EE 126 — Probability and Random Processes

March 06, 2018 — DTMCs III, Poisson Process I

1 Reading

W13.4. Poisson Process

A Poisson process, parameterized by the rate λ, is defined as $\{N_t,\ t \geq 0\}$ where

$$N_t = \begin{cases} \max\{n \geq 1 \mid T_n \leq t\} & \text{if } t \geq T_1 \\ 0 & \text{if } t < T_1 \end{cases}$$

$T_n = S_1 + ... + S_n$ where $S_1, ..., S_n$ (the times between jumps) are i.i.d. Exp(λ) random variables. We refer to T_n as the n**th jump time** of the Poisson process.

The Poisson process is memoryless, and N_t has a Poisson distribution with mean λt.

6.1. The Bernoulli Process

A Bernoulli process involves a sequence of independent Bernoulli trials (i.e. coin flips) with fixed parameter p across all flips. The binomial and geometric random variables are associated with the Bernoulli process.

- The Bernoulli process is **memoryless**: if we start at any point in the Bernoulli process, the rest of the sequence can be modeled by a Bernoulli process that is independent of the past.
- If $\bar{T}$ is the time of the first success after time n, then $\bar{T} - n$ has a geometric distribution that is independent of $X_1, ..., X_n$ (the past values).

We can alternatively describe the Bernoulli process in terms of arrivals. We call Y_k the time of the kth arrival, and $T_k = Y_k - Y_{k-1}$ (or $T_1 = Y_1$ for $k = 1$) the kth inter-arrival time. Then $Y_k = T_1 + T_2 + ... + T_k$. (Note: the T_is will be i.i.d. geometric random variables.)

The mean of Y_k is $\mathbb{E}[Y_k] = {}^{k}/_{p}$, the variance is $var(Y_k) = {}^{k(1-p)}/_{p^2}$, and the PMF is $p_{Y_k}(t) = \binom{t-1}{k-1} p^k (1-p)^{t-k}$.

6.2. The Poisson Process

The Poisson process is the continuous version of the Bernoulli process, and should be used when there is no natural way to divide time into discrete intervals (e.g. if we don't know how small to make the intervals).

Let $P(k, \tau) = P(\text{there are exactly } k \text{ arrivals during an interval of length } \tau)$, which should hold for all intervals of length τ. There is also $\lambda > 0$, the arrival rate of the process.

A Poisson process is defined by the following properties:

- **Time homogeneity:** the probability $P(k, \tau)$ of k arrivals is the same for all intervals of the same length τ.
- **Independence:** the number of arrivals during one interval is independent of all arrivals outside of this interval.
- **Small interval probabilities:** the probabilities $P(k, \tau)$ satisfy $P(0, \tau) = 1 - \lambda\tau + o(\tau)$, $P(1, \tau) = \lambda\tau + o_1(\tau)$, and $P(k, \tau) = o_k(\tau)$ (for $k = 2, 3, ...$) where $o(\tau)$ fulfills $\lim_{\tau \to 0} \frac{o(\tau)}{\tau} = 0$ and $o_k(\tau)$ fulfills $\lim_{\tau \to 0} \frac{o_k(\tau)}{\tau} = 0$.

The Poisson and exponential random variables are both associated with the Poisson process. The number N_τ of arrivals can be described by a Poisson random variable, and the time T until the first arrival can be described by an exponential random variable.

Like the Bernoulli process, the Poisson process is memoryless. Also, if $\bar{T}$ is the time of the first arrival after time t, then $(\bar{T} - t) \sim \text{Exp}(\lambda)$ and is independent of the history of the process up to time t.

We can describe the Poisson process as a sequence of independent exponential random variables $T_1, T_2, ...$ with common parameter λ. Then the kth arrival time Y_k will be equal to $T_1 + ... + T_k$. The mean of Y_k will be given by $\mathbb{E}[Y_k] = {}^k/_\lambda$, the variance will be given by $var(Y_k) = {}^k/_{\lambda^2}$, and the PDF is $f_{Y_k}(y) = \frac{\lambda^k y^{k-1} e^{-\lambda y}}{(k-1)!}$ for $y \geq 0$.

Finally, we can take two independent Poisson processes with rates λ_1 and λ_2 and **merge** them by recording an arrival whenever one occurs in either process. This will form a new Poisson process with rate $\lambda_1 + \lambda_2$. We can also take a single Poisson process and **split** it, by keeping each arrival with probability p and discarding it with probability $1 - p$. The result is a Poisson process with rate λp. (Note: analogous results hold for Bernoulli processes.)

