

Single Jump Approximation

Sean L. Wu (slwood89@gmail.com)

4/29/2021

Single-Jump Approximation Algorithm

The single jump algorithm is a particular way to numerically simulate approximately correct trajectories from models expressed as continuous-time Markov chains (CTMC). Bréto et al. (2009) call this an Euler-multinomial scheme but we use the term single jump because an Euler scheme brings to mind approximating intensities linearly (as in tau-leaping) whereas the single jump algorithm actually creates a new model (a discrete-time Markov chain, DTMC) that accumulates probability correctly (via an Exponential CDF) over a time step (Δt), but also converges to the true CTMC model as $\Delta t \rightarrow 0$.

We are going to assume that the model is expressed in terms of a stoichiometry matrix $\mathbf{S}_{u \times v}$ and a column vector of hazard function $\Lambda_{v \times 1}$, where u is the number of compartments (species, for chemical reaction networks) and v is the number of events (reactions) that can cause a change in system state. We deal exclusively with the class of Markov population processes (Kingman 1969), where each event may only belong to one of three types: a birth (a single compartment gains 1 member), a death (a single compartment loses 1 member), or transfers (one compartment loses 1 member and another compartment gains 1 member). While the algorithm can be extended to more complex models, we restrict ourselves to this framework, as it is sufficiently general to model a wide class of processes.

Under these assumptions, \mathbf{S} has only (0,-1,1) elements, where \mathbf{S}_{ij} gives the effect of the j^{th} event on the i^{th} compartment. The hazard function $\lambda_j(t, X(t))$ takes as arguments the current simulation time t , and the current system state $X(t)$, which is a u -dimensional column vector giving the state of the system at time t . More precisely, $X(t) \in \mathbb{Z}^u = (X_1(t), \dots, X_u(t))$ so each element $X_i(t)$ gives the number of individuals in state i at that point in time. The hazard function returns a non-negative number which is proportional to the probability of event j happening *at time* t . CTMC models constructed in this manner are known as Poisson systems and are discussed for population models in Matis & Kiffe (2000), and defined rigorously in Brémaud (2013).

So we have a concrete example in mind, consider the Poisson system in Figure 1. It has $u = 3$ so the state vector is $X(t) = (X_1(t), X_2(t), X_3(t))$. There are $v = 6$ events which are seen by the edges on the graph, which move individuals between compartments. Edges that originate in the empty set \emptyset correspond to “birth” events (there is a single 1 in their column of \mathbf{S}), and edges that lead to \emptyset are “death” events (there is a single -1 in their column of \mathbf{S}). All other events take one individual from the originating compartment and put them in the recipient compartment when they occur. The edges are labeled as N_{ij} , which is a random variable telling us how many times this event happens over a time step. In an exact simulation of the model, because multiple events are not allowed to occur simultaneously in CTMC models, over a time step of infinitesimal size δt , all N_{ij} except (possibly) one will be equal to 0 (the one that is not 0 is equal to 1 is the one that happened). If everything is 0 it means nothing happened on that step.

In the model, individuals are born at some rate $\lambda_{\emptyset \rightarrow 1}$ and then arrive at compartment 1. Once there, they can either die, or advance to compartment 2. Once in compartment 2, they can either die, go back to compartment 1, or go to compartment 3. Compartment 3 is a sink, so that anyone who doesn’t die ends up there and never leaves (such “sink” compartments are sometimes handy in simulation to denote, for example, the cumulative number of deaths from a certain disease).




Figure 1: Continuous Time System

This model has the following stoichiometry matrix:

$$\mathbf{S} = \begin{matrix} & \emptyset \rightarrow 1 & 1 \rightarrow \emptyset & 1 \rightarrow 2 & 2 \rightarrow \emptyset & 2 \rightarrow 1 & 2 \rightarrow 3 \\ \begin{matrix} X_1 \\ X_2 \\ X_3 \end{matrix} & \begin{pmatrix} 1 & -1 & -1 & 0 & 1 & 0 \\ 0 & 0 & 1 & -1 & -1 & -1 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \end{matrix}$$

Now that we have this picture in mind, let's think about ways to make an approximate algorithm that is faster than exact simulation because it can take larger time steps $\Delta t \gg \delta t$. This will be faster than, say Gillespie's method (or any other exact algorithm) when there are a lot of individuals in the system, meaning that for a given Δt we can simulate *a lot* of things happening in one go, rather than simulating each event individually as Gillespie's method would. One method, commonly used, is known as tau-leaping (Gillespie 2001). It assumes that over Δt none of the hazards change very much so $\lambda_j(t + \Delta t, X(t + \Delta t)) \approx \lambda_j(t, X(t))$. When that holds, the total number of j events that occur over $[t, t + \Delta t)$ is a Poisson random variable with parameter $\lambda_j(t, X(t))\Delta t$, which is an Euler type approximation (ie; the amount of things happening over the interval Δt is about equal to the rate they were happening at the start of the interval, times the width of the interval).

Another way to make an approximation, which may be more reasonable for models of interacting individuals than for models of chemical reactions, is to say that over Δt each individual in a compartment can only experience a single event (somewhat Zen-like, this also includes the event of "no event"). This is an approximation because no such requirement is placed on the exact CTMC model. If one labeled all individuals in a certain compartment, there is nothing except extreme improbability of occurrence preventing the same individual experiencing two consecutive events in a very short time, and all others experiencing no events. The single jump approximation restricts this. Looking at Figure 1. For an individual in, say, compartment 2, over $[t, t + \Delta t)$ they can either die (go to \emptyset), go to 3, go back to 1, or do nothing (stay in 2), *but* once one of those happens, that individual cannot experience any more events over the time step.

