

1 Poisson processes, and Compound (batch) Poisson processes

1.1 Point Processes

Definition 1.1 A simple point process $\psi = \{t_n : n \geq 1\}$ is a sequence of strictly increasing points

$$0 < t_1 < t_2 < \cdots, \quad (1)$$

with $t_n \rightarrow \infty$ as $n \rightarrow \infty$. With $N(0) \stackrel{\text{def}}{=} 0$ we let $N(t)$ denote the number of points that fall in the interval $(0, t]$; $N(t) = \max\{n : t_n \leq t\}$. $\{N(t) : t \geq 0\}$ is called the counting process for ψ . If the t_n are random variables then ψ is called a random point process. We sometimes allow a point t_0 at the origin and define $t_0 \stackrel{\text{def}}{=} 0$. $X_n = t_n - t_{n-1}$, $n \geq 1$, is called the n^{th} interarrival time.

We view t as time and view t_n as the n^{th} arrival time. The word *simple* refers to the fact that we are not allowing more than one arrival to occur at the same time (as is stated precisely in (1)). In many applications there is a “system” to which customers are arriving over time (classroom, bank, hospital, supermarket, airport, etc.), and $\{t_n\}$ denotes the arrival times of these customers to the system. But $\{t_n\}$ could also represent the times at which phone calls are made to a given phone or office, the times at which jobs are sent to a printer in a computer network, the times at which one receives or sends email, the times at which one sells or buys stock, the times at which a given web site receives hits, or the times at which subways arrive to a station. Note that

$$t_n = X_1 + \cdots + X_n, \quad n \geq 1,$$

the n^{th} arrival time is the sum of the first n interarrival times.

Also note that the event $\{N(t) = 0\}$ can be equivalently represented by the event $\{t_1 > t\}$, and more generally

$$\{N(t) = n\} = \{t_n \leq t, t_{n+1} > t\}, \quad n \geq 1.$$

In particular, for a random point process, $P(N(t) = 0) = P(t_1 > t)$.

1.2 Renewal process

A random point process $\psi = \{t_n\}$ for which the interarrival times $\{X_n\}$ form an i.i.d. sequence is called a *renewal process*. t_n is then called the n^{th} *renewal epoch* and $F(x) = P(X \leq x)$ denotes the common interarrival time distribution. The *rate* of the renewal process is defined as $\lambda \stackrel{\text{def}}{=} 1/E(X)$.

Such processes are very easy to simulate, assuming that one can generate samples from F which we shall assume is so here; the key point being that we can use the recursion $t_{n+1} = t_n + X_{n+1}$, $n \geq 0$, and thus simply generate the interarrival times sequentially:

Simulating a renewal process with interarrival time distribution F up to time T :

1. $t = 0$, $N = 0$
2. Generate an X distributed as F .
3. $t = t + X$. If $t > T$, then stop.

4. Set $N = N + 1$ and set $t_N = t$.
5. Go back to 2.

Note that when the algorithm stops, the value of N is $N(T)$ and we have sequentially simulated all the desired arrival times $t_1, t_2 \dots$ up to $t_{N(T)}$.

1.3 Poisson process

Definition 1.2 A Poisson process at rate λ is a renewal point process in which the interarrival time distribution is exponential with rate λ : interarrival times $\{X_n : n \geq 1\}$ are i.i.d. with common distribution $F(x) = P(X \leq x) = 1 - e^{-\lambda x}$, $x \geq 0$; $E(X) = 1/\lambda$.

Simulating a Poisson process at rate λ up to time T :

1. $t = 0$, $N = 0$
2. Generate U .
3. $t = t + \lceil -(1/\lambda) \ln(U) \rceil$. If $t > T$, then stop.
4. Set $N = N + 1$ and set $t_N = t$.
5. Go back to 2.

1.4 Further properties of the Poisson process; a different algorithm for simulating

Here we review known properties of the Poisson process and use them to obtain another algorithm for simulating such a process.

