Stochastic Processes Notes

gordonbchen

1 Chapter 1

1.1 Markov Property

Markov property: $\boxed{\mathbb{P}(X_{n+1}|X_n, X_{n-1} \dots X_1, X_0) = \mathbb{P}(X_{n+1}|X_n)}$.

$P_{i,j} = \mathbb{P}(X_{n+1} = j|X_n = i)$. Since $\sum_j \mathbb{P}(X_{n+1} = j|X_n = i) = 1$, then $\boxed{P\vec{1} = \vec{1}}$.

m-step transition probability: $\boxed{\mathbb{P}(X_{n+m} = j|X_n = i) = P^m(i, j)}$.

1.2 Classifications

1.2.1 Transient and Recurrent States

Let $\mathbb{P}_y(T_y < \infty) = \rho_{yy}$ be the probability of returning to y, where $T_y = \min\{n \geq 1 : X_n = y\}$ is the time of first return to y. Then the probability of returning to y, k times is $\mathbb{P}_y(T_y^k < \infty) = \rho_{yy}^k$.

If $\rho_{yy} < 1$, then $\lim_{k\to\infty} \rho_{yy}^k = 0$ and the state is transient. The probability for visiting the state more and more times goes to 0, and at some point, the state is never visited again.

Otherwise, if $\rho_{yy} = 1$, then $\lim_{k\to\infty} \rho_{yy}^k = 1$ and the state is recurrent. The state will be visited infinitely many times.

If a set is finite, closed (cannot leave set), and irreducible (all states communicate with each other), then all of its states are recurrent.

1.2.2 Period

Period of Markov Chain: greatest common divisor of all $I_x = \{n \geq 1 : p^n(x, x) > 0\}$, length of all cycles.

Aperiodic Markov Chain: period = 1.

1.3 Stationary Distribution

Stationary distribution π: $\boxed{\pi P = \pi}$

1.3.1 Spectral Theory

$A = \begin{pmatrix} a & 1-a \\ 1-b & b \end{pmatrix}$. A has eigenvalues $\lambda_1 = 1, \lambda_2 = a + b - 1$. $|\lambda_2| < 1$.

$\vec{1}$ is a right eigenvalue of A, $A\vec{1} = \vec{1}$. π is a left eigenvalue of A, $\pi A = \pi$.

$A = PDP^{-1}$, where P are the right eigenvectors of A, D is a diagonal matrix of eigenvalues, and P^{-1} is the matrix of left eigenvectors (rows). To see this last part, note that $A^T = (P^{-1})^T DP^T$ and left eigenvectors are eigenvectors of A^T.

$A = \begin{pmatrix} \vec{r}_1 & \vec{r}_2 \end{pmatrix} \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix} \begin{pmatrix} \vec{l}_1 \\ \vec{l}_2 \end{pmatrix} = \lambda_1 \vec{r}_1 \vec{l}_1 + \lambda_2 \vec{r}_2 \vec{l}_2 = \begin{pmatrix} \pi_1 & \pi_2 \\ \pi_1 & \pi_2 \end{pmatrix} + \lambda_2 \vec{r}_2 \vec{l}_2$

$\lim_{n\to\infty} A^n = PD^nP^{-1} = \begin{pmatrix} \pi_1 & \pi_2 \\ \pi_1 & \pi_2 \end{pmatrix} + \lambda_2^n \vec{r}_2 \vec{l}_2 = \begin{pmatrix} \pi_1 & \pi_2 \\ \pi_1 & \pi_2 \end{pmatrix}$ since $|\lambda_2| < 1$

1.3.2 Doubly Stochastic

Doubly Stochastic: $\vec{1}P = \vec{1}$, ie columns of P sum to 1. Then $\pi = \vec{1}/n$.

Symmetric: $P = P^T$. Since a Markov Chain requires $P\vec{1} = \vec{1}$, $\vec{1}^T P = \vec{1}^T$, and $\pi = \vec{1}/n$. Special case of Doubly Stochastic.

1.3.3 Reversibility: Detailed Balance, Kolmogorov Cycle, and Symmetrizable

Reversible = Detailed Balance Condition = Symmetrizable Matrix = Kolmogorov Cycle Condition.

Detailed Balance Condition: $\boxed{\pi_i P_{ij} = \pi_j P_{ji}}$.

Kolmogorov Cycle Condition: for every cycle $x_0, x_1, x_2, \ldots, x_n = x_0$, $\boxed{\prod_{i=0}^{n-1} p(x_i, x_{i+1}) = \prod_{i=0}^{n-1} p(x_{i+1}, x_i)}$. For a 3×3 matrix one only has to check that $P_{12}P_{23}P_{31} = P_{13}P_{32}P_{21}$.

Symmetrizable matrix P: $\boxed{S = DPD^{-1}}$ is symmetric for some diagonal matrix D. If P is symmetrizable, then $D = \text{diag}(\sqrt{\pi_i})$.

All irreducible linear markov chains ($p(i,j) = 0$ if $|i - j| > 1$) are reversible, since all irreducible 2-state chains are reversible ($P_{ij}P_{ji} = P_{ji}P_{ij}$).

1.4 Limit Distributions

If a Markov Chain is irreducible, $\boxed{\frac{1}{n} \lim_{n\to\infty} p^n(x, y) = \pi(y)}$.

Perron-Frobenius: for irreducible and aperiodic P, $\boxed{\lim_{n\to\infty} P^n = \vec{1}\pi}$ and $\boxed{\lim_{n\to\infty} \mu P^n = \pi}$.

