ELEC 548 Byron's Review of Poisson Processes

A Exponential Distribution

A random variable T is said to be exponentially distributed with rate $\lambda > 0$ if its probability density function (PDF) is

$$f_T(t) = \begin{cases} \lambda e^{-\lambda t} & \text{if} \quad t \geq 0 \\ 0 & \text{if} \quad t \leq 0. \end{cases} \tag{1}$$

As a shorthand, we can write $T \sim \exp(\lambda)$.

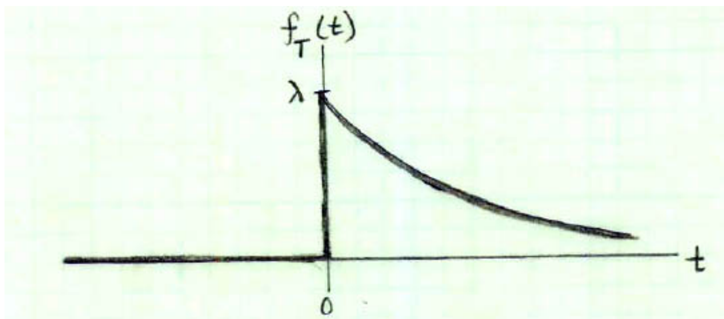

A peak ahead: The time between two consecutive spikes (a.k.a. the inter-spike interval or ISI) can be modeled by an exponential distribution.
Alternatively, we can describe T in terms of its cumulative distribution function (CDF).

$$F_T(t) = \Pr(T \leq t) = \begin{cases} 1 - e^{-\lambda t} & \text{if} \quad t \geq 0 \\ 0 & \text{if} \quad t \leq 0. \end{cases} \tag{2}$$

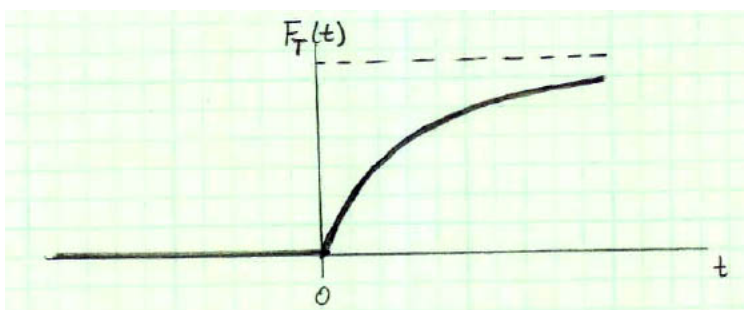

Note that the PDF and CDF (of any random variable) are related in the following way:

$$f_T(t) = \frac{dF_T(t)}{dt} \qquad F_T(t) = \int_{-\infty}^{t} f_T(t)\ dt \tag{3}$$

A.1 Mean and variance of the exponential

$$\begin{aligned} \mathrm{E}[T] &= \int t f_T(t)\,dt \\ &= \int_0^\infty t\lambda e^{-\lambda t}\,dt \end{aligned} \tag{4}$$

Integrating by parts; let $u = t$ and $\mathrm{d}v = \lambda e^{-\lambda t}$, so $\mathrm{d}u = \mathrm{d}t$ and $v = -e^{-\lambda t}$.

$$\begin{aligned}\mathsf{E}[T] &= uv|_0^\infty - \int_0^\infty v du \\ &= -t\, e^{-\lambda t}|_0^\infty + \int_0^\infty e^{-\lambda t}\,\mathrm{d}t \\ &= 0 - 0 + \left[-\frac{1}{\lambda}e^{-\lambda t}\right]_0^\infty \\ &= \frac{1}{\lambda}\end{aligned} \tag{5}$$

$$\begin{aligned}\mathsf{E}[T^2] &= \int t^2 f_T(t)\,\mathrm{d}t \\ &= \int_0^\infty t^2\,\lambda e^{-\lambda t}\,\mathrm{d}t\end{aligned} \tag{6}$$