Let's get rigorous about this approximation to see what it says in terms of probability. For each individual in compartment 2, there are $m = 4$ things that can happen to them, one for each outward edge plus the event of no event. Let's collect them in a probability vector $P(X_2)$. Because in the model, individuals are indistinguishable, this probability vector is the same for everyone in compartment 2. Let's also collect all the Bernoulli random variables for this individual that say which thing happens and what didn't happen in a binary vector $\delta(X_2)$. Again, because individuals are indistinguishable, by adding all these binary vectors for each person in compartment 2, we get the vector $\Delta(X_2)$, where each element is an integer that says how many

times that event (corresponding to the event with the probability in $P(X_2)$) happened over the time step. That means we know at the end of the time step where each individual in compartment 2 went (including back to 2).

$$P(X_2) = (P(22), P(21), P(23), P(2\emptyset))$$

$$\Delta(X_2) = (\tilde{N}_{22}, \tilde{N}_{21}, \tilde{N}_{23}, \tilde{N}_{2\emptyset})$$

We put the tilde on the elements of $\Delta(X_2)$ because they aren't the same random variables depicted in Figure 1. Those random variables are only allowed to be 0s or 1s, because the model is defined by infinitesimal time steps, so only nothing or at most one thing can happen at a time, whereas these ones can be any positive number less than or equal to the number of individuals in the compartment ($\tilde{N}_{2j} \leq X_2(t)$). One consequence of our approximation is that we don't know *when* in the interval $[t, t + \Delta t)$ each individual went somewhere, only that it happened, so the model turns from a continuous-time model to a discrete-time one, in fact it says that each individual calculates their probabilities to go somewhere at the start of the time step and then samples where they go at the end of the time step. For a rigorous proof that this discrete-time model converges to the continuous-time one as $\Delta t \rightarrow 0$, see the supplement to Bretó et al (2009). The model will have the compartmental diagram shown in Figure 2, where now self-transitions are allowed.




Figure 2: Discrete Time (Single Jump) System

Now let's work out the elements of $P(X_2)$. Let's begin by calculating $P(22)$, the probability that nothing happens. Let $\lambda_{2\bullet}(t, X(t)) = \lambda_{21}(t, X(t)) + \lambda_{23}(t, X(t)) + \lambda_{2\emptyset}(t, X(t))$ (where \bullet means "to anywhere").

$$P(22) = e^{-\sum_j \lambda_{2j}(t, X(t))\Delta t} = e^{-\lambda_{2\bullet}(t, X(t))\Delta t}$$

It is the survival function of an exponential random variate with rate parameter equal to $\lambda_{2\bullet}(t, X(t))$ evaluated at Δt . To calculate all the other probabilities, follow a formula that says the probability of the event happening is the probability of *anything* happening times the probability of that specific event happening. For example:

$$P(23) = [1 - P(22)] \left(\frac{\lambda_{23}(t, X(t))}{\lambda_{2\bullet}(t, X(t))} \right)$$

For some state i which has $j = 1, \dots, m$ possible destinations the general form is:

$$P(ii) = e^{-\sum_j \lambda_{ij}(t, X(t))}$$

$$P(ij) = \left[1 - e^{-\sum_j \lambda_{ij}(t, X(t))} \right] \left(\frac{\lambda_{ij}(t, X(t))}{\sum_j \lambda_{ij}(t, X(t))} \right)$$

That means that the random vector $\Delta(X_2)$ follows a Multinomial distribution:

$$\Delta(X_2) \sim \text{Multinomial}(X_2(t), P(X_2))$$

Or, put generally for the compartment i :

$$\Delta(X_i) \sim \text{Multinomial}(X_i(t), P(X_i))$$

These probabilities and random variables tell us how to deal with death and transfer events in Markov population processes, but we've neglected birth events. Luckily those are much easier. Like the rest, each birth event $N_{\emptyset i}$ has a function that gives the intensity it occurs as, $\lambda_{\emptyset i}(t, X(t))$. Because over the time step $X(t)$ doesn't change (it gets updated "all at once" at the end; think of synchronous versus asynchronous updating in cellular automata), the total "integrated" intensity is just $\lambda_{\emptyset i}(t, X(t))\Delta t$. Under the assumptions of the Markov model births are independent (in plain language, it means that if we *already* know lots of births are happening now due to the rate function, actually seeing one happen doesn't mean *even more* will happen), the random variable giving the number of birth events into i over the time step has the following distribution:

$$\tilde{N}_{\emptyset i} \sim \text{Poisson}(\lambda_{\emptyset i}(t, X(t))\Delta t)$$

So to simulate trajectories from the model we only need to use Poisson and Multinomial random variables (if a compartment only has one way out, then we just use a Binomial random variable, the "base case" of the Multinomial).

Implementation in R

What we want to be able to do in R is take any model expressed as a stoichiometry matrix \mathbf{S} and a vector of hazard functions $\Lambda(t, X(t))$ and then, given a desired Δt , return a function that takes as input $X(t)$ and returns the updated state vector at the end of the time step, $X(t + \Delta t)$.

References

- Bretó, Carles, et al. "Time series analysis via mechanistic models." *The Annals of Applied Statistics* 3.1 (2009): 319-348.
- Kingman, J. F. C. "Markov population processes." *Journal of Applied Probability* 6.1 (1969): 1-18.
- Matis, James H., et al. *Stochastic population models: a compartmental perspective*. Vol. 145. Taylor & Francis, 2000.
- Brémaud, Pierre. *Markov chains: Gibbs fields, Monte Carlo simulation, and queues*. Vol. 31. Springer Science & Business Media, 2013.
- Gillespie, Daniel T. "Approximate accelerated stochastic simulation of chemically reacting systems." *The Journal of chemical physics* 115.4 (2001): 1716-1733.