Lecture 6: Discrete-time Markov Chains III

Uffe Høgsbro Thygesen
Informatics and Mathematical Modelling
Technical University of Denmark
2800 Kgs. Lyngby – Denmark
Email: uht@imm.dtu.dk

Outline

Outline
The vector $\rho(k)$
The stationary distribution
Finite state spaces
Perron-Frobenius
Time to absorption
Hitting times
Detailed balance
Summary
Exercise

Recap the stationary distribution, and the vector $\rho(k)$.

Markov chains on finite state spaces

- What to learn from the eigenvalues/eigenvectors of $\mathbf{P}$
- How to analyse $\mathbf{P}$ numerically

Time-reversible Markov chains and “detailed balance”

The vector $\rho(k)$

Outline
The vector $\rho(k)$
The stationary distribution
Finite state spaces
Perron-Frobenius
Time to absorption
Hitting times
Detailed balance
Summary
Exercise

Let the chain start with $X_0 = k$, and stop monitoring the process at $T_k = \min\{n > 0 : X_n = k\}$.

Let N_i denote number of visits to state i

$$N_i = \#\{n : X_n = i, T_k \geq n\} = \sum_{n=1}^{\infty} \mathbf{1}(X_n = i, T_k \geq n)$$

Let $\rho_i(k) = \mathbb{E}^k N_i$. Taking expectations of the indicator functions

$$\rho_i(k) = \sum_{n=1}^{\infty} \mathbb{P}^k(X_n = i, T_k \geq n)$$

$\rho(k)$ for a two-state Markov chain

Outline
The vector $\rho(k)$
The stationary distribution
Finite state spaces
Perron-Frobenius
Time to absorption
Hitting times
Detailed balance
Summary
Exercise

$$\mathbf{P} = \begin{bmatrix} 1-p & p \\ q & 1-q \end{bmatrix} \text{ with } 0 < p, q \leq 1$$

We can write up the distribution of N_2 explicitly:

$$\mathbb{P}^1(N_2 = k) = \begin{cases} k = 0 & : & 1-p \\ k > 0 & : & p(1-q)^{k-1}q \end{cases}$$

Note the geometric tails. In particular

$$\rho(1) = (1, p/q)$$

so that $\rho(1) = \rho(1)\mathbf{P}$.

Compare the stationary distribution

$$\pi = \frac{1}{q+p}(q, p)$$

Lemma 6.4.5: $\rho(k) = \rho(k)\mathbf{P}$

During one cycle, we expect to spend $\rho_i(k)$ time steps at state i.

For each of these time steps, the probability that the next time step is at state j, is p_{ij}.

Thus

$$\rho_j(k) = \sum_i \rho_i(k) p_{ij}$$

The simple symmetric random walk

In this case

$$\rho_i(k) = 1$$

is a solution of $\rho(k) = \rho(k)\mathbf{P}$.

So: Start the walk in $k = 0$, and let i be arbitrary. The expected number of visits to i before returning to the origin is 1.

(We already saw this surprising result in theorem 3.10.8)

$\rho(k)$ and the mean recurrence time

The time until recurrence must be spent somewhere: $T_k = \sum_i N_i$.
Assuming positive recurrency, we take expectations w.r.t. $\mathbb{P}^k$:

$$\mu_k = \sum_i \rho_i(k)$$

This, together with the other properties of $\rho(k)$:

1. $\rho(k) = \rho(k)\mathbf{P}$,
2. $\rho_k(k) = 1$,

means that we can generate a stationary distribution from ρ, whenever $\sum_j \rho_j(k)$ is finite:

$$\pi = \rho(k)/\mu_k$$

Note that $\pi_k = 1/\mu_k$.

The stationary distribution

In general the distribution evolves according to the Chapman-Kolmogorov equations (lemma 6.1.8 with $n = 1$)

$$\mu^{(m+1)} = \mu^{(m)}\mathbf{P}$$

(Recall that $\mu_i^{(m)} = \mathbb{P}(X_m = i)$)
A stationary distribution is a constant (in time) solution to this recursion

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

In addition π must be a distribution - so

$$\pi_i \geq 0 \text{ and } \sum_{i=1}^{N} \pi_i = 1$$

Recap properties of the stationary distribution

1. If the initial state is distributed according to π, then any later state is, too.
2. The mean recurrence time μ_i is $1/\pi_i$.
3. If the chain is irreducible and aperiodic

$$p_{ij}(n) \to \pi_j \text{ as } n \to \infty \ .$$

4. If the chain is irreducible, the fraction of time spent in state i over a long time interval is π_i (see problem 7.11.32)

How to find the stationary distribution π?

(Simulate for a long time and plot a histogram - provided the chain is *ergodic*)

A simple technique is null spaces:

In Matlab

```
pivec = null(P'-eye(length(P)));
pivec = pivec/sum(pivec);
```

In R/ MASS

```
pivec <- Null(P - diag(nrow(P)))
pivec <- pivec/sum(pivec)
```


Finding π using eigenvectors

The stationarity condition

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

says that π is a left eigenvector of $\mathbf{P}$, with eigenvalue $\lambda = 1$.
In Matlab:

```
[V,D] = eig(P');
pivec = V(:, abs(diag(D) - 1)< 1e-6);
pivec = real(pivec/sum(pivec));
```

In R:

```
evs <- eigen(t(P))
pivec <- evs$vectors[, abs(evs$values - 1) < 1e-6]
pivec <- Re(pivec/sum(pivec))
```


The Perron-Frobenius theorem 6.6.1

We **know** that 1 is an eigenvalue of a stochastic matrix, because with $\mathbf{1} = (1, \dots, 1)'$:

$$\mathbf{P1} = \mathbf{1} \text{ , i.e. } \sum_j p_{ij} = 1$$

According to the Perron-Frobenius theorem: If the chain is aperiodic and irreducible, then the eigenvalue 1 is simple (i.e., multiplicity 1) and all other eigenvalues λ have $|\lambda| < 1$.

