MAS275 Probability Modelling Chapter 6: Poisson processes

Robin Stephenson

School of Mathematics and Statistics, University of Sheffield

Spring Semester, 2020

Introduction

Poisson processes are a particularly important topic in probability theory.

The one-dimensional Poisson process, which most of this section will be about, is a model for the random times of occurrences of instantaneous events.

Applications

There are many examples of things whose random occurrences in time can be modelled by Poisson processes, for example

- customers arriving in a queue
- incoming calls to a phone
- eruptions of a volcano
- and so on

Higher dimensional analogues of the basic Poisson process are also useful as models for random locations of objects in space.

Properties of the Poisson distribution

The Poisson distribution

Recall that the Poisson distribution with parameter λ, denoted $Po(\lambda)$, has probability function

$$p(x) = \frac{\lambda^x e^{-\lambda}}{x!}, x = 0, 1, 2, 3, \ldots$$

and a random variable with $Po(\lambda)$ distribution has $E(X) = \mathrm{Var}(X) = \lambda$.

Additivity

Proposition

The Poisson distribution is additive.

More precisely, if X and Y are independent random variables with distributions $Po(\lambda)$ and $Po(\mu)$ respectively, then $X + Y$ has the distribution $Po(\lambda + \mu)$

This property extends in an obvious way to more than two independent random variables.

Proof: See Exercise 17(b)(i).

Limit of binomials

Proposition

The Poisson distribution is a limit of binomial distributions.

Consider the sequence of binomial distributions $Bi(n, \mu/n)$ for fixed μ.

Then the number of trials n is increasing but the probability of success on a trial is μ/n which is decreasing in proportion, so that the mean μ remains fixed.

In the limit, the distribution is Poisson with parameter μ.

(See MAS113 for proof.)

The basic Poisson process

Time scale

We now move to a continuous time scale which is usually regarded as starting at time zero, so that it consists of the positive real numbers, and denote our time variable typically by t.

When we refer to a time interval, we adopt the convention that it excludes the left hand end point but includes the right hand end-point, say

$$(u, v] = \{t : u < t \leq v\}.$$

Basic Poisson process I

The Poisson process is described in terms of the random variables $N_{u,v}$ for $0 \leq u \leq v$, where $N_{u,v}$ is the number of occurrences in the time interval $(u, v]$.

The process has one parameter, which is a positive number λ known as the **rate** of the process.

λ is meant to measure the average or expected number of occurrences per unit time.

Basic Poisson process II

The basic Poisson process is then defined by the following two assumptions:

(a) For any $0 \leq u \leq v$, the distribution of $N_{u,v}$ is Poisson with parameter $\lambda(v - u)$.

(b) If $(u_1, v_1], (u_2, v_2], \ldots, (u_k, v_k]$ are disjoint time intervals then $N_{u_1,v_1}, N_{u_2,v_2}, \ldots, N_{u_k,v_k}$ are independent random variables.

Basic Poisson process III

Note that by assumption (a) and the mean of a Poisson random variable, $N_{u,v}$ has mean $\lambda(v - u)$, which gives the correct interpretation to the rate λ as described above.

Do these work?

To show that a process actually exists which satisfies these assumptions still requires a bit of work.

The special properties of Poisson distributions are important: if say $u < v < w$ then

$$N_{u,w} = N_{u,v} + N_{v,w}$$

Because all three of these random variables are to have Poisson distributions, and the two on the RHS are independent, the assumptions can only work because of the additivity property.

We will see a bit more on justifying the existence of a process satisfying the assumptions later.

Why Poisson?

We can also ask about why the assumption of Poisson distributions might make sense in a modelling context.

Divide $(0, t]$ up into n small intervals, and assume that there is the same small probability of an occurrence in each, assuming that the probability of more than one occurrence in an interval is negligible.

In order to fix the expected number of occurrences at λt, this probability must be set equal to $\lambda t/n$.

Binomial limit

Then the number of occurrences in the interval has the binomial distribution $Bi(n, \lambda t/n)$.

As we let n tend to infinity, the small intervals become smaller and so the approximation to the continuous time scale becomes closer

By the relationship between the Poisson and Binomial distributions (see Proposition 17) the distribution of the number of occurrences approaches the Poisson distribution with parameter λt.

This suggests the Poisson distribution as a sensible model in the genuinely continuous time setting.

Other distributions?

