Poisson Point Measures & Stochastic Geometry

Nicolas CURIEN

September 16, 2026

Contents

I Poisson Point Measures **5**

1 Construction of PPMs **6**
1.1 Review of Poisson random variables . . . 6
1.1.1 Definition and Laplace transform . . . 6
1.1.2 Infinite divisibility . . . 7
1.2 Poisson point measures . . . 9
1.2.1 Point measures . . . 9
1.2.2 Construction of PPM . . . 11
1.3 Basic operations . . . 14
1.4 Recovering the points . . . 15

2 Useful formulas **19**
2.1 Exponential formula . . . 19
2.2 The Mecke formula . . . 22
2.2.1 Multi-point Mecke . . . 25
2.3 Exercises . . . 25

3 Random measures on Polish spaces **27**
3.1 Radon measures . . . 27
3.1.1 Weak and vague convergence . . . 27
3.1.2 Random measures . . . 28
3.2 Point measures . . . 30
3.2.1 Convergence of atoms . . . 30
3.2.2 An example . . . 31

II Poissonization 33

4 Temporal Poisson processes 34
4.1 Poisson counting process . . . 34
4.2 Direct applications . . . 38
4.2.1 Order statistics . . . 38
4.2.2 Markov properties . . . 39

5 Inhomogeneous PPMs 41
5.1 Inhomogeneous Poisson point measures . . . 41
5.1.1 The Gumbel cloud . . . 41
5.1.2 The polynomial cloud . . . 42
5.1.3 The birthday process . . . 44
5.2 Hawkes process . . . 46

6 Poissonization 49
6.1 Choosing with clocks . . . 49
6.1.1 Properties of exponential laws . . . 49
6.1.2 Two de-Poissonization tools . . . 50
6.2 Examples without branching . . . 51
6.2.1 A matchbox warm-up . . . 51
6.2.2 Coupon collector . . . 53
6.2.3 The coupon collector's earworm . . . 54
6.2.4 Balls in bins . . . 55
6.2.5 Pill problem . . . 56
6.2.6 O.K. Corral . . . 57
6.2.7 Kingman's coalescent: splitting at the root . . . 60
6.3 Branching: Yule and binary search trees . . . 61
6.3.1 From the discrete binary search tree to the Yule tree . . . 62
6.3.2 Pólya urn and splitting at the root . . . 65

III Stochastic Geometry 67

7 Models based on points 68
7.1 Basic properties and intensities . . . 68
7.1.1 Intensities . . . 70
7.2 Neighbor models . . . 71
7.2.1 Boolean percolation . . . 72

7.2.2 The k-nearest-neighbor graph 74
7.3 Poisson–Voronoi tessellations 76
7.3.1 Definition and elementary geometry 76
7.3.2 The Delaunay dual 78
7.3.3 Stationarity and scaling 79
7.3.4 The typical Poisson–Voronoi cell 79
7.3.5 The zero cell 80
7.3.6 Planar mean values 81
7.3.7 Power and Laguerre tessellations 82
7.4 Random parking from a Poisson rain 83

8 Models based on lines **86**
8.1 The kinematic measure 86
8.1.1 Parametrizing the set of lines 86
8.1.2 The kinematic measure 88
8.2 Cauchy–Crofton formulas 90
8.2.1 The curve formula 90
8.2.2 Cauchy's perimeter formula 91
8.2.3 Integration of chord lengths 92
8.3 Poisson line tessellation 92
8.3.1 General position and the induced tessellation 93
8.3.2 Mean geometric characteristics 96
8.3.3 The zero cell 98
8.3.4 Extension to higher dimensions 99
8.4 Bibliographical remarks 99

Introduction

Since Kolmogorov's axiomatization, probability theory can, in principle, be regarded as a branch of analysis: it amounts to studying probability measures, while random variables are merely measurable functions. Yet every probabilist knows that probabilistic intuition becomes much more powerful once we stop thinking of a random variable as a deterministic function and instead regard it as a genuinely random object. Structural properties of a probability measure μ often reveal themselves through the random behavior of a variable X with law μ, and the probabilistic language makes it possible to uncover and express deeper properties. As a simple illustration, try writing the law of large numbers purely in analytic language!

When μ is an infinite measure, however, this correspondence with random variables seems to break down: there is no way to normalize μ and sample a single point according to it. Yet infinite measures are often more natural than finite ones, precisely because they enjoy stronger symmetry, invariance, or scaling properties. Think, for instance, of Lebesgue measure on $\mathbb{R}^d$. Poisson random measures provide a beautiful way to restore a probabilistic interpretation of such measures and, in a sense, to sample from an infinite measure μ. The price to pay is that one no longer samples a single point, but rather a random collection—a cloud—of points. Developed in its modern general form during the 1950s and 1960s, this remarkable object now permeates much of modern probability, from stochastic geometry and stochastic processes to stochastic analysis and statistical physics.

In these lecture notes, intended as an advanced course in master, we introduce the basic theory of Poisson random measures, their construction, and some of their structural properties, with particular emphasis on time-layered spaces. This provides the foundations for the Poissonization trick, a powerful method that yields quick and elegant proofs of discrete probability results by embedding discrete models into continuous time. In the second part, we turn to stochastic geometry and present several beautiful constructions based on Poisson processes of points and lines in $\mathbb{R}^d$.

Part I:

Poisson Point Measures

Given a measure μ on a space E, one would usually like to sample a *random element* $X \in E$ according to μ. This is obviously possible if μ is a probability measure or, more generally, if μ has finite mass, so that we can consider the normalized probability measure $\mu(\cdot)/\mu(E)$. This seems hopeless when μ is genuinely infinite, as is, for example, Lebesgue measure on Euclidean spaces. The concept of a Poisson random cloud provides a way to sample not just a single point $X \in E$ but also a random collection of points $X_i \in E$ for $i \in I$ in a uniform manner even in the presence of an infinite measure. This beautiful object, called a **Poisson point measure**, is a central tool in modern probability theory. Its construction relies on the infinite divisibility properties of the Poisson law, which we first recall.

Figure 1: An AI-generated image of a Poisson point measure

I: Construction of PPMs

Remark 1.1 (Terminology). In the literature, Poisson point measures are also called Poisson point processes, Poisson random measures, or Poisson point clouds. The most common term is perhaps "Poisson point process", but the word *process* can be misleading since it may suggest a time parametrization (which shall appear in Chapter 4). We shall mostly use *Poisson point measure* (PPM), which emphasizes the measure-valued nature of the object.

1.1 Review of Poisson random variables

In this section we recall the basic algebra of Poisson[1] random variables.

1.1.1 Definition and Laplace transform

Definition 1.1. *Let $\lambda \geqslant 0$. A random variable N, with values in $\mathbb{Z}_{\geqslant 0} = \{0,1,2,\dots\}$, is said to have the* Poisson distribution with parameter λ *if*

$$\mathbb{P}(N = k) = \mathrm{e}^{-\lambda}\frac{\lambda^k}{k!}, \qquad k \geqslant 0.$$

We write $N \sim \mathrm{Poi}(\lambda)$.