Integrating by parts; let $u = t^2$ and $\mathrm{d}v = \lambda e^{-\lambda t}$, so $\mathrm{d}u = 2t\,\mathrm{d}t$ and $v = -e^{-\lambda t}$.

$$\begin{aligned}\mathsf{E}[T^2] &= uv|_0^\infty - \int_0^\infty v du \\ &= -t^2\, e^{-\lambda t}|_0^\infty + \int_0^\infty 2te^{-\lambda t}\,\mathrm{d}t \\ &= 0 - 0 + \int_0^\infty 2te^{-\lambda t}\,\mathrm{d}t \\ &= \frac{2}{\lambda}\int_0^\infty \lambda te^{-\lambda t}\,\mathrm{d}t \\ &= \frac{2}{\lambda^2}\end{aligned} \tag{7}$$

and

$$\mathsf{var}(T) = \mathsf{E}(T^2) - (E[T])^2 = \frac{1}{\lambda^2} \tag{8}$$

A.2 Memoryless property of exponential random variables

In words: Say that the waiting time for a bus to arrive is exponentially distributed. If I've been waiting for t seconds, then the probability that I must wait s more seconds is the same as if I hadn't waited at all.

With math:

$$\mathsf{Pr}\,(T > t + s | T > t) = \mathsf{Pr}\,(T > s) \tag{9}$$

To show this,

$$\mathsf{Pr}\,(T > t + s | T > t) = \frac{\mathsf{Pr}\,(T > t + s)}{\mathsf{Pr}\,(T > t)} \tag{10}$$

Intuition: Conditioning the exponential is like starting at a point away from zero on the x axis of the PDF. Turning that new thing into a distribution implies renormalizing, but because of an exponential's shape, that stretching turns it into exactly the same thing as it was before.

B Defining the Poisson process

B.1 Constructing a Poisson process

Let $t_1, t_2, \ldots$ be independent exponential random variables with parameter λ. Let $T_n = t_1 + t_2 + \ldots + t_n$ for $n \geq 1$, where $T_0 = 0$. Define $N(s) = \max\{n : T_n \leq s\}$. $N(s)$ is a Poisson process.

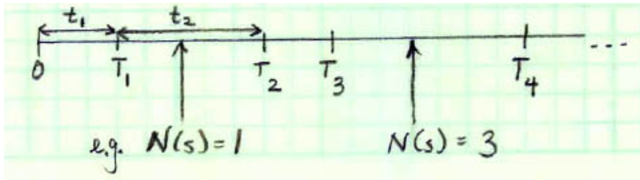

If a Poisson process is used to model a spie train, then:

— t_n is the n^{th} interspike interval (ISI).

— T_n is the time at which the n^{th} spike occurs.

— $N(s)$ is the number of spikes by time s.

B.2 Properties of the Poisson process

Why is $N(s)$ called a Poisson process rather than an exponential process?

Property 1: $N(s)$ has a Poisson distribution with mean λs.

First, recognize that $N(s) = n$ iff $T_n \leq s < T_{n+1}$. In other words, the n^{th} spike occurs *before* time s and the $(n+1)^{\text{th}}$ spike occurs *after* time s.

$$\begin{aligned}
\Pr(N(s) = n) &= \int_0^s \Pr(T_{n+1} > s | T_n = t) f_{T_n}(t)\, dt \\
&= \int_0^s \Pr(t_{n+1} > s - t) f_{T_n}(t)\, dt \\
&= \int_0^s e^{-\lambda(s-t)} f_{T_n}(t)\, dt
\end{aligned}$$

Recall that summing independent random variables implies convolving their PDF's. If we take Fourier transforms of the PDF's, then we can multiply rather than convolve.

$$\begin{aligned}
\mathfrak{F}\{f_{T_n}\} &= \prod_{i=1}^{n} \mathfrak{F}\{f_{T_i}\} \\
&= \left[\mathfrak{F}\{\lambda e^{-\lambda t} u(t)\}\right]^n \\
&= \left[\frac{\lambda}{\lambda + \jmath\omega}\right]^n
\end{aligned}$$