You might like to think about which other distributions you have encountered might have assumptions similar to (a) and (b) which work.

Simulation

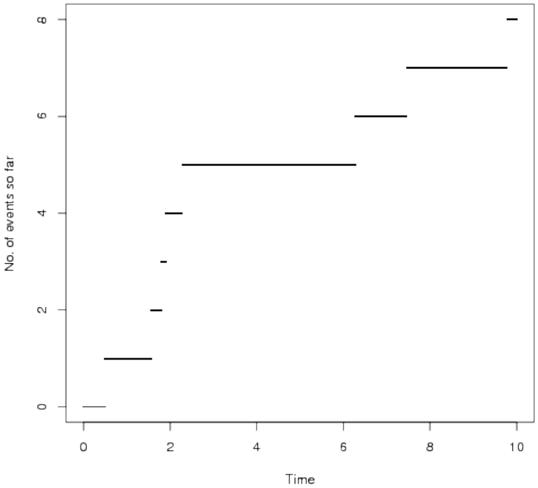

Example

Example

Volcanic eruptions

Multiple occurrences and inter-occurrence times

Inter-occurrence times

Let T_1 denote the length of time until the first occurrence, ...

... T_2 denote the length of time between the first and second occurrences, ...

... and so on, so that T_n represents the time between occurrences $n-1$ and n.

These random variables are called **inter-occurrence times**. We first show that these cannot be zero, and we will then show which distribution they have.

No multiple occurrences

Theorem

The probability that in $(0, t]$ two occurrences of a Poisson process with rate λ occur at exactly the same time is zero.

Exponential distribution

Theorem

Inter-occurrence times are independent of each other, and are exponentially distributed with parameter λ.

Variable rate Poisson process

Variable rate

We can make the basic Poisson process more flexible and realistic as a model by allowing the rate of the process to vary with time, $\lambda(t)$ say.

This can take into account, for example, the fact that traffic is heavier at rush hours, the rate of emission of particles from a radioactive isotope declines with time, and so on.

Assumptions

The only change to the definition of the Poisson process is that the assumption (a) is replaced by the following:

For any $0 \leq u \leq v$, the distribution of $N_{u,v}$ is Poisson with parameter $\int_u^v \lambda(t)dt$.

This assumption generalises that of the constant rate case, and gives the correct interpretation of "rate" when this rate is varying.

If $\lambda(t)$ is a constant we recover the basic Poisson process.

Simulation

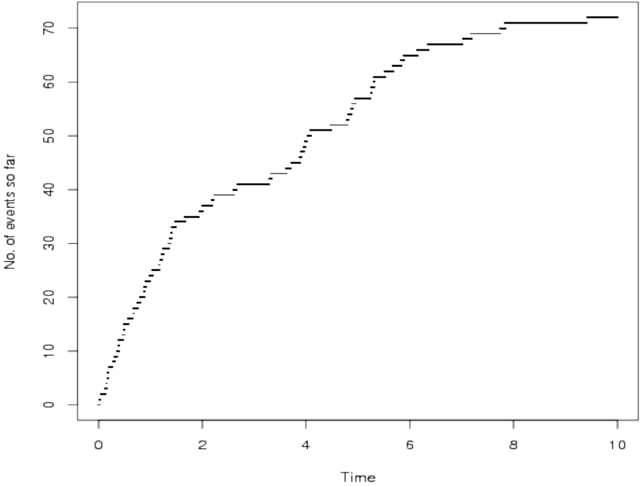

Independence

The independence assumption (b) still holds when the rate is variable.

This still works because of the additivity property of Poisson random variables, and also because of the additivity of integrals, namely that if $u < v < w$ then

$$\int_u^w \lambda(t)dt = \int_u^v \lambda(t)dt + \int_v^w \lambda(t)dt.$$

Example

Example

Email arrivals

Superposition, marking and thinning

Superposition

In some situations we have more than one Poisson process running.

If these are independent, then the process formed by combining them is also a Poisson process.

Theorem

Let $(N_{u,v})$ and $(M_{u,v})$ be independent Poisson processes with (possibly variable) rates $\lambda(t)$ and $\mu(t)$ respectively. Then $(N_{u,v} + M_{u,v})$ also forms a Poisson process, with rate $\lambda(t) + \mu(t)$.

Marking

Sometimes the occurrences in a Poisson process may be categorised as each belonging to one of a number of types.

