ECE276B: Planning & Learning in Robotics
Lecture 2: Markov Chains

Nikolay Atanasov

natanasov@ucsd.edu

UC San Diego
JACOBS SCHOOL OF ENGINEERING
Electrical and Computer Engineering

Outline

Markov Chain

- **Stochastic process**: indexed collection of random variables $\{x_0, x_1, \ldots\}$
- **Markov chain**: memoryless stochastic process $\{x_0, x_1, \ldots\}$:
  - x_0 has probability density function $p_0(\cdot)$
  - x_{t+1} conditioned on x_t has probability density function $p_f(\cdot \mid x_t)$ and is independent of the history $x_{0:t-1}$
- **Markov assumption**:

  "The future is independent of the past given the present"

Markov Chain

Stochastic process defined by a tuple $(\mathcal{X}, p_0, p_f)$:

- $\mathcal{X}$ is a discrete or continuous space
- $p_0(\cdot)$ is a prior pdf defined on $\mathcal{X}$
- $p_f(\cdot \mid \mathbf{x})$ is a conditional pdf defined on $\mathcal{X}$ for given $\mathbf{x} \in \mathcal{X}$ that specifies the stochastic process transitions

- When the state space is finite, $\mathcal{X} := \{1, \ldots, N\}$, the pdf p_f can be represented by an $N \times N$ transition matrix with elements:

$$P_{ij} := \mathbb{P}(x_{t+1} = j \mid x_t = i) = p_f(j \mid x_t = i)$$

Example: Student Markov Chain

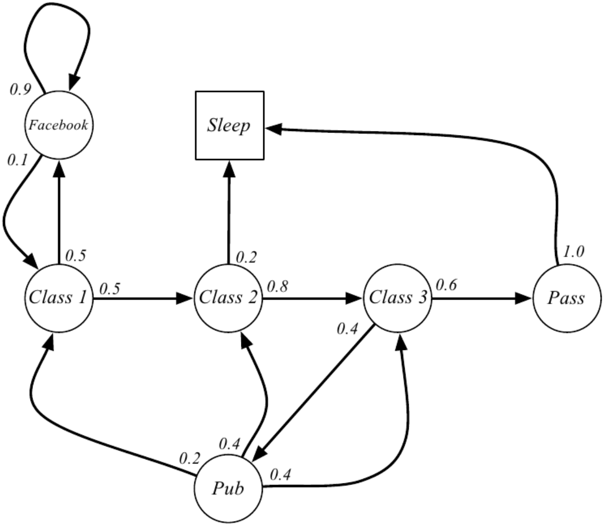

Example: Student Markov Chain

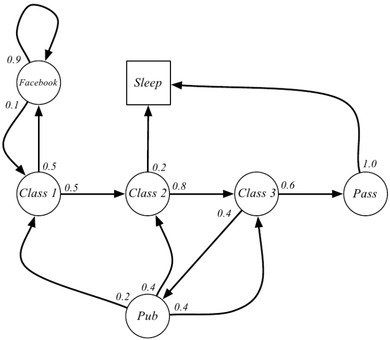

- Sample paths:
  - C1 C2 C3 Pass Sleep
  - C1 FB FB C1 C2 Sleep
  - C1 C2 C3 Pub C2 C3 Pass Sleep
  - C1 FB FB C1 C2 C3 Pub C1 FB FB FB C1 C2 Sleep

- Transition matrix:

$$P = \begin{matrix} FB \\ C1 \\ C2 \\ C3 \\ Pub \\ Pass \\ Sleep \end{matrix} \begin{bmatrix} 0.9 & 0.1 & 0 & 0 & 0 & 0 & 0 \\ 0.5 & 0 & 0.5 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.8 & 0 & 0 & 0.2 \\ 0 & 0 & 0 & 0 & 0.4 & 0.6 & 0 \\ 0 & 0.2 & 0.4 & 0.4 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

Chapman-Kolmogorov Equation

- ***n*-step transition probabilities** of Markov chain on $\mathcal{X} = \{1, \ldots, N\}$

$$P_{ij}^{(n)} := \mathbb{P}(x_{t+n} = j \mid x_t = i) = \mathbb{P}(x_n = j \mid x_0 = i)$$