The reason that the Poisson process is named so is because: For each fixed $t > 0$, the distribution of $N(t)$ is Poisson with mean λt :

$$P(N(t) = k) = e^{-\lambda t} \frac{(\lambda t)^k}{k!}, \quad k \geq 0.$$

In particular, $E(N(t)) = \lambda t$, $Var(N(t)) = \lambda t$, $t \geq 0$. In fact, the number of arrivals in ANY interval of length t , $N(s+t) - N(s)$ is also Poisson with mean λt :

$$P(N(s+t) - N(s) = k) = e^{-\lambda t} \frac{(\lambda t)^k}{k!}, \quad s > 0, \quad k \geq 0,$$

and $E(N(s+t) - N(s)) = \lambda t$, $Var(N(s+t) - N(s)) = \lambda t$, $t \geq 0$.

$N(s+t) - N(s)$ is called a length t *increment* of the counting process $\{N(t) : t \geq 0\}$; the above tells us that the Poisson counting process has increments that have a distribution that is Poisson and only depends on the length of the increment. Any increment of length t is distributed as Poisson with mean λt .

This is an example of a process having *stationary increments*: Any increment of length t has a distribution that only depends on the length t .

The Poisson process also has *independent increments*, meaning that non-overlapping increments are independent: If $0 \leq a < b < c < d$, then the two increments $N(b) - N(a)$, and $N(d) - N(c)$ are independent rvs.

Remarkable as it may seem, it turns out that the Poisson process is completely characterized by stationary and independent increments:

Theorem 1.1 *Suppose that ψ is a simple random point process that has both stationary and independent increments. Then in fact, ψ is a Poisson process. Thus the Poisson process is the only simple point process with stationary and independent increments.*

The practical consequences of this theorem: To check if some point process is Poisson, one need only verify that it has stationary and independent increments. (One would not need to check to see if it is a renewal process and then check that the interarrival time distribution is exponential.)

Remark 1.1 A Poisson process at rate λ can be viewed as the result of performing an independent Bernoulli trial with success probability $p = \lambda dt$ in each “infinitesimal” time interval of length dt , and placing a point there if the corresponding trial is a success (no point there otherwise). Intuitively, this would yield a point process with both stationary and independent increments; a Poisson process: The number of Bernoulli trials that can be fit in any interval only depends on the length of the interval and thus the distribution for the number of successes in that interval would also only depend on the length; stationary increments follows. For two non-overlapping intervals, the Bernoulli trials in each would be independent of one another since all the trials are i.i.d., thus the number of successes in one interval would be independent of the number of successes in the other interval; independent increments follows. This Bernoulli trials idea can be made mathematically rigorous.

1.4.1 Partitioning Theorems for Poisson processes and random variables

Given $X \sim \text{Poiss}(\alpha)$ (a Poisson rv with mean α) suppose that we imagine that X denotes some number of objects (arrivals during some fixed time interval for example), and that independent of one another, each such object is of type 1 or type 2 with probability p and $q = 1 - p$ respectively. This means that if $X = n$ then the number of those n that are of type 1 has a $\text{Bin}(n, p)$ distribution and the number of those n that are of type 2 has a $\text{Bin}(n, q)$ distribution. Let X_1 and X_2 denote the number of type 1 and type 2 objects respectively ; $X_1 + X_2 = X$. The following shows that in fact if we do this, then X_1 and X_2 are independent Poisson random variables with means $p\alpha$ and $q\alpha$ respectively.

Theorem 1.2 (Partitioning a Poisson r.v.) *If $X \sim \text{Poiss}(\alpha)$ and if each object of X is, independently, type 1 or type 2 with probability p and $q = 1 - p$, then in fact $X_1 \sim \text{Poiss}(p\alpha)$, $X_2 \sim \text{Poiss}(q\alpha)$ and they are independent.*

Proof : We must show that

$$P(X_1 = k, X_2 = m) = e^{-p\alpha} \frac{(p\alpha)^k}{k!} e^{-q\alpha} \frac{(q\alpha)^m}{m!}. \quad (2)$$

$$P(X_1 = k, X_2 = m) = P(X_1 = k, X = k + m) = P(X_1 = k | X = k + m)P(X = k + m).$$

But given $X = k + m$, $X_1 \sim \text{Bin}(k + m, p)$ yielding

$$P(X_1 = k | X = k + m)P(X = k + m) = \frac{(k + m)!}{k!m!} p^k q^m e^{-\alpha} \frac{\alpha^{k+m}}{(k + m)!}.$$