This is sometimes referred to as **marking**: think of each occurrence as being given a random "mark".

Specifically we will assume that each occurrence in a Poisson process with (possibly variable) rate $\lambda(t)$ is given, independently of everything else, one of k different marks with probabilities $p_1, p_2, \ldots, p_k$ respectively.

Write the total number of occurrences in $(u, v]$ as $N_{u,v}$ (as before) and write the number of occurrences of type i in $(u, v]$ as $N_{u,v}^{(i)}$.

Lemma

The following result about Poisson random variables will be useful.

Lemma

Let X be a Poisson random variable with parameter μ, and imagine that, conditional on $X = x$, we have x objects each of which is of one of k types.
Assume further that each of these objects is of type i with probability p_i, independently of the other objects.
Let the number of objects of type i be Y_i.
The (unconditional) joint distribution of $Y_1, Y_2, \ldots, Y_k$ is such that they are independent Poisson, with parameters $p_1\mu, p_2\mu, \ldots$ and $p_k\mu$ respectively.

Theorem

Theorem

For each i, the process given by $(N^{(i)}_{u,v})$ (counting the occurrences which are type i) is a Poisson process with rate $\lambda(t)p_i$, and the k processes for the different types are independent of each other.

Proof.

This essentially follows from Lemma 21 together with the independence properties of Poisson processes. □

Variable rates

It is even possible to allow the probabilities of the marks to be dependent on time, say $p_1(t), p_2(t), \ldots, p_k(t)$.

Then the generalised result is that the marked processes are independent Poisson processes with variable rates $p_1(t)\lambda(t), p_2(t)\lambda(t), \ldots, p_k(t)\lambda(t)$ respectively.

Thinning

One special case of marking is where $k = 2$ and the process of marking consists of either retaining the occurrence, with probability p, or deleting it, with probability $q = 1 - p$.

Then the process of retained points is Poisson with rate $p\lambda$, and in this context the property is often known as the **thinning** property.

Example

Example

University applications

Conditioning on the number of occurrences in an interval

Conditioning

Sometimes we know how many occurrences there are in a given interval, and are interested in how they are distributed within the interval.

Theorem

Assume that we have a Poisson process with constant rate λ. Given that there are n occurrences in the time interval $(0, t]$ say, the positions of these occurrences are distributed as a random sample of size n from the uniform distribution on that interval.

Binomial

Note that this implies that, conditional on there being n occurrences in $(0, t]$, the number of occurrences in any interval $(u, v] \subseteq (0, t]$ (so $0 \leq u < v \leq t$) has a $Bi(n, (v-u)/t)$ distribution,. . .

. . . as each of the n occurrences would have probability $(v-u)/t$ of being in $(u, v]$, independently of the others.

We will prove this latter version of the statement.

Variable rate version

This result generalises to the variable rate case, but the uniform distribution is replaced by the distribution which has p.d.f.

$$f(s) = \frac{\lambda(s)}{\int_0^t \lambda(x)dx},$$

which is the distribution on $(0, t]$ whose density is proportional to the rate of the original process.

So the number of occurrences in any interval $(u, v] \subseteq (0, t]$ has a

$$Bi\left(n, \frac{\int_u^v \lambda(s)ds}{\int_0^t \lambda(s)ds}\right)$$

distribution.

Example

Example

Conditioning on number of events

Simulation I

We can reverse this idea to construct a Poisson process, for example for simulation purposes, or to convince ourselves that Poisson processes really exist.

Simulation II

Assuming that we want to construct a variable rate Poisson process with rate $\lambda(t)$, we can do the following:

- Divide the positive real line up into intervals $(n-1, n]$ for each positive integer n.
- To each of these intervals $(n-1, n]$ assign a Poisson random variable X_n with parameter $\int_{n-1}^{n} \lambda(t)\, dt$. These Poisson random variables should be independent of each other.

Simulation III

- If $X_n = 0$, then there will be no occurrences in the interval $(n-1, n]$; if $X_n = x > 0$, then we create a random sample of x random variables on $(n-1, n]$ with probability density function $f(s) = \frac{\lambda(s)}{\int_{n-1}^{n} \lambda(x)dx}$, in a similar manner to above. The values of these random variables will give the times of occurrences in the interval.

Simulation IV

The second step assumes the integral is finite; if for one of the intervals it is not we will need to be more careful.

