EM 509: Stochastic Processes

Class Notes

Stochastic Models

(Lecture 6): Part 2

Dr. R. Palamakumbura

Dr. Rathnamali Palamakumbura, University of Peradeniya, Sri Lanka [rpalam@pdn.ac.lk]

Stochastic Models

- In part 1 of the lecture note we studied Markov chains in discrete time.
- In this part 2 of the lecture notes we will study continuous time Markov chains.
- This means that the chains stays in each state a random time, that is a continuous random variable with a distribution that may depend on the state.
- The state of the chain at time t is denoted X(t), where t ranges over the nonnegative real numbers.

Stochastic Models: Markov Processes

- **Continuous time Markov chains**

- Markov Property:

  If $\forall s, t \geq 0, \;\; i, j, x(u) \in \mathbb{Z}^{+} \;\; s.t. \;\; 0 \leq u < s,$

$$P\{X(t+s) = j \,|\, X(s) = i,\, X(u) = x(u)\} = P\{X(t+s) = j \,|\, X(s) = i\}$$

- We say that the process $\{X(t), t \geq 0\}$ is a continuous-time Markov chain if the process satisfies the above Markov property.

- That is, a continuous-time Markov chain is a stochastic process having the Markovian property that the conditional distribution of the future X(t + s) given the present X(s) and the past X(u), 0 u < s, depends only on the present and is independent of the past.

- In addition to having the Markov property for the jumps, we also want the jumps to be independent of the duration of time that is spent in a specific state.

Stochastic Models: Markov Processes

- **Continuous time Markov chains**

- For example suppose that a continuous-time Markov chain enters state i at some time, say, time 0 and suppose that the process does not leave state i (that is, a transition does not occur) during the next ten minutes. What is the probability that the process will not leave state i during the following five minutes?

  - Since the process is in state i at time 10 , by the Markovian property, the probability that it remains in that state during the interval [10, 15] is the (unconditional) probability that it stays in state i for at least five minutes.

  - That is, if T_i denotes the amount of time that the process stays in state i before making a transition into a different state, then

$$P\{T_i > 15 \mid T_i > 10\} = P\{T_i > 5\} .$$

  - In general, by the same reasoning, $P\{T_i > s + t \mid T_i > s\} = P\{T_i > t\} \forall s, t \geq 0.$ Hence, the random variable T_i is memoryless.

Stochastic Models: Markov Processes

- **Continuous time Markov chains**

- In order to achieve this, the exponential distribution is used since it is the only continuous distribution that would ensure this property.

  (Read Chapter 5 of Introduction to probability models by Ross, which I have already uploaded for details.)

- That is a continuous-time Markov chain is a stochastic process:

  - that moves from state to state in accordance with a discrete-time Markov chain. That is when the state i is left, a new state $j \neq i$ is chosen according to the transition probabilities of a discrete time Markov chain.

Stochastic Models: Markov Processes

- **Continuous time Markov chains**

  - the amount of time it spends in each state, before proceeding to the next state, is exponentially distributed. That is if the current state is i, the time until the state is changed has an exponential distribution with parameter say $\lambda(i)$.

    and

  - the amount of time the process spends in state i, and the next state visited, must be independent random variables since otherwise the information as to how long the process has already been in state i would be relevant to the prediction of the next state contradicting the Markov property.

Stochastic Models: Markov Processes

- **Continuous time Markov chains**

- Thus, a continuous-time Markov chain $\{X(t)\}$ is composed of a:

  1. discrete-time Markov chain $\{X_n\}$, the jump chain for the transitions and

  2. exponential random variables for the holding times .

- The $\lambda(i)'$s are called the holding-time parameters .

- Note that the state space is still finite or countably infinite. The discrete/continuous distinction refers to how the time is measured.

- Sometimes the term, Markov process is used in continuous time only and for discrete time the term Markov chain is used. But here we will use Markov process or Markov chain for both cases.

Stochastic Models: Markov Processes

- **Continuous time Markov chains**

- If, $P\{X(t+s)=j \,|\, X(s)=i\}$ is independent of s, then the continuous time Markov chain is said to have time homogeneous transition probabilities.

- That is the probability $P\{X(t+s)=j \,|\, X(s)=i\}$ depends only on time through the difference $(s+t)-s=t$.

- All Markov chains considered here will be assumed to have time homogeneous transition probabilities.

- We can define the transition probabilities, the probability that the chain is in state j, t time units after having been in state i as

$$p_{ij}(t) = P(X(t)=j \,|\, X(0)=i).$$

- For each t, we then get a transition matrix $P(t)$ with entries $p_{ij}(t), i, j \in \mathcal{S}$, which has the following properties.

