Renewal processes

Monday, November 17, 2008
12:05 PM

References: Karlin & Taylor Ch. 5
(also Resnick Ch. 3)

Note correction on Homework problem 3.1c

Point process on the positive real line: $\{\tau_j\}_{j=1}^{\infty}$

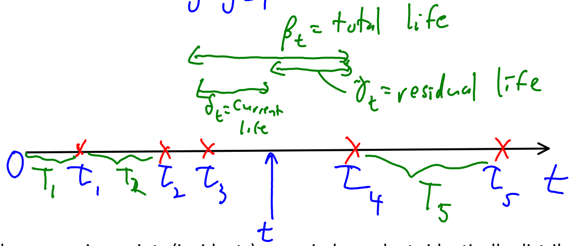

The time interval between the successive points (incidents) τ_j are independent, identically distributed random variables with a common cumulative distribution function (CDF)

$$F(t) = P(T_j \le t)$$

$$T_j = \tau_j - \tau_{j-1}$$

For the special case where $$F(t) = \begin{cases} 1-e^{-\lambda t} & \text{for } t \geq 0 \\ 0 & \text{for } t < 0 \end{cases}$$

This corresponds to an exponential distribution of the time intervals, and therefore this renewal process is just the Poisson point process (with rate λ) and mean time between incidents $= 1/\lambda$.

But sometimes it's useful to consider point processes where the probability distribution for the time between incidents is not exponentially distributed. Renewal processes are designed to allow this extension. We define a counting process:

$$N(t) = \sum_{j=0}^{\infty} I\{t > \tau_j\}$$

= # incidents before time t

This counting process is a Markov process if and only if the time intervals are exponentially distributed (meaning that we are describing a Poisson counting process). However, even though we don't have the Markov property, we do still have the property that at the special moments of time τ_j, the future is independent of the past, that is, we start anew.

Examples of renewal processes:

- Poisson process
- Equipment lifetime, including replacement
- Detectors or neurons with refractory periods
- Following distance in traffic
- Inventory demand and restocking
- Certain special events associated to Markov processes with the strong Markov property:
  - times of successive visits to a state
  - times of successive maxima (if the dynamics have some sort of translation invariance with respect to states)

Fundamental objects of study in renewal processes:

- statistics of the renewal counting function $N(t)$
- total life β_t
- current life δ_t
- residual life γ_t

We will see that one can derive explicit formulas for the properties of these statistics, after which application just requires an easy computation.

The more nontrivial aspect of renewal theory centers around a concept of recursion (renewal argument) and an associated renewal theorem that may need to be applied creatively.

To warm up, let's do some explicit calculations with the Poisson point process, which is a special case.

We know the CDF for the interval times:

$$F_T(t) = \begin{cases} 1-e^{-\lambda t} & \text{for } t \geq 0 \\ 0 & \text{for } t < 0 \end{cases}$$

$N(t)$ has a Poisson distribution with mean λt

By the Markov property underlying Poisson point process, the residual life

γ_t has exponential distribution with mean $1/\lambda$

$$F_{\gamma_t}(t') = P(\gamma_t \leq t') = \begin{cases} 1-e^{-\lambda t'} & \text{for } t \geq 0 \\ 0 & \text{for } t < 0 \end{cases}$$

The other times we have to calculate more carefully. Let's look at the current life.

$$F_{\delta_t}(t') = P(\delta_t \leq t') = 1 - P(\delta_t > t')$$

$$P(\delta_t > t') = P(N(t) - N(t-t') = 0)$$

(no incidents in time interval $[t-t', t]$)

$$= e^{-\lambda t'} \quad \text{provided } 0 \leq t' \leq t$$

$$P(\delta_t > t') = 0 \quad \text{for } t' \geq t$$

$$F_{\delta_t}(t') = \begin{cases} 1-e^{-\lambda t'} & 0 \leq t' \leq t \\ 1 & t' \geq t \\ 0 & t' < 0 \end{cases}$$

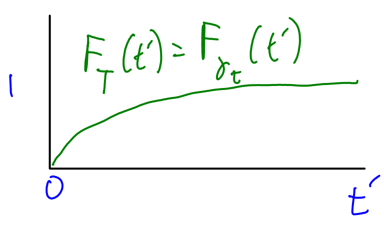

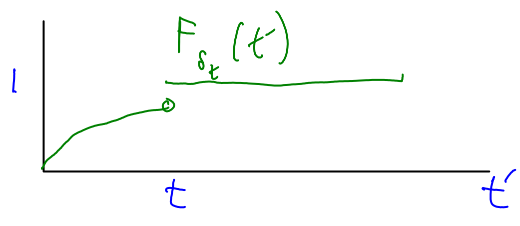

Total life $\beta_t = \gamma_t + \delta_t$

independent for Poisson process

This allows us to compute the probaiblity distribution for the total life explicitly, but we won't pause to do this since we'll have a more general formula later. Suffice to say for large t, the probability distribution for the total life will have the form of a gamma distribution with parameter 2.

