

Computational Methods in Stochastics

Lecture IV

Markov Chains and Stochastic Models

In what follows, the basis for doing **Markov chain Monte Carlo** simulations (**MCMC**) is given.

We introduce the notation used by Wilkinson, which is better suited for continuous state spaces (at the end of this lecture). Accordingly, some stuff is restated in this notation.

The concepts of the following Markov chain part are necessary for understanding and justifying e.g. Hamiltonian MC method.

In what follows, Markov processes and stochastic models are dealt with in (i) **discrete time and discrete state-space** (ii) **discrete time and continuous state-space** (iii) **continuous time and discrete state-space**.

Markov Processes and Stochastic Models in Discrete Time and Discrete State-space

Markov Chains

Stationary distributions

A distribution is said to be a *stationary distribution* of the homogeneous Markov chain governed by the matrix \mathbf{P} if

$$p = p\mathbf{P}.$$

So, p is a row eigenvector of the transition matrix with eigenvalue equal to 1.

If at some time n we have $p^{(n)} = p$, then $p^{(n+1)} = p^{(n)}\mathbf{P} = p\mathbf{P} = p$.

Similarly, $p^{(n+k)} = p, \forall k \geq 0$.

This is saying that, if a chain has a stationary distribution, it retains that distribution for all future time.

Markov Chains

Convergence

Convergence to a stationary distribution is a rather technical subject. We won't go into that. Suffice it to say that one hopes to find the *equilibrium distribution*, i.e. a stationary distribution to which p converges. Convincing oneself that a found stationary distribution is also the equilibrium distribution can be hard, to say the least.

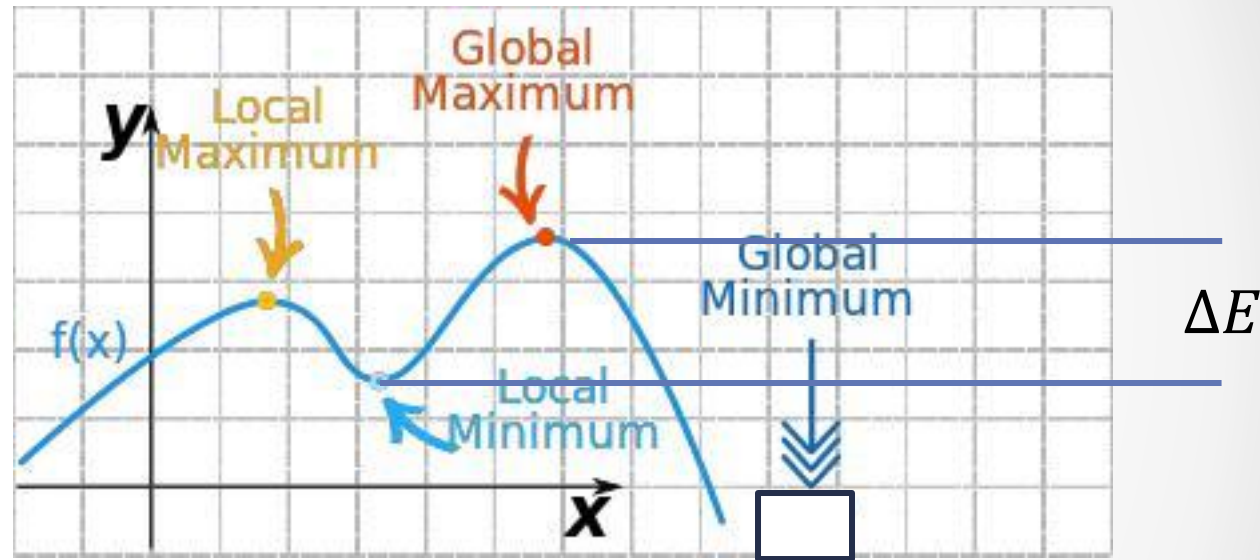
Markov Chains

Finding a stationary state by MC means to a physicist or chemist finding an *energy minimum* and to a statistician finding *maximum likelihood*.

If $y = f(x)$ is proportional to the (potential) energy at x , then the likelihood for corresponding stationary distributions (of parameters) has maxima at the energy minima.

In physics etc.:

$P_{ij} \propto e^{-\Delta E/k_B T}$ ← **Boltzmann weight** – everyone should know this.



Detailed balance

At a stationary state the detailed balance holds:

$$p_i P_{ij} = p_j P_{ji}$$


Markov Chains

In equilibrium there's a stationary distribution:
 $\Delta E(\mathbf{x}^A, \mathbf{x}^B)$ between certain states $\mathbf{x}^A = (x_1^A, x_2^A, \dots)$ and $\mathbf{x}^B = (x_1^B, x_2^B, \dots)$ is constant. The system (Markov chain) has reached a steady state and we are typically sampling states from a stationary distribution corresponding to the global or a local maximum likelihood / minimum energy.

The Boltzmann weight $e^{-\Delta E/k_B T}$ holds for all, i.e. also large, ΔE (see the pic on the previous slide). The distributions in the two states can be very different and changing in time.

Markov Chains and MC

Notation in Wilkinson's book:

The set $\{\theta^{(t)} | t = 0, 1, 2, \dots\}$ is a **discrete time stochastic process**. The **state space** is such that $\theta^{(t)} \in S, \forall t$, and may be discrete or continuous. (In lecture 3 and the other book θ were parameters and X_n were states.) ←  Yes, it's frustrating.

A Markov chain: For $A \subseteq S, t = 0, 1, 2, \dots,$

$$\begin{aligned} & P(\theta^{(t+1)} \in A | (\theta^{(t)} = x, \theta^{(t-1)} = x_{t-1}, \dots, \theta^{(0)} = x_0)) \\ &= P(\theta^{(t+1)} \in A | \theta^{(t)} = x), \forall x, x_{t-1}, \dots, x_0 \in S \end{aligned}$$

In the case of no t dependence, $P(\theta^{(t+1)} \in A | \theta^{(t)} = x) = P(x, A), \forall t$, and the Markov chain is homogeneous; the *transition kernel* $P(x, A)$ (previously one-step transition probability $P_{ij}^{n, n+1} = P_{ij}$) determines the behaviour of the chain.

Markov Chains and MC

Writing $P(x, y) = P(\theta^{(t+1)} = y | \theta^{(t)} = x)$, the **stochastic transfer matrix** becomes

$$P = \begin{pmatrix} P(x_1, x_1) & \cdots & P(x_1, x_r) \\ \vdots & \ddots & \vdots \\ P(x_r, x_1) & \cdots & P(x_r, x_r) \end{pmatrix}$$

The i th row is the probability distribution of the values of $\theta^{(t+1)}$ under the condition that $\theta^{(t)} = x_i$ (note that here the index i starts from 1 whereas it started from 0 in the previous strictly discrete notation).

