Lecture 5: Agent Based Simulations of Differential Equations

Radu Balan

Department of Mathematics, NWC
University of Maryland, College Park, MD

Version: February 5, 2024

Differential Equations

In this lecture we discuss *Agent Based Simulations* of differential equations $\frac{dx}{dt} = f(t,x)$. where $x \in \mathbb{R}^n$. We shall focus on the special class *Linear Systems of Differential Equations with constant coefficients* :

$$\frac{dx}{dt} = Rx \quad , \quad x(0) = x_0$$

that have an additional conservation law:

$$\frac{d}{dt}(x_1 + x_2 + \cdots + x_n) = 0.$$

This last condition is equivalent to asking that the matrix R^T has the constant vector 1 in its null space: $R^T 1 = 0$, or, equivalently, $1^T R = 0$, i.e., R has a left null-vector. Indeed this is the case because:

$$\frac{d}{dt}(x_1 + \cdots + x_n) = \frac{d}{dt}(1^T x) = 1^T Rx = 0$$

Linear Systems with Constant Coefficients

The analytic solution

The unique solution of this linear system of ODEs is given by

$$x(t) = e^{Rt}x(0)$$

where the matrix exponential $e^{Rt} = \sum_{k=0}^{\infty} \frac{t^k}{k!} R^k$ can be computed more easily when the matrix R diagonalizes: If $R = T \Lambda T^{-1}$ for some invertible matrix $T \in \mathbb{R}^{n \times n}$ and diagonal matrix Λ (that contains its eigenvalues), then

$$e^{Rt} = T e^{\Lambda t} T^{-1} \quad , \quad e^{\Lambda t} = diag(e^{\lambda_1 t}, \cdots, e^{\lambda_n t}).$$

An alternative way of computing the matrix exponential is to use the inverse Laplace transform: $e^{Rt} = \mathcal{L}^{-1}\{(sI - R)^{-1}\}|_t$.

Linear Systems with Constant Coefficients

The analytic solution (2)

Note the following result:

Lemma

If $1^T R = 0$ then $1^T e^{Rt} = 1^T$, for every t.
Conversely, if $1^T e^{Rt} = 1^T$, for every t then $1^T R = 0$.

This result is shown easily from the series definition of the matrix exponential: if $1^T R = 0$ then

$$1^T e^{Rt} = 1^T (I + Rt + \frac{t^2}{2} R^2 + \frac{t^3}{3!} + \cdots) = 1^T I + 0 + 0 + \cdots = 1^T$$

Conversely, if $1^T e^{Rt} = 1^T$, then taking the derivative at $t = 0$ yields $1^T R = 0$.

This lemma establishes the conservation law mentioned earlier: $1^T x(t) = 1^T x(0)$ for all t, if and only if $1^T R = 0$.

Systems of Differential Equations

Numerical Simulations: The Euler Method

Consider the general differential equation, or system of first order differential equations, $\frac{dx}{dt} = f(t,x)$ with initial condition, $x(0) = x_0$. Here $x = x(t) \in \mathbb{R}^n$ is a vector-valued function. The Euler method estimates numerically the solution at any time $t > 0$:

Algorithm (The Euler Method)

Inputs: *The function* $f = f(t,x)$, *step size* $h > 0$, *end time* T_{max}.

1. *Initialize* $k = 0$, $t_0 = 0$, $x_0 = x(0)$.
2. *Repeat the following steps until* $t \geq T_{max}$:
   1. *Compute:* $v = f(t_k, x_k)$,
   2. *Update:* $x_{k+1} = x_k + hv$,
   3. *Update:* $t_{k+1} = t_k + h$,
   4. *Increment:* $k = k + 1$.

Outputs: *Number of iterations: last value of* k *(it is also equal to* $\lceil \frac{T_{max}}{h} \rceil$ *);The estimated solution at* T_{max} *is the last value* x_k.

The Problem: Agents based Simulations

The problem we like to solve is to present a stochastic method that solves the same differential equation. We analyze in detail the IVP $\frac{dx}{dt} = Rx$, $x(0) = x_0 \in (\mathbb{R}^+)^n$ when $1^T R = 0$. This particular linear system of differential equations is very much related to the SIR and SEIR models.
Our problem is to find a Markov chain with matrix of transition probabilities Π so that it approximates well the solution $x(T_{max})$ at a future time T_{max}.
The strategy is to initialize a large number of "agents", say M, each in one of n compartments. The initial number in each compartment is propostional to the entries of $x(0)$. Then model is driven forward so that at each time step agents can move from one compartment to another according to transition probabilities in Π. After a number of steps, the distribution of agents is converted back into a solution $x(t)$.
Details follow in next slides.

