Note 14. Poisson processes

Alex Fu

Fall 2022

1 Constructions

Poisson processes are in some sense the simplest continuous-time Markov processes. They inherit many nice properties of the Poisson distribution, some of which we recall: the parameter or **rate** λ is equal to the mean and the variance; independent Poisson random variables can be summed, and those sums decomposed; and the Poisson distribution is a limit of binomial distributions — the *law of rare events*.

Definition 1 (Counting process*).

A **counting process** is a nondecreasing $\mathbb{N}$-valued random process $(N_t)_{t\geq 0}$. The *number* of occurrences at time t is $N(t) = N_t$. We also write $N(s,t) := N(t) - N(s)$ for the number of occurrences, or *increment*, between times s and t, $s < t$.

One of the simplest counting processes is the count at times $t = 0, 1, 2, \ldots$ of an event, or *arrival*, that occurs independently with probability p, such as the number of successes of a repeated trial. In other words, $N(t) \sim \text{Binomial}(t, p)$, and every increment $N(t, t+1)$ is i.i.d. as Bernoulli(p). Such a counting process is called a **binomial process**.

We now follow the same strategy of continuization as in the law of rare events: we take finer and finer discrete time steps, keeping the rate $p/\Delta t = \lambda$ as an invariant. The resulting continuous-time process thus describes rare occurrences that happen with a constant rate: customer arrivals, natural disasters, phone calls to a store, and much more.

Definition 2 (Poisson process I).

The **Poisson process** with rate $\lambda > 0$, whose distribution is denoted PP(λ), is the counting process $(N_t)_{t\geq 0}$ that satisfies the following.

1. $N(0) = 0$ and $N(t) \sim \text{Poisson}(\lambda t)$.

2. **Stationary increments**. For every $s, t \geq 0$, $N(s, s+t) \stackrel{\mathsf{d}}{=} N(t)$. The increment in any interval of length t is distributed as Poisson(λt).
3. **Independent increments**. For every $k \in \mathbb{N}$ and times $t_0 < t_1 < \cdots < t_k$, the random variables $N(t_0, t_1)$, …, $N(t_{k-1}, t_k)$ are independent.

The simple plot of t vs N_t brings a surprising amount of insight.

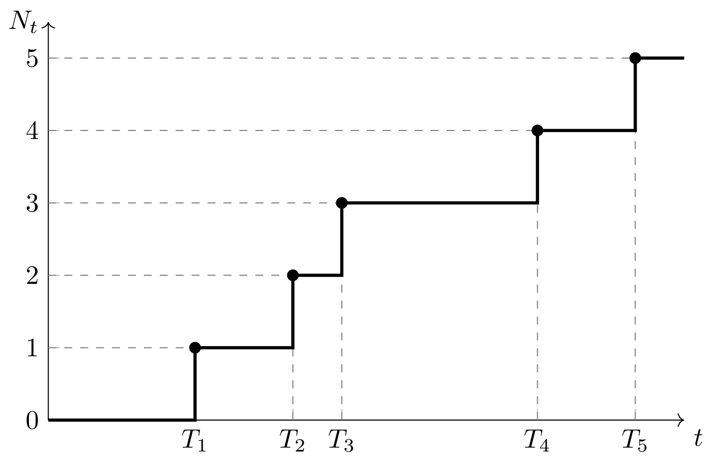

DTMCs have the basic simplifying assumption of *time-homogeneity*, which is incorporated by the fact that λ is constant. In continuous time, we impose a similar assumption of well-behavedness: **regularity**, that multiple consecutive increments or "double jumps" have negligible probability. If h is a small length of time, then

$$\begin{aligned}
\mathbb{P}(N(t, t+h) = 0) &= 1 - \lambda h - o(h) \\
\mathbb{P}(N(t, t+h) = 1) &= \lambda h + o(h) \\
\mathbb{P}(N(t, t+h) \geq 2) &= o(h)
\end{aligned}$$

where $o(h) \to 0$ as $h \to 0$. Regularity is a reasonable assumption: real-world occurrences mostly do not occur more than once in any infinitesimal interval of time. A simplified summary:

The key property of a Poisson process is intuitively $\mathbb{P}(N(t, t+dt) = 1) = \lambda\, dt$.

