IEOR 151 – Lecture 19
Markov Processes

1 Definition

A Markov process is a process in which the probability of being in a future state conditioned on the present state and past states is equal to the probability of being in a future state conditioned only on the present state. There are certain key features of Markov processes that can be used to classify different models:

1.1 State Space

The space to which the states of the model belong can be used to classify different models. The state describes the current configuration of the system, and it captures the important aspects of the system. The fact that the future evolution of a Markov process only depends on the current state is important because it means that for these classes of models the state completely characterizes the configuration of the system. In other words, there are no hidden variables that influence the evolution of the system. The notion of hidden states is particularly relevant in physics to quantum mechanics, where the question of whether the stochastic nature of quantum mechanics is due to hidden variables has been deeply studied.

Discrete state spaces are those in which the states are represented by a discrete set of objects. For example, we could describe a doctor in a hospital with a discrete state space in which some states could correspond to

- Meeting Patient
- Performing Procedure
- Idling
- Administration

Depending on the level of analysis, we might have more or less states to describe the system.

Continuous state spaces are those in which the states are represented by a continuous set of objects. An example is position of a ball we have thrown in space. The position of the ball can occupy different continuous values.

Hybrid state spaces are those in which the states are represented by a combination of discrete and continuous sets. For example, consider a bouncing ball. We need to use a discrete state to describe if the ball is in the air or if the ball is bouncing on the ground. Furthermore, we need continuous states to describe the position and velocity of the ball.

1.2 Initial Conditions

To define the evolution of a Markov process, we need to specify an *initial condition*, which represent the states of the system at the "start". The initial condition can be stochastic or deterministic. For instance, we can specify a distribution of states the system "starts" at.

1.3 Time

Two common classes of models are discrete-time and continuous-time models. In most Markov processes, time is a *privileged variable*, meaning that it is interpreted as a clock that keeps track of the duration between events; however, this is not the case in general. For instance, special and general relativity models from physics drop time from this privileged position in a precise way.

Discrete-time models are those in which time increases in discrete increments. Starting from $t = 0$, time advances as $t = t+1$, and actions in this class of models occur at every increment. Because of the discrete nature of time in this class of models, the time variable t is often used as an index for the state-variables. So if, say, the state is $x \in \mathbb{R}^p$, then the value of the state at time t is denoted x_t or $x[t]$. Because of the discrete nature of time, another convention is to use the variables n or k to denote time.

Continuous-time models are those in which time continuously increases. Starting from $t = 0$, time advances as $\frac{d}{dt}t = 1$, and actions in this class of models can occur at any point in time. Because of the continuous nature of time in this class of models, the time variable is used as the input into a function describing the states. So if the state is x, then the value of the state at time t is denoted as $x(t)$.

These are not the only classes of models for temporal evolution. For instance, hybrid-time models are those in which time increases in both discrete and continuous increments. These types of models occur when describing certain semi-autonomous, robotic, or embedded systems.

2 Markov Chains

A Markov chain is a Markov process in which the state space is discrete. These vertices will act as an abstraction for different quantities of the system. For instance, we will use vertices to represent the number of people waiting in the queue for a service system. The time in this class of models can be continuous, discrete, or hybrid. In this class, we will focus on continuous-time Markov chains.

Because of the discrete nature of the state space, the system is represented by a weighted directed graph $G = (V, E)$, where the vertices $v_i \in V$ represent states of the system and the edges $e_{ij} \in E$ denotes an edge from v_i going towards v_j. Furthermore, every vertex v_i has a self-loop, meaning an edge $e_{ii} \in E$ for all $v_i \in V$. By convention, these self-loops are not drawn because they have a fundamentally different characteristic.

The edges $e_{ij} \in E$ for $i \neq j$ have strictly positive weights w_{ij}, while the edges $e_{ii} \in E$ have non-positive weights. For a continuous-time Markov chain, the weights will represent transition rates. In particular, if edge e_{ij} has weight w_{ij} then this means that for a *small* increment of time h, we have the following probability of transitioning states for all $v_i, v_j \in V$ (even for $v_i = v_j$):

$$\mathbb{P}[x(t+h) = v_j | x(t) = v_i] \approx \mathbb{1}_{v_j = v_i} + w_{ij} h.$$

3 Poisson Process

A Poisson process is an integer-valued counting process $\{N(t), t \geq 0\}$ (with $N(t) \in \mathbb{N}$) in which the interarrival times are described by an exponential distribution. Recall that an exponential distribution with rate $\lambda > 0$ has distribution function $F(u) = 1 - \exp(\lambda u)$, and its mean is $1/\lambda$. Exponential distributions have an interesting *memoryless* property, meaning that if T has exponential distribution, then

$$\mathbb{P}[T > s + t | T > s] = \mathbb{P}(T > t).$$

What this intuitively means is that if you are waiting for an arrival for s units of time, then the probability of an arrival after t additional units of time does not depend on how long you have been waiting for. This model is realistic for some situations and unrealistic for other situations.

There are a number of important properties of Poisson processes:

1. $N(0) = 0$;
2. The number of arrivals in disjoint intervals are independent;
3. The number of arrivals in any time interval depends only on the length of the interval;
4. The distribution of $N(t)$ is given by a Poisson distribution;
5. Multiple arrivals cannot simultaneously occur;
6. Interarrival times have exponential distribution;
7. Arrivals are distributed uniformly on any interval of time.