When $\lambda = 0$, this means that $N = 0$ almost surely. We also use the convention $\mathrm{Poi}(\infty) = \delta_{+\infty}$. The **Laplace transform** is given by straightforward computation: for $N \sim \mathrm{Poi}(\lambda)$ and $t \geqslant 0$

$$\begin{aligned}\mathbb{E}[\mathrm{e}^{-tN}] &= \mathrm{e}^{-\lambda}\sum_{k\geqslant 0}\frac{\lambda^k \mathrm{e}^{-tk}}{k!} \\ &= \mathrm{e}^{-\lambda+\lambda\mathrm{e}^{-t}} = \exp\bigl(-\lambda(1-\mathrm{e}^{-t})\bigr). \end{aligned} \tag{1.1}$$

[1] Siméon-Denis Poisson (1781–1840), French mathematician. The Poisson law appeared in his 1837 book *Recherches sur la probabilité des jugements en matière criminelle et en matière civile*, in connection with a limiting regime for binomial probabilities.

Differentiating this identity at $t = 0$ gives the familiar formulas $\mathbb{E}[N] = \lambda$ and $\mathrm{Var}(N) = \lambda$. More generally, for every integer $r \geqslant 1$, the factorial moments (see Exercise 1.1) of the Poisson distribution are given by

$$\mathbb{E}[N(N-1)\cdots(N-r+1)] = \lambda^r.$$

It is easy to see that the Laplace transform $t \mapsto \mathbb{E}[\mathrm{e}^{-tN}]$ characterize the law of random variables N with values in $\{0,1,2,\ldots\} \cup \{\infty\}$, where we use the convention that $0 \cdot \infty = 0$. In connection with this convention, let us collect here a lemma regarding Laplace transform and possibly infinite random variables:

Lemma 1.1. *If $X \in [0,\infty]$ a non-negative possibly infinite random variable and $\phi(t) = \mathbb{E}[\mathrm{e}^{-tX}]$ for $t \geqslant 0$ is its Laplace transform. Then we have*

$$\mathbb{P}(X < \infty) = \lim_{t \downarrow 0} \phi(t).$$

Proof. For $x \in [0,\infty]$ we have $\mathrm{e}^{-tx} \to \mathbf{1}_{x<\infty}$ as $t \downarrow 0$ and so the the statement follows by dominated convergence theorem. □

1.1.2 Infinite divisibility

Proposition 1.2 (Stability under addition). *Let $N_1, N_2, \ldots$ be independent Poisson random variables with respective means $\lambda_1, \lambda_2, \ldots \in [0,\infty]$. Then*

$$N = \sum_{i=1}^{\infty} N_i \sim \mathrm{Poi}\left(\sum_{i=1}^{\infty} \lambda_i\right).$$

Proof. Let $\Lambda = \sum \lambda_i \in [0,\infty]$. We have the monotone almost sure convergence $\sum_{i=1}^{n} N_i \to N$ and so by the monotone convergence theorem we have for $s = \mathrm{e}^{-t} \in [0,1]$

$$\mathbb{E}[s^N] = \lim_{n\to\infty} \mathbb{E}\left[s^{N_1+\cdots+N_n}\right] = \lim_{n\to\infty} \prod_{i=1}^{n} \mathbb{E}[s^{N_i}] = \exp\left(-\Lambda(1-s)\right).$$

The right-hand side isthe Laplace transform (or probability-generating function) of a Poisson variable with mean Λ. Since the Laplace transform characterize the law of variables with values in $\{0,1,2,\ldots\} \cup \{\infty\}$ we are done. □

We also have the reverse operation, which allows us to split a Poisson random variable.

Proposition 1.3 (Splitting property). *Fix $\lambda \in [0,\infty)$. Let $N \sim \mathrm{Poi}(\lambda)$ and let $p_1, p_2, \ldots \geqslant 0$ satisfy $\sum_{i\geqslant 1} p_i = 1$. Conditionally on N, assign each of the N objects independently to category i with probability p_i. Let N_i be the number assigned to category i. Then the variables $N_1, N_2, \ldots$ are independent and*

$$N_i \sim \mathrm{Poi}(\lambda p_i).$$

Proof. It is enough to prove the result for two categories, assigned with probabilities p and $1-p$. The general case follows by iteration: for every $n \geqslant 1$, the $n+1$ variables

$$N_1, \dots, N_n, \sum_{j>n} N_j$$

are independent, with respective Poisson parameters $\lambda p_1, \dots, \lambda p_n, \lambda \sum_{j>n} p_j$. This determines all finite-dimensional distributions of $(N_i)_{i \geqslant 1}$.

For two categories, conditionally on N the vector (N_1, N_2) is binomially split, hence

$$\begin{aligned}
\mathbb{P}(N_1 = k_1, N_2 = k_2) &= \mathbb{P}(N = k_1 + k_2) \binom{k_1 + k_2}{k_1} p^{k_1} (1-p)^{k_2} \\
&= \mathrm{e}^{-\lambda} \frac{\lambda^{k_1+k_2}}{(k_1+k_2)!} \frac{(k_1+k_2)!}{k_1! k_2!} p^{k_1} (1-p)^{k_2} \\
&= \mathrm{e}^{-\lambda p} \frac{(\lambda p)^{k_1}}{k_1!} \cdot \mathrm{e}^{-\lambda(1-p)} \frac{(\lambda(1-p))^{k_2}}{k_2!}.
\end{aligned}$$

Thus N_1 and N_2 are independent Poisson variables with respective means λp and $\lambda(1-p)$. □

Exercise 1.1 (Factorial moments and the Poisson law). Let N be a non-negative integer-valued random variable. For $k \geqslant 0$, provided that $\mathbb{E}[N^k] < \infty$, the kth factorial moment of N is defined by

$$\mathbb{E}[(N)_k] = \mathbb{E}[N(N-1)\cdots(N-k+1)],$$

with the convention $(N)_0 = 1$.

1. Let $N \sim \mathrm{Poi}(\lambda)$, with $\lambda > 0$. Compute the factorial moments of N and show that the Poisson distribution is determined by its factorial moments.

   Hint: Prove that whenever the series is absolutely convergent we have

$$\mathbb{P}(N = r) = \frac{1}{r!} \sum_{j=0}^{\infty} \frac{(-1)^j}{j!} \mathbb{E}[(N)_{r+j}],$$

2. Suppose that

$$N = \sum_{i \in I} \mathbf{1}_{A_i},$$

   where I is finite or countable, $N < \infty$ almost surely, and $\mathbb{E}[N^k] < \infty$. Show that

$$\mathbb{E}[(N)_k] = \sum_{\substack{i_1, \dots, i_k \in I \\ \text{pairwise distinct}}} \mathbb{P}(A_{i_1} \cap \cdots \cap A_{i_k}).$$

Hint: Prove that

$$\left(\sum_{i\in I} y_i\right)_k = \sum_{\substack{i_1,\dots,i_k\in I\\ \text{pairwise distinct}}} y_{i_1}\cdots y_{i_k}, \qquad y_i \in \{0,1\}.$$

3. Use the previous questions and the method of moments to prove the following:

   (a) For every $\lambda > 0$,
   $$\mathrm{Bin}\left(n, \frac{\lambda}{n}\right) \xrightarrow[n\to\infty]{\text{law}} \mathrm{Poi}(\lambda).$$

   (b) Let σ_n be uniformly distributed on the symmetric group $\mathfrak{S}_n$. Then the number of fixed points of σ_n converges in distribution to $\mathrm{Poi}(1)$ as $n \to \infty$.

