Week #9: Counting Processes and Simulation

Notes by: Francisco Richter

May 15, 2025

Overview

Poisson processes are a fundamental concept in stochastic modeling, providing a rigorous mathematical framework for understanding events that occur randomly in time or space. Arrival times and counting processes (such as Poisson processes) find applications in a myriad of contexts, each with its own set of challenges and implications. For instance, in healthcare, modeling the arrival times of patients in an emergency room can be crucial for optimizing resource allocation and improving patient outcomes. Similarly, understanding the time intervals between bus arrivals at a specific stop can offer insights into public transportation scheduling and efficiency. In computer science, the arrival times of data packets in a network can be analyzed to optimize bandwidth and reduce latency. Businesses too can benefit; for example, modeling the arrival times of customers in a service queue, whether in a call center or a fast-food restaurant, can lead to enhanced service management.

Natural events like earthquakes, floods, and forest fires also exhibit arrival times that can be modeled to better understand and predict these phenomena. In retail, the time between customer arrivals at a checkout counter can inform decisions about staffing and service speed. Social media platforms often scrutinize the timing of posts or mentions to understand user engagement or to detect trending topics. In manufacturing, arrival times of components on an assembly line can be critical for identifying bottlenecks and optimizing production. Financial markets are another fertile ground where the arrival times of buy/sell orders can shed light on market dynamics. Finally, in ecology, monitoring the arrival times of different species at a watering hole or feeding station can offer invaluable data for conservation efforts and ecological research.

---

1 Counting Processes

Definition. Counting Process

A counting process is a stochastic process $\{N(t), t \geq 0\}$ that represents the total number of events that have occurred up to time t. The function $N(t)$ satisfies the following properties:

(1) $N(0) = 0$ (initial condition)

(2) $N(t)$ is integer-valued for all $t \geq 0$

(3) $N(t)$ is non-decreasing as t increases; that is, if $s < t$, then $N(s) \leq N(t)$

(4) The function $N(t)$ is right-continuous, meaning that for each t, $\lim_{s \to t^+} N(s) = N(t)$

In simpler terms, a counting process counts the number of times a certain event has occurred by any given time t. The count starts at zero and can only increase as time moves forward.

Consider a time interval T that we divide into n smaller intervals, each of length $\Delta t = \frac{T}{n}$. We are interested in counting the number of occurrences of a particular event within each small time interval Δt.

Initially, let us model this as a Bernoulli process. In each small time interval Δt, the event can either occur with probability p or not occur with probability $1 - p$:

$$P(\text{Event occurs in } \Delta t) = p, \quad P(\text{Event does not occur in } \Delta t) = 1 - p.$$

For large n and small Δt, we can relate p to a rate parameter λ by

$$p = \lambda \Delta t.$$

Let X be the number of events that occur in the entire interval $[0, T]$. The variable X is a sum of n independent Bernoulli random variables, each with success probability p. Thus,

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

The probability of observing exactly k events in T is

$$P(X = k) = \binom{n}{k} p^k (1 - p)^{n-k}.$$

Substituting $p = \lambda \Delta t$, we get

$$P(X = k) = \binom{n}{k} (\lambda \Delta t)^k \left(1 - \lambda \Delta t\right)^{n-k}.$$

As $n \to \infty$ (and hence $\Delta t \to 0$ while $n\Delta t = T$ remains constant), the binomial distribution converges to a Poisson distribution:

Theorem 1.1 (Convergence from Bernoulli to Poisson)**.** *Let X be the number of events in a time interval T divided into n intervals of length $\Delta t = \frac{T}{n}$. If each interval has a Bernoulli-distributed event occurrence with probability $p = \lambda \Delta t$, then as $n \to \infty$ with $n\Delta t = T$ fixed,*

$$\lim_{n \to \infty} P(X = k) = \frac{(\lambda T)^k}{k!} \exp(-\lambda T).$$

Hence, X follows a Poisson distribution with parameter λT in the limit.

---

2 Poisson Processes

A *Poisson process* holds a central position in stochastic processes because of its mathematical elegance and broad applicability. It models scenarios where events occur randomly in continuous time (or space) at a certain average rate.