The fact that increments or jumps occur one at a time also gives the following representation of a Poisson process as a chain of states drawn in a line:

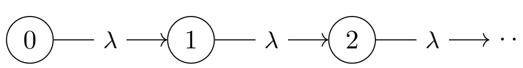

We will see that this in fact describes PP(λ) as a simple CTMC, parametrized by a single $\lambda > 0$. We now turn to leverage horizontal and vertical perspectives of the plot of N_t. The chain is the projection onto the vertical axis $\mathbb{N}$, the vertical distances are the increments of 1, and the height is always N_t. What about the projection down onto the horizontal axis?

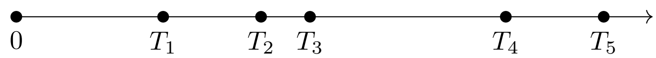

We have considered the Poisson process temporally so far, but by treating the time domain $\mathbb{R}_{\geq 0}$ as a space, in which points are randomly scattered, we gain a more general spatial interpretation often called the *Poisson point process*. It turns out that the PPP describes many common spatial distributions, such as wireless devices in a region, dust particles in air, or trees in a forest.

We can generalize Definition 1 as follows. The number of points inside any region with measure t is distributed as Poisson(λt); in 1 dimension, t is simply the length of the interval $(s, s+t]$. The number of points in disjoint regions is independent, which reduces to the property of independent increments.

A nice consequence of the spatial interpretation is that we find a uniformity of the points scattered in some fixed region $[0, N_t]$, which greatly simplifies our understanding of the arrival times. But, perhaps more importantly, by considering the positions of the points or arrival times T_i, and the horizontal distances or interarrival times $\tau_i = T_i - T_{i-1}$, we obtain a more natural construction of Poisson processes.

Definition 3 (Poisson process II).

Let $\lambda > 0$, let $(\tau_i)_{i=1}^{\infty}$ be i.i.d. Exponential(λ) **interarrival times** or *holding times*, and let $T_n := \sum_{i=1}^{n} \tau_i$ be **arrival times**, distributed as Erlang(n, λ). Then the **number of arrivals**

$$N_t := \sup\{n \geq 0 : T_n \leq t\}$$

for $t \geq 0$ form a **Poisson process** with rate λ.

So, Poisson processes are thus more easily modelled by a sum of i.i.d. samples of an Exponential distribution. The dual notations of N_t and T_n, number and time, also suggest some connections: $T_n = \inf\{t \geq 0 : N_t \geq n\}$ is the first time at which there are n arrivals. We wish to emphasize the following connection:

$N_t \geq n$ iff $T_n \leq t$.

In general, the sum $S_n = \sum_{i=1}^{n} X_i$ of i.i.d. samples of a fixed distribution are the arrival times, or *renewal times*, of a **renewal process**. We have already seen an application of renewal theory in the proof that $\pi(i) = 1/\mathbb{E}_i(T_i^+)$, where each revisit to i is an arrival or renewal.

Before our formal proof of the equivalence of Definition 1 and Definition 2, let us intuit why the holding times τ_i have an Exponential distribution. By the independence of increments, as $dt \to 0$,

$$\mathbb{P}(\tau_i > t) = (1 - \lambda\, dt)^{t/dt} = \left((1 - \lambda\, dt)^{1/d}\right)^t \to e^{-\lambda t},$$

which is precisely the ccdf of Exponential(λ).

2 Equivalency

We want to show that our definitions of a Poisson process coincide, which gives us the power of multiple perspectives, just as the four characterizations of DTMCs did. As a warmup,

Definition 4 (Erlang distribution).