   (c) (*) Consider the complete graph on the vertex set $\{1,\dots,n\}$. Thus every pair of distinct vertices is joined by an edge, and the graph has $n(n-1)/2$ edges. For $p \in [0,1]$, let $G(n,p)$ be the random graph obtained by retaining each edge independently with probability p. Let $T(n,p)$ be the number of triangles in $G(n,p)$. Show that, for every $\lambda > 0$,
   $$T\left(n, \frac{\lambda}{n}\right) \xrightarrow[n\to\infty]{\text{law}} \mathrm{Poi}\left(\frac{\lambda^3}{6}\right).$$

1.2 Poisson point measures

Let $(E,\mathcal{E})$ be a measurable space; no topology is needed. A natural way to describe a countable collection of points in E would be to index the points as $X_1, X_2, \dots$. Although this concrete approach works in many situations, it may be cumbersome on an abstract space, and repetitions must be allowed. A more symmetric point of view is to encode the collection as a counting measure.

1.2.1 Point measures

For $x \in E$, the Dirac measure at x is denoted by δ_x:

$$\delta_x(A) = \mathbf{1}_A(x), \qquad A \in \mathcal{E}.$$

Definition 1.2. *A* point measure *(or* counting measure*) on E is a measure of the form*

$$\nu = \sum_{i\in I} \delta_{x_i},$$

where I is finite or countable and $x_i \in E$. Repetitions among the x_i are allowed. The zero measure corresponds to $I = \varnothing$. We denote the set of point measures on E by $\mathcal{N}(E)$.

The same point may occur several times. For example, $3\delta_x + \delta_y$ assigns mass 3 to x and mass 1 to y. In an abstract measurable space it is preferable to allow multiplicities from the start.

We endow $\mathcal{N}(E)$ with the smallest σ-field $\mathfrak{N}$ making all evaluation maps

$$\nu \longmapsto \nu(A), \qquad A \in \mathcal{E},$$

measurable, where the range is $\mathbb{Z}_{\geqslant 0} \cup \{\infty\}$. A *random point measure* is then a measurable map

$$\mathbf{N} : (\Omega, \mathcal{F}, \mathbb{P}) \longrightarrow (\mathcal{N}(E), \mathfrak{N}). \tag{1.2}$$

Equivalently, $\mathbf{N}$ is a random point measure if it is a point measure and $\mathbf{N}(A)$ is a random variable for every $A \in \mathcal{E}$. Since $\mathfrak{N}$ is generated by the evaluation maps, the law of $\mathbf{N}$ is determined by the finite-dimensional laws of the family $(\mathbf{N}(A))_{A \in \mathcal{E}}$. The **mean intensity** of a random point measure is the measure (check it as an exercise) on $(E, \mathcal{E})$ defined by

$$\mu(A) = \mathbb{E}[\mathbf{N}(A)], \quad A \in \mathcal{E}.$$

Definition 1.3. *A* Poisson point measure *on $(E, \mathcal{E})$ with intensity μ is a random point measure Π such that*

(i) *for every $A \in \mathcal{E}$,*

$$\Pi(A) \sim \mathrm{Poi}(\mu(A));$$

(ii) *whenever $A_1, \ldots, A_r \in \mathcal{E}$ are pairwise disjoint, the random variables $\Pi(A_1), \ldots, \Pi(A_r)$ are independent.*

We write $\Pi \sim \mathrm{PPM}(E, \mathcal{E}, \mu)$, or simply $\mathrm{PPM}(\mu)$ when the underlying space is clear.

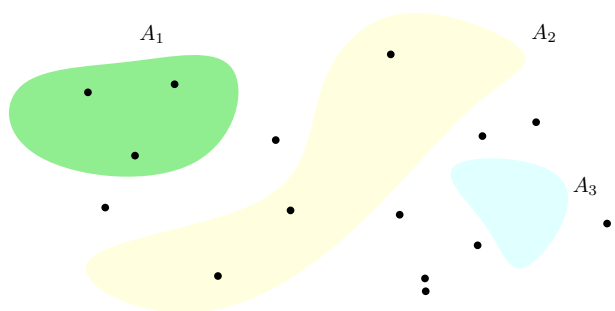

Figure 1.1: Illustration of the definition of a PPM: the numbers of atoms falling in disjoint regions are independent Poisson random variables.

A small point deserves to be checked: properties (i)–(ii) indeed determine the law among random point measures. For arbitrary $A_1, \dots, A_r \in \mathcal{E}$, let $B_1, \dots, B_m$ be the non-empty cells of the finite measurable partition generated by $A_1, \dots, A_r$. Each $\mathbf{N}(A_i)$ is a sum of some of the variables $\mathbf{N}(B_j)$. Since the B_j are pairwise disjoint, Definition 1.3 prescribes the joint law of $(\mathbf{N}(B_1), \dots, \mathbf{N}(B_m))$, and hence that of $(\mathbf{N}(A_1), \dots, \mathbf{N}(A_r))$. Thus all finite-dimensional laws of the evaluation family are determined.

Example 1.1. If μ has an atom at $x \in E$ then $\Pi(\{x\})$ is a non-trivial Poisson random variable, that is we may have repetitions in the atoms of a PPM.

Remark 1.2 (A subtlety). Let us stress a subtle point in Definition 1.3. A PPM is first of all a *random point measure* Π; its evaluations $(\Pi(A) : A \in \mathcal{E})$ are then required to have the Poisson finite-dimensional laws. A family of random variables $(\Pi_A : A \in \mathcal{E})$ with the same finite-dimensional laws does not by itself produce a random measure. Indeed, on a general measurable space, Kolmogorov's extension theorem ensures gives a family $(\Pi_A : A \in \mathcal{E})$ with the prescribed consistent finite-dimensional distributions. For each fixed disjoint sequence $A_1, A_2, \dots$, one then have

$$\sum_{i \geqslant 1} \Pi_{A_i} = \Pi_{\bigsqcup_{i \geqslant 1} A_i} \qquad \text{almost surely,}$$

but the exceptional null set can depend on the chosen sequence. This does not yield, on one common event of probability one, a countably additive map $A \mapsto \Pi_A$ simultaneously for all sequences. The point-measure requirement in the definition therefore carries a true content.

We can now state the main existence result of this chapter.

Theorem 1.4 (*Existence of PPMs*)

If μ is a σ-finite measure on a measurable space $(E, \mathcal{E})$, then there exists a Poisson point measure on E with intensity μ.

Example 1.2. A basic example is a Poisson point measure on $\mathbb{R}^d$ with Lebesgue intensity.

Remark 1.3 (History). Although the Poisson distribution dates back to the nineteenth century, Poisson random measures became a systematic tool in probability in the mid-twentieth century. They play a central role in the theory of Lévy processes and, through the work of Itô, in modern stochastic calculus.[2] They are also fundamental objects in stochastic geometry.

1.2.2 Construction of PPM

We prove Theorem 1.4 in two steps.

[2] 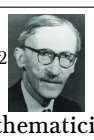 Paul Lévy (1886–1971), French mathematician. Kiyoshi Itô (1915–2008), Japanese mathematician.