6.2.1. Sums of Random Variables

Let $N, X_1, X_2, ...$ be independent random variables, where N takes nonnegative integer values. Let $Y = X_1 + ... + X_N$ for positive values of N, and let $Y = 0$ when $N = 0$.

- If $X_i \sim \text{Bernoulli}(p)$ and $N \sim \text{Binomial}(m, q)$, then $Y \sim \text{Binomial}(m, pq)$.
- If $X_i \sim \text{Bernoulli}(p)$ and $N \sim \text{Poisson}(\lambda)$, then $Y \sim \text{Poisson}(\lambda p)$.
- If $X_i \sim \text{Geometric}(p)$ and $N \sim \text{Geometric}(q)$, then $Y \sim \text{Geometric}(pq)$.
- If $X_i \sim \text{Exp}(\lambda)$ and $N \sim \text{Geometric}(q)$, then $Y \sim \text{Exp}(\lambda q)$.

2 Lecture

Hitting Times

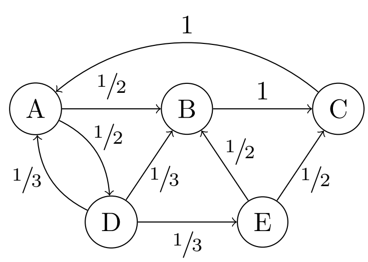

Hitting time: how many steps does it take to reach E for the first time, starting from A? Let $T_E = \min_{n \geq 0}\{X_n = E\}$. Then $\beta_E(A) = \mathbb{E}[T_E \mid X_0 = A]$. To solve for $\beta_E(A)$, we need to couple it with $\beta_E(i)$ (for all other states i) through the first-step equations.

More generally, we'll define $T_A = \min_{n \geq 0}\{X_n \in A\}$ where $A \subset \mathcal{X}$. Then $\beta(i) = \mathbb{E}[T_A \mid X_0 = i]$, and the first-step equations are

- $\beta(i) = 0, \quad i \in A, \;\; \forall i \in \mathcal{X}$
- $\beta(i) = 1 + \sum_{j \in \mathcal{X}} P_{ij} \beta(j), \quad i \notin A$

Note that we abbreviate $\beta_{\text{destination}}(i)$ as $\beta(i)$.

Let us define $P(\text{hitting } C \text{ before } E)$, starting at A, as

$$\alpha(A) = P(T_C < T_E \mid X_0 = A)$$

We again need to calculate $\alpha(i)$ for $i = A, B, C, D, E$.

$$\begin{aligned}
\alpha(A) &= \frac{1}{2}\alpha(B) + \frac{1}{2}\alpha(D) \\
\alpha(B) &= \alpha(C) \\
\alpha(C) &= 1 \\
\alpha(D) &= \frac{1}{3}\alpha(A) + \frac{1}{3}\alpha(B) + \frac{1}{3}\alpha(E) \\
\alpha(E) &= 0
\end{aligned}$$

If we solve this, we will find that $\alpha(A) = 4/5$, $\alpha(B) = 1$, $\alpha(C) = 1$, $\alpha(D) = 3/5$, and $\alpha(E) = 0$.

Accumulating Rewards

Let X_n be a Markov chain on $\mathcal{X}$ with transition probabilities P. Then define $A \subset \mathcal{X}$, $g : \mathcal{X} \mapsto \mathbb{R}$, and T_A the hitting time for A. If

$$\gamma(i) = \mathbb{E}\left[\sum_{n=0}^{T_A} g(X_n) \mid X_0 = i\right]$$

then

$$\gamma(i) = \begin{cases} g(i) & \text{if } i \in A \\ g(i) + \sum_j P_{ij}\gamma(j) & \text{if } i \notin A \end{cases}$$

Example. Flip a fair coin until we get two consecutive heads. What is the expected number of tails? (The number of tails corresponds to our expected payout.) The diagram for the Markov chain looks like this:

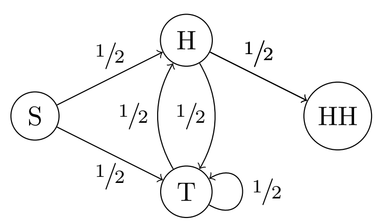

We have $g(S) = 0$, $g(T) = 1$, $g(H) = 0$, and $g(HH) = 0$. So

$$\begin{aligned}
\gamma(S) &= 0 + \frac{1}{2}\gamma(H) + \frac{1}{2}\gamma(T) \\
\gamma(H) &= 0 + \frac{1}{2}\gamma(T) + \frac{1}{2}\gamma(HH) \\
\gamma(T) &= 1 + \frac{1}{2}\gamma(H) + \frac{1}{2}\gamma(T) \\
\gamma(HH) &= 0
\end{aligned}$$

Solving yields $\gamma(S) = 3$. In other words, we expect to see 3 tails.