Table of Fourier Transforms
$e^{-at}u(t) \xrightarrow{\mathfrak{F}} \frac{1}{a + \jmath\omega}$
$t^n e^{-at}u(t) \xrightarrow{\mathfrak{F}} \frac{n!}{(a + \jmath\omega)^{n+1}}$

Taking the inverse transforms of both sides,

$$\begin{aligned}
f_{T_n}(t) &= \frac{\lambda^n}{(n-1)!} \cdot t^{n-1} e^{-\lambda t} u(t) \\
&= \lambda e^{-\lambda t} \frac{(\lambda t)^{n-1}}{(n-1)!} \quad \text{for } t \geq 0
\end{aligned} \tag{11}$$

This is called the **Erlang distribution** which is a special case of the **gamma distribution**. This will appear again when we try to model refractory periods.

$$\begin{aligned}
\Pr(N(s) = n) &= \int_0^s e^{-\lambda(s-t)} f_{T_n}(t)\,\mathrm{d}t \\
&= \int_0^s e^{-\lambda(s-t)} \lambda e^{-\lambda t} \frac{(\lambda t)^{n-1}}{(n-1)!}\,\mathrm{d}t \\
&= \frac{\lambda^n}{(n-1)!} e^{-\lambda s} \int_0^s t^{n-1}\,\mathrm{d}t \\
&= \frac{\lambda^n}{(n-1)!} e^{-\lambda s} \left[\frac{t^n}{n}\right]_0^s \\
&= e^{-\lambda s} \frac{(\lambda s)^n}{n!} = \text{Poisson}(\lambda s)
\end{aligned}$$

What does a Poisson distribution look like?

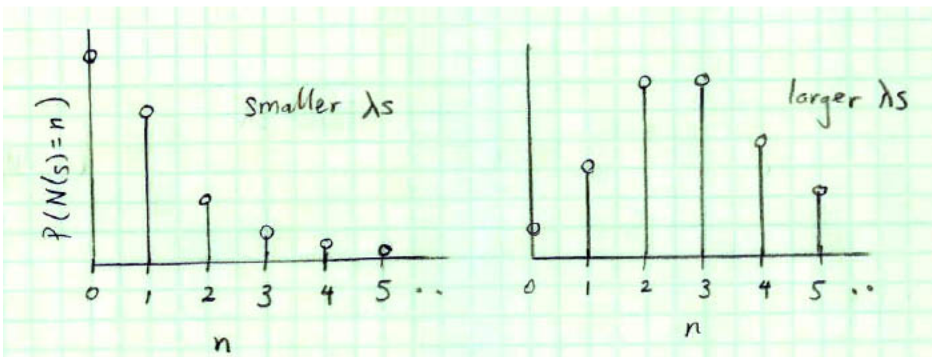

For smaller (λs), the exponential term dominates. For larger (λs), the polynomial term initially dominates.
What are the mean and variance of the Poisson distribution? First the mean:

$$\begin{aligned}
\mathrm{E}[N(s)] &= \sum_{n=0}^{\infty} n \cdot \Pr(N(s) = n) \\
&= \sum_{n=1}^{\infty} n \cdot e^{-\lambda s} \frac{(\lambda s)^n}{n!} \\
&= \lambda s \sum_{n=1}^{\infty} e^{-\lambda s} \frac{(\lambda s)^{n-1}}{(n-1)!} \\
&= \boxed{\lambda s}
\end{aligned} \tag{12}$$

For the variance, we can use a trick.

