COMPSCI 240: Reasoning Under Uncertainty

Nic Herndon and Andrew Lan

University of Massachusetts at Amherst

Spring 2019

Lecture 25: Markov Chains II

Recap: Discrete Markov Chain

- We consider discrete-time Markov chain, in which the state changes at certain discrete time instances, indexed by an integer variable t.
- A discrete Markov chain defines a series of random variables X_t, e.g., $\{X_0, X_1, X_2, \ldots\}$.
- A Markov Chain consists:
    - **State space**: a set of states in which the chain can be described at time t:

    $$\mathcal{S} = \{s_1, \ldots, s_k\}$$

    - **Transition probabilities** that describe the probability of transitioning from a state at $t-1$ to another state at t:

    $$P\left(X_t = s_j | X_{t-1} = s_i\right) = p_{ij} \text{ for all } 1 \leq i, j \leq k$$

    - An initial state X_0, in which the chain is initiated.

Markov Property

- The key assumption is that the transition probabilities (p_{ij}) for the state at time $t+1$ (state j) only depends on the state at time t (state i).
  - The value of X_{t+1} only depends on the value of X_t.
- Mathematically, the **Markov property** defines that

$$P(X_{t+1} = j | X_t = i, X_{t-1} = x_{t-1}, \cdots, X_0 = x_0)$$

$$= P(X_{t+1} = j | X_t = i)$$

$$= p_{ij}$$

- The transition probability p_{ij} must be **non-negative** and **sum to 1**:

$$\sum_{j=1}^{k} p_{ij} = 1, \text{ for all } i.$$

Markov Chain

- Write the **probability distribution** of each X_t as

$$\begin{aligned} v_t &= \langle v_t[1], v_t[2], \ldots, v_t[k] \rangle \\ &= \langle P\left(X_t = s_1\right), P\left(X_t = s_2\right), \ldots, P\left(X_t = s_k\right) \rangle. \end{aligned}$$

- If we happen to know the v_t, then we can compute v_{t+1} using the **Total Probability Law**.

$$P\left(X_{t+1} = s_j\right) = \sum_i P\left(X_{t+1} = s_j | X_t = s_i\right) P\left(X_t = s_i\right).$$

or

$$v_{t+1}[j] = \sum_i v_t[i] p_{ij}$$

States with Transition Probabilities

- Weight p_{ij} on arrow from state i to state j indicates the probability of transitioning to state j given we're in state i.

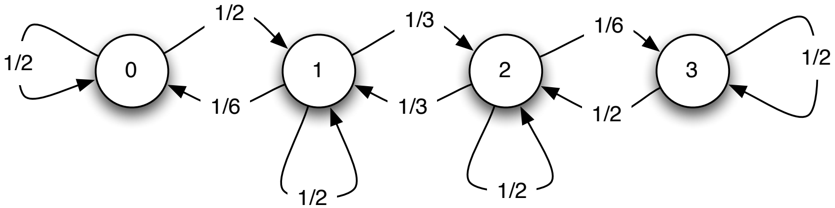

- Can work out things like "what's the probability we're in state 2 after two steps if we're currently in state 3."

Analyzing Markov Chains via Matrices

- Define **Transition probability matrix**:

$$A = \begin{pmatrix} p_{0,0} & p_{0,1} & p_{0,2} & p_{0,3} \\ p_{1,0} & p_{1,1} & p_{1,2} & p_{1,3} \\ p_{2,0} & p_{2,1} & p_{2,2} & p_{2,3} \\ p_{3,0} & p_{3,1} & p_{3,2} & p_{3,3} \end{pmatrix} = \begin{pmatrix} 1/2 & 1/2 & 0 & 0 \\ 1/6 & 1/2 & 1/3 & 0 \\ 0 & 1/3 & 1/2 & 1/6 \\ 0 & 0 & 1/2 & 1/2 \end{pmatrix}$$