We denote:

$$\begin{aligned} P(\theta^{(t)} = x_1) &= \pi^{(t)}(x_1) \\ P(\theta^{(t)} = x_2) &= \pi^{(t)}(x_2) \\ &\vdots \\ P(\theta^{(t)} = x_r) &= \pi^{(t)}(x_r) \end{aligned}$$

Markov Chains and MC

→ As an r -dimensional row vector

$$\pi^{(t)} = \left(\pi^{(t)}(x_1), \pi^{(t)}(x_2), \dots, \pi^{(t)}(x_r) \right).$$

The probability density at time $t + 1$ can be computed as a **total probability** (old notation: $\Pr\{\cdot\} = P(\cdot)$):

$$\begin{aligned} P(\theta^{(t+1)} = x_1) &= P(x_1, x_1) \pi^{(t)}(x_1) + P(x_2, x_1) \pi^{(t)}(x_2) + \dots \\ &\quad + P(x_r, x_1) \pi^{(t)}(x_r) \end{aligned}$$

Similarly for $P(\theta^{(t+1)} = x_2)$ etc. →

$$\begin{aligned} &\left(\pi^{(t+1)}(x_1), \pi^{(t+1)}(x_2), \dots, \pi^{(t+1)}(x_r) \right) = \\ &= \left(\pi^{(t)}(x_1), \pi^{(t)}(x_2), \dots, \pi^{(t)}(x_r) \right) \times \begin{pmatrix} P(x_1, x_1) & \dots & P(x_1, x_r) \\ \vdots & \ddots & \vdots \\ P(x_r, x_1) & \dots & P(x_r, x_r) \end{pmatrix} \end{aligned}$$

$$\Leftrightarrow \pi^{(t+1)} = \pi^{(t)} P$$

Markov Chains and MC

The way an **MC simulation** would proceed:

$$\pi^{(1)} = \pi^{(0)} P$$

$$\pi^{(2)} = \pi^{(1)} P = \pi^{(0)} P P = \pi^{(0)} P^2$$

⋮

$$\pi^{(t)} = \pi^{(0)} P^t$$

So, if the one-step transition matrix is P , then the m -step transition matrix is P^m .

Also, the $(m + n)$ -step transition matrix is $P^m P^n = P^{m+n}$.

The set of linear equations corresponding to the above statement are called the *Chapman-Kolmogorov equations*.

Markov Chains and MC

Derivation of Chapman-Kolmogorov equations

Define $P(n)$ to be the n -step transition matrix with elements $p_{ij}(n)$.

Due to the theorem of total probability,

$$\begin{aligned} p_{ij}(m+n) &= P(\theta^{(m+n)} = j | \theta^{(0)} = i) \\ &= \sum_{k=1}^r P(\theta^{(m+n)} = j | \theta^{(m)} = k) P(\theta^{(m)} = k | \theta^{(0)} = i) \end{aligned}$$

Chapman-Kolmogorov equations

$$\Leftrightarrow p_{ij}(m+n) = \sum_{k=1}^r p_{kj}(n) p_{ik}(m).$$

Two identical indices
are always summed
over.

Typically in the literature, Einstein's convention is used:

$$\Leftrightarrow p_{ij}(m+n) = \sum_{k=1}^r p_{ik}(m) p_{kj}(n) = p_{ik}(m) p_{kj}(n).$$

Markov Chains and MC

Stationary distributions

$\pi = \pi P$ is a row eigenvector of P with corresponding eigenvalue $\lambda = 1$ (and a fixed point of the linear map induced by P).

$$\pi^{(n+1)} = \pi^{(n)}P = \pi P = \pi, \text{ or } \pi^{(n+k)} = \pi, \forall k.$$

Since, $\pi = \pi P \iff \pi - \pi P = 0$,

the stationary distribution can be solved from $\pi(I - P) = 0$.

Markov Chains and MC

Convergence

... in this notation: $\lim_{t \rightarrow \infty} \pi^{(t)} = \pi$ (for t discrete or continuous).

Since $\pi = \pi P$ is a row eigenvector of P , there is a corresponding eigenvalue λ : $\pi P = \lambda \pi \Rightarrow \pi P^t = \lambda^t \pi$.

For a stochastic matrix P , $|\lambda| \leq 1$ (row elements sum to 1). \rightarrow Distributions $\pi^{(t)}$ converge as t grows (exponential propagation $\sim e^{|\lambda|t}$).

Markov Chains and MC

Some Important Properties and Special Cases

Reversibility and Detailed Balance

If $\theta^{(0)}, \theta^{(1)}, \dots, \theta^{(N)}$ is a Markov chain, then so is $\theta^{(N)}, \theta^{(N-1)}, \dots, \theta^{(0)}$. (For proof, see Wilkinson, 5.2.5.)

The transition kernel for the reversed chain:

$$\begin{aligned} P_t^*(x, y) &= P(\theta^{(t)} = y | \theta^{(t+1)} = x) \\ &= \frac{P(\theta^{(t+1)} = x | \theta^{(t)} = y) P(\theta^{(t)} = y)}{P(\theta^{(t+1)} = x)} \quad (\text{Bayes Theorem}) \\ &= \frac{P(y, x) \pi^{(t)}(y)}{\pi^{(t+1)}(x)}. \end{aligned}$$

→ **The reversed chain is, in general, not homogeneous**, that is, transition probabilities depend on time.

Markov Chains and MC

However, in equilibrium $\pi^{(t+1)}(x) = \pi^{(t)}(x) = \pi$, so

$$P^*(x, y) = \frac{P(y, x)\pi(y)}{\pi(x)}.$$

→ **In equilibrium** (the distribution has reached equilibrium) **the reversed chain is homogeneous.**

If $P^*(x, y) = P(x, y)$, $\forall x, y$, then the chain is said to be (time) *reversible*. In equilibrium (stationarity), for reversible chains *detailed balance* holds:

$$\pi(x)P(x, y) = \pi(y)P(y, x), \quad \forall x, y.$$

Detailed balance states the symmetry in the flow of probability between two states: Probability of a transition from x to y = probability of a transition from y to x $\forall(x, y)$.

Markov Chains and MC

If detailed balance holds, then stationarity and reversibility hold for the Markov chain.

Proof.

$$\pi(x)P(x, y) = \pi(y)P(y, x), \quad \forall x, y.$$

Summing over x ,

$$\sum_x \pi(x)P(x, y) = \sum_x \pi(y)P(y, x) = \pi(y) \sum_x P(y, x) = \pi(y), \forall y$$

$\Leftrightarrow \pi P = \pi$. (The chain has a stationary distribution.)

Then as $\pi(x) = \pi(y)$: $\pi(x)P(x, y) = \pi(y)P(y, x)$

$\Leftrightarrow P(x, y) = P(y, x), \forall x, y$. (The chain is reversible.)

Markov Chains and MC

So, summing up the things stated of the relation between the three properties of a Markov chain:

detailed balance \leftrightarrow stationarity & reversibility.

Think about these conditions a bit: for example, the order of strictness.

Markov Chains and MC

Stochastic simulation and analysis

Simulation of a Markov chain = simulation of a new state randomly, when probabilities are given by a probability vector p .

An MC simulation of a Markov chain proceeds as:

1. A transition matrix is given or deduced for the system.
2. The initial distribution $\pi^{(0)} = (\pi^{(0)}(x_1), \pi^{(0)}(x_2), \dots, \pi^{(0)}(x_r))$ is given.
3. Sample an initial state $\theta^{(0)}$ from $\pi^{(0)}$ using a lookup method.
4. Sample the state for $\theta^{(t)}$ using the set of probabilities from the $\theta^{(t-1)}$ th row of P for $t = 1, 2, \dots$
($P(\theta^{(t)} = x_i) = \pi^{(t)}(x_i)$, $i \in [1, r]$, and $\pi^{(t)} = \pi^{(t-1)}P$).