Finite intensity. Assume first that $\mu(E) < \infty$. If $\mu(E) = 0$, simply take $\Pi = 0$. Otherwise, let $(X_i)_{i\geqslant 1}$ be i.i.d. E-valued random variables[3] with law $\mu/\mu(E)$, and let $N \sim \mathrm{Poi}(\mu(E))$ be independent of the sequence $(X_i)_{i\geqslant 1}$.

Define

$$\Pi = \sum_{i=1}^{N} \delta_{X_i}.$$

This is a random point measure, since the measurability of the X_i's and N ensure that events of the form $\{\Pi(A) = k\}$ for $A \in \mathcal{E}$ and $k \in \mathbb{Z}_{\geqslant 0}$ are measurable. If $A_1, \dots, A_r$ are pairwise disjoint measurable sets, add $A_0 = E \setminus \bigcup_{j=1}^{r} A_j$ and apply the splitting property (Proposition 1.3) to the measurable partition $A_0, A_1, \dots, A_r$. We obtain that $\Pi(A_1), \dots, \Pi(A_r)$ are independent and

$$\Pi(A_i) \sim \mathrm{Poi}\left(\mu(E)\frac{\mu(A_i)}{\mu(E)}\right) = \mathrm{Poi}(\mu(A_i)).$$

Thus Π is a PPM with intensity μ.

σ-finite intensity. For a σ-finite intensity, partition $E = \bigsqcup_{n\geqslant 1} E_n$ with $\mu(E_n) < \infty$, construct independent PPMs (again using the countable product construction) Π_n with intensities $\mu(\,\cdot \cap E_n)$, and set

$$\Pi = \sum_{n\geqslant 1} \Pi_n.$$

The defining Poisson laws and independence follow immediately from the finite case and Proposition 1.2. □

Corollary 1.5 (Avoidance probability). *If* $\Pi \sim \mathrm{PPM}(E, \mathcal{E}, \mu)$ *then for every* $A \in \mathcal{E}$,

$$\mathbb{P}(\Pi(A) = 0) = \mathrm{e}^{-\mu(A)},$$

where the right-hand side is interpreted as 0 *when* $\mu(A) = \infty$.

The following exercise gives a useful converse to the defining properties of a PPM: for a simple locally finite random point measure on $\mathbb{R}_+$, the void probabilities already determine the law.

Exercise 1.2 (Rényi's characterization of Poisson point measures). Let Π be an almost surely simple point measure

$$\Pi = \sum_{i\in I} \delta_{x_i}$$

[3]The required probability space can be obtained from the standard countable product-measure theorem, applied to $E^{\mathbb{Z}_{>0}}$ with its product σ-field, and then taking a product with a space carrying the Poisson variable N. The existence of such a probability space is in fact the most complicated part of the proof, but usually the most overlooked. However, this construction is done in undergraduate probability courses.

on $\mathbb{R}_+$, where I is finite or countable and the atoms x_i are pairwise distinct. Assume that Π is almost surely finite on compact sets and that, for every Borel set $A \subseteq \mathbb{R}_+$,

$$\mathbb{P}\big(\Pi(A) = 0\big) = \mathrm{e}^{-\mu(A)},$$

where μ is a diffuse Borel measure that is finite on compact sets (with the convention $\mathrm{e}^{-\infty} = 0$). The goal is to prove that

$$\Pi \sim \mathrm{PPM}(\mathbb{R}_+, \mathcal{B}(\mathbb{R}_+), \mu).$$

1. Let $I_1, \dots, I_m$ be pairwise disjoint Borel sets. Show that the events

$$\{\Pi(I_j) = 0\}, \qquad 1 \leqslant j \leqslant m,$$

are independent.

2. Fix $0 \leqslant a < b < \infty$, and for $0 \leqslant k \leqslant 2^n - 1$ set

$$D_{n,k} = \left[a + \frac{k}{2^n}(b-a),\ a + \frac{k+1}{2^n}(b-a)\right).$$

Show that, almost surely,

$$\Pi([a,b)) = \lim_{n\to\infty} \sum_{k=0}^{2^n-1} \mathbf{1}_{\{\Pi(D_{n,k}) \neq 0\}}.$$

3. Deduce that, for every $z \in \mathbb{C}$ with $|z| < 1$,

$$\mathbb{E}\left[z^{\Pi([a,b))}\right] = \lim_{n\to\infty} \prod_{k=0}^{2^n-1} \left(z + (1-z)\mathrm{e}^{-\mu(D_{n,k})}\right).$$

4. Show that

$$\max_{0 \leqslant k \leqslant 2^n-1} \mu(D_{n,k}) \xrightarrow[n\to\infty]{} 0,$$

and use the product in the previous question to determine the law of $\Pi([a,b))$.

5. Show that, for pairwise disjoint half-open intervals $[a_i, b_i)$, the random variables

$$\Pi([a_i, b_i))$$

are independent.

6. Deduce that Π is a Poisson point measure on $\mathbb{R}_+$ with intensity μ.

1.3 Basic operations

The stability properties of Poisson random variables under addition and splitting translate into analogous properties for Poisson point measures. The proofs are left as exercises.

Proposition 1.6 (Additivity). *Let $\Pi_1, \Pi_2, \ldots$ be independent PPMs on $(E, \mathcal{E})$ with respective intensities $\mu_1, \mu_2, \ldots$. Then*

$$\Pi = \sum_{n \geqslant 1} \Pi_n$$

is a Poisson point measure with intensity $\mu = \sum_{n \geqslant 1} \mu_n$.

Proposition 1.7 (Restriction). *Let $\Pi \sim \mathrm{PPM}(E, \mathcal{E}, \mu)$ and let $B \in \mathcal{E}$. Then*

$$\Pi_B = \Pi(\,\cdot\, \cap B), \qquad \Pi_{B^c} = \Pi(\,\cdot\, \cap B^c)$$

are independent PPMs with respective intensities $\mu(\,\cdot\, \cap B)$ and $\mu(\,\cdot\, \cap B^c)$.

Proposition 1.8 (Mapping). *Let $\Pi \sim \mathrm{PPM}(E, \mathcal{E}, \mu)$ and let $\phi : (E, \mathcal{E}) \to (F, \mathcal{F})$ be measurable. Define, for $B \in \mathcal{F}$,*

$$(\phi_\# \Pi)(B) = \Pi(\phi^{-1}(B)), \qquad (\phi_\# \mu)(B) = \mu(\phi^{-1}(B)).$$

Then $\phi_\# \Pi$ is a PPM on $(F, \mathcal{F})$ with intensity $\phi_\# \mu$.

Even when μ is σ-finite, its image $\phi_\# \mu$ need not be σ-finite. For example, the projection $\phi : \mathbb{R}^2 \to \mathbb{R}$, $\phi(x, y) = x$, sends two-dimensional Lebesgue measure to the measure

$$(\phi_\# \mathrm{Leb}_2)(A) = \begin{cases} 0, & \mathrm{Leb}_1(A) = 0, \\ +\infty, & \mathrm{Leb}_1(A) > 0, \end{cases}$$

which is not σ-finite. Thus the mapping proposition also produces a PPM on $\mathbb{R}$ with this non-σ-finite intensity.

Remark 1.4 (σ-finite and s-finite measures). A measure which is a countable sum of finite measures is called *s-finite*. This condition is strictly weaker than σ-finiteness and is preserved under measurable mappings. The construction above extends to s-finite intensities by superposing independent finite-intensity PPMs. We shall nevertheless mostly assume σ-finiteness, since it is the standard hypothesis and it simplifies several later arguments, including Theorem 1.11.