The Compartment Model

Imagine the following scenario: There are n compartments (or, states) indexed from 1 to n containing a total of M agents. Initially, the compartment i has $M_i(0)$ agents, $i \in [n]$. Differently said, $M_i(0)$ agents are initialized to state i (that is, in compartment i), for each i from 1 to n. For each agent k, $1 \leq k \leq M$, we denote by $S_k(t)$ its state at time t.Thus $S_k(t) \in \{1, 2, \cdots, n\}$, for every k and t.
Assume we are given a *transition probability matrix* $\Pi \in [0,1]^{n \times n}$, known also as *probability matrix*, or *stochastic matrix*, or *transition matrix*. At every T_0 units of time (e.g., at every T_0 seconds) agents can change state (that is, move from one compartment to another) or stay in the same state randomly according to a certain transition probability: the probability for an agent to go from compartment i to compartment j is given by $\Pi_{j,i}$, $i,j \in [n]$,

$$Prob[S_k(t+T_0) = j | S_k(t) = i] = \Pi_{j,i} \quad , \quad i,j \in [n] \ , \ k \in [M]$$

The Compartment Model

For the matrix Π to define a probability matrix it is necessary and sufficient that: (1) each $\Pi_{i,j} \geq 0$ (i.e., non-negative entries), and (2) for each $i \in [n]$, $\sum_{j=1}^{n} \Pi_{j,i} = 1$ (i.e., each column sums to 1). The second condition reflects the fact that each agent has to arrive in one of existing compartments. Hence no agent is born, or dies, in this model. The transition diagram is rendered in Figure 1.

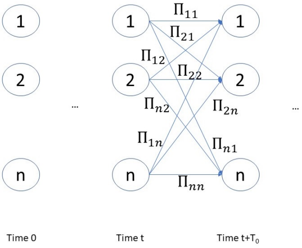

Figure: A rendition of the Markov chain associated to each agent: at time $t + T_0$

The Compartment Model (2)

Let $M_i(t)$ denote the number of agents in state i at time t. At time 0, the initial number of agents in state i is $M_i(0)$. The total number of agents is $M = \sum_{i=1}^{n} M_i(0)$ and is constant in time. Each of the M agents has a state variable $S_k(t)$, $1 \leq k \leq M$, that denotes the compartment (i.e., state) the agent k is in at time t. For instance, at time 0, the first $M_1(0)$ agents are in state 1, the subsequent $M_2(0)$ agents are in state 2, and so on:

$$\forall k \in [M] \; , \; S_k(0) = i \;\; \text{if and only if} \;\; \sum_{j=1}^{i-1} M_j(0) < k \leq \sum_{j=1}^{i} M_j(0)$$

As time evolves, agent states change and so do the number of agents in each compartment. These states and numbers become random variables: for each agent k, and compartment (state) i, $S_k(t)$ and $M_i(t)$ are random variables, $S_k(t) \in \{1, 2, \cdots, n\}$ and $M_i(t) \in \{1, 2, \cdots, M\}$.

The Compartment Model (3)

For simulation of time-invariant processes (differential equations or stochastic processes), the transition probability matrix Π is independent of t. In general, the compartment model framework allows for modeling and simulations of time-dependent processes.

Our first immediate goal is to find a relationship between the matrix R that characterizes the linear system of differential equations $\frac{dx}{dt} = Rx$ and the transition matrix Π of this Compartment model so that, at every time t that is multiple of T_0, $t = pT_0$ for some integer $p \geq 0$, the expectation of the number of agents in state i matches the exact solution at time t, when the linear system and the Markov chain process are initialized with the same vector $(M_1(0), M_2(0), \cdots, M_n(0))$:

$$\mathbb{E}\begin{bmatrix} M_1(pT_0) \\ \vdots \\ M_n(pT_0) \end{bmatrix} = e^{RpT_0}\begin{bmatrix} M_1(0) \\ \vdots \\ M_n(0) \end{bmatrix}. \quad (*)$$

The Compartment Model (4)

For the Markov chain depicted in Figure 1, the expectation is computed as follows. Let $I_{k,i}(t)$ denote the indicator function of whether agent k is in state (compartment) i at time t: $I_{k,i}(t) = 1$ if $S_k(t) = i$, and it is 0 otherwise. Thus

$$M_i(t) = \sum_{k=1}^{M} I_{k,i}(t) \quad , \quad \sum_{i=1}^{n} I_{k,i}(t) = 1$$

and:

$$\mathbb{E}[M_i(t)] = \sum_{k=1}^{M} \mathbb{E}[I_{k,i}(t)] = \sum_{k=1}^{M} Prob(I_{k,i}(t) = 1) = \sum_{k=1}^{M} Prob(S_k(t) = i).$$

Using the transition probability matrix Π and marginalizing the joint distribution $(S_k(t+T_0), S_k(t))$:

$$Prob(S_k(t+T_0) = i) = \sum_{j=1}^{n} Prob(S_k(t+T_0) = i, S_k(t) = j) =$$

The Compartment Model (5)

$$= \sum_{j=1}^{n} Prob(S_k(t+T_0)=i|S_k(t)=j)Prob(S_k(t)=j) = \sum_{j=1}^{n} \Pi_{i,j} Prob(S_k(t)=j)$$

$$\mathbb{E}[M_i(t+T_0)] = \sum_{k=1}^{M}\sum_{j=1}^{n} \Pi_{j,i} Prob(S_k(t)=j) = \sum_{j=1}^{n} \Pi_{i,j} \sum_{k=1}^{M} Prob(S_k(t)=j) = \sum_{j=1}^{n} \Pi_{i,j}\mathbb{E}[M_j(t)$$