$$\begin{aligned}
\mathrm{E}[N(s)(N(s)-1)] &= \sum_{n=0}^{\infty} n(n-1) \cdot \Pr(N(s) = n) \\
&= \sum_{n=2}^{\infty} n(n-1) \cdot e^{-\lambda s} \frac{(\lambda s)^n}{n!} \\
&= (\lambda s)^2 \sum_{n=2} \infty e^{-\lambda s} \frac{(\lambda s)^{n-2}}{(n-2)!} \\
&= (\lambda s)^2.
\end{aligned} \tag{13}$$

$$\begin{aligned}\text{var}\left[N(s)\right] &= \text{E}\left[N(s)^2\right] - \left(\text{E}\left[N(s)\right]\right)^2 \\ &= \text{E}\left[N(s)\left(N(s)-1\right)\right] + \text{E}\left[N(s)\right] - \left(\text{E}\left[N(s)\right]\right)^2 \\ &= (\lambda s)^2 + \lambda s - (\lambda s)^2 \\ &= \boxed{\lambda s}\end{aligned} \tag{14}$$

Property 2: $N(t+s) - N(s),\ t \geq 0,\ \sim$ Poisson(λt) and independent of $N(r),\ 0 \leq r \leq s$.

In other words, if you look forward from any time s, that is itself a Poisson process independent of anything that's already happened. This picture provides the intuition:

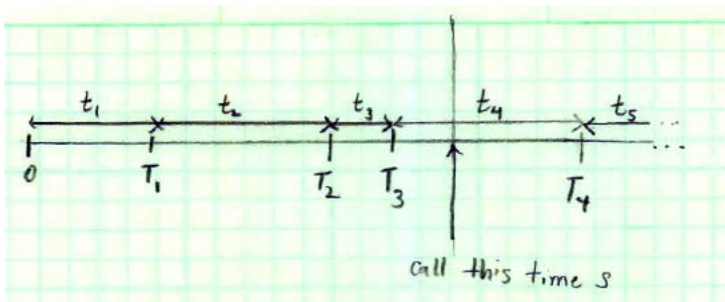

Looking forward from time s, the time until the first spike (at T_4) is distributed as an exponential with parameter λ and independent of anything that came before it, by the memoryless property of the exponential. Subsequent ISI's ($t_5, t_6, \ldots$) are $\sim$ exp(λ) and independent of anything before time s.

Property 3: $N(t)$ has <u>independent increments.</u>

If $s_0 < s_1 < \ldots < s_n$, then

$$N(s_1) - N(s_0), N(s_2) - N(s_1), \ldots, N(s_n) - N(s_{n-1}) \text{ are independent.}$$

In other words, if you take spike counts in non-overlapping windows, the spike counts are independent.

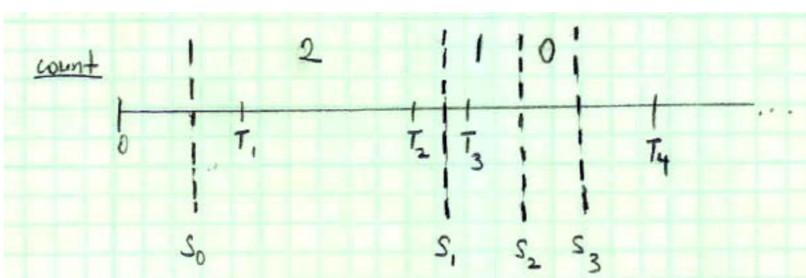

To summarize, if $N(s), s \geq 0$ is a Poisson process, then

(i) $N(0) = 0$

(ii) $N(t+s) - N(s) \sim$ Poisson(λt)

(iii) $N(t)$ has indepdent increments

Conversely if i, ii, and iii hold, then $N(s), s \geq 0$ is a Poisson process.

B.3 Another view of the Poisson process

So far, we have derived the Poisson process using i.i.d. exponential ISI's. Another very useful way of thinking about the Poisson process is using the Bernoulli process. The Poisson process is the continuous-time limit of the Bernoulli process, which is defined in discrete time.

Bernoulli Process

At each time step, flip a coin to decide whether the neuron spikes (1) or not (0). The coin flips are independent of each other.