The continuous random variable X follows the **Erlang distribution** Erlang(k, λ) for $k \in \mathbb{N}$ and $\lambda > 0$ if X is the sum of k i.i.d. Exponential(λ) r.v.s. The Erlang distribution is a special case of the *gamma* distribution for integer k.

Proposition 1 (Pdf of Erlang distribution).

The probability density function of $T_n \sim$ Erlang(n, λ) is

$$f_{T_n}(t) = \frac{(\lambda t)^{n-1}}{(n-1)!} \lambda e^{-\lambda t}.$$

Proof. The full joint distribution of $(T_1, \ldots, T_n)$ may be easier to work with, because it allows us to use the simpler interarrival times:

$$\begin{aligned} f_{T_1,\ldots,T_n}(t_1, \ldots, t_n) &= f_{\tau_1,\ldots,\tau_n}(t_1 - t_0, \ldots, t_n - t_{n-1}) \cdot \mathbb{1}\{t_1 \leq \cdots \leq t_n\} \\ &= \mathbb{1}_{t_1 \leq \cdots \leq t_n} \prod_{i=1}^{n} f_{\tau_i}(t_i - t_{i-1}) \\ &= \mathbb{1}_{t_1 \leq \cdots \leq t_n} \prod_{i=1}^{n} \lambda e^{-\lambda(t_i - t_{i-1})} \\ &= \mathbb{1}_{t_1 \leq \cdots \leq t_n} \lambda^n e^{-\lambda t_n}. \end{aligned}$$

Interestingly, there is no dependence on the values of $t_1, \ldots, t_{n-1}$, as we will find in another way in [Prop]. Now, we can integrate over the nuisance variables $t_1, \ldots, t_{n-1}$ to find

$$f_{T_n}(t) = \int_{t_1 \leq \cdots \leq t_{n-1} \leq t} f_{T_1,\ldots,T_{n-1},T_n}(t_1, \ldots, t_{n-1}, t)\, dt_1 \cdots dt_{n-1} = \frac{t^{n-1}}{(n-1)!} \cdot \lambda^n e^{-\lambda t}$$

by the symmetry of $t_1, \ldots, t_{n-1}$, as the hypercube $[0, t]^{n-1}$ has volume t^{n-1}, divided equally into $(n-1)!$ regions by permutations, so the support of the integral has volume $t^{n-1}/(n-1)!$. □

Proposition 2 (Poisson distribution).

N_t in Definition 2 is indeed distributed as Poisson(λt).

Proof. We recall the pmf of Poisson(λt), which is interestingly equal to $\frac{t}{n} \cdot f_{T_n}(n)$:

$$\mathbb{P}(\text{Poisson}(\lambda t) = n) = \frac{(\lambda t)^n}{n!} e^{-\lambda t}.$$

$N_t = n$ iff $T_n \le t$ and $T_{n+1} > t$, so $p_{N_t}(n)$ is equal to

$$\begin{aligned}
\mathbb{P}(T_1 \le t, \dots, T_n \le t, T_{n+1} > t) &= \int_{t_1 \le \dots \le t_n \le t < t_{n+1}} f_{T_1, \dots, T_{n+1}}(t_1, \dots, t_{n+1}) \, dt_1 \cdots dt_{n+1} \\
&= \int_{t_1 \le \dots \le t_n \le t} e^{-\lambda(t - t_n)} \prod_{i=1}^{n} \lambda e^{-\lambda(t_i - t_{i-1})} \, dt_1 \cdots dt_n \\
&= \frac{t^n}{n!} \lambda^n e^{-\lambda t}.
\end{aligned}$$

□

The main distributions: $\tau_i \sim$ Exponential(λ), $T_n \sim$ Erlang(n, λ), and $N_t \sim$ Poisson(λt).

As we have hinted at in the two proofs above,

Proposition 3 (Uniformity of past arrival times).

