Lecture-22: Reversibility

1 Introduction

Definition 1.1. A stochastic process $X : \Omega \to \mathcal{X}^{\mathbb{R}}$ is *time reversible* if the vector $(X_{t_1}, \dots, X_{t_n})$ has the same distribution as $(X_{\tau - t_1}, \dots, X_{\tau - t_n})$ for all finite positive integers $n \in \mathbb{N}$, time instants $t_1 < t_2 < \dots < t_n \in \mathbb{R}$ and shifts $\tau \in \mathbb{R}$.

Lemma 1.2. *A time reversible process is stationary.*

Proof. It suffices to show that for any shift $s \in \mathbb{R}$, finite $n \in \mathbb{N}$, and time instants $t_1 < \dots < t_n \in \mathbb{R}$, the random vectors $(X_{t_1}, \dots, X_{t_n})$ and $(X_{s+t_1}, \dots, X_{s+t_n})$ have identical distribution regardless of s. This follows from time reversibility of X, since both $(X_{t_1}, \dots, X_{t_n})$ and $(X_{s+t_1}, \dots, X_{s+t_n})$ have the same distribution as $(X_{-t_1}, \dots, X_{-t_n})$, by taking $\tau = 0$ and $\tau = -s$ respectively. □

Theorem 1.3. *A time-homogeneous Markov process $X : \Omega \to \mathcal{X}^{\mathbb{R}}$ with countable state space $\mathcal{X}$ and probability transition kernel $P : \mathbb{R}_+ \to \mathcal{M}(\mathcal{X})^{\mathcal{X}}$ is time reversible iff it is stationary and there exists $\pi \in \mathcal{M}(\mathcal{X})$ that satisfies the detailed balanced conditions for all states $x, y \in \mathcal{X}$ and times $t \in \mathbb{R}_+$,*

$$\pi_x P_{xy}(t) = \pi_y P_{yx}(t). \tag{1}$$

When such a distribution π exists, it is the invariant distribution of the process.

Proof. We assume that the process X is time reversible, and hence stationary. We denote the stationary distribution by $\pi \in \mathcal{M}(\mathcal{X})$, and by time reversibility of X for $\tau = 2s + t$, we have

$$P_\pi \{X_s = x, X_{s+t} = y\} = P_\pi \{X_s = y, X_{s+t} = x\}.$$

Hence, we obtain the detailed balanced conditions in Eq. (1).

Conversely, let π be the distribution that satisfies the detailed balanced conditions in Eq. (1), then summing up both sides over $y \in \mathcal{X}$, we see that π is the invariant distribution for Markov process X. Let $x \in \mathcal{X}^m$, then applying detailed balanced equations in Eq. (1) repeatedly, we can write

$$\pi(x_1) P_{x_1 x_2}(t_2 - t_1) \dots P_{x_{m-1} x_m}(t_m - t_{m-1}) = \pi(x_m) P_{x_m x_{m-1}}(t_m - t_{m-1}) \dots P_{x_2 x_1}(t_2 - t_1).$$

For the time homogeneous stationary Markov process X, it follows that for all $t_0 \in \mathbb{R}_+$

$$P_\pi\Big(\cap_{i=1}^m \{X_{t_i} = x_i\} \Big) = P_\pi\Big(\cap_{i=1}^m \{X_{t_0 + t_m - t_i} = x_i\} \Big).$$

Since $m \in \mathbb{N}$ and $t_0, t_1, \dots, t_m$ were arbitrary, the time reversibility follows for all $\tau = t_0 + t_m$. □

1.1 Reversible chains

Corollary 1.4. *A stationary time-homogeneous discrete time Markov chain $X : \Omega \to \mathcal{X}^{\mathbb{Z}}$ with transition matrix $P \in \mathcal{M}(\mathcal{X})^{\mathcal{X}}$ is time reversible iff there exists $\pi \in \mathcal{M}(\mathcal{X})$ that satisfies the detailed balanced conditions for all states $x, y \in \mathcal{X}$,*

$$\pi_x P_{xy} = \pi_y P_{yx}. \tag{2}$$

When such a distribution π exists, it is the invariant distribution of the process.