- **Chapman-Kolmogorov**: the n-step transition probabilities can be obtained recursively from the 1-step transition probabilities:

$$P_{ij}^{(n)} = \sum_{k=1}^{N} P_{ik}^{(m)} P_{kj}^{(n-m)}, \qquad \forall i,\ j,\ n,\ 0 \leq m \leq n$$

$$P^{(n)} = \underbrace{P \cdots P}_{n \text{ times}} = P^n$$

- Given the transition matrix P and a vector $\mathbf{p}_0 := [p_0(1), \ldots, p_0(N)]^\top$ of prior probabilities, the vector of probabilities $\mathbf{p}_n$ after n steps is:

$$\mathbf{p}_n^\top = \mathbf{p}_0^\top P^n$$

Example: Student Markov Chain

$$P = \begin{array}{c} FB \\ C1 \\ C2 \\ C3 \\ Pub \\ Pass \\ Sleep \end{array} \begin{bmatrix} 0.9 & 0.1 & 0 & 0 & 0 & 0 & 0 \\ 0.5 & 0 & 0.5 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.8 & 0 & 0 & 0.2 \\ 0 & 0 & 0 & 0 & 0.4 & 0.6 & 0 \\ 0 & 0.2 & 0.4 & 0.4 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

$$P^2 = \begin{array}{c} FB \\ C1 \\ C2 \\ C3 \\ Pub \\ Pass \\ Sleep \end{array} \begin{bmatrix} 0.86 & 0.09 & 0.05 & 0 & 0 & 0 & 0 \\ 0.45 & 0.05 & 0 & 0.4 & 0 & 0 & 0.1 \\ 0 & 0 & 0 & 0 & 0.32 & 0.48 & 0.2 \\ 0 & 0.08 & 0.16 & 0.16 & 0 & 0 & 0.6 \\ 0.1 & 0 & 0.1 & 0.32 & 0.16 & 0.24 & 0.08 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

$$P^{100} = \begin{array}{c} FB \\ C1 \\ C2 \\ C3 \\ Pub \\ Pass \\ Sleep \end{array} \begin{bmatrix} 0.01 & 0 & 0 & 0 & 0 & 0 & 0.99 \\ 0.01 & 0 & 0 & 0 & 0 & 0 & 0.99 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

First Passage Time

- **First passage time**: the number of transitions necessary to reach state j for the first time is a random variable:

$$\tau_j := \min\{t \geq 1 \mid x_t = j\}$$

- **Recurrence time**: the first passage time τ_i from $x_0 = i$ to $j = i$
- **Probability of first passage in n steps**: $\rho_{ij}^{(n)} := \mathbb{P}(\tau_j = n \mid x_0 = i)$

$$\begin{aligned}
\rho_{ij}^{(1)} &= P_{ij} \\
\rho_{ij}^{(2)} &= [P^2]_{ij} - \rho_{ij}^{(1)} P_{jj} \qquad \text{(first time we visit } j \text{ should not be 1!)} \\
&\vdots \\
\rho_{ij}^{(n)} &= [P^n]_{ij} - \rho_{ij}^{(1)} [P^{n-1}]_{jj} - \rho_{ij}^{(2)} [P^{n-2}]_{jj} - \cdots - \rho_{ij}^{(n-1)} P_{jj}
\end{aligned}$$

- **Probability of first passage**: $\rho_{ij} := \mathbb{P}(\tau_j < \infty \mid x_0 = i) = \sum_{n=1}^{\infty} \rho_{ij}^{(n)}$
- **Number of visits to j up to time n**:

$$v_j^{(n)} := \sum_{t=0}^{n} \mathbb{1}\{x_t = j\} \qquad v_j := \lim_{n \to \infty} v_j^{(n)}$$

Recurrence and Transience

- **Absorbing state**: a state j such that $P_{jj} = 1$
- **Transient state**: a state j such that $\rho_{jj} < 1$
- **Recurrent state**: a state j such that $\rho_{jj} = 1$
- **Positive recurrent state**: a recurrent state j with $\mathbb{E}[\tau_j \mid x_0 = j] < \infty$
- **Null recurrent state**: a recurrent state j with $\mathbb{E}[\tau_j \mid x_0 = j] = \infty$
- **Periodic state**: can only be visited at integer multiples of t
- **Ergodic state**: a positive recurrent state that is aperiodic

Recurrence and Transience

Total Number of Visits Lemma

$\mathbb{P}(v_j \geq k+1 \mid x_0 = j) = \rho_{jj}^k$ for all $k \geq 0$

Proof:
By Markov property and induction: $\mathbb{P}(v_j \geq k+1 \mid x_0 = j) = \rho_{jj}\mathbb{P}(v_j \geq k \mid x_0 = j)$.

