Problem Set #5

Due 19 May 2024

Note, you are submitting all your problems as Jupyter notebooks, so make your code, figures and text readable in a single file.

1) Download the `fit_spectra.ipynb` file, using the pulse example in class, find the best fit.

2) Use a Monte Carlo methods to compare your estimates for the volume of an N-dimensional unit ball to the exact formula in a plot (go from d = 1 to 7, and use N = 10000 samples). The exact formula is

$$V_d = \frac{\pi^{d/2}}{\Gamma(\frac{d}{2}+1)}$$

The Gamma function is available in `scipy.special.gamma`. Hint, compare the volume of a N-dimensional square to a N-dimensional sphere.

3) Consider the following infinite Markov chain.

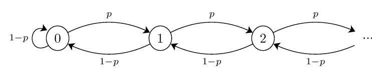

(a) Starting with a value of $p = 0.5$, show a histogram of the number of steps before returning to state 0 (assume once state 0 is reaching the simulation stops, using $N = 10000$ to start). Why are there no odd number of steps?
(b) How does this change is $p = 0.8$ (i.e., easier to go forward than back)?
(c) Returning to $p = 0.5$ and allowing state 0 to have the probabilities shown above, after $N = 10000$ steps, what is the distribution of states? How does this compare to the analytic expectation?

4) A typical problem arises when doctors are experimenting with two different treatments, without knowing which one is better. Call the treatments A and B. One of them is likely to be better, but we don't know which one. A series of patients will each be given one of the treatments. We aim to find a strategy that ensures that as many as possible of the patients are given the better treatment — though we don't know which one this is. Suppose that, for any patient, treatment A has P(success) = α, and treatment B has P(success) = β, and all patients are independent. Assume that $0 < \alpha < 1$ and $0 < \beta < 1$.
Consider two strategies the doctors might use:

- The random strategy for allocating patients to treatments A and B is to choose from the two treatments at random, each with probability 0.5, for each patient.
- The two-armed bandit strategy is more clever. For the first patient, we choose treatment A or B at random (probability 0.5 each). If patient n is given treatment A and it is successful, then we use treatment A again for patient n+1, for all n = 1, 2, 3, ... If A is a failure for patient n, we switch to treatment B for patient n + 1. A similar rule is applied if patient n is given treatment B: if it is successful, we keep B for patient n+1; if it fails, we switch to A for patient n + 1.

  Define the two-armed bandit process to be a Markov chain with state space (A, S), (A, F), (B, S), (B, F), where (A, S) means that patient n is given treatment A and it is successful, and so on.

(a) Draw the network for the Markov chain (A, S), (A, F), (B, S), (B, F) using probabilities α and β.

(b) Write the transition matrix.

(c) For $\alpha = 0.3$, 0.5 and 0.7, find the P(success) as a function of β for the random versus two-armed strategy.