Given $N_t = n$, the arrival times $T_1, \dots, T_n$ are jointly distributed as the order statistics of n i.i.d. Uniform$([0, t])$ random variables:

$$f_{T_1, \dots, T_n | N_t}(t_1, \dots, t_n \mid n) = \frac{n!}{t^n} \cdot \mathbb{1}_{t_1 < \dots < t_n < t}.$$

Proof. Indeed, we can simply check that

$$f_{T_1, \dots, T_n | N_t}(t_1, \dots, t_n \mid n) \cdot p_{N_t}(n) = \frac{n!}{t^n} \mathbb{1}_{t_1 < \dots < t_n < t} \cdot \frac{(\lambda t)^n}{n!} e^{-\lambda t} = \lambda^n e^{-\lambda t}.$$

□

There is one more key insight we can get from the plot of t vs N_t: a Poisson process is self-similar — the distribution of its shape is the same when any (t, N_t) is taken as the origin $(0, 0)$. Given the present $N_t = n$, the past is uniform, and the future carries on memoryless.

Proposition 4 (Memorylessness of Poisson process).

If $(N_t)_{t \ge 0}$ is a Poisson process with rate λ, then for any $s \ge 0$, the shifted process started at time s, $(N_{s+t} - N_s)_{t \ge 0}$, is also a Poisson process with rate λ.

Proof. Recall the memorylessness of the Exponential distribution. If $X \sim$ Exponential(λ), then $X \mid X > s$ is distributed as $s +$ Exponential(λ):

$$\mathbb{P}(X > s + t \mid X > s) = \frac{\mathbb{P}(X > s + t)}{\mathbb{P}(X > s)} = \frac{e^{-\lambda(s+t)}}{e^{-\lambda s}} = e^{-\lambda t} = \mathbb{P}(X > t).$$

This is the somewhat paradoxical result that no longer how long you have already waited for an Exponential bus, you are always expected to wait for another λ time.

$$\mathbb{P}(N(s+t) = m \mid N(s) = n) = \mathbb{P}(N(s, s+t) = m - n)$$

We could derive the above from the memorylessness of the Exponential, but it will be immediately implied by the property of stationary increments. □

The ability to look into an ongoing Poisson process at any arbitrary time $s \geq 0$, without needing to know any of its history except for $N(s)$, is an incredibly desirable property we want to show. Now that we have warmed up by seeing several techniques for translating between $N_t \leftrightarrow T_n$, we are ready to show the equivalency of the definitions of Poisson processes:

Definition 1 is a statement about N_t, but we only readily know T_n, so we can translate between the two by $N_t \geq n$ iff $T_n \leq t$. The distribution of T_n is easier to find from the joint distribution of $T_1, \ldots, T_n$, as the latter can involve the interarrival times $\tau_1, \ldots, \tau_n$, whose key properties of *independence* and *Exponential distribution* give us a product of simpler terms.

Proposition 5.

Definition 2 satisfies the property of stationary increments.

Proof. Fix $s > 0$, and let T be the time until the first arrival after time s.

$$\begin{aligned}
\mathbb{P}(T > \tau \mid N_s = n, T_n = t) &= \mathbb{P}(\tau_{n+1} > (s - t) + \tau \mid \tau_{n+1} > s - t,\, T_n = t) \\
&= \mathbb{P}(\tau_{n+1} > (s - t) + \tau \mid \tau_{n+1} > s - t) \\
&= \mathbb{P}(\tau_{n+1} > \tau) \\
&= e^{-\lambda \tau}.
\end{aligned}$$

The interarrival times are independent — thus so are τ_{n+1} and T_n — and memoryless by being Exponential. As the conditional distribution of T does not depend on N_t or T_n, it must also be the unconditional (marginal) distribution, which we can verify by the law of total probability. By conditioning on $T_1, \ldots, T_n$, we find that T is independent on the arrival times, interarrival times, and in fact $(N_t)_{0 \leq t \leq s}$. This is in fact memorylessness.