Using the fact that $e^\alpha = e^{p\alpha} e^{q\alpha}$ and other similar algebraic identities, the above reduces to (2) as was to be shown. ■

The above theorem generalizes to Poisson processes:

Theorem 1.3 (Partitioning a Poisson process) *If $\psi \sim PP(\lambda)$ and if each arrival of ψ is, independently, type 1 or type 2 with probability p and $q = 1 - p$ then in fact, letting ψ_i denote the point process of type i arrivals, $i = 1, 2$, the two resulting point processes are themselves Poisson and independent: $\psi_1 \sim PP(p\lambda)$, $\psi_2 \sim PP(q\lambda)$ and they are independent.*

The above generalizes in the obvious fashion to $k \geq 2$ types (type i with probability p_i) yielding independent Poisson processes with rates $\lambda_i = p_i\lambda$, $i \in \{1, 2, \dots, k\}$.

The above is quite interesting for it means that if Poisson arrivals at rate λ come to our lecture room, and upon each arrival we flip a coin (having probability p of landing heads), and route all those for which the coin lands tails (type 2) into a different room, only allowing those for which the coin lands heads (type 1) enter our room, then the arrival processes to the two room are independent and Poisson.

The above allows us to simulate two desired independent Poisson processes (e.g., two different “classes” of arrivals) by simulating only one and partitioning it:

Simulating two independent Poisson process at rates λ_1, λ_2 up to time T :

1. $t = 0, t_1 = 0, t_2 = 0, N_1 = 0, N_2 = 0$, set $\lambda = \lambda_1 + \lambda_2$, set $p = \lambda_1/\lambda$.
2. Generate U .
3. $t = t + \lceil -(1/\lambda) \ln(U) \rceil$. If $t > T$, then stop.
4. Generate U . If $U \leq p$, then set $N_1 = N_1 + 1$ and set $t_{N_1 1} = t$; otherwise ($U > p$) set $N_2 = N_2 + 1$ and set $t_{N_2 2} = t$
5. Go back to 2.

The above algorithm generalizes in a straightforward fashion to handle $k \geq 2$ such independent processes.

1.5 Uniform property of Poisson process arrival times: a new simulation algorithm

Suppose that for a Poisson process at rate λ , we condition on the event $\{N(t) = 1\}$, the event that exactly one arrival occurred during $(0, t]$. We might conjecture that under such conditioning, t_1 should be uniformly distributed over $(0, t)$. To see that this is in fact so, choose $s \in (0, t)$. Then

$$\begin{aligned} P(t_1 \leq s | N(t) = 1) &= \frac{P(t_1 \leq s, N(t) = 1)}{P(N(t) = 1)} \\ &= \frac{P(N(s) = 1, N(t) - N(s) = 0)}{P(N(t) = 1)} \\ &= \frac{e^{-\lambda s} \lambda s e^{-\lambda(t-s)}}{e^{-\lambda t} \lambda t} \\ &= \frac{s}{t}. \end{aligned}$$

It turns out that this result generalizes nicely to any number of arrivals, and we present this next.

Let U_1, U_2, \dots, U_n be n i.i.d uniformly distributed r.v.s. on the interval $(0, t)$. Let $U_{(1)} < U_{(2)} < \dots < U_{(n)}$ denote them placed in ascending order. Thus $U_{(1)}$ is the smallest of them ,

$U_{(2)}$ the second smallest and finally $U_{(n)}$ is the largest one. $U_{(i)}$ is called the i^{th} order statistic of U_1, \dots, U_n .