0–1 Law for the Total Number of Visits

j is recurrent iff $\mathbb{E}[v_j \mid x_0 = j] = \infty$

Proof: Since v_j is discrete, we can write $v_j = \sum_{k=0}^{\infty} \mathbb{1}\{v_j > k\}$ and

$$\mathbb{E}[v_j \mid x_0 = j] = \sum_{k=0}^{\infty} \mathbb{P}(v_j \geq k+1 \mid x_0 = j) = \sum_{k=0}^{\infty} \rho_{jj}^k = \frac{1}{1-\rho_{jj}}$$

Recurrence Is Contagious

i is recurrent and $\rho_{ij} > 0 \quad \Rightarrow \quad j$ is recurrent and $\rho_{ji} = 1$

Mean First Passage Time

- **Mean first passage time**: $M_{ij} := \mathbb{E}\left[\tau_j \mid x_0 = i\right]$
- By the law of total probability:

$$M_{ij} = P_{ij} + \sum_{k \neq j} P_{ik}(1 + M_{kj}) = 1 + \sum_{k \neq j} P_{ik} M_{kj}$$

- Let $M \in \mathbb{R}^{N \times N}$ with elements M_{ij} contain all mean first passage times
- The matrix of mean first passage times satisfies:

$$M = \mathbf{1}\mathbf{1}^\top + P(M - D)$$

where $D = \mathbf{diag}(M_{11}, \ldots, M_{NN})$ and $\mathbf{1} = \begin{bmatrix} 1 & \cdots & 1 \end{bmatrix}^\top$

Equivalence Classes

- $i \to j$: state j is **accessible** from state i if $P_{ij}^{(n)} > 0$ for some n
- Every state is accessible from itself since $P_{ii}^{(0)} = 1$
- $i \leftrightarrow j$: state i and j **communicate** if they are accessible from each other
- **Equivalence class**: a set of states which communicate with each other
- **Example**: find the equivalence classes for this Markov chain

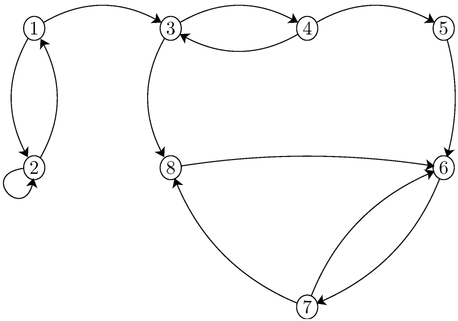

Classification of Markov Chains

- **Absorbing Markov Chain**: contains at least one absorbing state that can be reached from every other state (not necessarily in one step)
- **Irreducible Markov Chain**: all states communicate with each other
- **Ergodic Markov Chain**: an aperiodic, irreducible and positive recurrent Markov chain

Periodicity

- Periodicity has an important role in the long-term behavior of a Markov chain
- The **period** of a state i is the largest integer d_i such that $P_{ii}^{(n)} = 0$ whenever n is not divisible by d_i
  - If $d_i > 1$, then i is **periodic**
  - If $d_i = 1$, then i is **aperiodic**
- If $i \leftrightarrow j$, then $d_i = d_j$. Hence, all states of an irreducible Markov chain have the same period.
- Two integers are **co-prime** if their greatest common divisor (gcd) is 1
- If we can find co-prime l and m such that $P_{ii}^{(l)} > 0$ and $P_{ii}^{(m)} > 0$, then i is aperiodic
- Since 1 is co-prime to every integer, any state i with a self-transition is aperiodic