Markov Chains and MC

A layman's statement on simulations of a Markov process:

When the distribution changes in time, which means the system is out of equilibrium and the states are non-stationary, one can simulate starting from an initial distribution $\pi^{(0)}$ simply using $\pi^{(t)} = \pi^{(t-1)}P$ to explore how (and if) the system approaches equilibrium. In some cases there may be several equilibrium states and where one ends up depends on $\pi^{(0)}$. This is the “dynamical aspect”.

When we are in the stationary (that may be an equilibrium) state, we are interested in sampling different realisations of the distribution in that state. Then we apply the method outlined on the previous slide. In order to reach the stationary state, we may first do the above procedure sufficiently long so that $\pi^{(t)} \rightarrow \pi$. Then we can use $\pi^{(0)} = \pi$ to start sampling $\theta^{(t)}$. (Continues \rightarrow)

Markov Chains and MC

Sampling $\theta^{(t)}$ for a Markov process means doing Monte Carlo using the probabilities in the $\theta^{(t-1)}$ th row of P . (This means the row corresponding the value $\theta^{(t-1)} = x$; the x^{th} row includes transition probabilities from x to all states ($P(x, y) = P(\theta^{(t)} = y | \theta^{(t-1)} = x)$)).

Markov Processes and
Stochastic Models in
Discrete Time and
Continuous State-space

Markov Chains and MC

Markov chains with continuous state-spaces

(still discrete time)

When, for example, **sampling from continuous distributions** using a Markov chain (MCMC), one is working with a continuous state-space.

A commonly used model for such simulations in time-series analysis is the *first-order auto-regressive model*, **AR(1)**.

The model for an AR(1) process $\{Z_t | t = 1, 2, \dots\}$ can be summarised as

$$Z_t = \alpha Z_{t-1} + \epsilon_t, \epsilon_t \sim N(0, \sigma^2).$$

This means: The value of the stochastic process at time t , Z_t , depends on the value of the the stochastic process at time $t - 1$.

Markov Chains and MC

The non-deterministic part of the AR(1) process is given by the noise term ϵ_t . The noise process $\{\epsilon_t | t = 1, 2, \dots\}$ is assumed to be independent: the pair of random quantities ϵ_i and ϵ_j are independent of one another $\forall i \neq j$.

The conditional distribution of Z_t given $Z_{t-1} = z_{t-1}$:

$$Z_t | (Z_{t-1} = z_{t-1}) \sim N(\alpha z_{t-1}, \sigma^2) \quad (\text{The mean of the Gaussian changes.})$$

AR(1) is a Markov chain and its state-space $S = \mathbb{R}$ (continuous).

(Higher-order AR-models are **not** Markov chains, e.g. AR(2):
 $Z_t = \alpha Z_{t-1} + \beta Z_{t-2} + \epsilon_t, \epsilon_t \sim N(0, \sigma^2)$.)

Markov Chains and MC

Transition kernels

For a *homogeneous chain*, define $P(x, A) = P(\theta^{(t+1)} \in A | \theta^{(t)} = x)$.

$P(x, \{y\}) = 0$ for continuous state-spaces \rightarrow define

$$\begin{aligned} P(x, y) &= P(\theta^{(t+1)} \leq y | \theta^{(t)} = x) \\ &= P(\theta^{(1)} \leq y | \theta^{(0)} = x), \quad \forall x, y \in S. \end{aligned}$$

This is a **conditional cumulative distribution function (CDF)** and *distributional form of the transition kernel*.

The corresponding **conditional density** defines the *density form of the transition kernel* (useful in cases of multiple dimensions, e.g. $S \subseteq \mathbb{R}^n$):

$$p(x, y) = \frac{\partial}{\partial y} P(x, y) = p(y|x), \quad x, y \in S.$$

Markov Chains and MC

Example

$$\text{AR}(1): \theta^{(t+1)} = \alpha\theta^{(t)} + \epsilon_t, \quad \epsilon_t \sim N(0, \sigma^2)$$

$$\Rightarrow (\theta^{(t+1)} | \theta^{(t)} = x) \sim N(\alpha x, \sigma^2)$$

Transition kernel

$$\Rightarrow p(x, y) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left\{-\frac{1}{2}\left(\frac{y - \alpha x}{\sigma}\right)^2\right\}$$

(density for y for a given fixed value of x)

Markov Chains and MC

Stationarity and reversibility

The state $\theta^{(t)}$ is a realisation from the probability density function $\pi^{(t)}(x)$, $x \in S$. Total probability:

$$\pi^{(t+1)}(y) = \int_S p(x, y) \pi^{(t)}(x) dx$$

For the stationary distribution,

$$\pi(y) = \int_S p(x, y) \pi(x) dx \quad \leftarrow$$

referring back to Lecture 2, here x is the given variable over which we integrate to get marginal $\pi(y)$

The transition density for the reversed chain (from Bayes)

$$p_t^*(x, y) = \frac{p(y, x) \pi^{(t)}(y)}{\pi^{(t+1)}(x)}$$

Markov Chains and MC

... homogenises in the stationary limit

$$p^*(x, y) = \frac{p(y, x)\pi(y)}{\pi(x)}.$$

→ Again, for a (time) reversible chain, (the continuous form) of **the detailed balance** equations hold:

$$\pi(x)p(x, y) = \pi(y)p(y, x), \quad \forall x, y \in S.$$

Integrating both sides of the detailed balance condition gives $\pi(y) = \int_S p(x, y)\pi(x)dx$ → **Any homogeneous Markov chain satisfying detailed balance is reversible with stationary distribution $\pi(\cdot)$.**

Markov Chains and MC

Stochastic simulation and analysis

1. Sample $\theta^{(0)}$ from $\pi^{(0)}(\cdot)$ using an appropriate method.
2. Simulate (the realisation) $\theta^{(t+1)}$ from $p(\theta^{(t)}, \cdot)$ for $t = 1, 2, \dots$ using one of the standard techniques. (So, accept $\theta^{(t+1)}$ according to $P(x, y) = P(\theta^{(t+1)} \leq y | \theta^{(t)} = x)$.)

Computing $P(x, y)$ is not trivial. For the normal distribution, Python has in the library `scipy.stats.norm.cdf(x)` that computes the value of CDF for the value x for the *standard normal distribution* $N(0, 1)$. CDF for $N(\mu, \sigma^2)$ can then be computed using $F(x) = \Phi\left(\frac{x-\mu}{\sigma}\right)$, where $F(\cdot)$ is CDF of $N(\mu, \sigma^2)$ and $\Phi(\cdot)$ is CDF of $N(0, 1)$.

Markov Chains in Continuous Time and Discrete State-space

(the last variation of Markov chains)

A Markov Chain
Gang clearly with
no memory of how
they got to the
present state:



Markov Chains and MC

Markov Chains in Continuous Time

[Online book, Ch 11.3](#)

Let's state how the Markov chains are constructed in continuous time.

Markov process in continuous time:

$$\begin{aligned} P(X(t + dt) = x | \{X(t) = x(t) | t \in [0, t]\}) = \\ = P(X(t + dt) = x | X(t) = x(t)), \quad \forall t \in [0, \infty), x \in S. \end{aligned}$$

Finite state-space

The process can take on one of r states, $S = \{1, 2, \dots, r\}$.

