8.808/8.308 IAP 2026 Recitation 7-8: Simulation of Markov Jump Processes

Jessica Metzger

jessmetz@mit.edu | Office hours: 1/9, 1/14, 1/20, 1/27 11am-12pm (8-320)

January 21, 2026

Contents

1 Simple example: lightning strikes . . . 1
 1.1 Simple probabilities . . . 1
 1.2 Simulation strategy . . . 2
2 Multiple rates . . . 2
 2.1 Continuous time: the Gillespie algorithm . . . 2
 2.2 Many particles: discrete time . . . 3
3 Ising model . . . 3
 3.1 Detailed balance . . . 4

Consider a system with states $\{\varphi\}$, whose probability distribution over the states is described by the master equation

$$\partial_t P(\varphi, t) = \sum_{\varphi'} \Big[W(\varphi' \to \varphi) P(\varphi') - W(\varphi \to \varphi') P(\varphi) \Big] \, . \tag{1}$$

Today, we will show how to simulate these dynamics faithfully.

1 Simple example: lightning strikes

Consider, as a simple example, the number of lightning strikes k that a person has received in their lifetime. Suppose they are uncorrelated and occur randomly at a rate $w \sim 1$ per decade. A lightning strike represents a transition from state k to $k+1$. Thus, the master equation for the probability $P(k,t)$ of having received k lightning strikes by time t is

$$\partial_t P(k,t) = wP(k-1) - wP(k) \, . \tag{2}$$

Given a time interval τ, the number of strikes can be 0, 1, etc. On average, it is $w\tau$.

1.1 Simple probabilities

Let's break our interval τ into n intervals of length dt, i.e. $\tau = ndt$, with very small dt.

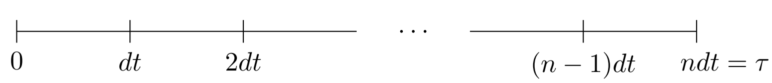

We can then exactly calculate the probabilities of having different numbers of strikes within τ:

0. The probability to have **0 strikes** during dt is $1 - wdt + \mathcal{O}(dt)$. Thus, the probability to have 0 strikes during τ is the probability to have 0 strikes during each of the $n = \tau/dt$ uncorrelated time-intervals of duration dt:

$$P[0 \text{ strikes during } \tau] = \lim_{dt \to 0} \big(1 - wdt\big)^{\tau/dt} = e^{-w\tau} \tag{3}$$

1. The probability to have **1 strike** during time dt is wdt. Thus, the probability to have a single strike at time $t < \tau$ is $e^{-wt} \cdot wdt \cdot e^{-w(\tau-t)} = we^{-w\tau}dt$. Integrating this over all possible times t gives us

$$P[1 \text{ strike during } \tau] = \int_0^\tau dtwe^{-w\tau} = w\tau e^{-w\tau} \ . \tag{4}$$

k. The probability to have **k strikes** during τ is found analogously:

$$P[k \text{ strikes during } \tau] = \frac{1}{k!}\int_0^\tau dt_1 \int_0^\tau dt_2 \ldots \int_0^\tau dt_k e^{-wt_1} \cdot w \cdot e^{-w(t_2-t_1)} \cdot w \cdot \ldots \cdot e^{-w(\tau-t_k)} \tag{5}$$

$$= \frac{1}{k!}\left(\int_0^\tau dtw\right)^k e^{-w\tau} \tag{6}$$

$$= \frac{1}{k!}(w\tau)^k e^{-w\tau} \ . \qquad \text{(Poisson distribution)} \tag{7}$$

This leads us to our simulation strategy:

1.2 Simulation strategy

There are two main classes of simulation:

- **Discrete time.** Choose a small time τ and
  1. Do 0 strikes with probability $1 - w\tau \sim e^{-w\tau}$
  2. Do 1 strikes with probability $w\tau \sim w\tau e^{-w\tau}$
  3. Ignore scenarios with more than 1 strike, since their probability is $\mathcal{O}(\tau^2)$.
- **Continuous time.** Because the probability to see zero strikes during time τ is $e^{-w\tau}$, we can infer that the probability distribution of *waiting times* is proportional to $e^{-w\tau}$. Thus, because we have only a single process, we can simply draw the waiting time τ from the distribution

$$P_\tau(\tau) = we^{-w\tau} \ . \tag{8}$$

  Then, we do 1 lightning strike, and increment time by τ.

The continuous time method is exact, and the discrete time method is only accurate to order τ.

2 Multiple rates

Let's now consider a system with more than 1 event. If we have m different possible events, with rates $w_1, w_2, \ldots, w_m$ respectively, we can again calculate the elementary probabilities:

0. The probability to have **0 events** during an interval τ is

$$P[0 \text{ events}] = e^{-w_1\tau}e^{-w_2\tau}\ldots e^{-w_m\tau} = \exp\left(-\sum_{i=1}^m w_i\tau\right) \ . \tag{9}$$

1. The probability to have **1 event** during τ is

$$P[1 \text{ event}] = w_1\tau e^{-w_1\tau}e^{-w_2\tau}\ldots e^{-w_m\tau} + w_2\tau e^{-w_1\tau}e^{-w_2\tau}\ldots e^{-w_m\tau} + \ldots + w_m\tau e^{-w_1\tau}e^{-w_2\tau}\ldots e^{-w_m\tau} \tag{10}$$

$$= \left(\sum_{i=1}^m w_i\tau\right)\exp\left(-\sum_{i=1}^m w_i\tau\right) \ . \tag{11}$$