Example 1.5 (Random walks on edge-weighted graphs). Consider an undirected graph $G = (\mathcal{X}, E)$ with the vertex set $\mathcal{X}$ and the edge set $E \subseteq \binom{\mathcal{X}}{2}$ being a subset of unordered pairs of elements from $\mathcal{X}$. We say that y is a neighbor of x, if $e = \{x, y\} \in E$ and denote $x \sim y$. We assume a function $w : E \to \mathbb{R}_+$, such that w_e is a positive number associated with each edge $e = \{x, y\} \in E$. Let $X_n \in \mathcal{X}$ denote the location of a particle on one of the graph vertices at the nth time-step. Consider the following random discrete time movement of a particle on this graph from one vertex to another. If the particle is currently at vertex x then it will next move to vertex y with probability

$$P^G_{xy} \triangleq P(\{X_{n+1} = y\} \mid \{X_n = x\}) = \frac{w_{\{x,y\}}}{\sum_{e \in E} w_e \mathbb{1}_e(x)} \mathbb{1}_E(\{x, y\}).$$

The Markov chain $X : \Omega \to \mathcal{X}^{\mathbb{N}}$ describing the sequence of vertices visited by the particle is a random walk on an undirected edge-weighted graph. Google's PageRank algorithm, to estimate the relative importance of webpages, is essentially a random walk on a directed graph!

Proposition 1.6. *Consider an irreducible time-homogeneous Markov chain $X : \Omega \to \mathcal{X}^{\mathbb{Z}}$ that describes the random walk on an edge weighted graph with a finite number of vertices. In steady state, this Markov chain is time reversible with stationary probability of being in a state $x \in \mathcal{X}$ given by*

$$\pi_x = \frac{\sum_{e \in E} w_e \mathbb{1}_e(x)}{2 \sum_{f \in E} w_f}. \tag{3}$$

Proof. Using the definition of transition probabilities for this Markov chain and the given distribution $\pi \in \mathcal{M}(\mathcal{X})$ defined in (3), we notice that

$$\pi_x P^G_{xy} = \frac{w_{\{x,y\}}}{2 \sum_{f \in E} w_f} \mathbb{1}_E(\{x, y\}), \qquad \pi_y P^G_{yx} = \frac{w_{\{x,y\}}}{2 \sum_{f \in E} w_f} \mathbb{1}_E(\{x, y\}).$$

Hence, the detailed balance equation for each pair of states $x, y \in \mathcal{X}$ is satisfied, and the result follows. □

We can also show the following *dual* result.

Lemma 1.7. *Consider a time reversible Markov chain $X : \Omega \to \mathcal{X}^{\mathbb{Z}}$ on a finite state space $\mathcal{X}$ and transition probability matrix $P \in \mathcal{M}(\mathcal{X})^{\mathcal{X}}$. Then, there exists a random walk on a weighted, undirected graph G with the same transition probability matrix P.*

Proof. Since X is time reversible, it is stationary and has a positive invariant distribution $\pi \in \mathcal{M}(\mathcal{X})$ such that $\pi_x P_{xy} = \pi_y P_{yx}$ for each $(x, y) \in \mathcal{X}^2$. This implies that $P_{xy} > 0$ iff $P_{yx} > 0$, and thus we can create a graph $G = (\mathcal{X}, E)$, where

$$E \triangleq \left\{ \{x, y\} \in \binom{\mathcal{X}}{2} : P_{xy} P_{yx} > 0 \right\}.$$

For each edge $\{x, y\} \in E$, we set the edge weights $w_{\{x,y\}} \triangleq \pi_x P_{xy} = \pi_y P_{yx}$. With this choice of weights, it is easy to check that $w_x \triangleq \sum_{e \in E} w_e \mathbb{1}_e(x) = \pi_x$, and the transition matrix associated with a random walk on this graph is exactly P with $P^G_{xy} = \frac{w_{\{x,y\}}}{w_x} = P_{xy}$. □

Is every Markov chain time reversible?
Let $X : \Omega \to \mathcal{X}^{\mathbb{Z}}$ be a time homogeneous discrete time Markov chain with probability transition matrix $P \in \mathcal{M}(\mathcal{X})^{\mathcal{X}}$.