Periodicity

- A matrix P is **non-negative** if all $P_{ij} \geq 0$
- A matrix P is **stochastic** if its rows sum to 1, i.e., $\sum_j P_{ij} = 1$ for all i
- A non-negative matrix P is **quasi-positive** if there exists a natural number $m \geq 1$ such that all entries of P^m are strictly positive
- If P is a stochastic matrix and is quasi-positive, i.e., all entries of P^m are positive, then for all $n \geq m$ all entries of P^n are positive
- **Aperiodicity Lemma**: A stochastic transition matrix P is irreducible and aperiodic if and only if P is **quasi-positive**.
- A finite Markov chain with transition matrix P is ergodic if and only if P is **quasi-positive**

Stationary and Limiting Distributions

- **Stationary distribution**: a vector $\mathbf{w} \in \{\mathbf{p} \in [0,1]^N \mid \mathbf{1}^\top \mathbf{p} = 1\}$ such that $\mathbf{w}^\top P = \mathbf{w}^\top$
- **Limiting distribution**: a vector $\mathbf{w} \in \{\mathbf{p} \in [0,1]^N \mid \mathbf{1}^\top \mathbf{p} = 1\}$ such that:

$$\lim_{t \to \infty} \mathbb{P}(x_t = j | x_0 = i) = \mathbf{w}_j$$

- If it exists, the limiting distribution of a Markov chain is stationary
- **Absorbing chains** have limiting distributions with nonzero elements only in absorbing states
- **Ergodic chains** have a unique limiting distribution (Perron-Frobenius Thm)
- **Periodic chains** may not have a limiting distribution; their stationary distribution has $w_j > 0$ only for recurrent states and w_j is the frequency $\frac{v_j^{(n)}}{n+1}$ of being in state j as $n \to \infty$

Example

- Consider a Markov chain with:
  - state space $\mathcal{X} = \{0, 1\}$
  - prior pmf $\mathbf{p}_0 = [\mathbb{P}(x_0 = 0),\ \mathbb{P}(x_0 = 1)]^\top = [\gamma,\ 1-\gamma]^\top$
  - transition matrix with $a, b \in [0, 1]$, $0 < a + b < 2$:

$$P = \begin{bmatrix} 1-a & a \\ b & 1-b \end{bmatrix}$$

- By induction: $P^n = \frac{1}{a+b}\begin{bmatrix} b & a \\ b & a \end{bmatrix} + \frac{(1-a-b)^n}{a+b}\begin{bmatrix} a & -a \\ -b & b \end{bmatrix}$
- Since $-1 < 1 - a - b < 1$: $\lim_{n\to\infty} P^n = \frac{1}{a+b}\begin{bmatrix} b & a \\ b & a \end{bmatrix}$
- Limiting distribution: exists and is not dependent on the initial pmf $\mathbf{p}_0$:

$$\lim_{t\to\infty} \mathbf{p}_t^\top = \lim_{t\to\infty} \mathbf{p}_0^\top P^t = \frac{1}{a+b}\mathbf{p}_0^\top \begin{bmatrix} b & a \\ b & a \end{bmatrix} = \left[\frac{b}{a+b},\ \frac{b}{a+b}\right]$$

Example

- If $a = b = 1$, the transition matrix is $P = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$
- This Markov chain is periodic:

$$x_t = \begin{cases} x_0 & \text{if } t \text{ is even} \\ x_1 & \text{if } t \text{ is odd} \end{cases}$$

- Stationary distribution: $\mathbf{w} = \left[\frac{1}{2},\ \frac{1}{2}\right]$
- Limiting distribution: does not exist. The pmf $\mathbf{p}_t$ does not converge as $t \to \infty$ and depends on $\mathbf{p}_0$

Outline

Absorbing Markov Chains

- Interesting questions:
  - Q1: On average, how many times is the process in state j?
  - Q2: What is the probability that the state will eventually be absorbed?
  - Q3: What is the expected absorption time?
  - Q4: What is the probability of being absorbed by j given that we started in i?