2.1 Homogeneous Poisson Process

Definition. Homogeneous Poisson Process

Let $(N(t) : t \geq 0)$ be a counting process. We say that $N(t)$ is a *Homogeneous Poisson Process* with rate $\lambda > 0$ if:

(1) $N(0) = 0$.

(2) The increments are independent.

(3) The number of events in any interval of length t follows a Poisson distribution with mean λt.

2.1.1 Waiting Times and Memoryless Property

A crucial insight is that the waiting times between successive events in a Homogeneous Poisson Process are *exponentially distributed* with parameter λ. This implies that the Poisson process also has the *memoryless property*: the time until the next event does not depend on how much time has already elapsed since the last event.

Theorem 2.1 (Exponential Waiting Times)**.** *In a Homogeneous Poisson Process with rate λ, the time T until the first event is exponentially distributed:*

$$P(T \leq t) = 1 - e^{-\lambda t}.$$

Theorem 2.2 (Memoryless Property)**.** *For an exponentially distributed waiting time T,*

$$P\big(T > s + t \,\big|\, T > s\big) = P(T > t),$$

and hence the underlying Poisson process is memoryless.

2.2 Non-Homogeneous Poisson Process

In many real-world settings, the rate of arrival $\lambda(t)$ varies over time. A *Non-Homogeneous Poisson Process* (NHPP) generalizes the basic Poisson framework by allowing $\lambda(t)$ to be a function of t. Then:

- $\{N(t)\}$ still has independent increments.
- The number of events in $[s, t]$ is Poisson with mean $\int_s^t \lambda(u)\, du$.

---

3 Birth–Death Processes

Birth–death processes form a classic family of continuous-time Markov chains often used to model population dynamics. Let $N(t)$ be the population at time t. Suppose:

- λ_n is the **birth rate** when the population is n.
- μ_n is the **death rate** when the population is n.

Define $P_n(t) = P\big(N(t) = n\big)$. The *master equation* (or Kolmogorov forward equation) reads:

$$\frac{dP_n(t)}{dt} = \lambda_{n-1}\, P_{n-1}(t) + \mu_{n+1}\, P_{n+1}(t) - (\lambda_n + \mu_n)\, P_n(t).$$

Such processes can be simulated directly (e.g., via the Gillespie algorithm described below) or analyzed theoretically for their steady-state distributions.

---

4 Simulations in Stochastic Processes

Example 4.1 (Birth–Death Simulation)**.** *If $\lambda_n = \lambda$ (constant) and $\mu_n = n\mu$, then births occur at a constant rate λ, while deaths occur at a rate proportional to the current population n. One can use a stochastic simulation to observe how the population fluctuates over time around a mean level near λ/μ.*

4.1 Gillespie Algorithm

The Gillespie algorithm (also known as the *stochastic simulation algorithm*) is widely used for simulating discrete-event systems—particularly chemical reaction networks. It can also be applied to birth–death processes, queueing systems, and any scenario where events occur randomly in continuous time with state-dependent rates.

Algorithm Overview

1. **Initialization:** Set $t = 0$ and the initial state $X(t)$.
2. **Compute Propensities:** For each possible event (e.g., reaction) j, calculate the propensity $a_j(X(t))$.
3. **Compute Total Propensity:**
$$a_0(X(t)) = \sum_j a_j(X(t)).$$
4. **Time to Next Event:** Draw $r_1 \sim U(0, 1)$ and compute
$$\Delta t = \frac{1}{a_0(X(t))} \ln\Big(\frac{1}{r_1}\Big).$$
5. **Select Event:** Draw another $r_2 \sim U(0, 1)$. Choose event μ such that
$$\sum_{j=1}^{\mu-1} a_j < r_2\, a_0 \leq \sum_{j=1}^{\mu} a_j.$$
6. **Update State:** Execute event μ, update $X(t)$, and set $t \leftarrow t + \Delta t$.
7. **Repeat:** Return to Step 2 until a designated stopping time or condition.