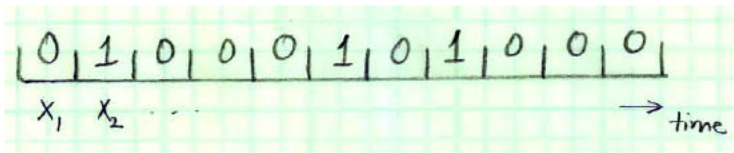

At the i^{th} time step,

$$X_i \sim \text{Bernoulli}(p) \text{ i.i.d.}$$

$$X_i = \begin{cases} 1 & \text{with probability } p \\ 0 & \text{with probability } (1-p) \end{cases}$$

Let S_n be the number of spikes up to and including the n^{th} time step.

$$S_n = \sum_{i=1}^{n} X_i$$

$$S_n \sim \text{Binomial}(n, p)$$

$$\Pr(S_n = k) = \binom{n}{k} p^k (1-p)^{n-k}$$

$$\text{E}[S_n] = np \Rightarrow \text{We expect to see } np \text{ spikes in } n \text{ time steps.}$$

Without proof here, as $n \to \infty$ and $p \to 0$, the Bernoulli process becomes the Poisson process, where $np = \lambda s$. So the Bernoulli process provides an intuitive way to think about the Poisson process.

We can also go in the other direction and consider the probability that a Poisson process gives a spike in a small time window of duration δ. The number of spikes in this window is $\sim$ Poisson$(\lambda\delta)$.

$$\Pr(0 \text{ spikes in } [t, t+\delta]) = e^{-\lambda\delta} = \boxed{1 - \lambda\delta} + \text{O}(\delta^2)$$

$$\Pr(1 \text{ spikes in } [t, t+\delta]) = e^{-\lambda\delta} \cdot \lambda\delta = \boxed{\lambda\delta} - \text{O}(\delta^2)$$

$$\Pr(> 1 \text{ spikes in } [t, t+\delta]) = e^{-\lambda\delta} = \text{O}(\delta^2)$$

If δ is small, O(δ^2) terms $\to 0$. Thus, whether or not a neuron spikes in this small window can be determined with a coin flip, where the probability of a spike is $\lambda\delta$.

B.4 Thinning

Suppose $N(s)$ is a Poisson process with rate λ. Each time a spike occurs, a coin is flipped. If the coin comes up heads (with probability p), the spike is assigned to output stream 1. Else, the spike is assigned to output stream 2. The two output streams are each an independent Poisson process with rates λp and $\lambda(1-p)$, respectively.

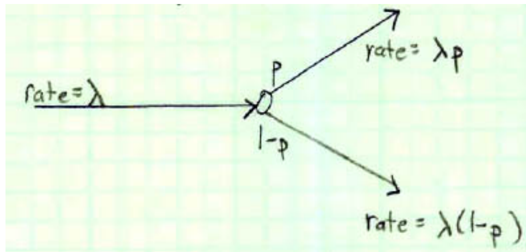

B.5 Superposition

Suppose $N_1(s)$ and $N_2(s)$ are independent Poisson processes with rates λ_1 and λ_2, respectively. Then $N_1(s) + N_2(s)$ is a Poisson process with rate $\lambda_1 + \lambda_2$.

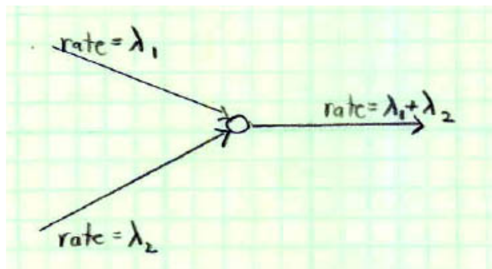

C Inhomogeneous Poisson Processes

So far, we've considered a Poisson process that is homogeneous – it's rate does not change as a function of time. However, the firing rates of neurons typically do change with time. To model the time-dependent activity of neurons, we need a non-stationary process, such as the *inhomogeneous* Poisson process.