Markov Chains and MC

If at time t the process is in state $x \in S$, its future behaviour can be characterised by *the transition kernel*

$$p(x, t, x', t') \equiv P(X(t + t') = x' | X(t) = x).$$

If this function does not depend explicitly on t , the process is said to be *homogeneous* and the kernel can be written $p(x, x', t')$. For each value of t' this kernel can be expressed as an $r \times r$ transition matrix, $P(t')$. ($P(0) = I$.)

Just as in the discrete time case

$$P(t + t') = P(t)P(t') = P(t')P(t)$$

Chapman-Kolmogorov equations:

$$p(i, j, t + t') = \sum_{k=1}^r p(i, k, t)p(k, j, t'), \quad i, j = 1, \dots, r.$$

Markov Chains and MC

Define the *transition rate matrix* Q as a derivative of $P(t')$ at t'

$$Q = \frac{d}{dt'} P(t')|_{t'=0} = \lim_{\delta t \rightarrow 0} \frac{P(\delta t) - P(0)}{\delta t} = \lim_{\delta t \rightarrow 0} \frac{P(\delta t) - I}{\delta t}.$$

The elements of Q give the “**hazards**” of moving to different states.

→ The infinitesimal transition matrix

$$P(dt) = I + Qdt. \Rightarrow P(t + dt) = P(dt)P(t) = [I + Qdt]P(t).$$

→ The stationary distribution π :

$$\pi P(dt) = \pi$$

$$\pi(I + Qdt) = \pi$$

$$\pi Q = 0.$$

Markov Chains and MC

$$\begin{aligned}\frac{d}{dt}P(t) &= \frac{P(t+dt) - P(t)}{dt} = \frac{P(dt)P(t) - P(t)}{dt} \\ &= \frac{P(dt) - I}{dt}P(t) = QP(t).\end{aligned}$$

$\Rightarrow \frac{d}{dt}P(t) = QP(t)$; the initial condition $P(0) = I$.

$\Rightarrow P(t) = \exp\{Qt\} = \sum_{k=0}^{\infty} \frac{1}{k!} (Qt)^k$, where $(Qt)^0 \equiv I$. ($\exp\{\cdot\}$ is the matrix exponential function, see e.g. Golub & van Loan: Matrix Computations. But forget about it, because...)

We can avoid working with differential equations in matrix form by writing in the component form *Kolmogorov's backward equations*:

$$\frac{d}{dt}p(i, j, t) = \sum_{k=1}^r q_{ik}p(k, j, t), \quad i, j = 1, 2, \dots, r.$$

Markov Chains and MC

Had we expanded $P(t + dt)$ as $P(t)P(dt)$ rather than $P(dt)P(t)$, we would have ended up with

$$\frac{d}{dt}P(t) = P(t)Q$$

\Rightarrow Kolmogorov's forward equations:

$$\frac{d}{dt}p(i, j, t) = \sum_{k=1}^r q_{kj}p(i, k, t), \quad i, j = 1, 2, \dots, r.$$

Markov Chains and MC

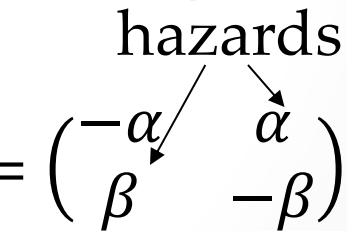
Example

A simple model for the activation of a single prokaryotic gene. The gene will be activated unless a repressor protein is bound to its regulatory region. Two states: state 0 (inactive) and state 1 (active). (\rightarrow The state vector is $(0 \ 1)$.)

In the inactive state there's a constant hazard of $\alpha > 0$ for activation.

In the active state there's a constant hazard of $\beta > 0$ for inactivation.

The rows of Q must sum to zero \rightarrow $Q = \begin{pmatrix} -\alpha & \alpha \\ \beta & -\beta \end{pmatrix}$



Solving $\pi Q = 0$ gives the stationary distribution

$$\pi = \left(\frac{\beta}{\alpha + \beta}, \frac{\alpha}{\alpha + \beta} \right). \quad (\pi \text{ is a PDF, so obviously } \sum_i \pi_i = 1.)$$

Markov Chains and MC

The infinitesimal transition matrix

$$P(dt) = I + Qdt = \begin{pmatrix} 1 - \alpha dt & \alpha dt \\ \beta dt & 1 - \beta dt \end{pmatrix}.$$

This process is also known as the *telegraph process* that can form a basis for a simple model of ion channels.

In summary: When solving for stationary distribution, use $\pi Q = 0$, and $\sum_i \pi_i = 1$, and that elements on each row of Q must sum to zero.

The last bit comes from the fact that the elements on row i of Q give all transition probabilities **per time dt** from state to i to all other states. In other words, the elements in Q are transition rates (velocities).

Markov Chains and MC

The following drill is just to show you how tedious even simple exact calculations for MC models are.

If you don't like it, put it in your pipe and smoke it (meaning, you can skip it).

To get the stationary distribution you need to also use Kolmogorov's (backward) eqs.

$\frac{d}{dt} P(t) = QP(t)$. Solving from this, e.g. , notation $p(i, j, t) = P_{ij}(t)$.

$$\frac{d}{dt} P_{00}(t) = -\alpha P_{00}(t) + \beta P_{01}(t) \text{ and } P_{01}(t) = 1 - P_{00}(t)$$

$$\Rightarrow \frac{d}{dt} P_{00}(t) = \beta - (\alpha + \beta)P_{00}(t).$$

Then let $Q_{00}(t) = e^{(\alpha+\beta)t} P_{00}(t) \dots$

Markov Chains and MC

Finally, $Q_{00}(t) = \left(\frac{\beta}{\alpha+\beta}\right) e^{(\alpha+\beta)t} + \left(\frac{\alpha}{\alpha+\beta}\right)$, and

$$P_{00}(t) = \frac{\beta}{\alpha+\beta} + \frac{\alpha}{\alpha+\beta} e^{-(\alpha+\beta)t}.$$

$$P_{01}(t) = 1 - P_{00}(t) = \frac{\alpha}{\alpha+\beta} - \frac{\alpha}{\alpha+\beta} e^{-(\alpha+\beta)t}.$$

By symmetry,

$$P_{11}(t) = \frac{\alpha}{\alpha+\beta} + \frac{\beta}{\alpha+\beta} e^{-(\alpha+\beta)t},$$

$$P_{10}(t) = \frac{\beta}{\alpha+\beta} - \frac{\beta}{\alpha+\beta} e^{-(\alpha+\beta)t}.$$

For details, see ISM,
Chapter 6.3.3. – if
interested.

From the previous eqs, the stationary state:

$$\begin{aligned} \lim_{t \rightarrow \infty} P_{11}(t) = \pi_1 &= \frac{\alpha}{\alpha + \beta} \\ \lim_{t \rightarrow \infty} P_{00}(t) = \pi_0 &= \frac{\beta}{\alpha + \beta} \end{aligned} \Rightarrow \pi = \left(\frac{\beta}{\alpha + \beta}, \frac{\alpha}{\alpha + \beta} \right)$$

Conclusion: it makes sense to simulate & compute the expectation values numerically.