In this section, we examine three illustrative simulation frameworks that showcase how stochastic processes are used in different modeling contexts. We first look at *opinion dynamics* (combining Markov chains and truncated exponential updates), then a *forest fire* cellular automaton, and finally an *agent-based* ant foraging simulation.

4.2 Cellular Automaton: The Forest Fire Model

Definition. Stochastic Cellular Automaton

A stochastic cellular automaton is defined by $(\mathcal{G}, S, \mathcal{N}, T)$, where $\mathcal{G}$ is a lattice, S is the set of cell states, $\mathcal{N}$ is the neighborhood structure, and T is a probabilistic transition rule.

In the *forest fire* model, each cell represents a patch of land and can be:

$$\{\text{Grass}, \text{Tree}, \text{Burning}, \text{Empty}\}.$$

Rules (with given probabilities) control growth from Grass to Tree, ignition from Tree to Burning, and consumption of Burning to Empty. A cell may also *spontaneously* ignite at a small probability. Repeated updates reveal patterns of fire spread and regrowth.

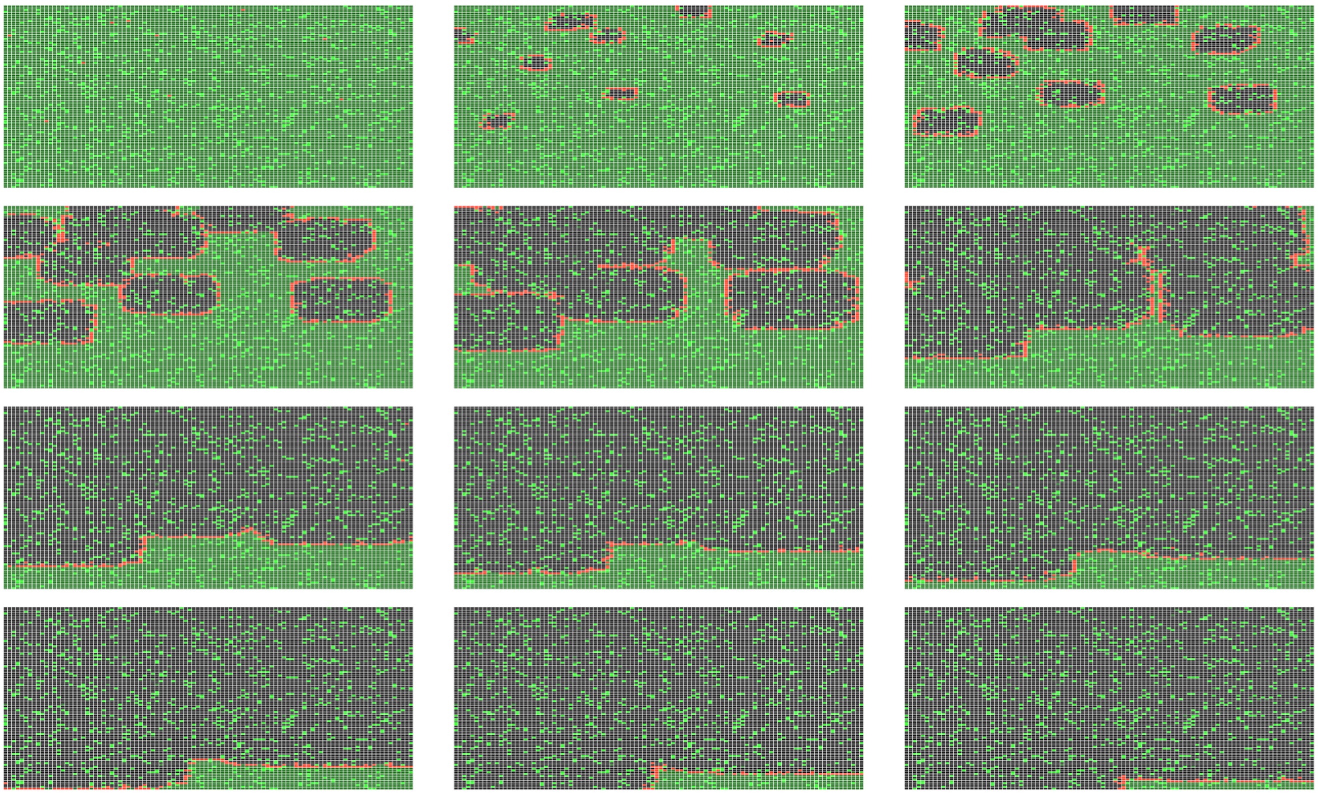

Figure 1: Snapshots from a forest fire simulation, illustrating stages of fire spread and eventual regrowth.