1. If the process is not stationary, then no. To see this, we observe that

   $$P\{X_{t_1} = x_1, X_{t_2} = x_2\} = \nu_{t_1}(x_1) P_{x_1 x_2}(t_2 - t_1), \quad P\{X_{\tau - t_2} = x_2, X_{\tau - t_1} = x_1\} = \nu_{\tau - t_2}(x_2) P_{x_2 x_1}(t_2 - t_1).$$

   If the process is not stationary, the two probabilities can't be equal for all times τ, t_1, t_2 and states $x_1, x_2 \in \mathcal{X}$.
2. If the process is stationary, then it is still not true in general. Suppose we want to find a stationary distribution $\alpha \in \mathcal{M}(\mathcal{X})$ that satisfies the detailed balance equations $\alpha_x P_{xy} = \alpha_y P_{yx}$ for all states $x, y \in \mathcal{X}$. For any arbitrary Markov chain X, one may not end up getting any solution. To see this

consider a path $x \to y \to z$ such that $P_{xy}P_{yz}P_{zx} > 0$. Time reversibility condition implies that

$$\alpha_x P_{xy} P_{yz} P_{zx} = \alpha_x P_{xz} P_{zy} P_{yx}.$$

However, this would imply that $\frac{P_{zy}P_{yx}}{P_{xy}P_{yz}} = \frac{P_{zx}}{P_{xz}}$, which is not true in general. Thus, we see that a necessary condition for time reversibility is $P_{xy}P_{yz}P_{zx} = P_{xz}P_{zy}P_{yx}$ for all $x, y, z \in \mathcal{X}$.

Theorem 1.8 (Kolmogorov's criterion for time reversibility of Markov chains). *A stationary Markov chain $X : \Omega \to \mathcal{X}^{\mathbb{Z}}$ is time reversible if and only if starting in state $x_0 \in \mathcal{X}$, any path back to state x_0 has the same probability as the time reversed path, for all initial states $x_0 \in \mathcal{X}$. That is, for any $n \in \mathbb{N}$ and state vector $x \in \mathcal{X}^n$*

$$P_{x_0x_1} P_{x_1x_2} \dots P_{x_nx_0} = P_{x_0x_n} P_{x_nx_{n-1}} \dots P_{x_1x_0}. \tag{4}$$

Proof. The detailed balance equation for a time reversible Markov process X implies that (4) holds for any finite set of states. Conversely, if (4) holds for any non-negative integer $n \in \mathbb{N}$, then for any states $x_0, y \in \mathcal{X}$, we have

$$(P^{n+1})_{x_0y} P_{yx_0} = \sum_{x_1,x_2,\dots x_n} P_{x_0x_1} \dots P_{x_ny} P_{yx_0} = \sum_{x_1,x_2,\dots x_n} P_{x_0y} P_{yx_n} \dots P_{x_1x_0} = P_{x_0y}(P^{n+1})_{yx_0}.$$

Taking the limit $n \to \infty$ and noticing that $\lim_{n\to\infty}(P^n)_{xy} = \pi_y$ for all $x, y \in \mathcal{X}$, we observe that X is a time-reversible process. □

1.2 Reversible Processes

Corollary 1.9. *A stationary Markov process $X : \Omega \to \mathcal{X}^{\mathbb{R}}$ with generator matrix $Q \in \mathbb{R}^{\mathcal{X}\times\mathcal{X}}$ is time reversible iff there exists a probability distribution $\pi \in \mathcal{M}(\mathcal{X})$, that satisfies the detailed balanced conditions*

$$\pi_x Q_{xy} = \pi_y Q_{yx}, \quad x, y \in \mathcal{X}. \tag{5}$$

When such a distribution π exists, it is the invariant distribution of the process.

Definition 1.10. Consider a stationary time-homogeneous Markov process $X : \Omega \to \mathcal{X}^{\mathbb{R}}$ with invariant distribution $\pi \in \mathcal{M}(\mathcal{X})$ and the generator matrix $Q \in \mathbb{R}^{\mathcal{X}\times\mathcal{X}}$. We denote the total number of transitions from state x to state y in the time duration $(0, t]$ by

$$N_t^{xy} \triangleq N^{xy}(0,t] \triangleq \sum_{n\in\mathbb{N}} \mathbb{1}_{(0,t]}(S_n) \mathbb{1}_{\{Z_{n-1}=x, Z_n=y\}}.$$

The *probability flux* from state x to state y is defined as $\Phi_{xy} \triangleq \lim_{t\to\infty} \frac{1}{t} N_t^{xy}$.