In matrix notation:

$$\mathbb{E}\begin{bmatrix} M_1(t+T_0) \\ \vdots \\ M_n(t+T_0) \end{bmatrix} = \Pi\,\mathbb{E}\begin{bmatrix} M_1(t) \\ \vdots \\ M_n(t) \end{bmatrix}.$$

Iterating over t, we obtain:

$$\mathbb{E}\begin{bmatrix} M_1(pT_0) \\ \vdots \\ M_n(pT_0) \end{bmatrix} = \Pi^p \begin{bmatrix} M_1(0) \\ \vdots \\ M_n(0) \end{bmatrix} \quad , \quad p = 0,1,2,\ldots \quad (**)$$

The Compartment Model (5)

Comparing (*) and (**) we observe that equation (*) is satisfied if:

$$\boxed{\Pi = e^{T_0 R}}.$$

One natural question: what properties of R imply that $e^{T_0 R}$ is a probability transition matrix?
The following result answers this question[1].

Proposition

Let R be a $n \times n$ matrix with real entries. The following are equivalent:

1. *For every $\tau \geq 0$, $e^{\tau R}$ is a transition probability matrix, i.e., has non-negative entries, and each column sums to one.*
2. *The matrix R satisfies two conditions:*
   1. *Every off-diagonal element of R is non-negative: $\forall i \neq j$, $R_{i,j} \geq 0$.*
   2. *Constant vector 1 is a right null-vector for R^T: $R^T 1 = 0$.*

[1] condition 2.2 was discussed earlier

The Compartment Model (6)

Why the Proposition?

$\Rightarrow$

Assume $e^{\tau R}$ defines a probability matrix for every $\tau \geq 0$. Thus: $1^T e^{\tau R} = 1^T$ (each columns sums to 1). Take derivative with respect to τ at $\tau = 0$ and obtain: $1^T R = 0$. Hence $R^T 1 = 0$.

On the other hand, for small $\tau > 0$, $e^{\tau R} \approx I + \tau R$. Hence, if $R_{i,j} < 0$ for some $i \neq j$, then the (i,j) entry of $e^{\tau R}$ will be negative for small enought τ. Contradiction. Hence all off-diagonal entries in R must be non-negative.

$\Leftarrow$

Assume $R^T 1 = 0$ and $R_{i,j} \geq 0$ for all $i \neq j$. First note $1^T R = 0$ and $1^T R^k = 0$ for all $k = 1, 2, 3, \ldots$. Using the power series for the matrix exponential, obtain that

$$1^T e^{\tau R} = 1^T \left(I + \sum_{k=1}^{\infty} \frac{\tau^k}{k!} R^k \right) = 1^T I = 1^T$$

Thus each column of $e^{\tau R}$ sums to one.

The Compartment Model (7)

Why the Proposition? (cont'ed)

Assume now that $R_{i,j} \geq 0$ for all $i \neq j$. Choose a real α large enough so that the matrix $R + \alpha I$ has all entries non-negative. Using the power series formula deduce that $e^{\tau(R+\alpha I)}$ has non-negative entries. But $R = (R + \alpha I) - \alpha I$ and

$$e^{\tau R} = e^{\tau(R+\alpha I)} e^{-\alpha \tau I} = e^{\tau(R+\alpha I)} diag(e^{-\alpha \tau}) = e^{-\alpha \tau} e^{\tau(R+\alpha I)}$$

where the entries in the last matrix are all non-negative. Thus $e^{\tau R}$ has non-negative entries.

□

How to simulate an Agent based Model

Assume you know the transition matrix Π and the initial population of agents $(M_1(0), \cdots, M_n(0))$.
Initialize agent states accordingly, $S_1(0), \cdots, S_M(0) \in [n]$.
For each $i \in [n]$, pre-compute and save in hash tables $a_{1,i} = 0$, $b_{1,i} = \Pi_{1,i}$, $a_{2,i} = b_{1,i}$, $b_{2,i} = b_{1,i} + \Pi_{2,i}$, $a_{3,i} = b_{2,i}$, $b_{3,i} = b_{2,i} + \Pi_{3,i}$ and so on until $a_{n,i} = 1 - \Pi_{n,i}$, $b_{n,i} = 1$.

At each time step $t = pT_0$, for each agent $k \in [M]$:

- Draw a random variable z uniformly distributed in $[0, 1)$.
- If $S_k(t) = i$, then find $j \in [n]$ so that $z \in [a_j, b_j)$.
- Assign $S_k(t + T_0) = j$.

Repeat until reach end of simulation time T_{max}.