2.1 Continuous time: the Gillespie algorithm

The analogue of the continuous-time algorithm mentioned above is known as the **Gillespie algorithm**. Observing that we again know the probability distribution of waiting time between events τ, we first simply sample τ:

$$P_\tau(\tau) \propto \exp\left(-\sum_{i=1}^m w_i\tau\right) \equiv e^{-W\tau} \ . \tag{12}$$

Then, we need to decide which event to execute. Given that one event must occur, the probability of event i occurring will be proportional to its rate w_i. We thus choose event i with probability

$$P[\text{event } i] = \frac{w_i}{\sum_{j=1}^{m} w_j} \equiv \frac{w_i}{W} \,. \tag{13}$$

Then, we execute the chosen event i and update time by τ. (We also must update our rates $\{w_i\}$ to reflect to change in state.)

These steps are summarized as follows:

Gillespie Algorithm

i) Sample waiting time τ from PDF $\propto \exp\left[-\sum_{i=1}^{m} w_i \tau\right]$

ii) Choose the event with probability $w_i / \sum_{j=1}^{m} w_j$

iii) Update time $t \to t + \tau$. Update rates $\{w_i\}$ to reflect the changed state. Start again at i).

As an example, we will simulate the following 3-state Markov process

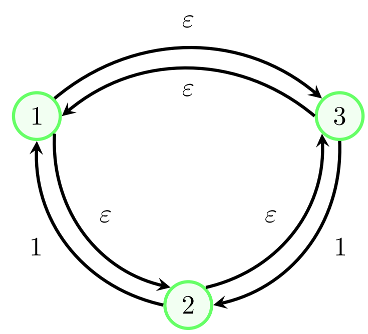

and measure the current. See `Rec7_3state.jl` and `Rec7_module.jl` for the code.

If you have time, you can calculate the current analytically, and compare that to your simulation.

2.2 Many particles: discrete time

Consider a system with a very large number N of particles. While the Gillespie algorithm is exact, it can be time-intensive for a large number of particles: the tower sampling takes $\mathcal{O}(N)$. In this case, it is better to use the following discrete time algorithm, where the timestep τ is fixed:

Discrete-time Algorithm (good for large N)

i) Randomly choose an agent n (particle, spin, etc.)

ii) Do something with probability $P_{\text{Do}} = \sum_{i \in A_n} w_i \tau$, where A_n is the set of possible events for agent n. Do nothing with probability $1 - P_{\text{Do}}$.

iii) Increase time by τ/N, update all rates, and repeat.

When N is very large, τ/N approaches zero. Thus it isn't necessary to make P_{Do} small. As long as it is ≤ 1, the error will decrease as the timestep τ/N decreases.

3 Ising model

You will simulate the Ising model in 2d. The Ising model describes a system of N spins $\{s_i\}$, with s_i taking on values $-1, +1$. The energy of the system is given by the Hamiltonian

$$H = -\sum_{i,j} J_{ij} s_i s_j + \sum_i h_i s_i \tag{14}$$

for couplings J_{ij} and the externally-applied field h_i. We will take $h_i = 0$ (no external field), and

$$J_{ij} = \begin{cases} J\ , & i,j \text{ nearest neighbors} \\ 0\ , & \text{otherwise} \end{cases} \quad . \tag{15}$$

We assume $J > 0$, i.e. the interactions are ferromagnetic (rather than antiferromagnetic).

Spin s_i experiences the energy $H_i(s_i) = -Js_i \sum_{j\sim i} s_j \equiv -Js_i m_i$. The possible changes of energy when flipping from s_i to $-s_i$ are $\Delta E \in \{-8J, \ldots, 8J\}$.

We need to construct a set of flipping rates that accurately sample the probability distribution of this system. To do so, we will take advantage of the detailed balance property of this equilibrium system.

3.1 Detailed balance

This is an equilibrium system; thus we know that in the steady state a transition from a state with energy φ to energy φ' will satisfy the detailed balance property

$$P(\varphi)W(\varphi \to \varphi') = P(\varphi')W(\varphi' \to \varphi) \tag{16}$$

where state φ has probability $\propto e^{-\beta H(\varphi)}$, and thus

$$\frac{P(\varphi)}{P(\varphi')} = e^{-\beta[H(\varphi)-H(\varphi')]}\ . \tag{17}$$

Thus, for our system to reach a state of detailed balance, its transition rates must satisfy

$$\frac{W(\varphi \to \varphi')}{W(\varphi' \to \varphi)} = e^{\beta[H(\varphi)-H(\varphi')]}\ . \tag{18}$$

Thus, for a spin i, we can choose the rate of flipping to be

$$W(s_i \to -s_i) = e^{\frac{1}{2}\beta[H_i(s_i)-H_i(-s_i)]}\ . \tag{19}$$

All that is left is to determine the appropriate τ. We need the probability to execute a flip to be ≤ 1. Thus, we choose $\tau = 1/[\max W(s_i \to -s_i)] = e^{-4\beta J}$, which is the largest possible τ. Choosing the largest possible τ ensures that we waste as little time as possible generating random numbers and doing nothing. Because we will take $N \to \infty$, the error will go to zero as explained above.

For the code, see `Rec7_2D_Ising.jl` and `Rec7_2D_Ising_module.jl`