But let's just look at the means:

Stieltjes integral

$$\lim_{\|P\| \to 0} \sum_j t_j' \left(F_{\delta_t}(t_{j+1}') - F_{\delta_t}(t_j')\right)$$

Expected residual life $\langle \gamma_t \rangle = \frac{1}{\lambda}$

Expected current life: $\boxed{\langle \delta_t \rangle = \int_0^\infty t' \, dF_{\delta_t}(t')}$

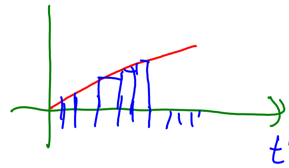

(generalizes the formula $\int_0^\infty t' \rho_{\delta_t}(t') \, dt'$

to cases where the PDF may not be well defined. When the PDF is well-defined then by definition $dF_{\delta_t}(t') = \rho_{\delta_t}(t') \, dt'$)

It's often useful to integrate this formula by parts:

$$\langle \delta_t \rangle = \int_0^\infty \overset{u}{t'} \, \overset{dv}{dF_{\delta_t}(t')} = \overset{u}{t'} \overset{v}{\left(F_{\delta_t}(t') - 1\right)} \Big|_{t'=0}^{\infty}$$

$$- \int_0^\infty \left(F_{\delta_t}(t') - 1\right) dt'$$

$$\boxed{\langle \delta_t \rangle = \int_0^\infty \left(1 - F_{\delta_t}(t')\right) dt'}$$

A generally useful formula for nonnegative random variables.

$$= \int_0^t e^{-\lambda t'} dt' = \frac{1}{\lambda}\left(1 - e^{-\lambda t}\right)$$

$$\to \frac{1}{\lambda} \text{ as } t \to \infty$$

$$\langle \beta_t \rangle = \langle \gamma_t + \delta_t \rangle = \langle \gamma_t \rangle + \langle \delta_t \rangle = \frac{1}{\lambda} + \frac{1}{\lambda}\left(1 - e^{-\lambda t}\right)$$

$$\to \frac{2}{\lambda} \text{ as } t \to \infty$$

Here we see the Poisson paradox: The expected time between two successive incidents has mean $1/\lambda$

But the expected time between the next incident and the previous incident about a given point t has mean $2/\lambda$

The resolution is to understand that the total life arises from sampling the time intervals straddling a given time point, and this biases the statistics of these time intervals to longer time intervals because longer time intervals cover more time! This explains why the mean total life is always greater than or equal to the average time between incidents.

Simple example: $T = \begin{cases} 1 & \text{w/prob } 1/2 \\ 1/2 & \text{w/prob } 1/2 \end{cases}$

The average time interval is $\frac{1}{2}(1) + \frac{1}{2}\left(\frac{1}{2}\right) = 3/4$

If I sample these intervals based on their straddling a given point

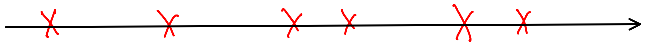

The sampling point is twice is likely to fall within a long interval than a short interval because the long intervals cover twice as much space. The average total life is then:

$$\frac{2}{3}(1) + \frac{1}{3}\left(\frac{1}{2}\right) = 5/6 > 3/4$$

This corresponds to the length bias inherent in sampling the time intervals based on whether they straddle a certain time point.

Now let's turn to developing the more general theory of renewal processes. Most calculations hinge upon a recursion concept (renewal argument) and a fundamental renewal theorem, which we now develop.

Renewal argument (Karlin and Taylor Sec. 5.4A)

- first-step analysis for renewal processes

First consider the mean number of incidents up to time t.

$dp_T(t')dt'$ when make sense

$$M(t) = \mathbb{E}\, N(t)$$

$$\mathbb{E}\, N(t) = \int_0^\infty \mathbb{E}[N(t) \mid T_1 = t']\, dF_T(t')$$

(law of total expectation, conditioning on continuously distributed r.v. T_1)

$$\mathbb{E}[N(t) \mid T_1 = t'] = \begin{cases} 0 & \text{for } t' > t \\ 1 + \mathbb{E}(N(t-t')) & 0 \le t' \le t \end{cases}$$

renewal property

$$M(t) = \int_0^t (1 + M(t-t'))\, dF_T(t')$$

Note this recursive expression for $M(t)$ in terms of its values at earlier times. Since time is continuous, the recursion isn't as straightforward as for discrete case, but similar analytical tools will still be helpful. But first we rewrite this equation in standard form.

$$M(t) = F_T(t) - \underbrace{F_T(0)}_{=0} + \int_0^t M(t-t')\, dF_T(t')$$

$$\boxed{M(t) = F_T(t) + \int_0^t M(t-t')\, dF_T(t')}$$

This is an example of a "renewal equation." You can use similar approach to calculate other moments of the renewal counting process. The general strategy for working with renewal processes is to derive a "renewal equation" for the statistic of interest, and then use the solution and asymptotic procedures which we develop below.

Solving the renewal equation:

Recall that for discrete-time stochastic models with recursive structures, the probability generating function was very useful (Z-transform). Here, for continuous time on the nonnegative real axis, the analogous tool is the Laplace transform.

$$\hat{M}(z) = \int_0^\infty e^{-zt} M(t)\,dt$$

where $z \in \mathbb{C}$ with large enough real part.

$$\hat{M}(z) = \hat{F}_T(z) + \hat{M}(z) \underbrace{z \hat{F}_T(z)}_{\int_0^\infty e^{-zt'} dF_T(t')}$$

$$\hat{M}(z) = \frac{\hat{F}_T(z)}{1 - z\hat{F}_T(z)}$$

This is pretty easy but the hard part is inverting the Laplace transform…sometimes it's easy (lookup in tables, etc.) but generally speaking the inversion formula is

$$M(t) = \frac{1}{2\pi i} \int_{\mathcal{C}} e^{zt}\, \hat{M}(z)\,dz$$

Bromwich contour

This complex contour integral formula is useful when the Laplace transform cannot be exactly inverted because one can express the solution in terms of sums of residues at the singularities of the Laplace transform.

For the purposes of long-time asymptotics, though, it is often helpful to simply refer to the Renewal Theorem which is essentially proved using these kinds of Laplace transform techniques.

To prepare to state this theorem, we need one technicality:

A probability distribution is said to be arithmetic if it is concentrated on a lattice of equally spaced points including the origin. Such probability distributions for interval times introduce some technical concerns -- see the text; we'll focus on nonarithmetic probability distributions.