Markov Chains and MC

Stochastic Simulation

Given the infinitesimal transition matrix $P(dt) = I + Qdt$, we have for small time steps Δt

$$P(\Delta t) \approx I + Q\Delta t.$$

$P(\Delta t)$ is regarded as **the transition matrix of a discrete time Markov chain** \rightarrow simulated sequence at times $0, \Delta t, 2\Delta t, 3\Delta t, \dots$

Precision can be improved by using the exact value

$$P(\Delta t) = \exp\{Q\Delta t\}. \quad (\text{We started from } \frac{d}{dt}P(t) = QP(t); P(0) = I.)$$

Δt must be chosen sufficiently small for not to miss interesting transitions.

Markov Chains and MC

The examples to follow show the principles on constructing efficient algorithms for stochastic simulation (of systems in continuous time).

Discrete event simulation

The process is currently in state $x \rightarrow$ the x th row of Q gives the hazards for the transitions to other states. The row sums to zero $\rightarrow -q_{xx}$ gives the combined hazard for moving away from the current state (q_{xx} is non-positive) \rightarrow the time to a transition event is exponential with rate $-q_{xx}$. ($-q_{xx}$ = the sum of the off-diagonal elements of the x th row.)

When this transition event occurs, the new state will be random with probabilities proportional to the x th row of Q .

Markov Chains and MC

Consider the probability that the next event will be in the time interval $(t + t', t + t' + dt]$ and will consist of a move to state j . Let this probability divided by dt be $f(t', j|t, i)$. The probability $f(t', j|t, i)dt$:

$$f(t', j|t, i)dt = P(\text{Next event in } (t + t', t + t' + dt]|t, i) \\ \times P(j|\text{Next event in } (t + t', t + t' + dt], t, i)$$

The first term: the hazards for the individual transitions are given by the off-diagonal elements of the i th row of Q ; the combined hazard is $-q_{ii}$. (The probability of two events occurring in the interval $(t, t + dt]$ is $O(dt^2)$ and can be neglected, $dt \ll 1$.)

Markov Chains and MC

The time to the next event is $Exp(-q_{ii}) \rightarrow$ the first term is $-q_{ii}e^{q_{ii}t'} dt$.

The second term is

This is saying that $j \neq i$, so it has no redundancy with $X(t + t' + dt) = j$

↓

$$\begin{aligned} P(X(t + t' + dt) = j | [X(t + t') = i] \cap [X(t + t' + dt) \neq i]) \\ = \frac{P(X(t + t' + dt) = j | X(t + t') = i)}{P(X(t + t' + dt) \neq i | X(t + t') = i)} = \frac{q_{ij} dt}{\sum_{k \neq i} q_{ik} dt} = \frac{q_{ij}}{-q_{ii}}. \end{aligned}$$

The first term times the second term:

$$f(t', j | t, i) = -q_{ii}e^{q_{ii}t'} \times \frac{q_{ij}}{-q_{ii}}$$

Markov Chains and MC

$f(t', j|t, i)$ factorises into the form of a probability density for the time to the next event and a probability mass function for the type of that event. → The next event can be simulated with the generation of two random variables. In addition the two random variates can be simulated independently. →

The standard discrete event simulation:

1. Initialise the process at $t = 0$ with initial state i .
2. Call the current state i . Simulate the time to the next event, t' , as an $Exp(-q_{ii})$ random quantity.
3. Put $t := t + t'$.
4. Simulate new state j as a discrete random quantity with PMF $-q_{ij}/q_{ii}$, $j \neq i$.
5. Output the time t and state j .
6. If $t < T_{max}$, return to step 2.

Markov Chains and MC

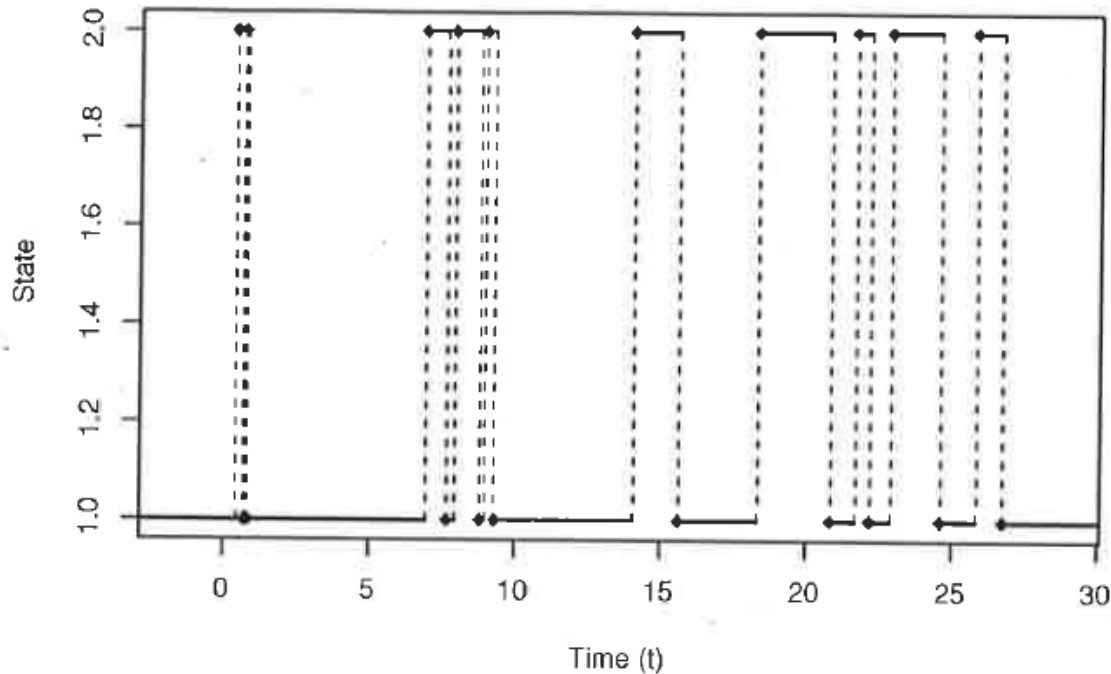


Figure 5.4 A simulated realisation of the simple gene activation process with $\alpha = 0.5$ and $\beta = 1$.

(In the figure, 1 corresponds to state 0 and 2 corresponds to state 1.)

Markov Chains and MC

Countable state-space; immigration-death process

The model: Individuals arrive into the population with constant hazard λ . Each individual dies independently with constant hazard μ . There is no reproduction.

The key transition equations:

(write down all possible transitions)

(To make sense, it must be that $\underline{\mu = c/x}$, where c is a constant. Only then is this a homogeneous Markov process.)

$$P(X(t + dt) = x + 1 | X(t) = x) = \lambda dt$$

$$P(X(t + dt) = x - 1 | X(t) = x) = x\mu dt$$

$$P(X(t + dt) = x | X(t) = x) = 1 - (\lambda + x\mu)dt$$

$$P(X(t + dt) = y | X(t) = x) = 0, \quad \forall y \notin \{x - 1, x, x + 1\}.$$

These equations define a homogeneous Markov process with infinite state-space $S = 0, 1, 2, \dots$

Markov Chains and MC

It is straightforward to write down the infinite-dimensional Q matrix (tridiagonal form):

$$Q = \begin{bmatrix} -\lambda & \lambda & 0 & 0 & 0 & \dots \\ \mu & -\lambda - \mu & \lambda & 0 & 0 & \dots \\ 0 & 2\mu & -\lambda - 2\mu & \lambda & 0 & \dots \\ 0 & 0 & 3\mu & -\lambda - 3\mu & \lambda & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

The (infinite dimensional) **stationary distribution** $\pi = (\pi_0, \pi_1, \pi_2, \dots)$:

$$\pi Q = 0 \Rightarrow \pi_k = \frac{\lambda^k}{k! \mu^k} \pi_0, \quad k = 1, 2, \dots$$

Markov Chains and MC

The terms are of the expansion $\pi_0 e^{\lambda/\mu}$. Imposing the unit-sum constraint, $\pi_0 = e^{-\lambda/\mu}$, so

$$\pi_k = \frac{(\lambda/\mu)^k}{k!} e^{-\lambda/\mu}, \quad k = 1, 2, \dots$$

So, the stationary distribution of this process is a Poisson random quantity with mean λ/μ .