4.3 Agent-Based Modeling: Ant Foraging Simulation

Definition. Agent

An agent A is a tuple

$$A = (S, R, G, T, M, D),$$

where S is its state space, R is its role, G its goal, T the transition rules, M its memory, and D its decision-making process.

In *ant foraging*:

- Each ant moves in a grid world with position (X, Y) and a flag F indicating if it is carrying food.
- If an ant finds food, it switches to $F = 1$ and heads back to the nest, depositing pheromones.
- Ants not carrying food follow pheromone trails or wander randomly.

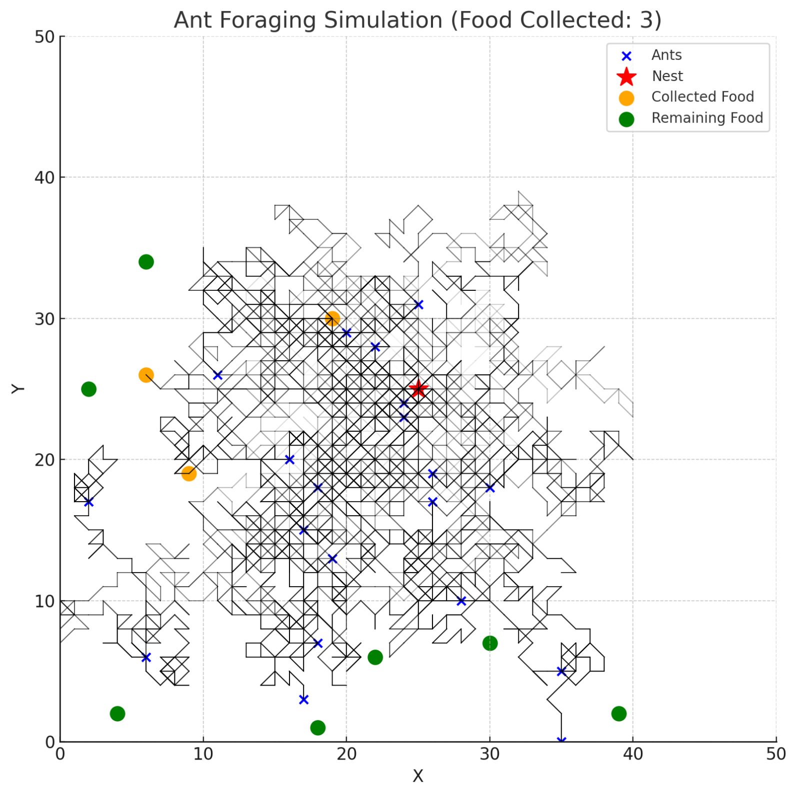

Figure 2: Ant foraging simulation on a 50×50 grid; nest at the red star, food sites in orange.

These agent-based simulations show how simple local rules can result in complex emergent behaviors, such as coordinated foraging.

4.4 Stochastic Simulation of Opinion Dynamics

We combine a Markov chain for who speaks (which person is chosen to talk at each step) with a truncated exponential model for how one individual's opinion influences another's. Let there be m individuals indexed by $i = 1, \ldots, m$. Each person has an opinion $o_i \in [0, 1]$. The *speaker selection* is governed by a discrete-time Markov chain X_n, which indicates which individual speaks at time n.

