Lecture 4

Last updated by Serik Sagitov: November 22, 2012

Abstract

A course based on the book Probabilities and Random Processes by Geoffrey Grimmett and David Stirzaker. Chapter 6. Markov chains (6.4-6.5, 6.8-6.9).

1 Stationary distributions

A vector of probabilities $\boldsymbol{\pi} = (\pi_j, j \in S)$ is a stationary distribution for the Markov chain X_n , if given X_0 has distribution $\boldsymbol{\pi}$, X_n has the same distribution $\boldsymbol{\pi}$ for any n, or in other words $\boldsymbol{\pi}$ is a left eigenvector of the transition matrix

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

Theorem 1 An irreducible chain (aperiodic or periodic) has a stationary distribution π iff the chain is positively recurrent; in this case π is the unique stationary distribution and is given by $\pi_i = 1/\mu_i$, where $\mu_i = \mathbb{E}(T_i|X_0 = i)$ and T_i is the time of first return to *i*.

Proof*. Let $\rho(k) = (\rho_j(k), j \in S)$ where $\rho_k(k) = 1$ and

$$\rho_j(k) = \sum_{n=1}^{\infty} \mathbb{P}(X_n = j, T_k \ge n | X_0 = k)$$

is the mean number of visits of the chain to the state j between two consecutive visits to state k. Then

$$\sum_{j \in S} \rho_j(k) = \sum_{j \in S} \sum_{n=1}^{\infty} \mathbb{P}(X_n = j, T_k \ge n | X_0 = k)$$
$$= \sum_{n=1}^{\infty} \mathbb{P}(T_k \ge n | X_0 = k) = \mathbb{E}(T_k | X_0 = k) = \mu_k$$

Main steps of the proof:

- 1. if the chain is irreducible recurrent, then $\rho_j(k) < \infty$ for any k and j, and furthermore, $\rho(k)\mathbf{P} = \rho(k)$,
- 2. if the chain is irreducible recurrent, there exists a positive root \mathbf{x} of the equation $\mathbf{xP} = \mathbf{x}$, which is unique up to a multiplicative constant; the chain is positively recurrent iff $\sum_{i \in S} x_i < \infty$.

Theorem 2 Let s be any state of an irreducible chain. The chain is transient iff there exists a non-zero bounded solution $(y_j : j \neq s)$ satisfying $|y_j| \leq 1$ for all j to the equations

$$y_i = \sum_{j \in S \setminus \{s\}} p_{ij} y_j, \quad i \in S \setminus \{s\}.$$

$$\tag{1}$$

Proof^{*}. Main step. Let τ_j be the probability of no visit to s ever for a chain started at state j. Then the vector $(\tau_j : j \neq s)$ satisfies (1).

Example 3 Random walk with retaining barrier. Transition probabilities

$$p_{00} = q$$
, $p_{i-1,i} = p$, $p_{i,i-1} = q$, $i \ge 1$.

Let $\rho = p/q$.

- If q < p, take s = 0 to see that $y_j = 1 \rho^{-j}$ satisfies (1). The chain is transient.
- Solve the equation $\pi \mathbf{P} = \pi$ to find that there exists a stationary distribution, with $\pi_j = \rho^j (1 \rho)$, if and only if q > p.
- If q > p, the chain is positively recurrent, and if q = p = 1/2, the chain is null recurrent.

Theorem 4 Ergodic theorem. For an irreducible aperiodic chain we have that

$$p_{ij}^{(n)} \to \frac{1}{\mu_j} \text{ as } n \to \infty \text{ for all } (i,j)$$

More generally, for an aperiodic state j and any state i we have that $p_{ij}^{(n)} \rightarrow \frac{f_{ij}}{\mu_j}$, where f_{ij} is the probability that the chain ever visits j starting at i.

2 Reversibility

Theorem 5 Put $Y_n = X_{N-n}$ for $0 \le n \le N$ where X_n is a stationary Markov chain. Then Y_n is a Markov chain with π :

$$\mathbb{P}(Y_{n+1} = j | Y_n = i) = \frac{\pi_j p_{ji}}{\pi_i}.$$

The chain Y_n is called the time-reversal of X_n . If π exists and $\frac{\pi_j p_{ji}}{\pi_i} = p_{ij}$, the chain X_n is called reversible (in equilibrium). The detailed balance equations

$$\pi_i p_{ij} = \pi_j p_{ji} \text{ for all } (i,j).$$
⁽²⁾

Theorem 6 Consider an irreducible chain and suppose there exists a distribution π such that (2) holds. Then π is a stationary distribution of the chain. Furthermore, the chain is reversible.

Proof. Using (2) we obtain

$$\sum_{i} \pi_i p_{ij} = \sum_{i} \pi_j p_{ji} = \pi_j.$$

Example. Ehrenfest model of diffusion: flow of m particles between two connected chambers. Pick a particle at random and move it to another chamber. Let X_n be the number of particles in the first chamber. State space $S = \{0, 1, ..., m\}$ and transition probabilities

$$p_{i,i+1} = \frac{m-i}{m}, \qquad p_{i,i-1} = \frac{i}{m}.$$

The detailed balance equations

$$\pi_i \frac{m-i}{m} = \pi_{i+1} \frac{i+1}{m}$$

imply

$$\pi_i = \frac{m-i+1}{i}\pi_{i-1} = \binom{m}{i}\pi_0.$$

Using $\sum_i \pi_i = 1$ we find that the stationary distribution $\pi_i = \binom{m}{i} 2^{-n}$ is a symmetric binomial.