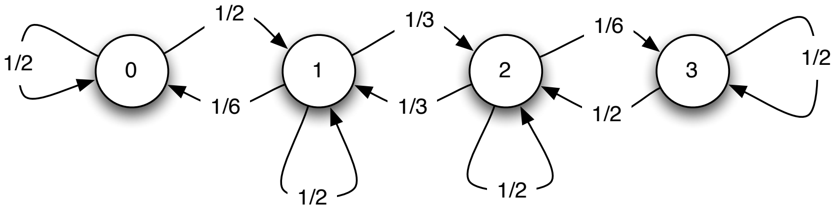

Simulation of the queue if there is initially one person

Given

$$v_t = \left\langle \sum_i p_{i1} v_{t-1}[i], \sum_i p_{i2} v_{t-1}[i], \cdots, \sum_i p_{ik} v_{t-1}[i] \right\rangle .$$

$$\begin{array}{rcl}
v_0 & = & \langle 0.000, 1.000, 0.000, 0.000\rangle \\
v_1 & = & \langle 0.167, 0.500, 0.333, 0.000\rangle \\
v_2 & = & \langle 0.167, 0.444, 0.333, 0.056\rangle \\
v_3 & = & \langle 0.158, 0.416, 0.342, 0.084\rangle \\
v_4 & = & \langle 0.148, 0.401, 0.352, 0.099\rangle \\
v_5 & = & \langle 0.142, 0.391, 0.359, 0.109\rangle \\
v_6 & = & \langle 0.136, 0.386, 0.364, 0.114\rangle \\
v_7 & = & \langle 0.133, 0.382, 0.368, 0.118\rangle \\
v_8 & = & \langle 0.130, 0.380, 0.370, 0.120\rangle \\
\vdots & \vdots & \vdots \\
v_\infty & = & \langle 0.125, 0.375, 0.375, 0.125\rangle
\end{array}$$

Simulation of the queue if there is initially three people

$$
\begin{array}{rcl}
v_0 & = & \langle 0.000, 0.000, 0.000, 1.000\rangle \\
v_1 & = & \langle 0.000, 0.000, 0.500, 0.500\rangle \\
v_2 & = & \langle 0.000, 0.167, 0.500, 0.333\rangle \\
v_3 & = & \langle 0.028, 0.250, 0.472, 0.251\rangle \\
v_4 & = & \langle 0.056, 0.296, 0.404, 0.204\rangle \\
v_5 & = & \langle 0.078, 0.324, 0.423, 0.177\rangle \\
v_6 & = & \langle 0.093, 0.341, 0.407, 0.159\rangle \\
v_7 & = & \langle 0.104, 0.353, 0.397, 0.148\rangle \\
v_8 & = & \langle 0.111, 0.360, 0.389, 0.140\rangle \\
\vdots & \vdots & \vdots \\
v_\infty & = & \langle 0.125, 0.375, 0.375, 0.125\rangle
\end{array}
$$

Steady State Distribution

Do all Markov chains have the property that eventually the distribution settles to the "same steady" state regardless of the initial state?

Definition

We have

$$v = \lim_{t \to \infty} v_t$$

$$\langle v[1], v[2], \ldots, v[k] \rangle = \lim_{t \to \infty} \langle v_t[1], v_t[2], \ldots, v_t[k] \rangle$$

If we have

$$v[j] = \sum_{i=1}^{k} p_{ij} v[i] \text{ for } j = 1, \cdots, k$$

and

$$\sum_{j=1}^{k} v[j] = 1$$

Then, we say v is a **steady state distribution for the Markov Chain.**

Queuing Example

For the queuing example, we had

$$P = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} & 0 & 0 \\ \frac{1}{6} & \frac{1}{2} & \frac{1}{3} & 0 \\ 0 & \frac{1}{3} & \frac{1}{2} & \frac{1}{6} \\ 0 & 0 & \frac{1}{2} & \frac{1}{2} \end{pmatrix}$$

if $\mathbf{v} = (v[1], v[2], v[3], v[4])$ then

$$v[1] = \frac{v[1]}{2} + \frac{v[2]}{6}$$

$$v[2] = \frac{v[1]}{2} + \frac{v[2]}{2} + \frac{v[3]}{3}$$

$$v[3] = \frac{v[2]}{3} + \frac{v[3]}{2} + \frac{v[4]}{2}$$

$$v[4] = \frac{v[3]}{6} + \frac{v[4]}{2}$$

Furthermore,

$$v[1] + v[2] + v[3] + v[4] = 1$$

Solving these gives us

$$\mathbf{v} = (0.125, 0.375, 0.375, 0.125)$$

Question

- "Most" Markov Chains have a unique steady state distribution regardless of initial state that is approached by successive iterations from **any starting distributions**.
- **Question**: Under what circumstances do we have a unique steady state distribution?