When person i (opinion o_i) influences person j (opinion o_j), assume w.l.o.g. $o_i < o_j$. Define

$$d = o_j - o_i,$$

and let $\Delta \in [0, d]$ be the shift bringing o_j partially toward o_i. Then

$$o_j' = o_j - \Delta,$$

with Δ drawn from a truncated exponential on $[0, d]$:

$$f_T(\Delta) = \frac{\lambda(T)\, \exp(-\lambda(T)\Delta)}{1 - \exp(-\lambda(T)\, d)},$$

where $\lambda(T)$ is a function of a *waiting time* T that encodes open-mindedness. A large T yields smaller $\lambda(T)$, hence bigger possible opinion changes. A small T yields minimal changes.

4.4.1 Simulation Procedure

1. **Initialization:** Assign initial opinions $o_i(0)$ and pick X_0.
2. **Markov Step:** Draw X_{n+1} from the transition probabilities of the Markov chain.
3. **Opinion Update:** If X_n influences X_{n+1}, update the latter's opinion via the truncated exponential rule.
4. **Iterate:** Repeat for a large number of steps.

Questions of interest include:

- Does the opinion set $\{o_i(n)\}$ converge over time?
- If convergence occurs, what is the limiting opinion?
- How do changes in $\lambda(T)$ affect the speed or possibility of consensus?

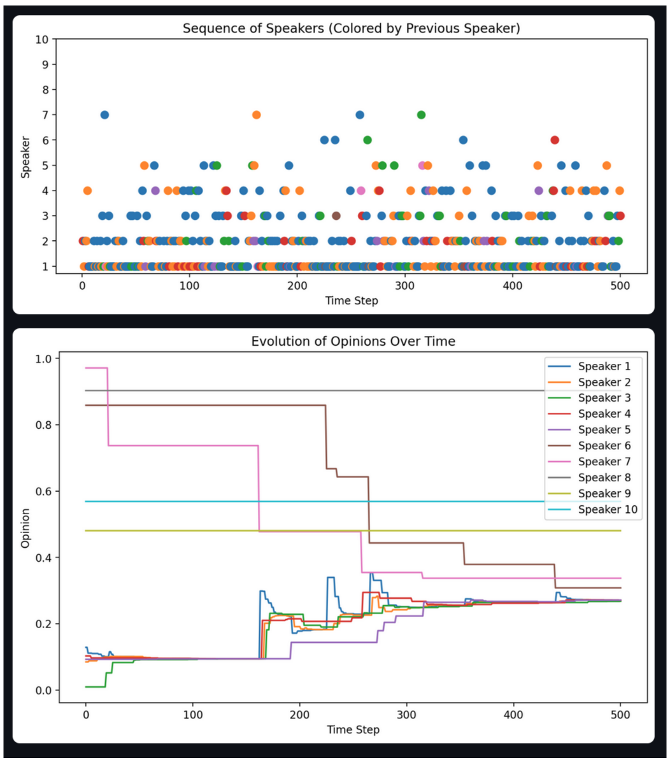

Figure 3: Conceptual illustration: person i (opinion o_i) influences person j (opinion o_j), shifting o_j closer to o_i.

Questions

1. Consider a Markov chain with three states, labeled 1, 2, and 3. The transition rates are

$$W_{1\to 2} = \alpha, \quad W_{2\to 1} = \beta, \quad W_{2\to 3} = \gamma, \quad W_{3\to 2} = \delta,$$

and all other rates are zero.

(a) Write down the system of first-order differential equations describing the time evolution of $P_1(t)$, $P_2(t)$, and $P_3(t)$.

(b) Use the fact that $P_1(t) + P_2(t) + P_3(t) = 1$ to reduce the system to fewer equations.

(c) Identify any stationary (steady-state) solution(s) $P_i(\infty)$ in terms of $\alpha, \beta, \gamma, \delta$.

Interpretation: Suppose at $t = 0$, the system is entirely in state 1: $P_1(0) = 1$, $P_2(0) = 0$, $P_3(0) = 0$. Describe qualitatively (without fully solving) how you expect the probability distribution among the states to evolve over time for different relative magnitudes of $\alpha, \beta, \gamma, \delta$.

