S_n = -M$. Note that

$$Z_i = \begin{cases} 1 & \text{with probability } P_1(1-P_2), \\ -1 & \text{with probability } P_2(1-P_1), \\ 0 & \text{otherwise}. \end{cases}$$

We now let

$$p = \mathbf{Pr}[Z_i = 1 \mid Z_i \neq 0] = \frac{P_1(1-P_2)}{P_1(1-P_2) + P_2(1-P_1)} \quad \text{and} \quad q = 1 - p.$$

Then the probability $\mathbf{Pr}[\text{test ends with } S_n = -M]$ is exactly the same as the probability of loss, namely, $\mathbf{Pr}[\text{lose} \mid X_0 = M]$ in the Gambler's ruin model of winning threshold $2M$.

Applying the result of the Gambler's ruin model, we conclude directly that

$$\begin{aligned} \mathbf{Pr}[\text{test ends with } S_n = -M] &= 1 - \mathbf{Pr}[\text{win} \mid X_0 = M] \\ &= 1 - \frac{1-(q/p)^M}{1-(q/p)^{2M}} = 1 - \frac{1}{1+(q/p)^M} \\ &= \frac{1}{1+(p/q)^M} = \frac{1}{1+\left(\frac{P_1(1-P_2)}{P_2(1-P_1)}\right)^M}. \end{aligned}$$

3 Another Example of Random Walks

Consider the following random walk, where $q = 1 - p$. The only difference between it and the Gambler's ruin is the successive state of state 0.

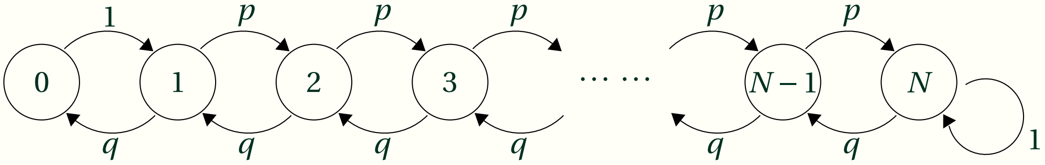

Figure 3: Another random walk similar to Gambler's ruin

Let H_i be the number of steps to reach state N for the first time when $X_0 = i$, and $h_i = \mathbb{E}[H_i]$. We first consider two simple cases:

- $h_0 = 1 + h_1$ since starting from 0 the only choice is moving to 1;
- $h_N = 0$.

Generally, for all $i \geq 1$, we have

$$h_i = q \cdot h_{i-1} + p \cdot h_{i+1} + 1\,,$$

which yields a *inhomogeneous second order linear recurrence* $h_{i+1} = h_i/p - h_{i-1} \cdot q/p - 1/p$. Let $s_i = -i/(2p-1)$. Then $s_{i+1} = s_i/p - s_{i-1} \cdot q/p - 1/p$, and thus

$$(h_{i+1} - s_{i+1}) = (h_i - s_i)/p - (h_{i-1} - s_{i-1}) \cdot q/p\,.$$

Applying the method in Section 2, we obtain that $h_i - s_i = a + b(q/p)^i$ for some a and b.

The key to solve an inhomogeneous linear recurrence is to find a *particular solution*, such as $s_i = -i/(2p-1)$ in our case. However, the way to find a particular solution is heuristic and a bit tricky. Now we would like to introduce another method.

Let Y_i be the number of steps to reach state $i+1$ for the first time when $X_0 = i$, and $y_i = \mathbb{E}[Y_i]$. Then it is clear that

$$H_i = \sum_{j=i}^{N-1} Y_j \quad \text{and} \quad h_i = \sum_{j=i}^{N-1} y_j\,.$$

Similarly to the case of h_i, we can calculate the solution to the boundary case and find the recurrence for y_i as follows:

- $y_0 = 1$;
- $y_i = 1 + q(y_{i-1} + y_i)$ for all $i \geq 1$.

Hence for all $i \geq 1$,

$$y_i = \frac{1}{p} + y_{i-1} \cdot \frac{q}{p}.$$

It implies that

$$\begin{aligned}
y_0 &= 1 \\
y_1 &= \frac{1}{p} + \frac{q}{p} = \frac{1}{p} + \alpha \\
y_2 &= \frac{1}{p} + \alpha\left(\frac{1}{p} + \alpha\right) = \frac{1}{p} + \frac{1}{p} \cdot \alpha + \alpha^2 \\
y_3 &= \frac{1}{p} + \alpha\left(\frac{1}{p} + \frac{1}{p} \cdot \alpha + \alpha^2\right) = \frac{1}{p} + \frac{1}{p} \cdot \alpha + \frac{1}{p} \cdot \alpha^2 + \alpha^3 \\
&\vdots \\
y_i &= \frac{1}{p} \sum_{j=0}^{i-1} \alpha^j + \alpha^i ,
\end{aligned}$$

where we use α to denote q/p. If $\alpha = 1$, then $y_i = 2i - 1$, otherwise

$$y_i = \frac{1}{p} \cdot \frac{1 - \alpha^i}{1 - \alpha} + \alpha^i .$$

In particular, if $\alpha = 1$, then $h_0 = \sum_{i=0}^{N-1}(2i+1) = N^2$.

Remark. The solution here is similar to the solution to the well-known *Coupon Collector Problem.*

The fact that $h_0 = N^2$ if $\alpha = 1$ has many important applications, such as the 2-SAT problem.
2-SAT is the problem of determining whether a CNF formula whose clauses consist of exact 2 literals has satisfying assignments. For example,

$$(x \vee y) \wedge (y \vee \neg z) \wedge (\neg x \vee z)$$

is a 2-CNF formula, and $x = y = z = \text{true}$ is one of its satisfying assignments. The problem of 2-SAT is to determine whether such formulas have satisfying assignments.
There exists a polynomial-time deterministic algorithm based on graph theory that can solve 2-SAT problem perfectly. But now we introduce a much simpler randomized algorithm that can solve this problem with high probability.

1. Let $V = \{v_1, v_2, \ldots, v_n\}$ be the set of variables. Pick an arbitrary assignment $V \to \{\text{true}, \text{false}\}$.
2. If there exists a clause c that has not been satisfied yet, then pick one of two variables incident to c uniformly at random and flip its value. Repeat this step $100n^2$ times, or until there does not exist an unsatisfied clause.

3. The algorithm outputs no solution if there still exists an unsatisfied clause after running $100n^2$ times.

We claim that our algorithm outputs the correct answer with probability at least $1 - 1/100$.

Proof. It is clear that if a 2-SAT instance has no solution then our algorithm will always gives the correct answer. So we consider the probability that our algorithm outputs no solution conditioned on that the instance indeed has a satisfying assignment.

Let $\sigma : V \to \{\mathrm{true}, \mathrm{false}\}$ be a satisfying assignment and our algorithm produces $100n^2$ assignments $\sigma_0, \sigma_1, \ldots, \sigma_{100n^2}$. We claim that the probability that there is no such k that $\sigma_k = \sigma$ is at most $1/100$.

Let X_i be a random variable that

$$X_i = |\{v \in V : \sigma_i(v) \neq \sigma(v)\}|.$$

In fact, $\{X_i\}$ is not a Markov chain. But we can still analyze it. The whole proof will be given in the next lecture. □

(To be continued...)