Proposition 1.9 (Marking). *Let $\Pi \sim \mathrm{PPM}(E, \mathcal{E}, \mu)$ and suppose that we are given a measurable enumeration*

$$\Pi = \sum_{i=1}^{N} \delta_{X_i}, \qquad N \in \mathbb{Z}_{\geqslant 0} \cup \{\infty\}.$$

Let $(Y_i)_{i\geqslant 1}$ be i.i.d. random variables with values in $(F,\mathcal{F})$ and common law ν, independent of N and of the sequence $(X_i)_{i\geqslant 1}$. Then

$$\Pi^F := \sum_{i=1}^{N} \delta_{(X_i,Y_i)} \sim \mathrm{PPM}(E\times F, \mathcal{E}\otimes\mathcal{F}, \mu\otimes\nu).$$

The proof of the above proposition is slightly more complicated than the additivity, mapping and restriction properties. We will see a slick proof using the exponential formula in the next chapter, but in the mean time we encourage the reader to find their own proof.

1.4 Recovering the points ☠

Theorem 1.4 is very general: it only assumes σ-finiteness of the intensity. On an arbitrary measurable space, however, it is not automatic that the atoms of a PPM can be recovered as measurable random variables. The following example shows why an additional measurability assumption is needed. Let $E = [0,1]$ and let $\mathcal{E}$ be the countable–cocountable σ-field. Define

$$\mu(A) = \begin{cases} 0, & A \text{ countable}, \\ 1, & A^c \text{ countable}. \end{cases}$$

Let $(X_i)_{i\geqslant 1}$ be i.i.d. E-valued variables with law μ, let $N \sim \mathrm{Poi}(1)$ be independent of them, and set

$$\Pi = \sum_{i=1}^{N} \delta_{X_i}.$$

Also define the random measure $\widetilde{\Pi} = N\mu$, i.e. $\widetilde{\Pi}(A) = N\mu(A)$. The random measures Π and $\widetilde{\Pi}$ have the same finite-dimensional laws for all their evaluations on $\mathcal{E}$, and these laws satisfy the Poisson count and independence formulas. However, Π is a point measure whereas, on $\{N > 0\}$, $\widetilde{\Pi}$ is not a point measure. Thus the Poisson formulas alone do not characterize a point measure among all random measures; the point-measure requirement in Definition 1.3 is genuine.

To recover individual points we shall use the following assumption:

$$(\Delta) \qquad \Delta_E := \{(x,x) : x \in E\} \in \mathcal{E}\otimes\mathcal{E}. \tag{1.3}$$

If $A \in \mathcal{E}\otimes\mathcal{E}$, then every section $A_x = \{y : (x,y) \in A\}$ belongs to $\mathcal{E}$. In particular, under (Δ) every singleton $\{x\}$ is measurable. The assumption holds in all standard Borel spaces, and in particular for the Borel σ-field of a Polish space.

Exercise 1.3. Show that the countable–cocountable σ-field in the preceding example does not satisfy (Δ).

A useful equivalent formulation of (Δ) is the existence of a countable family $(C_n)_{n\geqslant 1} \subset \mathcal{E}$ separating points[4]. Indeed, if such a family exists, then

$$\Delta_E = \bigcap_{n\geqslant 1} \big((C_n \times C_n) \cup (C_n^c \times C_n^c)\big).$$

Conversely, every element of the σ-field generated by measurable rectangles belongs to the σ-field generated by countably many of those rectangles. Apply this to Δ_E: the countably many sides of these rectangles must separate points, otherwise (x,x) and (x,y) could not be distinguished for some $x \neq y$.

We start with a simpler question than recovering the atoms, namely deciding whether two atoms coincide. By Example 1.1, this is the case with positive probability as soon as μ has an atom, the following proposition proves the converse:

Proposition 1.10 (Simplicity). *Suppose that Π is a PPM on $(E,\mathcal{E})$ satisfying (Δ) with a σ-finite intensity μ. If μ is diffuse, then Π is almost surely simple.*

Proof. We first treat the case $\mu(E) < \infty$. Let $(C_j)_{j\geqslant 1}$ be a countable family separating points, and let $\mathcal{P}_n$ be the finite partition generated by $C_1, \dots, C_n$. We claim that

$$\delta_n := \max_{D\in\mathcal{P}_n} \mu(D) \xrightarrow[n\to\infty]{} 0. \tag{1.4}$$

If not, there exsits $\varepsilon > 0$ and cells of arbitrarily fine partitions with mass at least ε. Since the partitions are finite and refining, König's lemma[5] yields a nested sequence

$$D_1 \supset D_2 \supset \cdots, \qquad D_n \in \mathcal{P}_n, \qquad \mu(D_n) \geqslant \varepsilon.$$

Because (C_j) separates points, $\bigcap_n D_n$ contains at most one point. Continuity from above, using $\mu(D_1) < \infty$, gives

$$\mu\left(\bigcap_{n\geqslant 1} D_n\right) = \lim_{n\to\infty} \mu(D_n) \geqslant \varepsilon,$$

contradicting diffuseness. Thus $\delta_n \to 0$.

Now let

$$A_n = \{\exists D \in \mathcal{P}_n : \ \Pi(D) \geqslant 2\},$$

and notice that $\{\Pi \text{ has an atom of multiplicity at least two}\} \subset \cap_n A_n$ so it suffices to prove that $\mathbb{P}(A_n) \to 0$ as $n \to \infty$. Before doing that by a variation of the first-moment method, if $N \sim \mathrm{Poi}(\lambda)$, then

$$\mathbb{P}(N \geqslant 2) \leqslant \frac{1}{2}\,\mathbb{E}[N(N-1)] = \frac{\lambda^2}{2}.$$

[4]Equivalently, $(E,\mathcal{E})$ admits an injective measurable map into [0,1]. This is strictly weaker than being a standard Borel space.

[5]if you have never seen this lemma, make the proof yourself!

Consequently,

$$\begin{aligned}\mathbb{P}(A_n) &\leqslant \sum_{D\in\mathcal{P}_n} \mathbb{P}(\Pi(D) \geqslant 2)\\ &\leqslant \frac{1}{2}\sum_{D\in\mathcal{P}_n} \mu(D)^2\\ &\leqslant \frac{1}{2}\,\delta_n \sum_{D\in\mathcal{P}_n} \mu(D) = \frac{1}{2}\,\delta_n \mu(E) \xrightarrow[n\to\infty]{(1.4)} 0.\end{aligned}$$

This proves the finite-intensity case.

For a general σ-finite diffuse intensity, partition E into countably many sets of finite μ-mass and apply the finite case to each restriction of Π. A countable intersection of the resulting probability-one events gives the conclusion. □

Theorem 1.11 (*Recovering the points*)

Assume that $E \neq \varnothing$, and suppose that Π is a PPM on $(E,\mathcal{E})$ satisfying (Δ) and having a σ-finite diffuse intensity μ. Then the atoms of Π can be enumerated measurably: there exist E-valued random variables $X_1, X_2, \ldots$ (with arbitrary values after the last atom when $\Pi(E) < \infty$) such that, with $N = \Pi(E) \in \mathbb{Z}_{\geqslant 0} \cup \{\infty\}$,

$$\Pi = \sum_{i=1}^{N} \delta_{X_i}.$$

Moreover, each X_i may be chosen as a measurable function of the point measure Π, i.e. using only the evaluation maps $\nu \mapsto \nu(A)$, $A \in \mathcal{E}$. By Proposition 1.10, these atoms are almost surely distinct.