Note that the joint density function of (U_1, U_2, \dots, U_n) is given by

$$g(s_1, s_2, \dots, s_n) = \frac{1}{t^n}, \quad s_i \in (0, t),$$

because each U_i has density function $1/t$ and they are independent. Now given any ascending sequence $0 < s_1 < s_2 < \dots < s_n < t$ it follows that the joint distribution $f(s_1, s_2, \dots, s_n)$ of the order statistics $(U_{(1)}, U_{(2)}, \dots, U_{(n)})$ is given by

$$f(s_1, s_2, \dots, s_n) = \frac{n!}{t^n},$$

because there are $n!$ permutations of the sequence (s_1, s_2, \dots, s_n) each of which leads to the same order statistics. For example if $(s_1, s_2, s_3) = (1, 2, 3)$ then there are $3! = 6$ permutations all yielding the same order statistics $(1, 2, 3)$: $(1, 2, 3)$, $(1, 3, 2)$, $(2, 1, 3)$, $(2, 3, 1)$, $(3, 1, 2)$, $(3, 2, 1)$.

Theorem 1.4 For any Poisson process, if $n \geq 1$, then: conditional on the event $\{N(t) = n\}$, the joint distribution of the n arrival times t_1, \dots, t_n is the same as the joint distribution of $U_{(1)}, \dots, U_{(n)}$, the order statistics of n i.i.d. $\text{unif}(0, t)$ r.v.s., that is, it is given by

$$f(s_1, s_2, \dots, s_n) = \frac{n!}{t^n}, \quad 0 < s_1 < s_2 < \dots < s_n < t.$$

Thus the following algorithm works:

Simulating a Poisson process at rate λ up to time T :

1. Generate N distributed as Poisson with mean λT . If $N = 0$ stop.
2. Set $n = N$. Generate n iid uniforms on $(0, 1)$, U_1, \dots, U_n and reset $U_i = TU_i$, $i \in \{1, \dots, n\}$ (this makes the U_i be uniform on $(0, T)$ as required).
3. Place the U_i in ascending order to obtain the order statistics $U_{(1)} < U_{(2)} < \dots < U_{(n)}$.
4. Set $t_i = U_{(i)}$, $i \in \{1, \dots, n\}$

The importance of the above is this: If you want to simulate a Poisson process up to time T , you need only first simulate the value of $N(T)$, then if $N(T) = n$ generate n i.i.d. $\text{Unif}(0, T)$ numbers (U_1, U_2, \dots, U_n) , place them in ascending order $(U_{(1)}, U_{(2)}, \dots, U_{(n)})$ and finally define $t_i = U_{(i)}$, $1 \leq i \leq n$. Simulating the value $N(T)$ is easy since $N(T)$ has a Poisson distribution with mean λT and we have several methods already to handle this.

Uniform numbers are very easy to generate on a computer and so this method can have computational advantages over simply generating exponential r.v.s. for interarrival times X_n , and defining $t_n = X_1 + \dots + X_n$. Exponential r.v.s. require taking logarithms to generate:

$$X_i = -\frac{1}{\lambda} \log(U_i),$$

where $U_i \sim \text{Unif}(0, 1)$ and this can be computationally time consuming.

Proof :[Theorem 1.4:] We will compute the density for

$$P(t_1 = s_1, \dots, t_n = s_n | N(t) = n) = \frac{P(t_1 = s_1, \dots, t_n = s_n, N(t) = n)}{P(N(t) = n)}$$

and see that it is precisely $\frac{n!}{t^n}$. To this end, we re-write the event $\{t_1 = s_1, \dots, t_n = s_n, N(t) = n\}$ in terms of the i.i.d. interarrival times as $\{X_1 = s_1, \dots, X_n = s_n - s_{n-1}, X_{n+1} > t - s_n\}$. For example if $N(t) = 2$, then $\{t_1 = s_1, t_2 = s_2, N(t) = 2\} = \{X_1 = s_1, X_2 = s_2 - s_1, X_3 > t - s_2\}$ and thus has density $\lambda e^{-\lambda s_1} \lambda e^{-\lambda(s_2 - s_1)} e^{-\lambda(t - s_2)} = \lambda^2 e^{-\lambda t}$ due to the independence of the r.v.s. X_1, X_2, X_3 , and the algebraic cancellations in the exponents.