According to Farkas' theorem (exercise 6.6.2), the left eigenvector can be taken to be a distribution.

Finding the time to absorption

Consider a chain with one absorbing state, and transition matrix

$$\begin{bmatrix} \mathbf{P} & \mathbf{p} \\ \mathbf{0} & 1 \end{bmatrix}$$

Let π be the initial distribution on $\{1, \ldots, N-1\}$. Let $\mathbf{e} = (1, \ldots, 1)'$.

The probability of not being absorbed by time n is $\pi \mathbf{P}^n \mathbf{e}$. This is the **survival function**.

The expected time to absorption is $\sum_{n=0}^{\infty} \pi \mathbf{P}^n \mathbf{e} = \pi (I - \mathbf{P})^{-1} \mathbf{e}$.

Finding the hitting time distribution $f_{ij}(n)$

Consider an irreducible chain.

Let T_j be the stopping time $\min\{n \geq 1 : X_n = j\}$.

$f_{ij}(\cdot)$ is the distribution of T_j, when starting in $X_0 = i$.

Modify $\mathbf{P}$ so state j is absorbing - this does not change f_{ij} when $i \neq j$.

Iterate $\mu^{(n+1)} = \mu^{(n)} \mathbf{P}$ starting with an $\mu^{(0)}$ which has a 1 at i and zeros elsewhere. Then

$$\mathbb{P}^i(T_j \leq n) = \mu_j^{(n)}$$

is the **distribution function** of T_j, and

$$f_{ij}(n) = \mathbb{P}^i(T_j \leq n) - \mathbb{P}^i(T_j \leq n-1)$$

The mean hitting time

(as in the gamble for the Jaguar)

Define ϕ so that $\phi_i = \mathbb{E}^i T_j$

Condition on the first time step!

$$\mathbb{E}^i T_j = \mathbb{E}^i (E^i T_j | X_1) = \sum_k p_{ik} (\phi_k + 1)$$

whenever $i \neq j$. When $i = j$ we have $\phi_j = 0$. In matrix-vector form:

$$\phi = D(\mathbf{P}\phi + \mathbf{1})$$

where $\mathbf{1} = (1, \ldots, 1)'$, and D is a diagonal matrix with ones in the diagonal *except* one zero at position (j, j).

Steady-state and detailed balance

The steady-state criterion

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

says that net flow away from i is zero. This is **global balance**.

A stronger criterion is that the net exchange between any two states i and j is zero:

$$\pi_i p_{ij} = \pi_j p_{ji}$$

This is the criterion of **detailed balance**.

Sum over i to see that this implies global balance.

Reversible chains and detailed balance

Detailed balance implies that the number of jumps $i \mapsto j$ is the same as the number of jumps $j \mapsto i$.

So we can't distinguish a **forward run** from a **backward run.**

Theorem 6.5.1 says that the reversed chain

$$Y_n = X_{N-n}$$

is Markov with transition probabilities

$$q_{ij} = \frac{\pi_j}{\pi_i} p_{ji}$$

Detailed balance is equivalent to $q_{ij} = p_{ij}$, so the original chain and the reversed chain have same statistics.

Stationary Markov chains on graphs

If the graph has no loops, then the Markov chain is reversible (exercise 6.5.9)

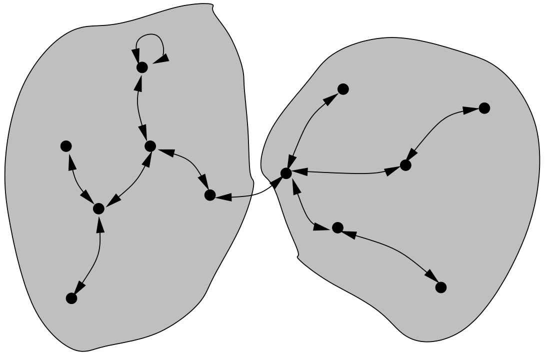

The Ehrenfest model of diffusion

M gas molecules distributed in two chambers. At each time step a random molecule switches chamber.
X_n is the number of molecules in the one chamber.
Transition probabilities:

$$p_{i(i+1)} = \frac{M-i}{M} \ , \ p_{i(i-1)} = \frac{i}{M}$$

Guess the steady-state distribution: binomial$(M, 1/2)$

$$\pi_i = \binom{M}{i} 2^{-M}$$

You may verify that this satisfies detailed balance:

$$\pi_i p_{i(i+1)} = \pi_{i+1} p_{(i+1)i}$$

Summary

- Numerical techniques for analysing finite Markov chains
- The eigenstructure of $\mathbf{P}$
- The Perron-Frobenius theorem
- Finding hitting time distributions using "Forward" iteration.
- "Backward" equations for mean hitting times
- Detailed balance and reversibility

Exercise

Computer exercise uht-02.

Next week

Section 6.8 on Poisson processes.

Section 6.9 on continuous-time Markov chains.