Proof. Fix a point $x_* \in E$, to be used as a dummy value when necessary, and let $(C_n)_{n\geqslant 1}$ be a countable family separating points. The binary codes $(\mathbf{1}_{C_n}(x))_{n\geqslant 1}$ induce a lexicographic order $\prec$ on E. The idea is simply to enumerate the atoms according to this order.

There is one small nuisance: an arbitrary countable subset of E need not have a $\prec$-smallest element. We therefore begin with the finite-intensity case. Then $\Pi(E) < \infty$ almost surely and, by Proposition 1.10, we may write its atoms as

$$X_1 \prec X_2 \prec \cdots \prec X_N, \qquad N = \Pi(E).$$

The only point to check is that these atoms can be chosen measurably from Π. We explain this only for X_1, which contains the main idea.

For $n \geqslant 1$, let

$$\mathcal{P}_n = \{D_1^n, \ldots, D_{r_n}^n\}$$

be the finite partition generated by $C_1, \dots, C_n$, with its cells listed in lexicographic order of their binary labels. The partitions refine one another, and since Π has only finitely many atoms, for all sufficiently large n distinct atoms belong to distinct cells. Hence, on $\{N \geqslant 1\}$ and up to a null event, for every $A \in \mathcal{E}$,

$$\{X_1 \in A\} = \bigcup_{n \geqslant 1} \bigcap_{m \geqslant n} \bigcup_{1 \leqslant i \leqslant r_m} \Big(\{\Pi(A \cap D_i^m) = 1\} \\ \cap \{\Pi(D_1^m \cup \cdots \cup D_{i-1}^m) = 0\} \Big).$$

The right-hand side is measurable in terms of the evaluation maps. Extending X_1 by x_* on $\{N = 0\}$ therefore makes X_1 measurable. The other selectors are obtained in the same way by looking successively for the second, third, *etc.* occupied cell; the notation is heavier but no new idea is involved.

Finally, in the σ-finite case, partition $E = \bigsqcup_{m \geqslant 1} E_m$ with $\mu(E_m) < \infty$, apply the finite construction to each restriction of Π to E_m, and concatenate the resulting finite lists. The counts $\Pi(E_m)$ are measurable, so this concatenation is measurable as well. This gives the desired enumeration. □

The take-away message of this section is that PPMs on well-behaved spaces, such as Polish spaces, with diffuse intensities can be viewed as finite or countable collections of distinct random points. There is then no issue with measurability, and one can manipulate these objects without undue concern. We shall be less strict in the coming chapters on measurability checks.

II: Useful formulas

Throughout this chapter, unless explicitly stated otherwise, we consider a Poisson point measure Π on a measurable space $(E, \mathcal{E})$ with a σ-finite intensity measure μ. Recall that if $f : E \to [0, \infty]$ is measurable and ν is a positive measure on E, we use the notation

$$\nu(f) = \int_E f \, d\nu = \int_E f(x) \, \nu(dx).$$

In particular, if $\nu = \sum_i \delta_{x_i}$ is a point measure, then

$$\nu(f) = \sum_i f(x_i) \quad \text{and takes values in } [0, \infty].$$

2.1 Exponential formula

Theorem 2.1 (*Exponential formula*)

For every nonnegative measurable function $f : E \to [0, \infty]$,

$$\mathbb{E}\big[\exp(-\Pi(f))\big] \;=\; \exp\Big(-\int_E (1 - \mathrm{e}^{-f}) \, d\mu\Big), \tag{2.1}$$

where as usual we use the convention $\mathrm{e}^{-\infty} = 0$.

A mnemonic way to remember the formula is to notice that there are only minus signs.

Proof. First suppose that f is a simple function

$$f = \sum_{j=1}^{r} t_j \mathbf{1}_{A_j},$$

where $t_j \geqslant 0$, the sets A_j are pairwise disjoint (non-necessarily of finite μ-measure). Then

$$\Pi(f) = \sum_{j=1}^{r} t_j \Pi(A_j).$$

By the defining properties of a Poisson point measure,

$$\begin{aligned}\mathbb{E}[\mathrm{e}^{-\Pi(f)}] &= \prod_{j=1}^{r} \mathbb{E}[\mathrm{e}^{-t_j \Pi(A_j)}] \\ &= \prod_{j=1}^{r} \exp\big(\mu(A_j)(\mathrm{e}^{-t_j} - 1)\big) \\ &= \exp\left(-\sum_{j=1}^{r} \mu(A_j)(1 - \mathrm{e}^{-t_j})\right) \\ &= \exp\left(-\int_E (1 - \mathrm{e}^{-f})\, d\mu\right).\end{aligned}$$

Thus the formula holds for such simple functions.

Now let $f : E \to [0, \infty]$ be arbitrary, then there exists a sequence of simple positive function f_n such that $f_n \uparrow f$ and we can then apply monotone (or even bounded) convergence theorem to complete the proof. □

The exponential formula determines the law of a random point measure in the sense of Section 1.2.1 (and even in the sense of a random measure when (E, d) is Polish, see Chapter 3). Indeed, taking

$$f = \sum_{j=1}^{r} t_j \mathbf{1}_{A_j}, \qquad t_j \geqslant 0,$$

gives the joint Laplace transform, hence the law of every finite family $\Pi(A_1), \ldots, \Pi(A_r)$. Indeed, it is easy to see that the Laplace transform $(t_1, \ldots, t_n) \mapsto \mathbb{E}[\exp(-\sum t_i N_i)]$ characterize the law of a random integer vector $\mathrm{N} = (N_1, \ldots, N_n) \in (\mathbb{Z}_{\geqslant 0} \cup \{\infty\})^n$. In particular, we can now prove the marking property stated in Proposition 1.9.

Proof of Proposition 1.9. Use the notation of that proposition and set

$$\mathcal{G} = \sigma(N, X_1, X_2, \ldots).$$

If $f : E \times F \to [0, \infty]$ be measurable, we introduce the notation

$$q(x) = \int_F \mathrm{e}^{-f(x,y)}\, \nu(dy), \qquad g(x) = -\log q(x),$$

with the convention $-\log 0 = +\infty$. Conditionally on $\mathcal{G}$, the marks are independent with law ν, and therefore

$$\mathbb{E}\Big[\mathrm{e}^{-\Pi^F(f)} \mid \mathcal{G}\Big] = \prod_{i=1}^{N} q(X_i) = \exp\left(-\sum_{i=1}^{N} g(X_i)\right) = \mathrm{e}^{-\Pi(g)}.$$

Taking expectations and applying the usual exponential formula to Π gives

$$\begin{aligned}\mathbb{E}\Big[\mathrm{e}^{-\Pi^F(f)}\Big] &= \exp\Big(-\int_E (1-\mathrm{e}^{-g(x)})\,\mu(dx)\Big)\\ &= \exp\Big(-\int_E\int_F (1-\mathrm{e}^{-f(x,y)})\,\nu(dy)\,\mu(dx)\Big)\\ &= \exp\Big(-\int_{E\times F} (1-\mathrm{e}^{-f})\,d(\mu\otimes\nu)\Big).\end{aligned}$$

This is the exponential formula for a Poisson point measure on $E\times F$ with intensity $\mu\otimes\nu$. Since the exponential formula determines the law of a random point measure, the conclusion of Proposition 1.9 follows. □

The exponential formula gives access to the law of $\Pi(f)$ for every nonnegative measurable f. In particular, it gives its first moment; this identity is a special case of Campbell's[1] formula:

Theorem 2.2 (*Campbell formula*)

For every nonnegative measurable function $f : E \to [0,\infty]$,

$$\mathbb{E}\Big[\int_E f(x)\,\Pi(dx)\Big] = \int_E f(x)\,\mu(dx).$$

In particular, if $\int \mathrm{d}\mu f < \infty$ then $\Pi(f)$ is almost surely finite. Conversely, if $\mu(E)=\infty$ and

$$\int_E \mu(\mathrm{d}x)\,1\wedge f(x) = \infty,$$

then $\Pi(f)=\infty$ almost surely.