We conclude that

$$\begin{aligned} P(t_1 = s_1, \dots, t_n = s_n | N(t) = n) &= \frac{P(t_1 = s_1, \dots, t_n = s_n, N(t) = n)}{P(N(t) = n)} \\ &= \frac{P(X_1 = s_1, \dots, X_n = s_n - s_{n-1}, X_{n+1} > t - s_n)}{P(N(t) = n)} \\ &= \frac{\lambda^n e^{-\lambda t}}{P(N(t) = n)} \\ &= \frac{n!}{t^n}, \end{aligned}$$

where the last equality follows since $P(N(t) = n) = e^{-\lambda t} (\lambda t)^n / n!$. ■

1.6 Non-stationary Poisson processes

Assuming that a Poisson process has a fixed and constant rate λ over all time limits its applicability. (This is known as a time-stationary or time-homogenous Poisson process, or just simply a stationary Poisson process.) For example, during rush hours, the arrivals/departures of vehicles into/out of Manhattan is at a higher rate than at (say) 2:00AM. To accommodate this, we can allow the rate $\lambda = \lambda(t)$ to be a deterministic function of time $t \geq 0$. For example, consider time in hours and suppose $\lambda(t) = 100$ per hour except during the time interval (morning rush hour) $(8, 9)$ when $\lambda(t) = 200$, that is

$$\lambda(t) = 200, t \in (8, 9), \lambda(t) = 100, t \notin (8, 9).$$

In such a case, for a given rate function $\lambda(t)$, the expected number of arrivals by time t is thus given by

$$m(t) \stackrel{\text{def}}{=} E(N(t)) = \int_0^t \lambda(s) ds. \quad (3)$$

We have already learned how to simulate a stationary Poisson process up to any desired time t , and next we will learn how to do so for a non-stationary Poisson process.

1.7 The non-stationary case: Thinning

In general the function $\lambda(t)$ is called the *intensity* of the Poisson process, and the following two conditions must hold (e.g. these two conditions define a non-stationary Poisson process):

1. For each $t > 0$, the counting random variable $N(t)$ is Poisson distributed with mean

$$m(t) = \int_0^t \lambda(s) ds.$$

$$E(N(t)) = m(t)$$

$$P(N(t) = k) = e^{-m(t)} \frac{m(t)^k}{k!}, \quad k \geq 0.$$

More generally, the increment $N(t+h) - N(t)$ has a Poisson distribution with mean $m(t+h) - m(t) = \int_t^{t+h} \lambda(s) ds$.

2. $\{N(t)\}$ has independent increments; if $0 \leq a < b < c < d$, then $N(b) - N(a)$ is independent of $N(d) - N(c)$.

We shall assume that the intensity function is bounded from above: There exists a $\lambda^* > 0$ such that

$$\lambda(t) \leq \lambda^*, \quad t \geq 0.$$

(In practice, we would want to use the smallest such upper bound.)

Then the simulation of the Poisson process is accomplished by a “thinning” method: First simulate a stationary Poisson process at rate λ^* , denote the arrival times of this process by $\{v_n : n \geq 1\}$. For example, sequentially generate iid exponential rate λ^* interarrival times and use the recursion $v_{n+1} = v_n + [-(1/\lambda^*) \ln(U_{n+1})]$, to obtain the stationary process arrival times. The rate λ^* is larger than needed for our actual process, so for each arrival time v_n , we independently perform a Bernoulli trial with success probability $\lambda(v_n)/\lambda^*$ to decide whether to keep it (success) or reject it (failure). The sequence of accepted times we denote by $\{t_n\}$ and forms our desired non-stationary Poisson process. To make this precise: for each arrival time v_n , we accept it with probability $p_n = \lambda(v_n)/\lambda^*$, and reject it with probability $1 - p_n$. Thus for each v_n we generate a uniform U_n and if $U_n \leq p_n$, then we accept v_n as a point, otherwise we reject it.

The thinning algorithm for simulating a non-stationary Poisson process with intensity $\lambda(t)$ that is bounded by λ^*

Here is the algorithm for generating our non-stationary Poisson process up to a desired time T to get the $N(T)$ arrival times $t_1, \dots, t_{N(T)}$.

1. $t = 0, N = 0$
2. Generate a U
3. $t = t + [-(1/\lambda^*) \ln(U)]$. If $t > T$, then stop.
4. Generate a U .
5. If $U \leq \lambda(t)/\lambda^*$, then set $N = N + 1$ and set $t_N = t$.
6. Go back to 2.