Reversible Markov Chains

Assume we have an irreducible Markov chain which is started at its invariant distribution π. Suppose that for every n, $(X_0, X_1, ..., X_n)$ has the same distribution as $(X_n, X_{n-1}, ..., X_0)$ (i.e. the time-reversal version). Then we call the chain **reversible.**

"If we take a sequence and play it forward or backward, it should be produced by *this* chain with *this* probability."

Fact: reversible or not, if we start with the Markov chain at π, the "time-reversed" sequence is a Markov chain. Why? Well,

$$\begin{aligned} P(X_k = i \mid X_{k+1} = j, X_{k+2} = i_{k+2}, ..., X_n = i_n) &= \frac{P(X_k = i, X_{k+1} = j, X_{k+2} = i_{k+2}, ..., X_n = i_n)}{P(X_{k+1} = j, X_{k+2} = i_{k+2}, ..., X_n = i_n)} \\ &= \frac{\pi(i) \cdot P_{ij} \cdot P_{j,i_{k+2}} \cdot P_{i_{k+2},i_{k+3}} \cdots P_{i_{n-1},i_n}}{\pi(j) \cdot P_{j,i_{k+2}} P_{i_{k+2},i_{k+3}} \cdots P_{i_{n-1},i_n}} \\ &= \frac{\pi(i) P_{ij}}{\pi(j)} \\ &= \tilde{P}_{ji} \text{ by definition.} \end{aligned}$$

This is a Markov chain since it depends only on i and j. The chain is reversible if $P_{ji} = \tilde{P}_{ji}$. (For a reversed Markov chain, $\tilde{P}_{ji} = P_{ji}$ for all $i, j \in \mathcal{X}$. Both the forward and reverse sequences should be described by the same machinery.)

For a reversible Markov chain, $\pi(j)P_{ji} = \pi(i)P_{ij}$ for all $i, j \in \mathcal{X}$. **This is the condition for reversibility.** (These pairwise "FLOW IN" = "FLOW OUT" equations are known as the *detailed balance equations.*)

Note: reversibility implies invariance.

Poisson Process

The *geometric distribution* (discrete time to success) is to the *exponential distribution* (continuous time to success) as the *Bernoulli coin-flip process* is to the *Poisson process.* The Poisson process is a continuous-time counting process.

To reiterate, the Poisson process is the continuous-time analogue of the "coin-flip" process (where the arrival time is continuous). It's a good model for arrival and departures, customers at a cashier, or photons hitting a detector.

We're just counting arrivals.

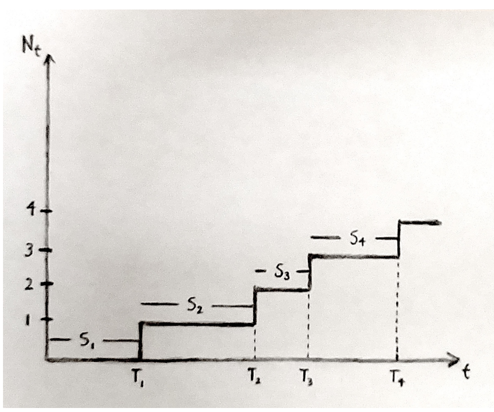

N_t represents how many things have arrived. T_i is the "arrival time," i.e. the absolute time at which an arrival happens. S_i is the "inter-arrival time," i.e. the time between arrivals (represented as i.i.d. Exp(λ) random variables with $\lambda > 0$). Note that $T_1 = S_1$, $T_2 = S_1 + S_2$, and $T_i = \sum_{k=1}^{i} S_k$.

$$N_t = \begin{cases} \max_{n \geq 1}\{n \mid T_n \leq t\} & \text{for } t \geq 0 \\ 0 & \text{for } t < T_1 \end{cases}$$

By defining our model in terms of i.i.d. exponentially random arrivals, we get a process which is memoryless! If we were to translate the origin, we would still have a Poisson process.

References

[1] D.P. Bertsekas and J.N. Tsitsiklis. *Introduction to Probability.* Athena Scientific books. Athena Scientific, 2002.

[2] Jean Walrand. *Probability in Electrical Engineering and Computer Science: An Application-Driven Course.* Quorum Books, Westport, CT, USA, 2014.