CS 70 Spring 2017 — Discrete Mathematics and Probability Theory — Rao — DIS 12a

1 Markov Chain Basics

A Markov chain is a sequence of random variables X_n, $n = 0, 1, 2, \ldots$. Here is one interpretation of a Markov chain: X_n is the state of a particle at time n. At each time step, the particle can jump to another state. Formally, a Markov chain satisfies the Markov property:

$$\Pr(X_{n+1} = j \mid X_n = i, X_{n-1} = i_{n-1}, \ldots, X_0 = i_0) = \Pr(X_{n+1} = j \mid X_n = i), \tag{1}$$

for all n, and for all sequences of states $i_0, \ldots, i_{n-1}, i, j$. In other words, the Markov chain does not have any memory; the transition probability only depends on the current state, and not the history of states that have been visited in the past.

(a) In lecture, we learned that we can specify Markov chains by providing three ingredients: $\mathscr{X}$, P, and π_0. What do these represent, and what properties must they satisfy?

(b) If we specify $\mathscr{X}$, P, and π_0, we are implicitly defining a sequence of random variables X_n, $n = 0, 1, 2, \ldots$, that satisfies (1). Explain why this is true.

(c) Calculate $\Pr(X_1 = j)$ in terms of π_0 and P. Then, express your answer in matrix notation. What is the formula for $\Pr(X_n = j)$ in matrix form?

2 Markov Conversation

Alice is hosting a party. As she's talking to her guests, she notices that conversations naturally transition between casual and more interesting topics. Consider the following simple model of conversations: Each type of topic takes a certain amount of time, and can transition to different topics as specified.

1. A Topics: These take 5 minutes. At the end, they can transition into a B topic (w.p. 25%), or the conversation can terminate (w.p. 75%).

2. *B* Topics: These take 16 minutes. At the end, they are always followed by an *A* Topic.

The following diagram illustrates the conversation flow, where "∅" means the conversation has terminated, and "*A*", "*B*" correspond to the conversation topics.

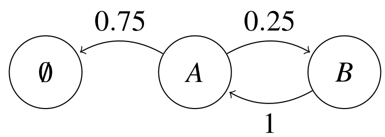

Using the above model:

(a) What is the expected length of a conversation that starts on an *A* topic?

As the party progresses, Alice revises her model of conversation to include three types of topics:

1. *A* Topics: These take 5 minutes. At the end, they can transition into a *B* topic (w.p. 25%), or the conversation can terminate (w.p. 75%).
2. *B* Topics: These take 15 minutes. At the end, they can be followed by a *C* topic (w.p. 50%), or can go back to an *A* topic (w.p. 50%).
3. *C* Topics: These take 25 minutes. They are always followed by a *B* topic.

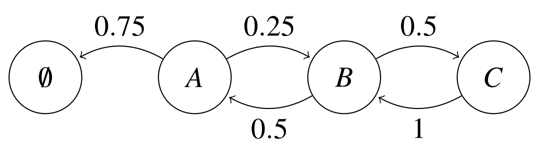

Alice starts to wonder how the expected length of her conversations depend on who she talks to. Assume the following model: If she talks to acquaintances, they start on an *A* topic. With close friends, they start on a *B* topic (w.p. 50%), or on a *C* topic (w.p. 50%). Using the revised model:

(b) Alice starts talking to her acquaintance Bob. What is the expected length of their conversation?

(c) Alice starts talking to her close friend Charlie. What is the expected length of their conversation?

(d) Assume people at the party are equally likely to be close friends or acquaintances. But Eve noticed that Alice and Dave talked for 45 minutes (i.e. they reached state ∅ after 45 minutes)! What is the probability that Dave is a close friend of Alice?