Definition $N(s), s \geq 0$ is an inhomogeneous Poisson process with rate $\lambda(r)$ if

(i) $N(0) = 0$

(ii) $N(t+s) - N(s) \sim \text{Poisson}(\int_s^{t+s} \lambda(r)\, dr)$

(iii) $N(t)$ has indepdent increments

Comparing with the previous definition of a homogeneous Poisson process, the only difference is that the Poisson mean is now $\int_s^{t+s} \lambda(r)\, dr$ rather than λt.
Note that if $\lambda(r)$ is flat, then the two definitions are equivalent!

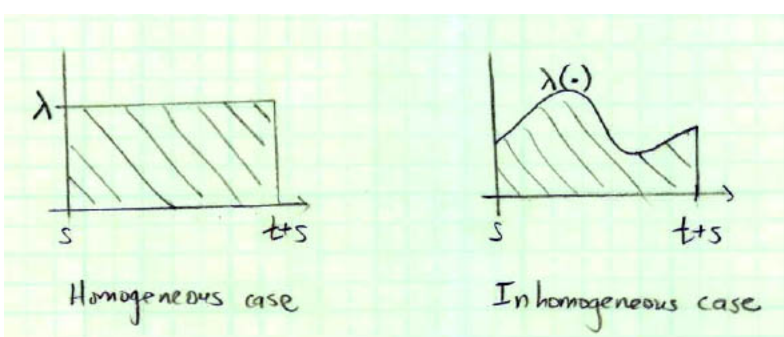

For an inhomogeneous Poisson process, the ISI's are no longer exponentially distributed or independent. Let's show this:

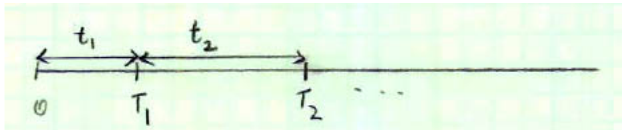

$$\text{Let } \mu(t) = \int_0^t \lambda(r)\, \mathrm{d}r$$

$$\Pr(t_1 > t) = \Pr\left(N(t) = 0\right) = e^{-\int_0^t \lambda(r)\,\mathrm{d}r} = e^{-\mu(t)}$$

$$f_{t_1} = -\frac{\mathrm{d}}{\mathrm{d}t}\Pr(t_1 > t) = \lambda(t)e^{-\mu(t)}, \text{ which is not exponential!}$$

Now, look forward in time from T_1.

$$\begin{aligned}\Pr(t_2 > s | t_1 = t) &= \Pr\left(N(t+s) - N(t) = 0\right)\\ &= e^{-\int_t^{t+s}\lambda(r)\,\mathrm{d}r}\\ &= e^{-(\mu(s+t)-\mu(t))}\end{aligned}$$

$$f_{t_2|t_1}(s) = -\frac{\mathrm{d}}{\mathrm{d}s}\Pr(t_s > s|t_1 = t) = \lambda(s+t)e^{-(\mu(s+t)-\mu(t))}$$

Since t_2 depends on t_1, the ISIs are not independent.
The joint distribution of ISI's is

$$\begin{aligned}f_{t_1,t_2}(t,s) &= f_{t_2|t_1}(s)\cdot f_{t_1}(t)\\ &= \lambda(t)\lambda(s+t)e^{-\mu(s+t)}.\end{aligned}$$

Changing variables from ISI's to spike times (i.e., $\nu_1 = t$, $\nu_2 = s + t$),

$$f_{T_1,T_2}(\nu_1,\nu_2) = \lambda(\nu_1)\lambda(\nu_2)e^{-\mu(\nu_2)}.$$

For more than two spikes, we get

$$\boxed{f_{T_1,\dots,T_n}(\nu_1,\dots,\nu_n) = \lambda(\nu_1)\dots\lambda(\nu_n)e^{-\mu(\nu_n)}}. \tag{15}$$