3 Poisson process and continuous-time Markov chains

A Poisson process N(t) with intensity λ is the number of events observed up to time t given that the inter-arrival times are independent exponentials with parameter λ :

$$\mathbb{P}(N(t) = k) = \frac{(\lambda t)^k}{k!} e^{-\lambda t}, \quad k \ge 0.$$

More generally, a pure birth process X(t) with intensities $\{\lambda_i\}_{i=0}^{\infty}$

- holds at state *i* an exponential time with parameter λ_i ,
- after the holding time it jumps up from i to i + 1.

Exponential holding times has no memory and therefore imply the Markov property in the continuous time setting.

Explosion: $\mathbb{P}(X(t) = \infty) > 0$ for a finite t. It is possible iff $\sum 1/\lambda_i < \infty$.

Definition 7 A continuous-time process X(t) with a countable state space S satisfies the Markov property if

$$\mathbb{P}(X(t_n) = j | X(t_1) = i_1, \dots, X(t_{n-1}) = i_{n-1}) = \mathbb{P}(X(t_n) = j | X(t_{n-1}) = i_{n-1})$$

for any states $j, i_1, \ldots, i_{n-1} \in S$ and any times $t_1 < \ldots < t_n$.

In the time homogeneous case compared to the discrete time case instead of transition matrices \mathbf{P}^n with elements $p_{ij}^{(n)}$ we have transition matrices \mathbf{P}_t with elements

$$p_{ij}(t) = \mathbb{P}(X(u+t) = j | X(u) = i).$$

Chapman-Kolmogorov: $\mathbf{P}_{t+s} = \mathbf{P}_t \mathbf{P}_s$ for all $t \ge 0$ and $s \ge 0$. Here $\mathbf{P}_0 = \mathbf{I}$ is the identity matrix.

Example. For the Poisson process we have

$$p_{ij}(t) = \frac{(\lambda t)^{j-i}}{(j-i)!} e^{-\lambda t}$$

and

$$\sum_{k} p_{ik}(t) p_{kj}(s) = \sum_{k} \frac{(\lambda t)^{k-i}}{(k-i)!} e^{-\lambda t} \frac{(\lambda s)^{j-k}}{(j-k)!} e^{-\lambda s} = \frac{(\lambda t + \lambda s)^{j-i}}{(j-i)!} e^{-\lambda (t+s)} = p_{ij}(t+s).$$

4 The generator of a continuous-time Markov chain

A generator $\mathbf{G} = (g_{ij})$ is a matrix with non-negative off-diagonal elements such that $\sum_j g_{ij} = 0$. A Markov chain X(t) with generator \mathbf{G}

- holds at state *i* an exponential time with parameter $\lambda_i = -g_{ii}$,
- after the holding time it jumps from i to $j \neq i$ with probability $h_{ij} = \frac{g_{ij}}{\lambda_i}$.

The embedded discrete Markov chain is governed by transition matrix $\mathbf{H} = (h_{ij})$ satisfying $h_{ii} = 0$. A continuous-time MC is a discrete MC plus holding intensities (λ_i) .

Example. The Poisson process and birth process have the same embedded MC with $h_{i,i+1} = 1$. For the birth process $g_{ii} = -\lambda_i$, $g_{i,i+1} = \lambda_i$ and all other $g_{ij} = 0$.

Kolmogorov equations. Forward equation: for any $i, j \in S$

$$p_{ij}'(t) = \sum_{k} p_{ik}(t)g_{kj}$$

or in the matrix form $\mathbf{P}'_t = \mathbf{P}_t \mathbf{G}$. It is obtained from $\mathbf{P}_{t+\epsilon} - \mathbf{P}_t = \mathbf{P}_t (\mathbf{P}_{\epsilon} - \mathbf{P}_0)$ watching for the last change. Backward equation $\mathbf{P}'_t = \mathbf{G}\mathbf{P}_t$ is obtained from $\mathbf{P}_{t+\epsilon} - \mathbf{P}_t = (\mathbf{P}_{\epsilon} - \mathbf{P}_0)\mathbf{P}_t$ watching for the initial change. These equations often have a unique solution

$$\mathbf{P}_t = e^{t\mathbf{G}} := \sum_{n=0}^{\infty} \frac{t^n}{n!} \mathbf{G}^n.$$

Theorem 8 Stationary distribution: $\pi \mathbf{P}_t = \pi$ for all t iff (a counterpart of the discrete time equation $\pi \mathbf{P} = \pi$)

$$\pi \mathbf{G} = \mathbf{0}.$$

Proof:

$$\boldsymbol{\pi} \mathbf{P}_t \stackrel{\forall t}{=} \boldsymbol{\pi} \quad \Leftrightarrow \quad \sum_{n=0}^{\infty} \frac{t^n}{n!} \boldsymbol{\pi} \mathbf{G}^n \stackrel{\forall t}{=} \boldsymbol{\pi} \quad \Leftrightarrow \quad \sum_{n=1}^{\infty} \frac{t^n}{n!} \boldsymbol{\pi} \mathbf{G}^n \stackrel{\forall t}{=} \mathbf{0} \quad \Leftrightarrow \quad \boldsymbol{\pi} \mathbf{G}^n \stackrel{\forall n}{=} \mathbf{0}$$

Example. Check that the birth process has no stationary distribution.

Theorem 9 Let X(t) be irreducible with generator **G**. If there exists a stationary distribution π , then it is unique and for all (i, j)

$$p_{ij}(t) \to \pi_j, \quad t \to \infty.$$

If there is no stationary distribution, then $p_{ij}(t) \rightarrow 0$ as $t \rightarrow \infty$.

Example. Poisson process holding times $\lambda_i = \lambda$. Then $\mathbf{G} = \lambda(\mathbf{H} - \mathbf{I})$ and $\mathbf{P}_t = e^{\lambda t(\mathbf{H} - \mathbf{I})}$.