The process can be simulated by discrete event simulation.

Markov Chains and MC

Inhomogeneous Poisson process

In the previous **homogeneous** (stationary) Poisson process the event **hazard λ is constant**.

Previously, N_t was defined as the number of events of the Poisson process in the interval $(0, t]$; $N_t \sim Po(\lambda t)$.

→ The process N_t is a *counting process*.

This counting process is a Markov process governed by the *homogeneous Markovian transition equations*

$$P(N_{t+dt} = x + 1 | N_t = x) = \lambda dt$$

$$P(N_{t+dt} = x | N_t = x) = 1 - \lambda dt$$

$$P(N_{t+dt} = y | N_t = x) = 0, \forall y \notin \{x, x + 1\}.$$

Markov Chains and MC

Poisson process can be generalised such that the event **hazard** is not a constant but a **function** $\lambda(t)$. \rightarrow **The probability of an event in the interval** $(t, t + dt]$ is $\lambda(t)dt$.

The *inhomogeneous Markovian transition equations* for the associated counting process N_t are

$$P(N_{t+dt} = x + 1 | N_t = x) = \lambda(t)dt$$

$$P(N_{t+dt} = x | N_t = x) = 1 - \lambda(t)dt$$

$$P(N_{t+dt} = y | N_t = x) = 0, \forall y \notin \{x, x + 1\}.$$

Markov Chains and MC

Since for a sufficiently small interval $(t, t + dt]$ $\lambda(t)$ is constant, the number of events in that interval is approximately Poisson $\sim Po(\lambda(t)dt)$, independent of all other intervals.

→ The number of events in the interval $(0, t]$, N_t is the sum (integral) over all intervals.

The sum of independent Poisson processes is a Poisson process:

$$N_t \sim Po\left(\int_0^t \lambda(s)ds\right) = Po(\Lambda(t)),$$

where $\Lambda(t) = \int_0^t \lambda(s)ds$ is the *cumulative hazard*.

The number of events in the interval $(s, t]$, $0 < s < t$ is $Po(\Lambda(t) - \Lambda(s))$.

Markov Chains and MC

Time-change of a Poisson process

To illustrate, take a unit Poisson process (whose rate is 1) having a counting process $M(t)$. Then define a new process that speeds up time by a factor of λ : $N(t) = M(\lambda t)$.

Obviously, $N(t)$ is - and so the time-change $t' = \lambda t$ gives - the counting process of a homogeneous Poisson process of rate λ . This generalises to non-linear time-changes.

Time-changes must be monotonic; to ensure this we write

$$t' = \int_0^t \lambda(s) ds, \quad \text{where } \lambda(t) > 0, \forall t > 0.$$

Markov Chains and MC

Again starting with a unit process $M(t)$, define

$$N(t) = M(t') = M\left(\int_0^t \lambda(s) ds\right) \sim Po\left(\int_0^t \lambda(s) ds\right).$$

$N(t)$ is an *inhomogeneous Poisson process* with rate $\lambda(t)$.

Slogan: A unit Poisson process can be transformed to an arbitrary Poisson process by distorting the time appropriately.

Markov Chains and MC

Proposition. For the inhomogeneous Poisson process with rate function $\lambda(t)$ **the time T to the *first* event** has distribution function

$F(t) = 1 - \exp\{-\Lambda(t)\}$, and hence the density function

$$f(t) = \lambda(t) \exp\{-\Lambda(t)\}, \quad t > 0.$$

Proof.

No events before
time $t < T$.

$$\begin{aligned} F(t) &= P(T \leq t) = 1 - P(T > t) = 1 - P(N_t = 0) \\ &= 1 - \frac{\Lambda(t)^0 \exp\{-\Lambda(t)\}}{0!} = 1 - \exp\{-\Lambda(t)\} \end{aligned}$$

Markov Chains and MC

In *stochastic simulation* the motivation often is to simulate the first/next event of an inhomogeneous process. One way to do this is to use the **transformation method**. We can simulate $u \sim U(0,1)$ and solve $F(t) = u$. Rearranging gives $\Lambda(t) = -\log(1 - u)$. Since u and $1 - u$ have the same distribution, we can solve $\Lambda(t) = -\log u$ for t .

If $\Lambda(t)$ is analytically invertible, then $t = \Lambda^{-1}(-\log u)$.

Since $\Lambda(t)$ is by construction monotonically increasing, $\Lambda(t) = -\log u$ will have at most one solution for $u \in (0, 1)$. If $\Lambda(t)$ does not tend to infinity as $t \rightarrow \infty$, there may not be a solution; in other words, the first event may not happen at all. To avoid this, in practise the event hazard $\lambda(t)$ is not allowed to decay faster than $1/t$.

Markov Chains and MC

Example

Consider the inhomogeneous Poisson process with rate function $\lambda(t) = \lambda t$ for some constant $\lambda > 0$.

The cumulative hazard $\Lambda(t) = \lambda t^2 / 2$.

$\Rightarrow N_t \sim Po(\lambda t^2 / 2) \Rightarrow$ the number of events in the interval $(s, t]$ is $Po(\lambda(t^2 - s^2) / 2)$.

The PDF for the time to the first event

$$f(t) = \lambda t \exp\left\{-\frac{\lambda t^2}{2}\right\}, \quad t > 0.$$

To simulate:

1. Sample $u \sim U(0,1)$.

2. Obtain the time to the first event as $t = \sqrt{-\frac{2 \log u}{\lambda}}$. Return to 1.

Solve

$$\lambda t^2 / 2 = -\log u \text{ for } t.$$

↓

Markov Chains and MC

Simulation of inhomogeneous Poisson processes

In the previous example, where transformation could be made analytically, simulation is effective: no wasted random numbers.

When $\lambda(t)$ is not this nice we need to do rejection-type simulation. The derivation of this method is typical in stochastics/statistics: Theoretically the justification of the method is not that straightforward, but intuitively – as an algorithm – it seems obvious.

An inhomogeneous Poisson process can be simulated by first simulating from a unit homogeneous Poisson process defined on particular region \mathbb{R}^2 .

Markov Chains and MC

The **2d Poisson process** is defined on a set $\Omega \subseteq \mathbb{R}^2$ and has the property that for any measurable subset $A \subseteq \Omega$ with area $\mu(A)$ the number of events in A will have a $Po(\mu(A))$ distribution.

→ The number of events in an infinitesimal square $(x, x + dx] \times (y, y + dy]$ will have a $Po(dx dy)$ distribution; the probability of an event in this region will be $dx dy$, independently of all other regions.

The *uniformity property* of the (2d) homogeneous Poisson process: Conditional on the number of points in a region Ω , those points are uniformly distributed over that region.