Sanity check: What does the spike train probability density 15 reduce down to for a homogeneous Poisson process?
For a homogeneous Poisson process, $\lambda(r) = \lambda_0 \forall r$. So,

$$f_{T_1,\dots,T_n}(\nu_1,\dots,\nu_n) = \lambda_0^n e^{-\lambda_0\nu_n}. \tag{16}$$

Note that this does not depend on the spike times $\nu_1,\dots,\nu_{n-1}$. Given that n spikes occured and the time of the last spike ν_n, all spike trains have the sam probability. Equation 16 could also have been obtained by multiplying exponential distributions, since ISI's are i.i.d.

$$\begin{aligned}f_{t_1,\dots,t_n}(u_1,\dots,u_n) &= \prod_{i=1} n\lambda_0 e^{-\lambda_0 u_i}\\ &= \lambda_0^n e^{-\lambda_0\left(\sum_{i=1}^n u_i\right)}\end{aligned}$$

where $\nu_n = \sum_{i=1}^n u_i$.

D Generating Poisson processes

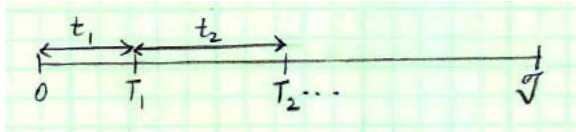

D.1 Homogeneous Poisson process with rate λ

Method 0 (Note - this is the most inefficient method!)

- Pick a bin size Δt such that $\left(1 - (1 + \lambda \Delta t)e^{-\lambda \Delta t}\right) \approx 0$ (i.e., the probability that more than one spike is emitted in a bin is very small).
- Generate a vector of Uniform$([01])$ random variables, $U_1, \dots, U_M$, where $M = \lceil \frac{\mathcal{T}}{\Delta t} \rceil$. (In Matlab use "`rand`".)
- Emit a spike in bin m if $U_m < \lambda \Delta t$, with corresponding spike time $T_k = m \Delta t$.

Method 1

- Generate i.i.d. exponential random variables $t_1, t_2, \dots$ with parameter λ. (In Matlab use "`exprnd`".)
- The spike times are $T_n = \sum_{i=1}^{n} t_i$
- If $T_n > \mathcal{T}$, stop

Method 2

- Draw $N \sim$ Poisson$(\lambda \mathcal{T})$, the number of spikes on the interval $[0, \mathcal{T}]$. (In Matlab use "`poissrnd`".)
- Draw $T_1, \dots, T_N \sim$ Uniform$([0, \mathcal{T})$. (In Matlab use "`rand`".) The $T_1, \dots, T_n$ are the spike times. (Cool!)

Why does Method 2 work?
The intuition is that a spike should not be more likely to occur at one time compared to another time (think of a Bernoulli process). More formally, Method 2 is based on the following (not-proved-here) theorem:

Theorem: If we conditionon $N(\mathcal{T}) = M$, then the set of spike times $\{T_1, \dots, T_M\}$ has the same distribution as $\{U_1, \dots, U_M\}$, where $U_1, \dots, U_M \sim$ Uniform $\left([0, \mathcal{T}]\right)$ i.i.d.

D.2 Inhomogeneous Poisson process with rate $\lambda(t)$

Method 0 See Method 0 above.

Method 1

- Let $\lambda_{\text{max}} = \max_t \lambda(t)$. Generate a homogeneous Poisson process with rate λ_{max} using one of the methods above.
- For $n = 1, \dots, N$

$$\left.\begin{array}{l} \text{Draw } U \sim \text{Uniform}([0,1]) \\ \text{If } U > \frac{\lambda(T_n)}{\lambda_{\text{max}}} \text{ reject the spike at } T_n \\ \text{Else, retain the spike at } T_n. \end{array}\right\} \text{thinning}$$

The spikes that are retained at the end of this procedure represent an inhomogeneous Poisson process with rate $\lambda(t)$.