11 Markov Chains

- Consider a sequence of discrete random variables $X_0, X_1, X_2, \dots$ where the index represents discrete time. Recall that an infinite sequence of discrete random variables is described by a collection of joint PMFs $P_{X_{t_1},X_{t_2},\dots,X_{t_m}}(x_{t_1}, x_{t_2}, \dots, x_{t_m})$ for every possible choice of indices $t_1 < t_2 < \cdots < t_m$ and any positive integer m.

- If the sequence $X_0, X_1, X_2, \dots$ is independent and identically distributed (i.i.d.), then

$$P_{X_{t_1},X_{t_2},\dots,X_{t_m}}(x_{t_1}, x_{t_2}, \dots, x_{t_m}) = P_X(x_{t_1})\, P_X(x_{t_2}) \cdots P_X(x_{t_m}) \ ,$$

  i.e., there is no dependence across time.

- The sequence $X_0, X_1, X_2, \dots$ has the **Markov property** if, any choice of m and indices $t_1 < t_2 < \cdots < t_m$, X_{t_m} conditioned on $X_{t_{m-1}}, \dots, X_{t_1}$ only depends on the most recent random variable $X_{t_{m-1}}$,

$$P_{X_{t_m}|X_{t_{m-1}},\dots,X_{t_1}}(x_{t_m}|x_{t_{m-1}}, \dots, x_{t_1}) = P_{X_{t_m}|X_{t_{m-1}}}(x_{t_m}|x_{t_{m-1}}) \ ,$$

- A **discrete-time Markov chain** is a sequence of discrete random variables $X_0, X_1, X_2, \dots$ satisfying the Markov property.
  - The random variable X_i is often called the **state** at time i.
  - The Markov property in this context is equivalent to the following statement: The state X_{t+1}, conditioned on the full history $X_0, \dots, X_t$, only depends on the current state X_t,

$$P_{X_{t+1}|X_t,\dots,X_0}(x_{t+1}|x_t, \dots, x_0) = P_{X_{t+1}|X_t}(x_{t+1}|x_t) \ .$$

- We focus on discrete-time Markov chains with the following properties:
  - Finite Range (i.e., Finite State Space): To simplify notation, we always label the range as $R_X = \{1, 2, \dots, K\}$.
  - Homogeneous: The conditional PMF $P_{X_{t+1}|X_t}(x_{t+1}|x_t)$ only depends on the values of x_{t+1} and x_t, not the time index t, and is described by the transition probabilities defined below.

- The **transition probabilities** P_{jk} are the probabilities of moving from state j to state k in one time step for $j, k \in R_X$,

$$\mathbb{P}\big[X_{t+1} = k \big| X_t = j\big] = P_{X_{t+1}|X_t}(k|j) = P_{jk} \quad \text{for all } t.$$

- The **n-step transition probabilities** $P_{jk}(n)$ are the probabilities of moving from state j to state k in exactly n time steps,

$$\mathbb{P}\big[X_{t+n} = k \big| X_t = j\big] = P_{X_{t+n}|X_t}(k|j) = P_{jk}(n) \quad \text{for all } t,$$

  and can be determined via the **Chapman-Kolmogorov equations,**

$$P_{jk}(n+m) = \sum_{i=1}^{K} P_{ji}(n) P_{ik}(m).$$

- It is often more convenient to write out all of the transition probabilities as matrix. Specifically, the **state transition matrix** is

$$\mathbf{P} = \begin{bmatrix} P_{11} & P_{12} & \cdots & P_{1K} \\ P_{21} & P_{22} & \cdots & P_{2K} \\ \vdots & \vdots & \ddots & \vdots \\ P_{K1} & P_{K2} & \cdots & P_{KK} \end{bmatrix}$$

  The row index j is for the current state, and the column index k is for the next state. To satisfy the normalization property, each row must sum to 1. We can also write the **Chapman-Kolmogorov equations in matrix form,** $\mathbf{P}(n+m) = \mathbf{P}(n)\mathbf{P}(m)$ where $\mathbf{P}(n) = \mathbf{P}^n$.

- The **state probability vector at time** t is

$$\underline{p}_t = \begin{bmatrix} P_{X_t}(1) \\ \vdots \\ P_{X_t}(K) \end{bmatrix}$$

  where the j^{th} entry $P_{X_t}(j) = \mathbb{P}[X_t = j]$ is the probability of occupying state j at time t. By normalization, the entries must sum to 1.