2. Consider a birth-death process on the non-negative integers $\{0, 1, 2, \dots\}$ with transitions

$$n \to n+1 \quad \text{at rate } \lambda_n, \qquad n \to n-1 \quad \text{at rate } \mu_n \text{ (for } n \geq 1),$$

and no other transitions are allowed.

(a) Write the system of first-order differential equations for $P_n(t)$ (the probability that the system is in state n at time t). In particular, clarify how the boundary condition at $n = 0$ should be handled.

(b) Assume $\lambda_n = \lambda$ and $\mu_n = \mu$ (constants) for $n \geq 1$. Show that a stationary distribution exists only if $\lambda < \mu$. Find that stationary distribution $P_n(\infty)$.

(c) If $\lambda \geq \mu$, argue what happens to $P_n(t)$ as $t \to \infty$. Does the system have a proper stationary distribution in that scenario?

3. Consider the tree in Figure 4, representing a birth-death process in continuous time, except that the extinction (death) rate μ is zero. The speciation (birth) rate λ depends on the color of the branch. There are only two possible colors: green and blue. Specifically,

$$\lambda = \begin{cases} \lambda_1 & \text{if the branch is green,} \\ \lambda_2 & \text{if the branch is blue,} \end{cases}$$

and whenever a branching (speciation) event occurs, the newly created descendant branch inherits its parent's color with probability $\frac{2}{3}$ or switches to the other color with probability $\frac{1}{3}$.

(a) Based on the observed branch colors and branching times, what is the probability (likelihood) of the evolutionary process depicted in Figure 4?

Hint: The overall probability factors into products of (i) waiting times to speciation events according to the branch-specific rate λ_1 or λ_2, and (ii) color-transition probabilities at each branching event.

(b) Outline how one could simulate such processes via a *Gillespie-type algorithm.*

References

[1] Ross, S. M. *Stochastic Processes.* John Wiley & Sons.

[2] Grimmett, G. R., & Stirzaker, D. R. *Probability and Random Processes.* Oxford University Press.

[3] Allen, L. J. S. *An Introduction to Stochastic Processes with Applications to Biology.* CRC Press.

[4] Ross, S. M. *Introduction to Probability Models.* Academic Press.

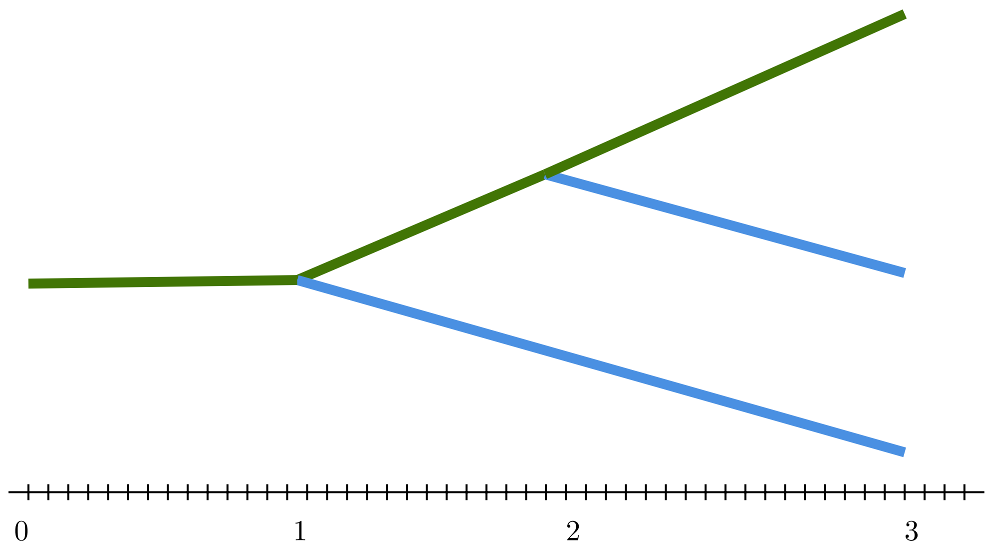

Figure 4: A schematic phylogenetic tree from time 0 to time 3, with branch colors green (dark) and blue (light).