Stochastic Models: Markov Processes

- **Continuous time Markov chain**

- Let $P(t)$ be the transition matrix for a continuous-time Markov chain with state space $\mathcal{S}$. Then

1. At $t = 0, \; P(0) = I$.

2. For all $i \in \mathcal{S}, \; t \geq 0, \; \sum_{j \in \mathcal{S}} p_{ij} = 1$

3. Chapman–Kolmogorov equations: $P(s+t) = P(s)P(t)$

With k an intermediate state at time s,

$$p_{ij}(s+t) = \sum_{k \in \mathcal{S}} P_i(X(s+t) = j \mid X(s) = k) P_i(X(s) = k \mid X(0) = i)$$

$$= \sum_{k \in \mathcal{S}} p_{ik}(s) p_{kj}(t)$$

which is the $(i,j)^{th}$ entry in the matrix $P(s)P(t)$.

Stochastic Models: Markov Processes

- **Continuous time Markov chains**

Note:

- Usually $P(t)$ is difficult or impossible to compute.
- In the discrete case it was convenient since $P^{(n)} = P^n$, so all the information needed is contained in the one-step transition matrix P.
- In the continuous case there is no analog of "one step," so we need to proceed differently in search of a more compact description.
- Let the jump chain have transition probabilities $p_{ij},\ i \neq j$ and consider the chain in a state i .
- The holding time is exp(λ(i)) and when it leaves, the chain jumps to state j with probability p_{ij}.

Stochastic Models

- **Continuous time Markov chains**

  - Now, if we consider the chain, only when it is in state i and disregard everything else, we can view the jumps from i as a **Poisson process** with rate $\lambda(i)$.

    (Read Chapter 5 of Introduction to probability models by Ross, which I have already uploaded for more details.)

  - For any other state j , the jumps from i to j is then a **thinned Poisson process with rate** $\lambda(i)p_{ij}$.

  - Thus, for any pair of states i and j , we can define the transition rate between i and j as $\gamma_{ij} = \lambda(i)p_{ij}$.

  - Let $\gamma_{ii} = -\sum_{i \neq j} \gamma_{ij}$ and define the generator as the matrix G, **with** $(i, j)^{th}$ **entry is** γ_{ij} .

Stochastic Models: Markov Processes

Continuous time Markov chains

- Note that once the γ_{ij} have been inserted, the diagonal elements γ_{ii} are chosen such that G has row sums equal to 0.

- The generator completely describes the Markov chain, since if we are given G, we can retrieve the holding-time parameters as

$$\lambda(i) = -\gamma_{ii},\ i \in \mathcal{S}$$

and the jump probabilities as

$$p_{ij} = -\frac{\gamma_{ij}}{\gamma_{ii}},\ j \neq i\,.$$

- Note that $p_{ii} \forall i \in \mathcal{S}$ since the $p'_{ij}s$ give the probability distribution when the chain leaves a state and there can be no jumps from a state to itself.

Stochastic Models: Markov Processes

- **Continuous time Markov chains**

Examples:

1. An ON/OFF system stays OFF for a time that is exp(λ) and OFF for a time exp(μ) (μ does not denote the mean here). Describe the system as a continuous time Markov chain.

- The holding-time parameters are λ and μ, and the only possible jumps are from 0(OFF) to 1 (ON) and vice versa. Note that $p_{ij} = 1, i \neq j$.

- Thus we have $\gamma_{01} = \lambda, \gamma_{10} = \mu, \gamma_{00} = -\lambda, \gamma_{11} = -\mu$.

- The generator $G = \begin{pmatrix} -\lambda & \lambda \\ \mu & -\mu \end{pmatrix}$.

Stochastic Models: Markov Processes

- **Discrete time Markov chains**

  - We can also describe the system in a graph as follows:

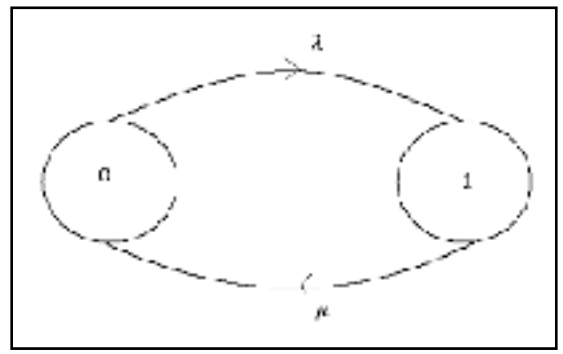

  - This is similar to how we described discrete-time Markov chains but the numbers on the arrows are now rates, not probabilities.