To simulate a 2d Poisson process on a finite region Ω :

1. Sample the *number of points* from $Po(\mu(\Omega))$.
2. Scatter these points uniformly over Ω .

Markov Chains and MC

If it is possible to simulate a unit Poisson process on a region $\bar{\Omega} \supseteq \Omega$, then the points of the originally sampled Poisson process that lie within Ω form a Poisson process on Ω . \rightarrow Rejection sampling can be applied on the bounding region $\bar{\Omega}$.

1d inhomogeneous Poisson process out of 2d (unit) Poisson process:

Consider an inhomogeneous Poisson process with bounded, non-negative hazard function $\lambda(t)$ satisfying $0 \leq \lambda(t) \leq U_\lambda$ for some upper-bound U_λ . We wish to simulate from this process on $(0, T]$ for some $T > 0$.

Let Ω be the region of \mathbb{R}^2 bounded by $x = 0, x = T, y = 0$, and $y = \lambda(x)$ and let $\bar{\Omega}$ be the rectangle $[0, T] \times [0, U_\lambda]$. \rightarrow By construction, $\Omega \subseteq \bar{\Omega} \subseteq \mathbb{R}^2$.

Markov Chains and MC

- To sample an inhomogeneous 1d Poisson process with $\lambda(t)$:
- sample a Poisson process on $\bar{\Omega}$ by first simulating the number of points $m \sim Po(U_\lambda T)$
 - for $i = 1, \dots, m$ independently sample an x -coordinate from $U(0, T)$ and y -coordinate from $U(0, U_\lambda)$
 - keep only the points in Ω , i.e. the points (x, y) satisfying $y \leq \lambda(x)$
 - from the above-obtained Poisson process on Ω marginally the x -coordinates give the event times of the 1d inhomogeneous Poisson process

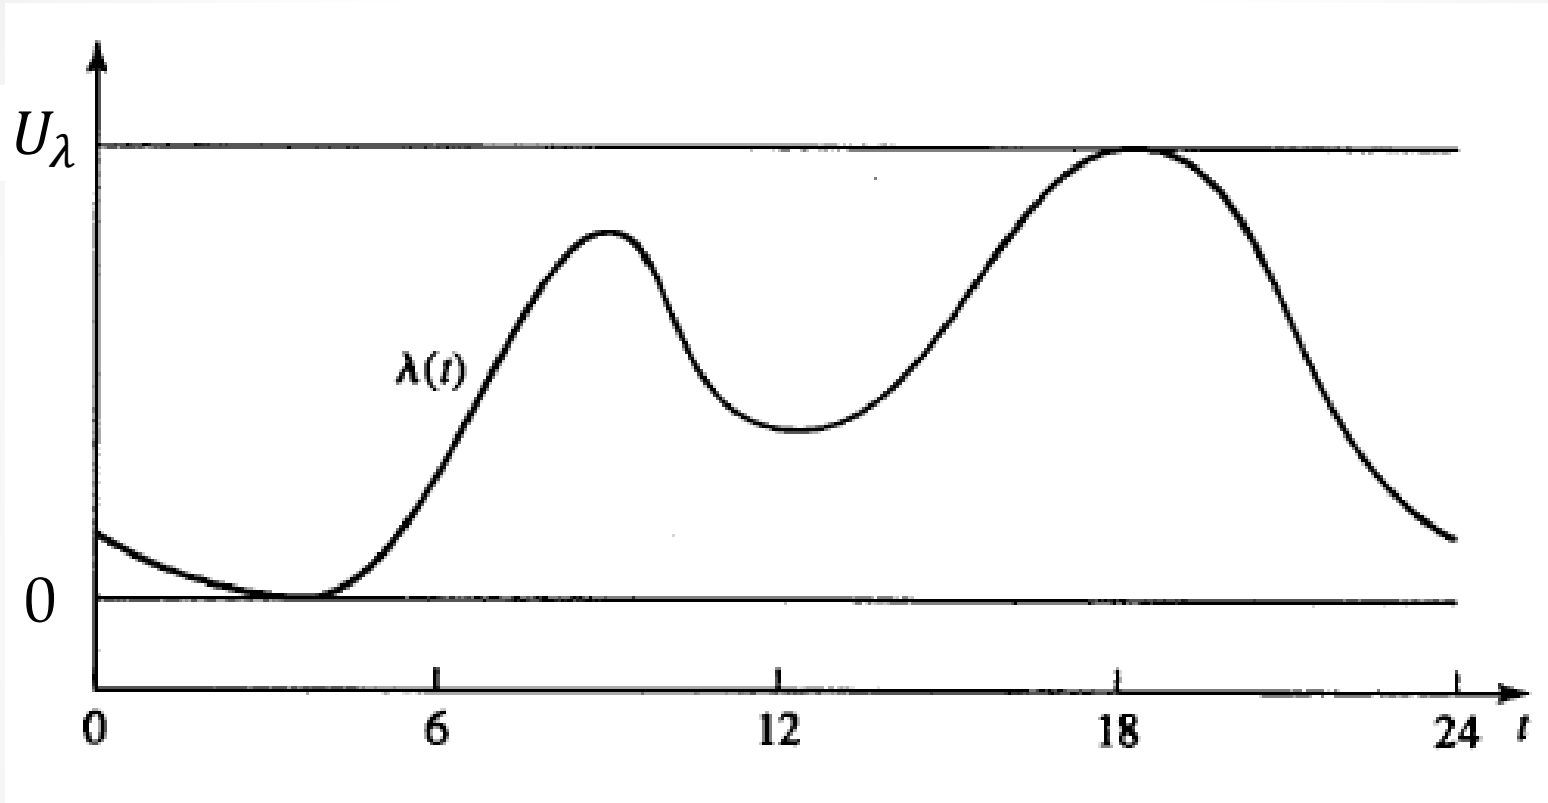
Markov Chains and MC

The algorithm for sampling event times from an inhomogeneous Poisson process:

1. Sample $m \sim Po(U_\lambda T)$
2. For $i := 1, \dots, m$
 - (a) Sample $x \sim U(0, T)$
 - (b) Sample $y \sim U(0, U_\lambda)$
 - (c) If $y \leq \lambda(x)$, print x

The printed x (m of them) represent a **single realisation** of the inhomogeneous Poisson process on the time interval $(0, T]$ **in any order**, i.e. not chronologically.

Markov Chains and MC



Markov Chains and MC

So, the point of drawing m from the Poisson distribution is this: If one does not do it in the beginning, the number of simulated events in the interval T is not Poissonian. So, one needs to do it. After the simulation we will have a number of events at different x proportional to $\left(\frac{\lambda(x)}{U_\lambda}\right) * Po(U_\lambda T)$ as we should. In the basic rejection without this part the number of events would not have any relation to Poissonian.

So, m drawn from Poissonian for the time interval makes sure this is a realisation of a Poisson process. In order to get complete statistics, in fact, several batches over different numbers m of runs drawn from $m \sim Po(U_\lambda T)$ should be done and results averaged over them. Only this way we can obtain **statistics** for the process.

Markov Chains and MC

If we are looking for **the time of the first event**, the previous algorithm is inefficient: multiple realisations have to be run until the end, find the minimum times and average over them.

