CS177, Homework 5
Due Date: Wednesday, May 14th

Reading

Olofsson Section 7.2.3 and your notes from class on estimating stationary distributions.

Summary

In this assignment you will use MATLAB to explore the stationary distributions of a small Markov chain. You should submit your solutions electronically to the EEE dropbox (MATLAB files, plots and a brief writeup with answers to questions). Where numerical answers are required, such as vectors of probabilities, please provide answers to 4 decimal places.

Problem 1: Implementing the power method

Implement a function in MATLAB called `powermethod.m` that uses the power method as described in class to compute the steady-state probabilities for a Markov chain with `M` possible states. Your function should take as inputs the `MxM` state transition matrix `P`, a `1xM` initial probability distribution vector `initial_pi` and a convergence threshold `epsilon`.

Your function should start with the initial vector and repeatedly apply the state transition matrix until the vector converges to a constant. In order to check if the vector is converging, at each iteration, you should should compare the previous value of the vector to the new value. This can be done by summing up the absolute value of the differences between the elements of $\pi^{(t)}$ for the current iteration and the value of $\pi^{(t-1)}$ from the previous iteration. If the sum of differences is less than the parameter `epsilon` you should stop and return the current vector. A typical value for epsilon might be 0.00001.

In order to prevent runaway execution, (e.g. if the matrix doesn't have a steady state distribution), your code should also stop if the number of iterations exceeds some constant number `MAX_ITERATIONS=10000`

Problem 2: Testing the power method

Let P be the transition matrix for a four state Markov chain where P is given by

$$\begin{array}{cccc} 0 & 1 & 0 & 0 \\ 0.5 & 0 & 0 & 0.5 \\ 0 & 0.6 & 0 & 0.4 \\ 0 & 0.6 & 0.4 & 0 \end{array}$$

1. Use your code to calculate the steady state probabilities for the Markov transition matrix P above. Initialize your algorithm with the uniform vector $\pi^{(0)} = [0.25, 0.25, 0.25, 0.25]$. What steady state probabilities does the vector converge to?

2. Modify your code to plot out the estimate of the steady state probability for state 2, $\pi_2^{(t)}$ as a function of the iteration. Your graph should display the value of $\pi_2^{(t)}$ on the x-axis as a function of $t = 1, \ldots, 50$ for the y-axis.

3. Generate the plot of $\pi_2^{(t)}$ for two different settings of the initial vector $\pi^{(0)} = [1, 0, 0, 0]$ and $\pi^{(0)} = [0, 0, 0, 1]$. Ideally these should be displayed on the same axis using the `hold on` command. Comment briefly on the results and save the resulting plot as a graphic file to submit with your assignment.

4. Run your code using the transition matrix

$$\begin{array}{cccc} 0.2 & 0 & 0 & 0.8 \\ 0 & 0.9 & 0.1 & 0 \\ 0 & 0.6 & 0.4 & 0 \\ 0.7 & 0 & 0 & 0.3 \end{array}$$

   starting with the initial vectors $\pi^{(0)} = [1, 0, 0, 0]$ and $\pi^{(0)} = [0.5, 0.5, 0, 0]$. Do these two different initializations converge to the same stationary distribution? If not, then why not?

Problem 3: Simulating a Markov Chain

Write a MATLAB function called `mc_simulate.m` which simulates a Markov chain. It should take as input the state transition matrix P the initial state pmf vector p_0 and a number of time steps T. It should return a sequence of states of length T sampled from the given Markov chain e.g. if $T = 12$ then it should return an array of integers like $[3, 3, 3, 1, 1, 2, 1, 3, 1, 2, 2, 3]$. Your function should choose an initial state based on the probability vector p_0. For each additional time step, it should use the appropriate row of P containing the conditional probabilities in order to select the next state conditioned on the current state.

Write a second MATLAB function called `compute_frequencies.m` which takes as input a sequence of states and returns the relative frequency with which each state occurs in the resulting sequence. For the example sequence above, your function would return $[0.3333, 0.2500, 0.4167]$ since the relative frequency of state 1 is $(4/12) = 0.3333$ etc.

Problem 4: Testing the simulator

1. Use your simulator to simulate a sequence of length $T = 20$ from the same state transition matrix used in problem 2.1. For the initial distribution use the uniform distribution $p_0 = [0.25, 0.25, 0.25, 0.25]$. What are the relative frequencies of the different states in this sequence?

2. Use your simulator to simulate sequences of length $T = 20, 50, 100, 1000, 10000, 100000$. What is the empirical frequency of state 2 for each of these sequences? Comment on the similarities and differences between this result and the power method from problem 2.