Absorbing Markov Chains

- **Canonical form**: reorder states so that transient come first: $P = \begin{bmatrix} Q & R \\ 0 & I \end{bmatrix}$
- One can show that $P^n = \begin{bmatrix} Q^n & * \\ 0 & I \end{bmatrix}$ and $Q^n \to 0$ as $n \to \infty$

  Proof: If j is transient, then $\rho_{ij} < 1$ and from the 0-1 Law:

$$\infty > \mathbb{E}\left[v_j \mid x_0 = i\right] = \mathbb{E}\left[\sum_{n=0}^{\infty} \mathbb{1}\{x_n = j\} \,\middle|\, x_0 = i\right] = \sum_{n=0}^{\infty} [P^n]_{ij}$$

- **Fundamental matrix**: $Z^A = (I - Q)^{-1} = \sum_{n=0}^{\infty} Q^n$
  - Expected number of times the chain is in state j: $Z^A_{ij} = \mathbb{E}\left[v_j \mid x_0 = i\right]$
  - Expected absorption time when starting from state i: $\sum_j Z^A_{ij}$
- **Absorption probability**: let B_{ij} be the the probability of reaching absorbing state j starting from transient state i:

$$B_{ij} = P_{ij} + \sum_{k \in \text{Transient}} P_{ik} B_{kj} \quad \Rightarrow \quad B = R + QB \quad \Rightarrow \quad B = Z^A R$$

Example: Drunkard's Walk

- Transition matrix:

$$P = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0.5 & 0 & 0.5 & 0 & 0 \\ 0 & 0.5 & 0 & 0.5 & 0 \\ 0 & 0 & 0.5 & 0 & 0.5 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

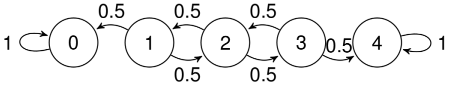

- Canonical form:

$$P = \begin{bmatrix} 0 & 0.5 & 0 & 0.5 & 0 \\ 0.5 & 0 & 0.5 & 0 & 0 \\ 0 & 0.5 & 0 & 0 & 0.5 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

- Fundamental matrix:

$$Z^A = (I - Q)^{-1} = \begin{bmatrix} 1.5 & 1 & 0.5 \\ 1 & 2 & 1 \\ 0.5 & 1 & 1.5 \end{bmatrix}$$

Outline

Markov Chains

Absorbing Markov Chains

Ergodic Markov Chains

General Finite Markov Chain

- A finite Markov chain might have several transient and recurrent classes
- As t increases, the chain is absorbed in one of the recurrent classes
- We can replace each recurrent class with an absorbing state to obtain a chain with only transient and absorbing states
- We can obtain the absorbtion probabilities from $B = Z^A R$
- Each recurrent class can then be analyzed separately

Perron-Frobenius Theorem (Finite Ergodic Markov Chain)

Theorem

Consider an irreducible, aperiodic, finite Markov chain with transition matrix P. Then, the following hold:

- 1 is the eigenvalue of max modulus, i.e., $|\lambda| < 1$ for all other eigenvalues
- 1 is a simple eigenvalue, i.e., the associated eigenspace and left-eigenspace have dimension 1
- The eigenvector associated with 1 is $\mathbf{1}$
- The unique left eigenvector $\mathbf{w}$ is nonnegative and $\lim_{n\to\infty} P^n = \mathbf{1}\mathbf{w}^\top$. Hence, the unique stationary distribution $\mathbf{w}$ is a limiting distribution for the Markov chain, i.e., any initial distribution converges to $\mathbf{w}$.

Perron-Frobenius Theorem (Ergodic Markov Chain)

Theorem

Consider an irreducible, aperiodic, countably infinite Markov chain. Then, one of the following holds.

- All states are transient and $\lim_{t\to\infty} \mathbb{P}(x_t = j | x_0 = i) = 0$, $\forall i, j$.
- All states are null-recurrent and $\lim_{t\to\infty} \mathbb{P}(x_t = j | x_0 = i) = 0$, $\forall i, j$.
- All states are positive-recurrent and there exists a limiting distribution $\mathbf{w}_j = \sum_i \mathbf{w}_i P_{ij}$, $\sum_j \mathbf{w}_j = 1$ such that:

