ECE276B: Planning & Learning in Robotics
Lecture 1: Markov Chains

Instructor:

Nikolay Atanasov: natanasov@ucsd.edu

Teaching Assistants:

Zhichao Li: zhl355@eng.ucsd.edu

Jinzhao Li: jil016@eng.ucsd.edu

UC San Diego

JACOBS SCHOOL OF ENGINEERING

Electrical and Computer Engineering

Definition of Markov Chain

- **Stochastic process**: an indexed collection of random variables $\{x_0, x_1, \ldots\}$ whose range is a measurable space $(\mathcal{X}, \mathcal{F})$
  - example: time series of weekly demands for a product
- A temporally homogeneous **Markov chain** is a stochastic process $\{x_0, x_1, \ldots\}$ such that:
  - $x_0 \sim p_0(\cdot)$ for a prior probability density function p on $(\mathcal{X}, \mathcal{F})$
  - $\mathbb{P}(x_{t+1} \in A \mid x_{0:t}) = \mathbb{P}(x_{t+1} \in A \mid x_t) = \int_A p_f(x \mid x_t) dx$ for $A \in \mathcal{F}$ and a conditional pdf $p_f(\cdot \mid x_t)$ on $(\mathcal{X}, \mathcal{F})$
- Intuitive definition:
  - In a Markov Chain the distribution of $x_{t+1} \mid x_{0:t}$ depends only on x_t (a memoryless stochastic process)
  - The state captures all information about the history, i.e., once the state is known, the history may be thrown away
  - "The future is independent of the past given the present" (**Markov Assumption**)

Markov Chain

A **Markov Chain** is a stochastic process defined by a tuple $(\mathcal{X}, p_0, p_f, T)$:

- $\mathcal{X}$ is discrete/continuous set of states
- p_0 is a prior pmf/pdf defined on $\mathcal{X}$
- $p_f(\cdot \mid \mathbf{x})$ is a conditional pmf/pdf defined on $\mathcal{X}$ for given $\mathbf{x} \in \mathcal{X}$ that specifies the stochastic process transitions.
- T is a finite/infinite time horizon

- In a finite-dimensional case, the transition pmf is summarized by a matrix $P_{ij} := \mathbb{P}(x_{t+1} = j \mid x_t = i) = p_f(j \mid x_t = i)$

Example: Student Markov Chain

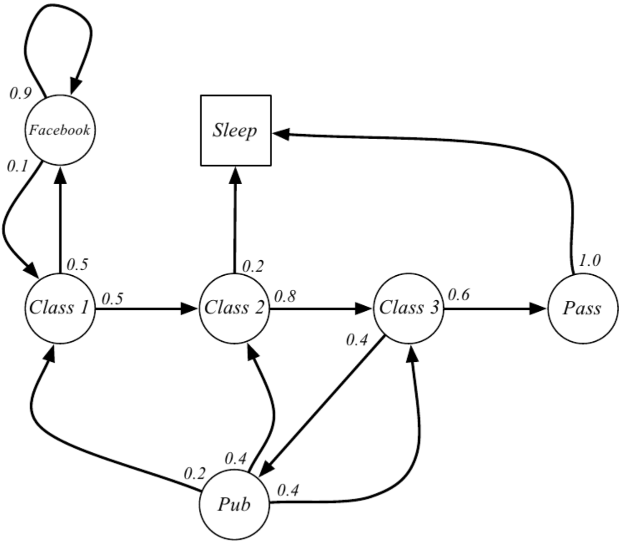

Example: Student Markov Chain

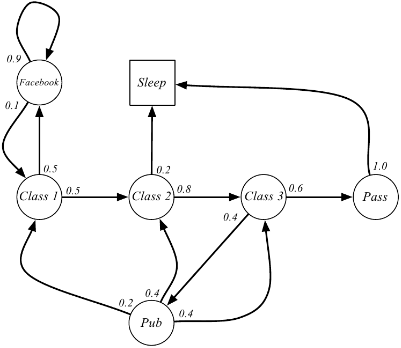

- Sample paths:
  - C1 C2 C3 Pass Sleep
  - C1 FB FB C1 C2 Sleep
  - C1 C2 C3 Pub C2 C3 Pass Sleep
  - C1 FB FB C1 C2 C3 Pub C1 FB FB FB C1 C2 Sleep

- Transition matrix:

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

Chapman-Kolmogorov Equation

- ***n*-step transition probabilities** of a time-homogeneous 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:

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

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

- 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 $\mathbf{p}_{t|t}$ after t steps is:

$$\boxed{\mathbf{p}_{t|t}^\top = \mathbf{p}_0^\top P^t}$$

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 go from $x_0 = i$ to state j for the first time is a random variable

$$\tau_{ij} := \inf\{t \geq 1 \mid x_t = j, x_0 = i\}$$

- **Recurrence Time**: the first passage time τ_{ii} to go 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

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

Proof: By the (strong) 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 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

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

Theorem: Recurrence is contagious

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

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**: it is possible to go from every state to every state (not necessarily in one step)
- **Ergodic Markov Chain**: an aperiodic, irreducible and positive recurrent Markov chain
- **Stationary distribution**: a vector $\mathbf{w} \in \{\mathbf{p} \in [0,1]^N \mid \mathbf{1}^\top \mathbf{p} = 1\}$ such that $\mathbf{w}^\top P = \mathbf{w}^\top$
  - Absorbing chains have stationary distributions with nonzero elements only in absorbing states
  - Ergodic chains have a unique stationary distribution (Perron-Frobenius Theorem)
  - Some periodic chains only satisfy a weaker condition, where $w_j > 0$ only for recurrent states and w_j is the frequency $\frac{v_j^{(n)}}{n+1}$ of being in state j as $n \to \infty$

Absorbing Markov Chains

- Interesting questions:
  - Q1: On average, how mant 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 the states so that the transient ones 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**: $Z^A = (I - Q)^{-1} = \sum_{n=0}^{\infty} Q^n$ exists for an absorbing Markov chain
  - 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: $\sum_j Z^A_{ij}$
  - Let $B = Z^A R$. The probability of reaching absorbing state j starting from state i is B_{ij}

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:

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

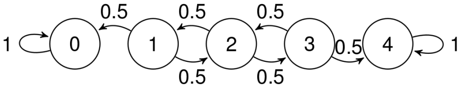

- Fundamental matrix:

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

Perron-Frobenius Theorem

Theorem

Let P be the transition matrix of an irreducible, aperiodic, finite, time-homogeneous Markov chain with stationary distribution $\mathbf{w}$. Then

- 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, $\mathbf{w}$ is the unique stationary distribution for the Markov chain and any initial distribution converges to it.

Fundamental Matrix for Ergodic Chains

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

$$
\begin{aligned}
(P - \mathbf{1}\mathbf{w}^\top)^n &= \sum_{i=0}^{n} (-1)^i \binom{n}{i} P^{n-i} (\mathbf{1}\mathbf{w}^\top)^i = P^n + \sum_{i=1}^{n} (-1)^i \binom{n}{i} (\mathbf{1}\mathbf{w}^\top)^i \\
&= P^n + \underbrace{\left[\sum_{i=1}^{n} (-1)^i \binom{n}{i}\right]}_{(1-1)^n - 1} (\mathbf{1}\mathbf{w}^\top) = P^n - \mathbf{1}\mathbf{w}^\top
\end{aligned}
$$

- 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

- **Fundamental matrix**: $Z^E := (I - P + \mathbf{1}\mathbf{w}^\top)^{-1}$ where P is the transition matrix and $\mathbf{w}$ is the stationary distribution.
- **Properties**: $w^T Z^E = \mathbf{w}^\top$, $Z^E \mathbf{1} = \mathbf{1}$, and $Z^E(I - P) = I - \mathbf{1}\mathbf{w}^\top$
- **Mean first passage time**: $m_{ij} := \mathbb{E}\left[\tau_j \mid x_0 = i\right] = \dfrac{Z^E_{jj} - Z^E_{ij}}{w_j}$

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:

$$w^T = \begin{bmatrix} 0.4 & 0.2 & 0.4 \end{bmatrix}$$

- Fundamental matrix:

$$I - P + \mathbf{1}w^T = \begin{bmatrix} 0.9 & -0.05 & 0.15 \\ -0.1 & 1.2 & -0.1 \\ 0.15 & -0.05 & 0.9 \end{bmatrix}$$

$$Z^E = \begin{bmatrix} 1.147 & 0.04 & -0.187 \\ 0.08 & 0.84 & 0.08 \\ -0.187 & 0.04 & 1.147 \end{bmatrix}$$

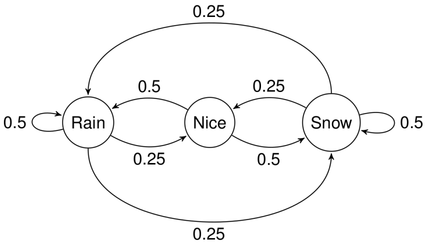

- Mean first passage time:

$$m_{12} = \frac{Z^E_{22} - Z^E_{12}}{w_2} = \frac{0.84 - 0.04}{0.2} = 4$$