One approach to **sample the event times in order** is to use properties of uniform order statistics: Denoting the sorted values $X_{(i)}$ for $X_i \sim U(0, T), i = 1, \dots, m$, so that $X_{(1)} \leq X_{(2)} \leq \dots \leq X_{(m)}$,

$$X_{(1)} = \min\{X_i\},$$

THIS IS AN ERROR MADE IN THE BOOK:

which has a CDF $F_{(1)}(x) = (x/T)^m$; simulate it by sampling $u \sim U(0, 1)$ and setting $x_{(1)} = Tu^{1/m}$. **This part is correct:**

→ Conditional on $x_{(i-1)}$, $X_{(i)}$ is the minimum of $m - i + 1$ uniforms on $[x_{(i-1)}, T]$. This can be simulated by sampling $u \sim U(0, 1)$ and setting $x_{(i)} = (T - x_{(i-1)})u^{1/(m-i+1)}$.

Markov Chains and MC

DO NOT USE!

This revised algorithm for sampling events in order as in the book by Wilkinson. **THIS IS WRONG:**

1. Set $x_{(0)} := 0$
2. Sample $m \sim Po(U_\lambda T)$
3. For $i := 1, \dots, m$
 - (a) Sample $u \sim U(0, 1)$
 - (b) Set $x_{(i)} := (T - x_{(i-1)})u^{1/(m-i+1)}$
 - (c) Sample $y \sim U(0, U_\lambda)$
 - (d) If $y \leq \lambda(x_{(i)})$, print $x_{(i)}$

Markov Chains and MC

THE PREVIOUS DERIVATION of the revised algorithm from the Wilkinson book IS COMPLETELY WRONG!

In Wilkinson's derivation there are two errors, which happily rectify each other so that every other sample appears in order and the whole range $(0, T]$ is sampled.

The first mistake: Wilkinson used CDF for maximum values instead of minimum values.

The second mistake: Wilkinson did not take into account the fact that for each recursive interval $(x_{(i-1)}, x_{(i)}]$, the starting point $x_{(i-1)}$ needs to be taken into account.

As a result, some ordering was obtained due to the correct condition $X_{(1)} \leq X_{(2)} \leq \dots \leq X_{(m)}$. However, this condition is not completely fulfilled due to the first mistake. Due to the second mistake the whole range $(0, T]$ is sampled despite the first mistake.

Markov Chains and MC

THE FOLLOWING IS CORRECT.

The CDF's for the first (and for reference, the last) event (see [Online Statistics Course](#)); here, $n = m$:

For X_1, X_2, \dots, X_n iid continuous random variables with pdf f and cdf F the density of the k th order statistic is

$$\begin{aligned} F_{(1)}(x) &= P(X_{(1)} < x) = 1 - P(X_{(1)} > x) \\ &= 1 - P(X_1 > x, \dots, X_n > x) = 1 - P(X_1 > x) \cdots P(X_n > x) \\ &= 1 - (1 - F(x))^n \end{aligned}$$

$$\begin{aligned} F_{(n)}(x) &= P(X_{(n)} < x) = 1 - P(X_{(n)} > x) \\ &= P(X_1 < x, \dots, X_n < x) = P(X_1 < x) \cdots P(X_n < x) \\ &= F(x)^n \end{aligned}$$

Markov Chains and MC

In the book, CDF for the last event was picked by mistake.
When we take the first-event CDF:

$$F_{(1)}(x) = 1 - [1 - F_{U(0,T]}(x)]^m = 1 - \left[1 - \frac{x}{T}\right]^m$$

Calculate **the inverse for the first event** in the interval $(0, T]$:
(the 1st event)

$$u = 1 - \left[1 - \frac{x}{T}\right]^m \Leftrightarrow x = [1 - (1 - u)^{1/m}]T \equiv x_{(1)} \swarrow$$

Then we calculate **the time for the i th event** $x_{(i)}$ conditional on $x_{(i-1)}$ - meaning that $i - 1$ events have occurred before.

Clearly, $x_i \sim U(x_{(i-1)}, T]$, and there will be $m - i + 1$ events in the interval $(x_{(i-1)}, T]$.

Markov Chains and MC

$$\text{CDF for } U(x_{(i-1)}, T]: F_{U(x_{(i-1)}, T]} = \frac{x_{(i)} - x_{(i-1)}}{T - x_{(i-1)}}$$

Accordingly, CDF for $X(i)$:

$$F_{X(i)} = 1 - \left[1 - F_{U(x_{(i-1)}, T)}\right]^{m-i+1} = 1 - \left[1 - \frac{x_{(i)} - x_{(i-1)}}{T - x_{(i-1)}}\right]^{m-i+1}$$

The inverse transformation:

$$\begin{aligned} u &= 1 - \left[1 - \frac{x_{(i)} - x_{(i-1)}}{T - x_{(i-1)}}\right]^{m-i+1} \Leftrightarrow \\ x_{(i)} &= x_{(i-1)} + \left[1 - (1 - u)^{\frac{1}{m-i+1}}\right] (T - x_{(i-1)}) = \\ &= x_{(i-1)} + \left[1 - u^{\frac{1}{m-i+1}}\right] (T - x_{(i-1)}) \end{aligned}$$

Markov Chains and MC

From the previous derivation, **the revised algorithm** for sampling events in order reads as:

1. Set $x_{(0)} := 0$
2. Sample $m \sim Po(U_\lambda T)$ ← This makes the thing Poissonian.
3. For $i := 1, \dots, m$
 - (a) Sample $u \sim U(0, 1)$
 - (b) Set $x_{(i)} := x_{(i-1)} + (T - x_{(i-1)})(1 - u^{1/(m-i+1)})$
 - (c) Sample $y \sim U(0, U_\lambda)$
 - (d) If $y \leq \lambda(x_{(i)})$, print $x_{(i)}$

This algorithm gives all events in order.

Markov Chains and MC

Thinning approach to jump time simulation

In simulating the process described above one can sample the *inter-event times* from an $Exp(U_\lambda)$ distribution to get the x -coordinates, then sample the corresponding y -coordinates from a $U(0, U_\lambda)$ distribution and reject the points outside Ω .

The thinning algorithm:

1. Set $t_0 := 0$
2. Set $i := 1$
3. Repeat
 - (a) Sample $t \sim Exp(U_\lambda)$ ← This makes the thing Poissonian.
 - (b) Set $t_i := t_{i-1} + t$
 - (c) If $t_i > T$, stop
 - (d) Sample $u \sim U(0, 1)$
 - (e) If $u \leq \lambda(t_i)/U_\lambda$, print t_i
 - (f) Set $i := i + 1$

Markov Chains and MC

The thinning algorithm only requires knowledge of the function $\lambda(t)$ – and only at a finite number of times (although time is continuous here) – and not on its integral or inverse.

Hence, the algorithm can be used when a hazard function is not analytically convenient or even deterministic; for example $\lambda(t)$ may be determined by a stochastic process. Only the upper bound has to be available.

The name comes from the fact that we are using rejection to *thin* a homogeneous Poisson process to give a target inhomogeneous process.

The algorithm can be modified to use any inhomogeneous Poisson process with a rate function $\nu(t)$ such that $\nu(t) \geq \lambda(t) \forall t \in [0, T]$ and which can be directly simulated. Often linear and quadratic forms for $\nu(t)$ are used. This is analogous to *envelope rejection*.

WHEW!

A lot to take in - learn the basic principles. For now, the methods (implementations) are important.