1 Continuous LLSE

Suppose that X and Y are uniformly distributed on the shaded region in the figure below.

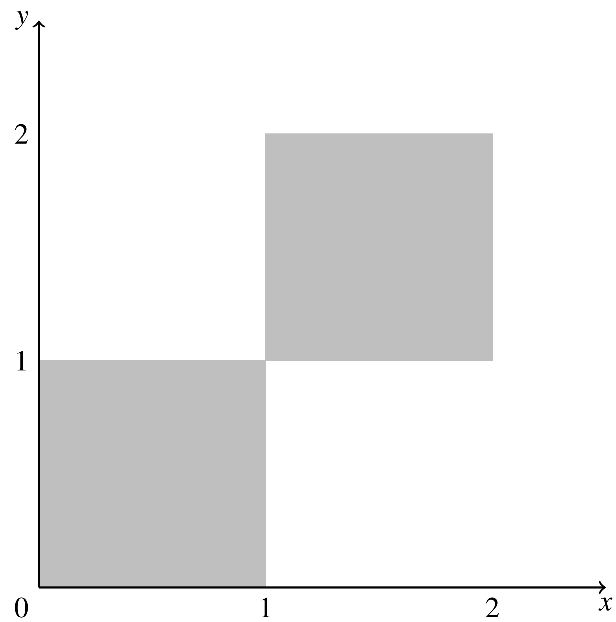

Figure 1: The joint density of (X,Y) is uniform over the shaded region.

That is, X and Y have the joint distribution:

$$f_{X,Y}(x,y) = \begin{cases} 1/2, & 0 \le x \le 1, 0 \le y \le 1 \\ 1/2, & 1 \le x \le 2, 1 \le y \le 2 \end{cases}$$

(a) Do you expect X and Y to be positively correlated, negatively correlated, or neither?

(b) Compute the marginal distribution of X.

(c) Compute $L[Y \mid X]$.

(d) What is $\mathbb{E}[Y \mid X]$?

2 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:

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

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 (**??**). Explain why this is true.

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

3 The Dwinelle Labyrinth

You have decided to take a humanities class this semester, a French class to be specific. Instead of a final exam, your professor has issued a final paper. You must turn in this paper *before* noon to the professor's office on floor 3 in Dwinelle, and it's currently 11:48 a.m.

Let Dwinelle be modeled by the following Markov chain. Instead of rushing to turn it in, we will spend valuable time computing whether or not we *could have* made it. Suppose walking between floors takes 1 minute.

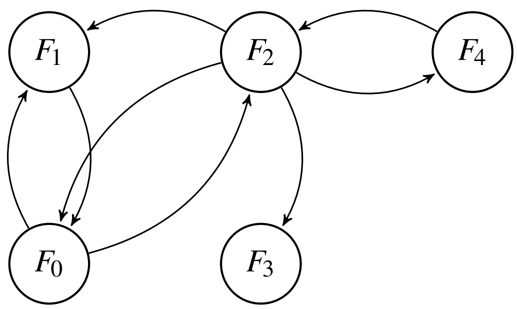

(a) Will you make it in time if you choose a floor to transition to uniformly at random? (If T_i is the number of steps needed to get to F_3 starting from F_i, where $i \in \{0,1,2,3,4\}$, is $\mathbb{E}[T_0] < 12$?)

(b) Will you make it in time, if for every floor, you order all accessible floors and are twice as likely to take higher floors? (If you are considering 1, 2, or 3, you will take each with probabilities $1/7$, $2/7$, $4/7$, respectively.)