Moreover, given $N_t = n$ and $T_n = t$, the subsequent interarrival times $\tau_{n+2}, \tau_{n+3}, \ldots$ after T are i.i.d. Exponential(λ), and by the same argument independent of $(N_t)_{0 \leq t \leq s}$. As these uniquely characterize $N(s, s+t)$ for $t \geq 0$, and $(T, \tau_{n+2}, \ldots) \overset{d}{=} (\tau_1, \tau_2, \ldots)$, we are done. □

Proposition 6.

Definition 2 satisfies the property of independent increments.

Proof. Let $t_0 < t_1 < \cdots < t_k$. We proceed by induction on $k \in \mathbb{N}$, where the base case is trivial. Supposing that $N(t_{i-1}, t_i)$ for $0 \leq i \leq k-1$ are independent, we consider the previous proof with $s = t_{k-1}$. Then the subsequent interarrival times, which determine $N(t_{k-1}, t_k)$, are independent of $(N_t)_{0 \leq t \leq s}$ and thus $N(t_{i-1}, t_i)$ for $0 \leq i \leq k-1$ as well. □

Definition 2 begets Definition 1, but we also want to show the reverse direction:

Proposition 7.

Let $(\tau_i)_{i=1}^{\infty}$ be a sequence of almost surely positive *interarrival times*, let $T_n := \sum_{i=1}^{n} \tau_i$, and let $N_t := \sup\{n \geq 0 : T_n \leq t\}$. If $N_t \sim$ Poisson(λt) for all $t \geq 0$ and has stationary and independent increments, then the τ_i are i.i.d. Exponential(λ).

Proof. We proceed by induction on $i \in \mathbb{Z}^+$. For the base case $i = 1$,

$$\mathbb{P}(\tau_1 > s) = \mathbb{P}(N(s) = 0) = e^{-\lambda s}$$

implies that $\tau_1 \sim$ Exponential(λ). Now, assuming that $\tau_1, \ldots, \tau_{n-1}$ are i.i.d. Exponential(λ), we let $t_{n-1} = \sum_{i=1}^{n-1} s_i$ and find that

$$\begin{aligned} \mathbb{P}(\tau_n > s \mid \tau_1 = s_1, \ldots, \tau_{n-1} = s_{n-1}) &= \mathbb{P}(N(t_{n-1}, t_{n-1} + s) = 0 \mid (N_t)_{0 \leq t \leq t_{n-1}}) \\ &= \mathbb{P}(N(t_{n-1}, t_{n-1} + s) = 0) \\ &= \mathbb{P}(N(s) = 0) \end{aligned}$$

by independent increments, stationary increments, and the fact that $\tau_1, \ldots, \tau_{n-1}$ uniquely specify and are specified by $(N_t)_{0 \leq t \leq t_{n-1}}$. As the conditional distribution of τ_n does not at all depend on the $\tau_1, \ldots, \tau_{n-1}$, we can check that this is also its unconditional (marginal) distribution, so τ_n is distributed as Exponential(λ) and jointly independent of $(\tau_1, \ldots, \tau_{n-1})$. □

The point of this section was not in the details of the proofs, but in the value of an equivalency: Definition 1 gives us properties of the Poisson distribution and the language of random processes more suited for proofs, while Definition 2 gives the properties of the Exponential distribution and an easier form for computation.

3 Merging and splitting

The applications of Poisson processes may seem limited by arrivals needing to be a "homogeneous" type, such as buses, customers, or packets, but two properties of the Poisson distribution allow us to also model "multityped" data: cars and trucks, different kinds of orders, packets from different sources, combining multiple independent processes into one.

Furthermore, amazingly, the reverse is also possible: given a Poisson stream of random arrivals that can be classified into types, the "thinned" or split streams of the arrivals of each individual type are themselves Poisson. For instance, if each scattered point is randomly colored, then the points of a particular color form another, smaller Poisson point process.