$\boxed{\frac{1}{n} \sum_{n=0}^{\infty} f(X_n) = \sum_x f(x)\pi(x)}$.

1.5 Exit Distributions

$V_F = \min\{n \geq 0 : X_n \in F\}$ is the Hitting Time of set F, the first time X takes on a state in F.

$\mathbb{P}_x(V_A < V_B)$ is the probability that X, starting in state $X_0 = x$, takes on a state in A (wins) before taking on a state in B (loses).

Let $\mathbb{P}_x(V_A < V_B) = h(x)$. $h(a) = 1$ for all $a \in A$, and $h(b) = 0$ for all $b \in B$. $\boxed{h(x) = \sum_y p(x, y)h(y)}$.

1.5.1 Example

A gambler wins \$1 with $p = 1/2$ and loses \$1 with $p = 1/2$. If he ever reaches \$0, the casino makes him stop. If he ever reaches \$4, he is happy and goes home. Let $X \in \{0, 1, 2, 3, 4\}$ be the amount of money the gambler has. Note that $X = 0$ and $X = 4$ are absorbing states. Let $A = \{4\}$ and $B = \{0\}$, so $h(x) = \mathbb{P}_x(V_A < V_B)$ is the probability that the gambler ends up going home with \$4 given that he currently has \$$x$.

$$
\begin{aligned}
h(0) &= 0 \\
h(1) &= \tfrac{1}{2}h(2) \\
h(2) &= \tfrac{1}{2}h(1) + \tfrac{1}{2}h(3) \\
h(3) &= \tfrac{1}{2}h(2) + \tfrac{1}{2} \\
h(4) &= 1
\end{aligned}
\qquad
\left(\begin{array}{c|ccc|c}
1 & 0 & 0 & 0 & 0 \\
\hline
-.5 & 1 & -.5 & 0 & 0 \\
0 & -.5 & 1 & -.5 & 0 \\
0 & 0 & -.5 & 1 & 0 \\
\hline
0 & 0 & 0 & 0 & 1
\end{array}\right)
\begin{pmatrix} h(0) \\ h(1) \\ h(2) \\ h(3) \\ h(4) \end{pmatrix}
=
\begin{pmatrix} 0 \\ 0 \\ 0 \\ .5 \\ 1 \end{pmatrix}
\qquad
\begin{pmatrix} 1 & -.5 & 0 \\ -.5 & 1 & -.5 \\ 0 & -.5 & 1 \end{pmatrix} h = \begin{pmatrix} 0 \\ 0 \\ .5 \end{pmatrix}
$$

Let $C = S - A - B$, $R = P_{x \in C, y \in C}$, and $v_i = p(i, a)$.

$\boxed{h = Rh + v}$. then $(I - R)h = v$ and $h = (I - R)^{-1}v$.

1.5.2 Limit distributions with transient states

$P =$

	1	2	3	4	5	6	7
1	.2	.3	.1	0	.4	0	0
2	0	.5	0	.2	.3	0	0
3	0	0	.7	.3	0	0	0
4	0	0	.6	.4	0	0	0
5	0	0	0	0	.5	.5	0
6	0	0	0	0	0	.2	.8
6	0	0	0	0	1	0	0

$P =$

	1	2	A	B
1	.2	.3	.1	.4
2	0	.5	.2	.3

$h = Rh + v$, solving gives $h = (11/40, 2/5)$.

$\{1, 2\}$ are transient states. $A = \{3, 4\}$ and $B = \{5, 6, 7\}$ are the 2 irreducible, recurrent components.

$\pi_A = (2/3, 1/3)$, and $\pi_B = (8/17, 5/17, 4/17)$.

$\lim_{n\to\infty} P^n =$

	T	A	B
T	0	$h\pi_A$	$(1-h)\pi_B$
A	0	$\vec{1}\pi_A$	0
B	0	0	$\vec{1}\pi_B$

The limit probability of going from a transient state to a recurrent state is the probability of exiting into that recurrent state times the stationary probability of that recurrent state inside the component.

1.6 Exit times

$\mathbb{E}_x[V_A]$: expected number of steps we take before hitting the exit set A if we start from state x.

Let $\mathbb{E}_x[V_A] = g(x)$. $g(a) = 0$ for all $a \in A$.

$g(x) = \sum_y p(x, y)(1 + g(y)) = \sum_y p(x, y) + \sum_y p(x, y)g(y)$, so $\boxed{g(x) = 1 + \sum_y p(x, y)g(y)}$.

1.6.1 Example

We continue with the gambler who has a 50/50 chance of winning/losing (1.5.1). $S = \{0, 1, 2, 3, 4\}$, and $A = \{0, 1\}$.

$g(0) = 0$
$g(1) = 1 + \frac{1}{2}g(2)$
$g(2) = 1 + \frac{1}{2}g(1) + \frac{1}{2}g(3)$
$g(3) = 1 + \frac{1}{2}g(2)$
$g(4) = 0$

$$\left(\begin{array}{c|ccc|c} 1 & 0 & 0 & 0 & 0 \\ \hline -.5 & 1 & -.5 & 0 & 0 \\ 0 & -.5 & 1 & -.5 & 0 \\ 0 & 0 & -.5 & 1 & -.5 \\ \hline 0 & 0 & 0 & 0 & 1 \end{array}\right) \begin{pmatrix} g(0) \\ g(1) \\ g(2) \\ g(3) \\ g(4) \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \\ 1 \\ 1 \\ 0 \end{pmatrix}$$

$\boxed{g = \vec{1} + Rg}$

$(I - R)g = \vec{1}$

$g = (I - R)^{-1}\,\vec{1}$