$$\lim_{t\to\infty} \mathbb{P}(x_t = j | x_0 = i) = \mathbf{w}_j > 0.$$

Fundamental Matrix for Ergodic Chains

- We can try to define a fundamental matrix as in the absorbing case but $(I - P)^{-1}$ does not exist because $P\mathbf{1} = \mathbf{1}$ (Perron-Frobenius)
- For absorbing chain, $I + Q + Q^2 + \ldots = (I - Q)^{-1}$ converges because $Q^n \to 0$
- For ergodic chain, $I + (P - \mathbf{1}\mathbf{w}^\top) + (P^2 - \mathbf{1}\mathbf{w}^\top) + \ldots$ converges because $P^n \to \mathbf{1}\mathbf{w}^\top$ (Perron-Frobenius)
- Note that $P\mathbf{1}\mathbf{w}^\top = \mathbf{1}\mathbf{w}^\top$ and $(\mathbf{1}\mathbf{w}^\top)^2 = \mathbf{1}\mathbf{w}^\top\mathbf{1}\mathbf{w}^\top = \mathbf{1}\mathbf{w}^\top$

$$(P - \mathbf{1}\mathbf{w}^\top)^n = \sum_{i=0}^{n}(-1)^i\binom{n}{i}P^{n-i}(\mathbf{1}\mathbf{w}^\top)^i = P^n + \sum_{i=1}^{n}(-1)^i\binom{n}{i}(\mathbf{1}\mathbf{w}^\top)^i$$

$$= P^n + \underbrace{\left[\sum_{i=1}^{n}(-1)^i\binom{n}{i}\right]}_{(1-1)^n-1}(\mathbf{1}\mathbf{w}^\top) = P^n - \mathbf{1}\mathbf{w}^\top$$

- Thus, the following inverse exists:

$$I + \sum_{n=1}^{\infty}(P^n - \mathbf{1}\mathbf{w}^\top) = I + \sum_{n=1}^{\infty}(P - \mathbf{1}\mathbf{w}^\top)^n = (I - P + \mathbf{1}\mathbf{w}^\top)^{-1}$$

Fundamental Matrix for Ergodic Chains

- Consider an ergodic Markov chain with transition matrix P and stationary distribution $\mathbf{w}$
- **Fundamental matrix**: $Z^E := (I - P + \mathbf{1}\mathbf{w}^\top)^{-1}$
  - $\mathbf{w}^\top Z^E = \mathbf{w}^\top$
  - $Z^E \mathbf{1} = \mathbf{1}$
  - $Z^E (I - P) = I - \mathbf{1}\mathbf{w}^\top$
- **Mean first passage time**:
  - $M_{ij} = \mathbb{E}[\tau_j \mid x_0 = i] = \dfrac{Z^E_{jj} - Z^E_{ij}}{w_j}, \; i \neq j$
  - $M_{ii} = \mathbb{E}[\tau_i \mid x_0 = i] = \dfrac{1}{w_i}$

Example: Land of Oz

- Transition matrix:
$$P = \begin{bmatrix} 0.5 & 0.25 & 0.25 \\ 0.5 & 0 & 0.5 \\ 0.25 & 0.25 & 0.5 \end{bmatrix}$$

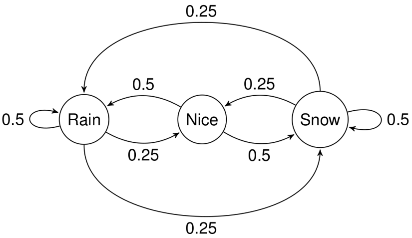

- Stationary distribution:
$$\mathbf{w}^\top = \begin{bmatrix} 0.4 & 0.2 & 0.4 \end{bmatrix}$$

- Fundamental matrix:
$$I - P + \mathbf{1}\mathbf{w}^\top = \begin{bmatrix} 0.9 & -0.05 & 0.15 \\ -0.1 & 1.2 & -0.1 \\ 0.15 & -0.05 & 0.9 \end{bmatrix}$$
$$Z^E = \begin{bmatrix} 1.147 & 0.04 & -0.187 \\ 0.08 & 0.84 & 0.08 \\ -0.187 & 0.04 & 1.147 \end{bmatrix}$$

- Mean first passage time:
$$M_{12} = \frac{Z^E_{22} - Z^E_{12}}{w_2} = \frac{0.84 - 0.04}{0.2} = 4$$