Proof. This could be deduced from the exponential formula, but the proof from scratch is quick: For $f=\mathbf{1}_A$ with $\mu(A)<\infty$, this is simply

$$\mathbb{E}[\Pi(A)] = \mu(A).$$

By linearity, the formula follows for nonnegative simple functions supported on a set of finite μ-mass. For arbitrary nonnegative measurable f, choose $f_m \uparrow f$ as in the proof of Theorem 2.1 and apply monotone convergence on both sides. This proves the Campbell formula and we deduce that $\Pi(f)<\infty$ as soon as f is μ-integrable. Conversely, recall from Lemma 1.1 that $\mathbb{P}(\Pi(f)<\infty) = \lim_{t\downarrow 0}\mathbb{E}[\mathrm{e}^{-t\Pi(f)}]$. We can then use the exponential formula

[1] Norman Robert Campbell (1880-1949), a British physicist.

to get

$$\mathbb{E}[\mathrm{e}^{-t\Pi(f)}] = \exp\left(-\int \mu(\mathrm{d}x)(1-\mathrm{e}^{-tf(x)})\right).$$

When $t \to 0$ we have $(1-\mathrm{e}^{-tf(x)}) \to \mathbf{1}_{f(x)=\infty}$ so in case of finite measure, by dominated convergence theorem we have

$$\mathbb{P}(\Pi(f) < \infty) = \exp\left(-\int \mu(\mathrm{d}x)\mathbf{1}_{f(x)=\infty}\right).$$

However, if μ is infinite and $\int \mathrm{d}\mu\ (1 \wedge f) = \infty$, using $1-\mathrm{e}^{-u} \geqslant 1 \wedge u$ we have

$$\int \mu(\mathrm{d}x)(1-\mathrm{e}^{-tf(x)}) \geqslant \int \mu(\mathrm{d}x)\ \big(1 \wedge (tf(x))\big) \geqslant t\int \mu(\mathrm{d}x)\ (1 \wedge f(x)) = \infty,$$

and so even when letting $t \to 0$ we get $\mathbb{P}(\Pi(f) < \infty) = 0$. □

If $f : E \to \mathbb{R}$ is measurable and

$$\int_E |f|\, d\mu < \infty,$$

then Campbell's formula applied to $|f|$ gives

$$\mathbb{E}[\Pi(|f|)] = \mu(|f|) < \infty.$$

Hence $\Pi(|f|) < \infty$ almost surely, so the signed integral $\Pi(f)$ is well defined almost surely, belongs to L^1, and

$$\mathbb{E}[\Pi(f)] = \int_E f\, d\mu.$$

2.2 The Mecke formula

Campbell's formula computes the expectation of a sum whose contribution at a point depends only on that point. The Mecke formula allows the contribution to depend on the entire configuration. The key of the formula sits on the size-biasing property of Poisson random variables:

Exercise 2.1. Let $N \sim \mathrm{Poi}(\lambda)$. For every $f : \mathbb{Z}_{\geqslant 0} \to \mathbb{R}_+$ we have

$$\frac{1}{\lambda}\mathbb{E}[Nf(N)] = \mathbb{E}[f(N+1)].$$

In words, the variable N biaised by N itself has the same distribution as $N+1$.

Theorem 2.3 (*Mecke formula*)

Let $\Pi \sim \mathrm{PPM}(\mu)$, where μ is σ-finite. Let $H : E \times \mathcal{N}(E) \to [0,\infty]$ be $\mathcal{E} \otimes \mathfrak{N}$-measurable.

Then

$$\mathbb{E}\left[\int_E H(x,\Pi)\,\Pi(dx)\right] = \int_E \mathbb{E}[H(x,\Pi+\delta_x)]\ \mu(dx).$$

Both sides may be infinite.

If H is real-valued rather than nonnegative, the same identity holds as soon as one side is absolutely integrable; equivalently it is enough to apply the nonnegative formula to $|H|$ and then subtract positive and negative parts.

Remark 2.1. The above formula has several names in the literature, sometimes depending on the field: Mecke[2] formula, Slivnyak–Mecke formula, refined Campbell formula, or even Palm formula. This is surely another application of Stigler's law of eponymy.

Proof. We leave to the reader the technical exercise of checking that all quantities are measurable and well-defined. We first assume that $\mu(E) = \lambda < \infty$. If $\lambda = 0$ the result is immediate, so suppose $\lambda > 0$. Since both sides depend measurably only on the point measure Π, whose law is fixed by its intensity, it is enough to verify the formula for the construction of Theorem 1.4. Let $N \sim \mathrm{Poi}(\lambda)$ and, conditionally on N, let $X_1,\dots,X_N$ be i.i.d. with law μ/λ. Then

$$\Pi = \sum_{j=1}^{N} \delta_{X_j}.$$

For brevity, write $\mathbf{x}_n = (x_1,\dots,x_n)$. By Tonelli's theorem,

$$\begin{aligned}\mathbb{E}&\left[\int_E H(x,\Pi)\,\Pi(dx)\right]\\ &= \mathrm{e}^{-\lambda}\sum_{n=1}^{\infty}\frac{1}{n!}\int_{E^n}\sum_{i=1}^{n} H\left(x_i,\sum_{j=1}^{n}\delta_{x_j}\right)\mu^{\otimes n}(d\mathbf{x}_n).\end{aligned}$$

The product measure $\mu^{\otimes n}$ is invariant under permutations of the coordinates. Therefore all n terms in the inner sum have the same integral, and

$$\begin{aligned}\mathbb{E}&\left[\int_E H(x,\Pi)\,\Pi(dx)\right]\\ &= \mathrm{e}^{-\lambda}\sum_{n=1}^{\infty}\frac{n}{n!}\int_{E^n} H\left(x_1,\sum_{j=1}^{n}\delta_{x_j}\right)\mu^{\otimes n}(d\mathbf{x}_n).\end{aligned}$$

[2] Joseph Mecke (1938-2014), a German mathematician. Izrail Mikhailovich Slivnyak (????- <2012), a Soviet mathematician.