Lemma 1.11. *For a time-homogeneous CTMC X, the probability flux from state x to state y is $\pi_x Q_{xy}$.*

Proof. Let $X_0 = x$ and $\tau_x^+(k)$ be the kth visiting time to state x. It follows that $\tau_x^+ : \Omega \to \mathbb{R}_+^{\mathbb{N}}$ is a renewal sequence. We consider the reward process $N^{xy} : \Omega \to \mathbb{Z}_+^{\mathbb{R}_+}$ where N_t^{xy} is the number of transitions from state x to y in the duration $(0, t]$. We denote the total number of transitions from state x to state y in the kth inter-renewal duration by

$$N^{xy}(k) \triangleq N^{xy}(\tau_x^+(k-1), \tau_x^+(k)] \triangleq \sum_{n\in\mathbb{N}} \mathbb{1}_{(\tau_x^+(k-1),\tau_x^+(k)]}(S_n) \mathbb{1}_{\{Z_{n-1}=x, Z_n=y\}}.$$

The number of visit to all states $y \in \mathcal{X}$ during kth successive visit to state $x \in \mathcal{X}$ is the number of transitions during $(\tau_x^+(k-1), \tau_x^+(k)]$, and we denote this number as $N^x(k) \triangleq \sum_{n\in\mathbb{N}} \mathbb{1}_{(\tau_x^+(k-1),\tau_x^+(k)]}(S_n)$. From the renewal reward theorem for the embedded DTMC $Z : \Omega \to \mathcal{X}^{\mathbb{Z}_+}$ with invariant distribution $u \in \mathcal{M}(\mathcal{X})$, we can write the average number of one-step transitions from state x to y as

$$u_x p_{xy} = \lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} \mathbb{1}_{\{Z_{n-1}=x, Z_n=y\}} = \frac{\mathbb{E}_x N^{xy}(k)}{\mathbb{E}_x N^x(k)} = u_x \mathbb{E}_x N^{xy}(k).$$

It follows that $\mathbb{E}_x N^{xy}(k) = p_{xy}$ and recall that $\mathbb{E}_x \tau_x^+(1) = \frac{1}{\pi_x \nu_x}$. From the renewal reward theorem applied to reward process N^{xy} and renewal sequence τ_x^+, we obtain

$$\lim_{t\to\infty} \frac{N_t^{xy}}{t} = \frac{\mathbb{E}_x N^{xy}(1)}{\mathbb{E}_x \tau_x^+(1)} = \pi_x \nu_x p_{xy} = \pi_x Q_{xy}.$$

□

Lemma 1.12. *For a stationary time-homogeneous Markov process $X : \Omega \to \mathcal{X}^{\mathbb{R}}$, probability flux balances across a cut $A \subseteq \mathcal{X}$, that is*

$$\sum_{y \notin A} \sum_{x \in A} \pi_x Q_{xy} = \sum_{x \in A} \sum_{y \notin A} \pi_y Q_{yx}.$$

Proof. Let $A \subseteq \mathcal{X}$, and denote the number of visits from states in A to states in A^c in the interval $(0, t]$ and probability flux from $A \to A^c$ as

$$N_t^{A,A^c} \triangleq \sum_{y \notin A} \sum_{x \in A} N_t^{xy}, \qquad \Phi^{A,A^c} = \sum_{y \notin A} \sum_{x \in A} \Phi_{xy} = \lim_{t \to \infty} \frac{1}{t} N_t^{A,A^c}.$$

By definition of probability flux across cut A, it suffice to show that $\left| N_t^{A,A^c} - N_t^{A^c,A} \right| \leqslant 1$, which follows from the observe that the difference $N_t^{A,A^c} - N_t^{A^c,A} = \mathbb{1}_{\{X_0 \in A\}} - \mathbb{1}_{\{X_t \notin A\}}$ for any time $t \in \mathbb{R}_+$. □