Note that when the algorithm stops, the value of N is $N(T)$ and we have sequentially simulated all the desired arrival times $t_1, t_2 \dots$ up to $t_{N(T)}$.

Here is a proof that this thinning algorithm works:

Proof: [Thinning works] Let $\{M(t)\}$ denote the counting process of the rate λ^* Poisson process, and let $\{N(t)\}$ the counting process of the resulting thinned process. We need to show that $\{N(t)\}$ has independent increments and that the increments are Poisson distributed with the correct $m(t)$ function. First note that $\{N(t)\}$ has independent increments since $\{M(t)\}$ does

and the thinning is done independently; $\{N(t)\}$ inherits independent increments from $\{M(t)\}$. So what is left to prove is that for each $t > 0$, $N(t)$ constructed by this thinning has a Poisson distribution with mean $m(t) = \int_0^t \lambda(s) ds$. We know that for each $t > 0$, $M(t)$ has a Poisson distribution with mean $\lambda^* t$. We will partition $M(t)$ into two types: $N(t)$ (the accepted ones), and $R(t)$ (the rejected ones), and conclude that $N(t)$ has the desired Poisson distribution (recall Theorem 1.2). To this end recall that conditional on $M(t) = n$, we can treat the n unordered arrival times as iid unif $(0, t)$ rvs. Thus a typical arrival, denoted by $V \sim Unif(0, t)$, will be accepted with conditional probability $\lambda(v)/\lambda^*$, conditional on $V = v$. Thus the unconditional probability of acceptance is $p = p(t) = E[\lambda(V)/\lambda^*] = (1/\lambda^*)(1/t) \int_0^t \lambda(s) ds$, and we conclude from partitioning (Theorem 1.2) that $N(t)$ has a Poisson distribution with mean $\lambda^* t p = m(t)$, as was to be shown. ■

1.8 Simulating a compound Poisson process

Another restriction we want to relax is the assumption that arrivals occur strictly one after the other. Buses and airplanes, for example, hold more than one passenger. This is easily remedied by assuming the bus arrival times form a Poisson process, but independently each bus holds a random number of passengers (generically denoted by B) with some probability mass function $P(k) = P(B = k)$, $k \geq 0$, and mean $E(B)$. Letting B_1, B_2, \dots denote the iid sequential bus sizes, the number of passengers to arrive by time t , $X(t)$ is given by

$$X(t) = \sum_{n=1}^{N(t)} B_n, \tag{4}$$

where $N(t)$ is the counting process for the Poisson process; $N(t) =$ the number of buses to arrive by time t . This is known as a *compound* or *batch* Poisson arrival process.

It is easily shown (using Wald's equation) that $E(X(t)) = E(N(t))E(B) = \lambda t E(B)$.

We could also take the bus arrival times as a non-stationary Poisson process, in which case we would have a non-stationary compound Poisson process, and then $E(X(t)) = E(N(t))E(B) = m(t)E(B)$. Finally, there are applications where we want to allow B to be a continuous rv with a density. For example, consider an insurance company to which claims are made at times forming a Poisson process, and iid claim sizes B_1, B_2, \dots . The claim sizes are amounts of money, and (4) denotes the total amount of money paid out by the insurance company during the time interval $(0, t]$.

Suppose that we wish to simulate a stationary compound Poisson process at rate λ with iid B_i distributed as (say) G (could be continuous or discrete). Suppose that we already have an algorithm for generating from G .

Algorithm for generating a compound Poisson process up to a desired time T to get $X(T)$

1. $t = 0$, $N = 0$, $X = 0$.
2. Generate a U
3. $t = t + [-(1/\lambda) \ln(U)]$. If $t > T$, then stop.
4. Generate B distributed as G .
5. Set $N = N + 1$ and set $X = X + B$
6. Go back to 2.

In the non-stationary case the algorithm becomes (using thinning):

1. $t = 0, N = 0, X = 0$.
2. Generate a U
3. $t = t + [-(1/\lambda^*) \ln(U)]$. If $t > T$, then stop.
4. Generate a U .
5. If $U \leq \lambda(t)/\lambda^*$, then set $N = N + 1$ and generate B distributed as G and set $X = X + B$.
6. Go back to 2.