- We can determine how the state probabilities change in one time step using either the

$$\text{Transition Probabilities: } P_{X_{t+1}}(k) = \sum_{j=1}^{K} P_{X_t}(j) P_{jk} \quad \text{or,}$$
$$\text{State Transition Matrix: } \underline{p}_{t+1} = \mathbf{P}^{\mathsf{T}} \underline{p}_t \ ,$$

- We can determine how the state probabilities change in n time steps using either the

$$n\text{-Step Transition Probabilities: } P_{X_{t+n}}(k) = \sum_{j=1}^{K} P_{X_t}(j) P_{jk}(n) \quad \text{or,}$$
$$\text{State Transition Matrix: } \underline{p}_{t+n} = (\mathbf{P}(n))^{\mathsf{T}} \underline{p}_t \ .$$

11.1 State Classification

- **State classification** is a systematic way to classify Markov chains, and is useful for determining which Markov chains have certain properties.

- State k is **accessible** from state j if it is possible to reach state k starting from state j in one or more time steps, $P_{jk}(n) > 0$ for some $n \geq 0$. Notation: $j \to k$
  - $P_{jk}(0)$ is the probability of going from state j to state k in exactly 0 time steps. Thus, $P_{jk}(0) = \begin{cases} 1 & j = k \\ 0 & j \neq k \end{cases}$ and we always have that j is accessible from itself, $j \to j$.

- States j and k **communicate** if $j \to k$ and $k \to j$. Notation: $j \leftrightarrow k$
  - Since we always have $j \to j$, we also always have that j communicates with itself, $j \leftrightarrow j$.

- A **communicating class** C is a subset of the states such that all states that belong to C communicate with each other. That is, if $j \in C$, then $k \in C$ if and only if $j \leftrightarrow k$.

  - A finite-state Markov chain can always be partitioned into disjoint communicating classes.

- A Markov chain is **irreducible** if all of its states belong to a single communicating class.
- A state j is **transient** if there is a state k such that k is accessible from j but j is not accessible from k, i.e., $j \to k$ but $k \nrightarrow j$.
  - Intuitively, once we reach state k from state j, we can never return to state j.
  - If a state j is not transient, then it is **recurrent.**
  - The states in a communicating classes are either all transient or all recurrent.
  - At least one communicating class is recurrent.
- The **period** d of a state j is the greatest common divisor of the length of all cycles from j back to itself.
  - A state is **aperiodic** if it has period 1.
  - If there are no cycles from a state back to itself, then its period is set to 1 by default.
  - All states in a communicating class have the same period. A communicating class is aperiodic if all of its states have period 1.
  - Shortcut: If a communicating class contains a cycle of length 1, then it is aperiodic.
  - A Markov chain is aperiodic if all its states are aperiodic.

11.2 Limiting State Probability Vector

- Intuitively, if we let a Markov chain run for a long time, we might expect the state probability vector to stabilize.
- Mathematically, are interested in the limit $\lim_{t\to\infty} \underline{p}_t$, when it exists.
- If a finite-state, homogeneous, discrete-time Markov chain is irreducible and aperiodic, then it has a **unique limiting probability state vector** $\underline{\pi} = \lim_{t\to\infty} \underline{p}_t$.
- The limiting state probability vector satisfies the following properties:
  - Normalization: $\displaystyle\sum_{j=1}^{K} \pi_j = 1$
  - Any initial state probability vector $\underline{p}_0$ will converge to $\underline{\pi}$ as $t \to \infty$.
  - **Steady-State Distribution:** $\underline{\pi} = \mathbf{P}^{\mathsf{T}}\underline{\pi}$.
  - $\underline{\pi}$ is an eigenvector of $\mathbf{P}^{\mathsf{T}}$ with eigenvalue 1.
- To solve for $\underline{\pi}$, we use a system of K linear equations obtained from $\underline{\pi} = \mathbf{P}^{\mathsf{T}}\underline{\pi}$ and $\displaystyle\sum_{j=1}^{K} \pi_j = 1$.
- We can also handle Markov chains that have a single recurrent communicating class along with additional transient communicating classes. Specifically, if the Markov chain has only one recurrent communicating class, there is still a unique limiting state probability vector. To calculate it, we set the probabilities of the transient states to 0, and then solve for the remaining values as in the irreducible case.