  - The jump chain has transition matrix

$$P = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

Stochastic Models: Markov Processes

- **Discrete time Markov chains**

2. A continuous-time Markov chain on state space {1, 2, 3} has generator

$$G = \begin{pmatrix} -6 & 2 & 4 \\ 1 & -2 & 1 \\ 3 & 1 & -4 \end{pmatrix}.$$

Suppose that the chain is in state 1. What is the expected time until it leaves, and what is the probability that it next jumps to state 2?

- The holding-time parameter in state 1 is $\lambda(1) = -\gamma_{11} = 6$, so the expected holding time is 1/6(mean of the exponential distribution).

- The probability to jump to state 2 is $p_{12} = -\frac{\gamma_{12}}{\gamma_{11}} = -\frac{2}{-6} = \frac{1}{3}$.

Stochastic Models: Markov Processes

- **Discrete time Markov chains**

- The generator now plays the role that the transition matrix did in the discrete case, and a logical question is how G relates to $P(t)$.

- The following gives the relationship, where $P'(t)$ denotes the matrix of the derivatives $p'_{ij}(t)$.

- The transition matrix $P(t)$ and generator G satisfy the backward equations

$$P'(t) = GP(t), t \geq 0$$

  and forward equations $P'(t) = P(t)G, t \geq 0.$

- It is usually difficult to solve the backward and forward equations and only in simple cases can we easily find the explicit form of $P(t)$.

- Since $P(0) = I$, the backward and forward equations also suggest a way to obtain the generator from $P(t)$ according to $G = P'(0)$.

Stochastic Models: Markov Processes

- **Continuous time Markov chains**

- **Stationary Distributions**

  - Consider a continuous time Markov chain with transition matrix $P(t)$.
    A probability distribution π which is such that

    $$\pi P(t) = \pi \ \forall t \geq 0$$

    is called a **stationary distribution** of the chain.

  - The intuition is the same as in the discrete case; the probability π is the proportion of time spent in state j in the long run.

  - Since we have pointed out how difficult it typically is to find $P(t),$ the definition does not give a computational recipe.

  - Instead, first differentiate with respect to t on both sides in the definition.

Stochastic Models: Markov Processes

- **Continuous time Markov chains**
- **Stationary Distributions**
  - Since π does not depend on t we have

$$\frac{d}{dt}(\pi P(t)) = \pi P'(t) = \frac{d}{dt}(\pi) = 0,\ t \geq 0.$$

  - In particular, with $t = 0$, $P'(0) = G$ and we have $\pi G = 0$ and as in the discrete case, we have the additional condition that the entries in π sum to 1.

  - Example: In the ON/OFF system $\pi G = 0 \implies -\pi_0 \lambda + \pi_1 \mu = 0.$

    With this and $\pi_0 + \pi_1 = 1$ gives $(\pi_0\,,\pi_1) = \left(\frac{\mu}{\mu+\lambda}\,,\frac{\lambda}{\mu+\lambda}\right).$

Stochastic Models: Markov Processes

Continuous time Markov chains

- The existence of a stationary distribution is again closely related to the concepts of irreducibility and recurrence.
- Irreducibility is only a property of how the states communicate and has nothing to do with holding times, so we call a continuous time Markov chain irreducible if its jump chain is irreducible.
- As for recurrence and transience, they are defined in the analogous way.

Let $S_i = \inf\{t \mid X(t) = i\}$ and if $S_i = \infty$ then i is never visited.

The only difference from the discrete case is that S_i is now a continuous random variable.

If $P_i(S_i < \infty) = 1$, state i is called recurrent and if $P_i(S_i < \infty) < 1$, state i is called transient .

Stochastic Models: Markov Processes

- **Continuous time Markov chains**

  - Here we use the notation S for "sum," since S_i is the sum of the holding times in all states visited before reaching i.

  - We will keep the notation τ_i ($\tau_i = \textbf{min}\{n \geq 1 \,|\, X_n = i\}$) for the return times in the jump chain.

  - Thus, if the holding time in state k is T_k, then S_i, τ_i are related by

$$S_i = \sum_{n=0}^{\tau_i - 1} T_{X_n}.$$

  - Now suppose that the state i is recurrent ($P(\tau_i < \infty) = 1$) in the jump chain $\{X_n\}$.

  - This means that τ_i is finite, and since also the $T_k' s$ are finite, $S_i' s$ must be finite and the state i is also recurrent in $\{X(t)\}$.

  - Thus, if the jump chain is recurrent, so is the continuous-time chain $\{X(t)\}$.