STATS 217: Introduction to Stochastic Processes I

Lecture 8

Random variables and stochastic processes

- A **random variable** is a function

$$X : \Omega \to \mathbb{R},$$

where Ω is a probability space (think of this as the space of outcomes of a random experiment).

- A **stochastic process** is a collection of random variables

$$(X_t)_{t \in \mathcal{T}}.$$

- The most common choices for us will be

$$\begin{aligned} \mathcal{T} &= \mathbb{Z}^{\geq 0} = \{0, 1, 2, \ldots, \}, \\ \mathcal{T} &= \mathbb{Z}, \\ \mathcal{T} &= \mathbb{R}. \end{aligned}$$

Markov chains

- A **discrete time Markov chain (DTMC)** is a stochastic process $(X_t)_{t\in\mathbb{Z}^{\geq 0}}$ satisfying the **Markov property**

$$\mathbb{P}[X_{n+1} = x_{n+1} \mid X_n = x_n, \dots, X_0 = x_0] = \mathbb{P}[X_{n+1} = x_{n+1} \mid X_n = x_n]$$

  for all $n \geq 1$ and $x_0, \dots, x_{n+1}$.
- In other words, conditioned on the present, the future is independent of the past.
- A DTMC is **time homogeneous** if

$$\mathbb{P}[X_{n+1} = j \mid X_n = i] = \mathbb{P}[X_{m+1} = j \mid X_m = i]$$

  for all i, j and all times n, m.
- From now on, unless specified otherwise, a DTMC is assumed to be time homogeneous.

Transition matrix

A DTMC is completely specified by the following pieces of information.

- The **state space** S, which is the collection of all possible values that $X_0, X_1, \ldots,$ could take.
- The **initial state** X_0.
- The **transition probabilities**

$$p_{ij} := \mathbb{P}[X_{n+1} = j \mid X_n = i] \quad \forall i, j \in S.$$

- By time homogeneity, the right hand side depends only on i, j and not on n.
- It will be useful to combine all of the transition probabilities into an $|S| \times |S|$ **transition matrix** P,

$$(P)_{ij} := p_{ij}.$$

- Note that P is **row-stochastic** i.e.

$$\sum_{j \in S} p_{ij} = 1 \quad \forall i \in S.$$

Examples

Two state Markov chain.

- State space: $S = \{0, 1\}$.
- Transition matrix:

$$\begin{pmatrix} 1-p & p \\ q & 1-q, \end{pmatrix}$$

for some $p, q \in [0, 1]$.

Examples

Symmetric simple random walk on the integers.

- What is the state space?
- What are the transition probabilities?

Examples

Gambler's ruin stopped at -$100 or $200.

- What is the state space?
- What are the transition probabilities?

- A state $i \in S$ is called **absorbing** if $p_{ii} = 1$.
- What are the absorbing states, if any?

Examples

Branching process with $Z_0 = 1$ and offspring distribution ξ.

- What is the state space?
- What are the transition probabilities?
- What are the absorbing states, if any?

Examples

Coupon collector. There are n different types of coupons, say, $\{1, \ldots, n\}$. Each day you get a uniformly random coupon (repetitions allowed). You stop once you've collected all n types of coupons.

Let X_i denote the number of different types of coupons you've collected by the end of day i. You start with $X_0 = 0$ coupons.

- What is the state space?
- What are the transition probabilities?
- Are there any absorbing states?
- On the problem set, you will study the time it takes to collect all n types of coupons.

Examples

Random walk on a graph. Let $G = (V, E)$ be an undirected graph on vertices $V = \{1, \ldots, n\}$ and edges E. We start at the vertex v_0 and at every time, move to a uniformly random neighbor of the current vertex.

Let X_i denote our position at time i.

- What is the state space?
- What are the transition probabilities?
- Are there any absorbing states?

Examples

Simple random walk on the *n*-dimensional hypercube. The n-dimensional hypercube is the undirected graph on $V = \{0,1\}^n$ where $u, v \in V$ are connected by an edge $e \in E$ if and only if u and v differ in exactly one coordinate.

- What is the state space?
- What are the transition probabilities?
- Starting from $(0, 0, \ldots, 0)$, can the random walk hit $(1, 1, \ldots, 1)$ in an even number of steps?

Examples

Lazy random walk on the n-dimensional hypercube. Transition matrix

$$P_{\text{lazy}} = \frac{1}{2} I + \frac{1}{2} P_{\text{simple}},$$

where I is the $2^n \times 2^n$ identity matrix and P_{simple} is the transition matrix of the simple random walk on the n-dimensional hypercube.

- What is this Markov chain doing?
- Starting at $(0, 0, \ldots, 0)$, can the random walk hit $(1, 1, \ldots, 1)$ in an even number of steps?

Examples

The Ehrenfest urn. n balls are distributed among two urns, urn A and urn B. At each time, we select a ball uniformly at random and move it from its current urn to the other urn.

- How can we model this as a Markov chain?

Examples

Polya's urn. We start with a single urn containing a red ball and a white ball. At each time, we select a ball uniformly at random and return it to the urn along with a new ball of the same color.

- How can we model this as a Markov chain?
- Let R_k denote the number of red balls in the urn after k new balls have been added. What are the possible values that R_k can take?
- On the homework, you will find the distribution of R_k.

Examples

Free throws. Consider a basketball player who makes free throws with the following probabilities

1/2 if she missed the last two times

2/3 if she made one of the last two throws

3/4 if she made both of her last two throws.

- Can this be modelled as a Markov chain?
- What is the state space?
- What are the transition probabilities?

Multi-step transition probabilities

- The transition probability p_{ij} tells us the probability of going from i to j in *one* step, i.e.

$$p_{ij} = \mathbb{P}[X_1 = j \mid X_0 = i].$$

- What about the probability of going from i to j in *two* steps i.e. what is

$$p_{ij}^2 := \mathbb{P}[X_2 = j \mid X_0 = i]?$$

- Well, to go from i to j in two steps, we must go from i to some state $k \in \Omega$ in one step and then from k to j in one step.

Multi-step transition probabilities

Using the law of total probability, we have

$$\begin{aligned}
\mathbb{P}[X_2 = j \mid X_0 = i] &= \sum_{k \in \Omega} \mathbb{P}[X_1 = k \wedge X_2 = j \mid X_0 = i] \\
&= \sum_{k \in \Omega} \mathbb{P}[X_1 = k \mid X_0 = i]\mathbb{P}[X_2 = j \mid X_0 = i \wedge X_1 = k] \\
&= \sum_{k \in \Omega} \mathbb{P}[X_1 = k \mid X_0 = i]\mathbb{P}[X_2 = j \mid X_1 = k] \\
&= \sum_{k \in \Omega} p_{ik} p_{kj} \\
&= (P^2)_{ij}.
\end{aligned}$$

Multi-step transition probabilities

- There is nothing special about two steps here and you should check that the same argument gives

$$p_{ij}^n := \mathbb{P}[X_n = j \mid X_0 = i] = (P^n)_{ij} \quad \forall n \geq 1.$$

- Since for any non-negative integers ℓ, m,

$$P^{\ell+m} = P^\ell P^m,$$

we obtain the **Chapman-Kolmogorov equations**

$$p_{ij}^{\ell+m} = \sum_{k \in \Omega} p_{ik}^\ell p_{kj}^m.$$