It is not too hard to show, using the marking and additivity properties of the Poisson distribution, that a process of occurrences constructed in this way will satisfy the assumptions with which we defined the Poisson process.

The spatial Poisson process

The spatial Poisson process

An important generalisation of the basic Poisson process is to replace the time scale with a space, and the aim is to model a random scattering of points in this space.

The space may be one-dimensional – for example if we wish to consider defects on a length of cable – and in that case it looks like the time scale, . . .

. . . or it may be in a higher dimension – two dimensions for positions of spots of rain on a pavement, three dimensions for positions of stars in space, for example.

Generalising length

To generalise the assumptions we made for the basic Poisson process, we need an analogue of the length of a time interval.

The natural way to do this is to consider length in one dimension, area in two dimensions, volume in three dimensions, and so on.

We will refer to length, area or volume, as appropriate, as **measure**, and denote the measure of the set A by $|A|$.

Number of points in a subset

The behaviour of the process can be described by random variables $N(A)$ for subsets A with finite measure:

$N(A)$ represents the number of points of the process which fall inside the set A.

The parameter λ is in this context called the **density** of the process.

Definition

A **spatial Poisson process** is now defined to be a process which satisfies the following assumptions, which are generalisations of those we used in the time setting.

1. For any set A of finite measure, $N(A)$ has the Poisson distribution with parameter $\lambda|A|$.
2. If $A_1, A_2, \ldots, A_k$ are disjoint sets of finite measure, then $N(A_1), N(A_2), \ldots, N(A_k)$ are independent random variables.

Properties

The assumptions work for essentially the same reasons as before, notably the additivity of independent Poisson random variables.

Most of the properties of the basic Poisson process have analogues in more than one dimension.

Superposition

Two independent spatial Poisson processes with rates λ and μ can be combined to form a spatial Poisson process with rate $\lambda + \mu$.

Marking

If the points of a spatial Poisson process with rate λ are given independent marks (from $1, 2, \ldots, k$) with probabilities $p_1, p_2, \ldots, p_k$ then the points with mark i form a spatial Poisson process with rate λp_i, and the processes corresponding to the different marks are independent.

Conditioning

If we know that there are n points of the process in A, then the conditional distribution of the location of the points is that of a random sample of size n from the uniform distribution on A.

In particular, if $B \subseteq A$, then the number of points in B is Binomial with parameters n and $|B|/|A|$.

Proofs

The proofs of all of these properties are natural generalisations of the proofs of the one-dimensional versions.

It is also possible to simulate from a spatial Poisson process in a similar way to the one described for the variable rate time Poisson process.

Simulation

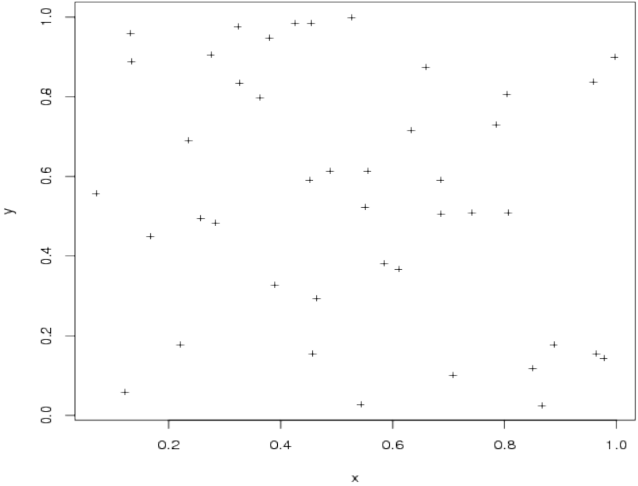

Inter-occurrence times?

One property of the basic Poisson process which does not naturally carry over is the joint distribution of inter-occurrence times.

There is no natural ordering of points in two or more dimensions, and so the analogue of inter-occurrence times does not exist.

However, it is possible to use a similar idea to calculate the distribution of the distance to the nearest point in the process from a given point.

Example

Example

Trees in a forest

Variable rate

It is also possible to define variable rate spatial Poisson processes.

The parameter of the Poisson distribution giving the number of points in a set A will be the integral over A of the rate function, just as for the one-dimensional case, but the integral is now a multidimensional one.

Clustering and regularity

An actual scatter of points may be **clustered** relative to the true randomness of the Poisson process –

for example, positions of plants each of which self-propagates within a local area –

or **regular** relative to the Poisson process –

for example, positions of birds' nests when there is a territorial effect inhibiting nests from being too close together.