Corollary 1.13. *For $A = \{x\}$, the above equation reduces to the full balance equation for state x, i.e.,*

$$\sum_{y : y \neq x} \pi_x Q_{xy} = \sum_{y : y \neq x} \pi_y Q_{yx}.$$

Definition 1.14. A time-homogeneous Markov process $X : \Omega \to \mathbb{Z}_+^{\mathbb{R}_+}$ is called a *birth-death process* if its generator matrix satisfies $Q_{x,y} = 0$ for all states $x, y \in \mathbb{Z}$ such that $|y - x| > 1$. We define two non-negative sequences birth and death rates denoted by $\lambda \in \mathbb{R}_+^{\mathbb{Z}}$ and $\mu \in \mathbb{R}_+^{\mathbb{N}}$, such that for all $n \in \mathbb{N}$

$$\lambda_n \triangleq Q_{n-1,n}, \qquad \mu_n \triangleq Q_{n,n-1}.$$

Proposition 1.15. *An ergodic birth-death process in steady-state is time-reversible.*

Proof. Since the process is stationary, the probability flux must balance across any cut of the form $A = \{0, 1, 2, \dots, n\}$, for $n \in \mathbb{Z}_+$. Since there are no other transitions possible across the cut, this is precisely the set of detailed balance equations $\pi_n \lambda_n = \pi_{n+1} \mu_{n+1}$ for each state $n \in \mathbb{Z}_+$, and hence the process is time-reversible. □

In fact, the following, more general, statement can be proven using similar ideas.

Proposition 1.16. *Consider an irreducible and ergodic Markov process $X : \Omega \to \mathcal{X}^{\mathbb{R}}$ on a countable state space $\mathcal{X}$ with generator matrix $Q \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ having the following property. For any pair of states $x \neq y \in \mathcal{X}$, the transition graph has a* unique *path $x = x_0 \to x_1 \to \cdots \to x_{n(x,y)} = y$ and $y = x_{n(x,y)} \to x_{n-1} \to \cdots \to x_0 = x$ of distinct states. Then the process X is time reversible at stationarity.*

Proof. Let the stationary distribution of X be $\pi \in \mathcal{M}(\mathcal{X})$, such that $\pi Q = 0$. We fix a state $x \in \mathcal{X}$, and define the set of states connected to x as $B_x \triangleq \{y \in \mathcal{X} : Q_{xy} > 0\}$. By theorem hypothesis, for each $y \in B_x$ we have a unique path $x \to y$ and $y \to x$, and thus we have $Q_{yx} > 0$ as well. For any $y \notin B_x$, the detailed balance equation is satisfied trivially for each pair (x, y). Let $y \in B_x$, then we can define

$$A_{xy} \triangleq \{z \in \mathcal{X} : z \text{ connected to } x \text{ via } y\}.$$

By definition of A_{xy}, we have a path $x \to y \to z$ for any $z \in A_{xy}$. From the hypothesis of unique paths, we have $x \in A_{xy}$. Further, since self transitions are not possible, $y \notin A_{xy}$. Since Q is irreducible, each state x is connected to every other state $z \in \mathcal{X} \setminus \{x\}$. Therefore, we observe that

$$A_{xy}^c = \{w \in \mathcal{X} : w \text{ connected to } x \text{ not via } y\}.$$

We observe that $x \notin A_{xy}^c$ and $y \in A_{xy}^c$. Next, we consider a pair of states (z, w) such that $z \in A_{xy} \setminus \{x\}$ and $w \in A_{xy}^c \setminus \{y\}$. In this case, if $Q_{zw} > 0$, then we have two paths $x \to y \to z \to w$ and another path from x to w without going via y, and that contradicts the hypothesis. It follows that $Q_{zw} = Q_{wz} = 0$ for all such pairs (z, w). This implies that there are no paths between $A_{xy} \setminus \{x\}$ and $A_{xy}^c \setminus \{y\}$. From the probability flux balance across cuts, we obtain the detailed balance equation

$$\pi_x Q_{xy} = \sum_{w \notin A_{xy}} \sum_{z \in A_{xy}} \pi_z Q_{zw} = \sum_{z \in A_{xy}} \sum_{w \notin A_{xy}} \pi_w Q_{wz} = \pi_y Q_{yx}.$$

Since the choice of states $x, y \in \mathcal{X}$ was arbitrary, the result follows. □

Exercise 1.17. Prove Corollary 1.4 and Corollary 1.9 from Theorem 1.3.