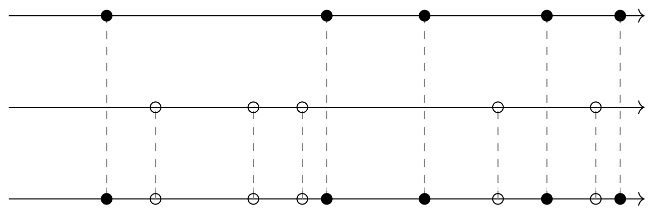

Or, if every incoming packet is independently routed to different servers with some probabilities for the purpose of load balancing, then the number of arrivals at each server can be easily analyzed as another Poisson process with rate proportional to its probability. Poisson processes thus play an integral part in *queueing theory* with generalized arrivals and departures.

Proposition 8 (Poisson merging).

If $X \sim$ Poisson(λ) and $Y \sim$ Poisson(μ) are independent, then $X + Y \sim$ Poisson$(\lambda + \mu)$.

Proposition 9 (Poisson splitting).

If $Z \sim$ Poisson(λ), $X :=$ Binomial$(Z, 1-p)$, and $Y := Z - X$, so that $Y \sim$ Binomial(Z, p), then $X \sim$ Poisson$(\lambda(1-p))$ and $Y \sim$ Poisson(λp) are independent.

These proofs should be familiar exercises. Note that it is not enough to find the expectations or rates, as we still need to show that the *distributions* are Poisson.

Theorem 1 (Merging of Poisson processes).

If $(N_t)_{t\geq 0}$ and $(M_t)_{t\geq 0}$ are independent Poisson processes with rates λ and μ respectively, then their sum $(N_t + M_t)_{t\geq 0}$ is also a Poisson process with rate $\lambda + \mu$.

Proof. Let $L := N + M$. To show independent increments, let $t_0 < \cdots < t_k$, and let us consider the joint distribution of the increments:

$$\mathbb{P}\left(\bigcap_{i=1}^{k}\{L(t_{i-1}, t_i) = \ell_i\}\right) = \sum_{n_1,\ldots,n_k : n_i \leq \ell_i} \mathbb{P}\left(\bigcap_{i=1}^{k}\{N(t_{i-1}, t_i) = n_i,\ M(t_{i-1}, t_i) = \ell_i - n_i\}\right)$$

$$= \sum_{n_1,\ldots,n_k : n_i \leq \ell_i} \prod_{i=1}^{k} \mathbb{P}(N(t_{i-1}, t_i) = n_i)\, \mathbb{P}(M(t_{i-1}, t_i) = \ell_i - n_i)$$

$$= \prod_{i=1}^{k} \sum_{n_i : n_i \le \ell_i} \mathbb{P}(N(t_{i-1}, t_i) = n_i) \, \mathbb{P}(M(t_{i-1}, t_i) = \ell_i - n_i)$$

$$= \prod_{i=1}^{k} \mathbb{P}(L(t_{i-1}, t_i) = \ell_i).$$

Here, we have simply used convolution, the independent increments of the independent processes N_t and M_t, then convolution once more. To show stationary increments, let $s \ge 0$:

$$\begin{aligned}
\mathbb{P}(L(s, s+t) = \ell) &= \sum_{n=0}^{\ell} \mathbb{P}(N(s, s+t) = n) \cdot \mathbb{P}(M(s, s+t) = \ell - n) \\
&= \sum_{n=0}^{\ell} \mathbb{P}(N(t) = n) \, \mathbb{P}(M(t) = \ell - n) \\
&= \mathbb{P}(L(t) = \ell)
\end{aligned}$$

by the stationary increments of N_t and M_t. Lastly, the Poisson distribution and rate of L_t are given by Proposition 8. □

Theorem 2 (Splitting of Poisson processes).