Regularity

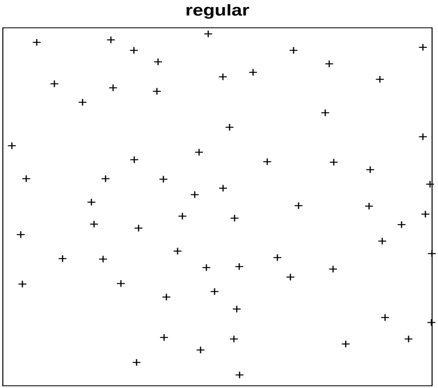

Clustering

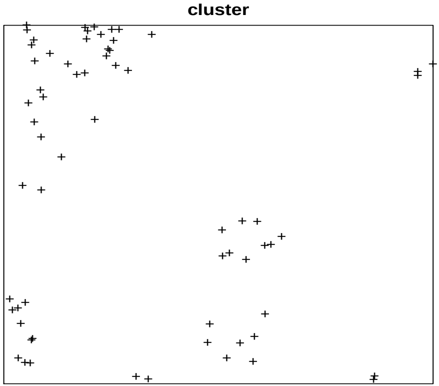

Compound Poisson processes*

Compound Poisson process

Suppose that events occur at random times but also each event carries with it some numerical value, and the chief interest is in the sum of these numerical values over a period of time.

Examples might be

- claims on an insurance company, which occur at random times but they differ in size
- fatalities in road accidents, where the accidents occur at random times but each accident may incur a number of deaths.

Model

The simplest model for such situations is to take the times of occurrences as a basic Poisson process and . . .

. . . then to assume that the sizes of the occurrences are random variables each with some known distribution, which may be discrete or continuous.

These random variables are independent of each other and of the times of the occurrences

This gives what is known as a **compound Poisson process**.

Notation

We will use the following notation.

$N(t)$ denotes the number of occurrences in the time interval $(0, t]$, previously written as $N_{0,t}$.

The sizes of the occurrences in chronological order are denoted by $Y_1, Y_2, \ldots$.

Then the sum of these over the time interval $(0, t]$ may be written

$$X(t) = \sum_{i=1}^{N(t)} Y_i.$$

Graphs

If we think graphically of plotting $N(t)$ and $X(t)$ against t, then

- the graph of $N(t)$ jumps upwards by 1 at each occurrence and stays constant in between
- the graph of $X(t)$ jumps (upwards or downwards, since Y_i could be negative) by the random quantities $Y_1, Y_2, \ldots$ at the times of the occurrences, and stays constant in between.

The process is completely specified by the rate λ of the Poisson process and the common distribution of the random variables $Y_1, Y_2, \ldots$.

Simulation

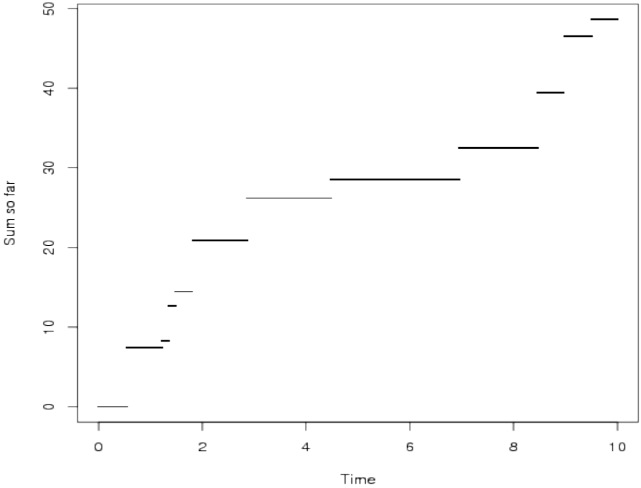

Further generalisations

We can combine the ideas of compound Poisson processes with variable rate and spatial Poisson processes as well.

For example, we might be modelling the locations of nests of some species of bird within some region.

We could treat the locations as points in two dimensional space to be modelled by a spatial Poisson process.

Further generalisations II

If some parts of the region are more favourable to the species than others, then we would expect a higher density of nests in these areas, so the model would have a variable rate which is higher in favourable areas and lower elsewhere.

If we wanted to model the total number of offspring raised, then a compound spatial Poisson model might be appropriate.