Writing $m = n - 1$, separating the first coordinate, and applying Tonelli again, we obtain

$$\begin{aligned}
&\mathbb{E}\left[\int_E H(x,\Pi)\,\Pi(dx)\right] \\
&\quad = \int_E \left[\mathrm{e}^{-\lambda} \sum_{m=0}^{\infty} \frac{1}{m!} \int_{E^m} H\left(x, \delta_x + \sum_{j=1}^{m} \delta_{y_j}\right) \mu^{\otimes m}(d(y_1,\ldots,y_m)) \right] \mu(dx).
\end{aligned}$$

The expression in square brackets is precisely $\mathbb{E}[H(x,\Pi+\delta_x)]$. This proves the formula when $\mu(E) < \infty$.

We now treat the σ-finite case. Choose $F_m \uparrow E$ with $\mu(F_m) < \infty$, and set

$$\Pi_m^{\mathrm{in}} = \Pi|_{F_m}, \qquad \Pi_m^{\mathrm{out}} = \Pi|_{F_m^c}.$$

By Proposition 1.7, these two random point measures are independent. Conditioning on Π_m^{out}, we apply the finite-intensity Mecke formula to Π_m^{in} with the measurable function

$$H_m(x,\Pi) = H_m(x, \Pi_m^{\mathrm{out}} + \Pi_m^{\mathrm{in}}) = \mathbf{1}_{F_m}(x) \cdot H(x, \Pi_m^{\mathrm{out}} + \Pi_m^{\mathrm{in}})$$

and obtain

$$\begin{aligned}
\mathbb{E}\left[\int_{F_m} H(x,\Pi)\,\Pi(dx)\right] &= \mathbb{E}\left[\int_{F_m} H_m(x, \Pi_m^{\mathrm{in}} + \Pi_m^{\mathrm{out}})\,\Pi_m^{\mathrm{in}}(dx)\right] \\
&= \mathbb{E}^{\mathrm{out}}\left[\int_{F_m} \mathbb{E}^{\mathrm{in}}\left[H_m(x, \Pi_m^{\mathrm{in}} + \Pi_m^{\mathrm{out}} + \delta_x)\right]\,\mu(dx)\right] \\
&= \int_{F_m} \mathbb{E}^{\mathrm{in+out}}\left[H_m(x, \Pi_m^{\mathrm{in}} + \Pi_m^{\mathrm{out}} + \delta_x)\right]\,\mu(dx).
\end{aligned}$$

Finally, both sides increase as $m \to \infty$, so monotone convergence proves the formula on E.

□

Remark 2.2. Taking $H(x,\nu) = f(x)$ in the Mecke formula gives Campbell's formula. Thus Mecke's identity is a genuine extension of Campbell's identity.

A common interpretation of the Mecke formula is as follows: a typical point X of Π is distributed according to μ, and, conditionally on X, the remaining point measure has the same distribution as Π. The proper meaning of this statement is precisely the Mecke formula. One can also consider "Π conditioned to have an atom at x". Even if this event has zero probability, the resulting point measure is obtained simply by adding δ_x to an independent sample of Π. For a general random point measure, the corresponding conditional law is called the Palm distribution[3].

[3] Conny Palm (1907-1951), a Swedish teletraffic engineer

2.2.1 Multi-point Mecke

The one-point Mecke formula has a very useful multi-point version. To get nice formulas, the correct object on the left-hand side is not the ordinary product measure $\Pi^{\otimes k}$ and one must exclude the possibility of selecting the same *occurrence* of a point more than once. The proof is left to the reader, and can be obtained by recursing from Theorem 2.3:

Theorem 2.4 (*Multi-point Mecke formula*)

Let $k \geqslant 1$ and let $H : E^k \times \mathcal{N}(E) \longrightarrow [0,\infty]$ be $\mathcal{E}^{\otimes k} \otimes \mathfrak{N}$-measurable. Let $\Pi \sim \mathrm{PPM}(E,\mu)$ written as

$$\Pi = \sum_{i \in I} \delta_{X_i}.$$

Then

$$\mathbb{E}\Big[\sum_{\substack{i_1,\dots,i_k \in I \\ \text{pairwise distinct}}} H(X_{i_1},\dots,X_{i_k},\Pi)\Big] = \int_{E^k} \mathbb{E}\left[H\left(x_1,\dots,x_k,\Pi + \sum_{r=1}^{k} \delta_{x_r}\right)\right] \mu^{\otimes k}(d(x_1,\dots,x_k)).$$

Specializing to functions $H(\mathbf{x},\Pi) = h(\mathbf{x})$, we obtain the multi-point Campbell formula:

Corollary 2.5 (Factorial Campbell formula). *For every $k \geqslant 1$ and every nonnegative measurable $h : E^k \to [0,\infty]$,*

$$\mathbb{E}\Big[\sum_{\substack{i_1,\dots,i_k \in I \\ \text{pairwise distinct}}} h(X_{i_1},\dots,X_{i_k})\Big] = \int_{E^k} h(\mathbf{x})\, \mu^{\otimes k}(d\mathbf{x})..$$

2.3 Exercises

Exercise 2.2 (Simplicity from the second factorial measure). Assume (Δ) and suppose that μ is diffuse. Use the multi-point Mecke formula to prove (again) that Π is almost surely simple.

Exercise 2.3 (Mecke's identity characterizes the Poisson law). Let $\mathbf{N}$ be a random point measure with σ-finite intensity measure μ and assume that, for every nonnegative measurable H,

$$\mathbb{E}\left[\int_E H(x,\mathbf{N})\, \mathbf{N}(dx)\right] = \int_E \mathbb{E}[H(x,\mathbf{N}+\delta_x)]\, \mu(dx).$$

First suppose that f is bounded, nonnegative, and vanishes outside a set of finite μ-mass. Set

$$L(t) = \mathbb{E}[\mathrm{e}^{-t\mathbf{N}(f)}], \qquad t \geqslant 0.$$

1. Apply Mecke's identity with

$$H(x,\nu) = f(x)\mathrm{e}^{-t\nu(f)}$$

to derive a differential equation for L.

2. Deduce that $L(t) = \exp\left(-\int_E (1-\mathrm{e}^{-tf})\,d\mu\right)$.

3. Extend the result by monotone approximation and deduce that $\mathbf{N}$ is a Poisson point measure with intensity μ.

Exercise 2.4 (Dickman distribution). Let Π be a PPM on $(0,1]$ with intensity

$$\mu(dx) = \frac{dx}{x},$$

and set

$$X = \int_{(0,1]} x\,\Pi(dx).$$

1. Show that $X < \infty$ almost surely and that $\mathbb{E}[X] = 1$.

2. Show that, for $\lambda \geqslant 0$,

$$\mathbb{E}[\mathrm{e}^{-\lambda X}] = \exp\left(\int_0^1 \frac{\mathrm{e}^{-\lambda x}-1}{x}\,dx\right) = \exp\big(I(-\lambda)\big),$$

where

$$I(s) = \int_0^s \frac{\mathrm{e}^t - 1}{t}\,dt.$$

This law is called the Dickman distribution.

3. Under the change of variables $x = \mathrm{e}^{-t}$, identify the image of Π with a homogeneous Poisson process of rate 1 on $\mathbb{R}_+$. If T_1 denotes its first point, show that $U = \mathrm{e}^{-T_1}$ is uniform on $(0,1)$ and deduce the distributional identity

$$X \overset{d}{=} U(1+X'),$$

where X' is an independent copy of X, independent of U.