Let $(N_t)_{t \ge 0} \sim \mathsf{PP}(\lambda)$, and let $(B_n)_{n \in \mathbb{N}}$ be i.i.d. Bernoulli(p) independent of (N_t). If

$$N_a(t) := |\{n \le N(t) : B_n = 0\}| \sim \text{Binomial}(N(t), 1-p)$$
$$N_b(t) := |\{n \le N(t) : B_n = 1\}| \sim \text{Binomial}(N(t), p)$$

so that $N_a(t) + N_b(t) = N(t)$, then $(N_a(t))_{t \ge 0}$ and $(N_b(t))_{t \ge 0}$ are also independent Poisson processes with rates $\lambda(1-p)$ and λp respectively.

Proof. By symmetry, we only have to show that $(N_a(t))_{t \ge 0}$ is a Poisson process independent of $N_b(t)$. To show independent increments, let $t_0 < \cdots < t_k$:

$$\begin{aligned}
\mathbb{P}\left(\bigcap_{i=1}^{k} \{N_a(t_{i-1}, t_i) = a_i\} \right) &= \sum_{n_1, \ldots, n_k : a_i \le n_i} \mathbb{P}\left(\bigcap_{i=1}^{k} \{N_a(t_{i-1}, t_i) = a_i\} \,\middle|\, \bigcap_{i=1}^{k} \{N(t_{i-1}, t_i) = n_i\} \right) \\
&\qquad \cdot \mathbb{P}\left(\bigcap_{i=1}^{k} \{N(t_{i-1}, t_i) = n_i\} \right) \\
&= \sum_{n_1, \ldots, n_k : a_i \le n_i} \left[\prod_{i=1}^{k} \binom{n_i}{a_i} (1-p)^{a_i} p^{n_i - a_i} \cdot \mathbb{P}(N(t_{i-1}, t_i) = n_i) \right]
\end{aligned}$$

$$= \prod_{i=1}^{k} \left[\sum_{n_i : a_i \le n_i} \binom{n_i}{a_i} (1-p)^{a_i} p^{n_i - a_i} \cdot \mathbb{P}(N(t_{i-1}, t_i) = n_i) \right]$$

$$= \prod_{i=1}^{k} \mathbb{P}(N_a(t_{i-1}, t_i) = a_i).$$

Like previously, we have used the law of total probability to find the joint distribution of $N_a(t)$ in terms of the distribution of $N(t)$, which we know has independent increments. We can then written the joint probabilities as a product, then transform back using the law of total probability once more. For the sum $N_t + M_t$, the law gives a convolution; for the conditional $N_a \mid N$, the law gives the corresponding form we have used above.

To show stationary increments, we will similarly leverage the stationary increments of $N(t)$:

$$\begin{aligned} \mathbb{P}(N_a(s, s+t) = a) &= \sum_{n : a \le n} \binom{n}{a} (1-p)^a p^{n-a} \cdot \mathbb{P}(N(s, s+t) = n) \\ &= \sum_{n : a \le n} \binom{n}{a} (1-p)^a p^{n-a} \cdot \mathbb{P}(N(t) = n) \\ &= \mathbb{P}(N_a(t) = a). \end{aligned}$$

Lastly, $N_a(t) \sim \text{Poisson}(\lambda(1-p)t)$ by Proposition 9. For the independence of the split processes, we show that $\{N_a(t_{i-1}, t_i)\}_{i=1}^k$ and $\{N_b(t_{j-1}, t_j)\}_{j=1}^k$ are independent sets for $t_0 < \cdots < t_k$. For $i \neq j$, independence follows from the independent increments of $N(t)$, in particular $N(t_{i-1}, t_i)$ and $N(t_{j-1}, t_j)$. For $i = j$, it follows from Poisson splitting for $N(t_{i-1}, t_i)$. □

4 Random incidence paradox

We will end our discussion of Poisson processes with the interesting random incidence property or **random incidence paradox** (RIP).

Unlike the Poisson point process of points scattered over all of $\mathbb{R}$, the Poisson process $(N_t)_{t \ge 0}$ is not actually translation-invariant, even though it is memoryless, because there is a fixed beginning of time: for instance, $N(2, 4)$ is not equal in distribution to $N(-1, 1) = N(1)$. In some sense, "the past is not infinite like the future." With this in mind, let us consider a Poisson process that has been running for a long time.

Proposition 10 (Existence of a past arrival).

As $t \to \infty$, there is at least one arrival almost surely:

$$\mathbb{P}\Big(\lim_{t \to \infty} N_t \ge 1 \Big) = 1.$$

Proof. By the monotonicity of the event $\{N_t = 0\}$, as $s \le t$ implies $\{N_s = 0\} \supseteq \{N_t = 0\}$,

$$\mathbb{P}\Big(\lim_{t\to\infty} N_t = 0\Big) = \lim_{t\to\infty} \mathbb{P}(N_t = 0) = \lim_{t\to\infty} e^{-\lambda t} = 0.$$

□

We can thus ask the following questions at time t after the process has ran for a long time.

1. What is the expected length of time before the next arrival?
2. What is the expected interarrival time of the interarrival interval t is in?

These two seem like the same question. The answer to the first is $\frac{1}{\lambda}$, the mean of an interarrival time, by memorylessness, yet the answer to the second is different, contrary to intuition.

Suppose that $N_t = n$, so that $T_n \le t < T_{n+1}$. Then

$$T_{n+1} - T_n = (t - T_n) + (T_{n+1} - t).$$

$T_{n+1} - t$ is distributed as Exponential(λ) by memorylessness, but the problematic term is the nonnegative length $t - T_n$, which has ccdf

$$\mathbb{P}(t - T_n > \tau) = \mathbb{P}(N(t-\tau, t) = 0) = \mathbb{P}(N(\tau) = 0) = e^{-\lambda\tau}$$

for $0 \le \tau \le t$. So $t - T_n$ is almost Exponential(λ), "truncated" to the interval $[0, t]$. If we imagine running the process in reverse, then $t - T_n$ is time before the "next," or previous, arrival, except if there is no previous arrival, then $t - T_n = t$. The reversed process is "stopped" by the time 0, so Poisson processes are not quite reversible.

In any case, by the tail-sum formula,

$$\begin{aligned}\mathbb{E}(T_{n+1} - T_n) = \mathbb{E}(t - T_n) + \mathbb{E}(T_{n+1} - t) &= \int_0^\infty \mathbb{P}(t - T_n > \tau)\, d\tau + \frac{1}{\lambda} \\ &= \int_0^t e^{-\lambda\tau}\, d\tau + \frac{1}{\lambda} \\ &= \frac{1 - e^{-\lambda t}}{\lambda} + \frac{1}{\lambda}.\end{aligned}$$

We notice that $e^{-\lambda t}$ is precisely $\mathbb{P}(N_t = 0)$, the probability there are no past arrivals, which forces the truncation. By Proposition 10, as we take $t \to \infty$, we have

$$\mathbb{E}(T_{n+1} - T_n) = \frac{2}{\lambda}.$$

A past arrival almost surely exists, just as a future arrival almost surely exists; the past is infinitely long, the time $t = 0$ infinitely far away, just like the future and $t = +\infty$, so the past and future are symmetric in distribution, the process the same run forwards and backwards, which explains the doubled expected interarrival time: $\frac{1}{\lambda}$ from t to T_{n+1}, and $\frac{1}{\lambda}$ from t to T_n.

While the answer of $\frac{2}{\lambda}$ makes sense, it may still seem as if we have two conflicting answers to the same question of finding an expected interarrival time. However, the second question is not the same as uniformly selecting an interarrival interval: each interval is weighted by its length, the number of points $t \in [T_n, T_{n+1})$. Thus t has a greater chance to fall in a longer interval, so we expect the chosen interval to be longer than a typical interarrival interval.

Our last remark: Poisson also means *fish* in French, so if you ever encounter any Poisson processes problems involving poisonous fish, that's why.

■