Simulation of Continuous-Time Markov Chains

Thursday, November 13, 2008
12:04 PM

Homework 3 due Thursday, November 20

Office Hours for Tuesday, November 18 are cancelled.

Markov time (stopping time) τ with respect to a filtration $\mathcal{A}_t$

is a random variable defined on the probability space underlying a stochastic process such that for every time $t \geq 0$ the event

$$\{\tau \leq t\} \in \mathcal{A}_t$$

Intuitively, this means that the random time τ has the property that you can tell whether or not it has already occurred just by looking at the information available up to the given moment.

Examples:

- the first time that a state of a Markov chain is visited
- the fifth time a certain state of a Markov chain is visited
- time to accumulate a certain reward
- time at which a Markov chain changes state
- time at which an optional decision is exercised

What is a random time which is not a Markov time?

- the last time to visit a state of a Markov chain
- the time at which a stochastic process or function thereof reaches its global maximum

Optimal stopping theory for choosing the best strategy based on a Markov time to maximize expected return

- Lawler Ch. 4

Strong Markov property:

For Markov times τ:

$$P(X(t+\tau)=j \mid \mathcal{A}_\tau) = P(X(t+\tau)=j \mid X(\tau)) \quad \text{for } t > 0$$

This generalizes the Markov property to allow the "present time" to be a random Markov time, so that the computation of statistics about the future of a Markov time depend only on the state of the system at the Markov time, and all other past information is redundant.

Discrete-time Markov chains have this strong Markov property automatically. (Resnick Sec. 1.8)

Continuous-time Markov chains also have the strong Markov property provided they are defined to be right-continuous, as we have been doing (Friedman, Stochastic Differential Equations, Sec. 2.2)

This concept now allows us to develop an efficient event-based simulation method for continuous-time Markov chains.

- Given the current state, determine how long the Markov chain stays in that state
- When the Markov chain changes state, determine the new state.

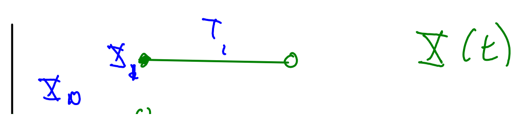

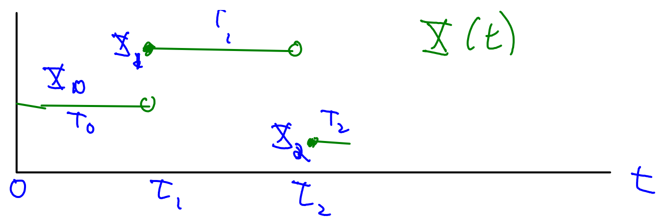

Suppose we start in a state i:

$X(s) = i$

First event simulation

1) Then we determine the time at which the Markov chain leaves this state by generating an exponentially distributed random time T with mean $1/\lambda_i$ where $A_{ii} = -\lambda_i$

This determines the time τ at which the Markov chain leaves state i through:

$$\tau = s + T = \inf_{t > s} \{ X(t) \neq i \}$$

The reason for this is the same as for the Poisson counting process.

2) Determine the next state of the continuous-time Markov chain by the following rule:

X_+ = state after the transition

D

X_- = state prior to jump

$$P(X_+ = j \mid X_- = i) = \frac{A_{ij}}{-A_{ii}} = \frac{A_{ij}}{\lambda_i} \quad \text{for } j \neq i$$

$$\left(\text{Recall } A_{ii} = -\sum_{j \neq i} A_{ij}\right)$$

One repeats this process for each new state visited.

Rigorous justification Karlin and Taylor Ch. 14

An alternative procedure to this is the following:

Next event simulation

Start in state i. Calculate random transition times corresponding to each possible transition out of state i.

Exponentially distributed random times $\{T_{ij}\}_{j \neq i}$ with $\langle T_{ij} \rangle = \frac{1}{A_{ij}}$

These correspond intuitively to the amount of time to wait until a transition from state i to state j would occur, were all other transitions out of state i to be ignored.

Take the shortest of these times, and this determines the time at which the transition happens, and the new state is exactly the one corresponding to the shortest transition time.

The equivalence of this procedure to the previous one discussed is shown in Lawler Ch. 3.

It seems like the first method is more efficient since fewer random variables have to be generated, but sometimes in high-dimensional, particularly network simulations, the second approach is easier to manage.

The next event simulation approach is much more amenable to parallelization than the first event approach.

G. Korniss (physics): using small world network for communication across computational models to improve the synchronization across nodes

Long-time properties of continuous-time Markov chains

Transience vs. recurrence?

- communication classes are determined in exactly the same way as for discrete-time Markov chains, based on the topology
- determining whether or not a communication class is transient or recurrent is done by simply looking at the associated discrete time Markov chain determined by the states visited by the continuous-time Markov chain (see the graph above)

$X_n = X(T_n)$ where T_n = nth transition time

The probability transition matrix associated to this discrete-time Markov chain is

$$\tilde{P}_{ij} = P(X_1 = j \mid X_0 = i) = \frac{A_{ij}}{-A_{ii}} = \frac{A_{ij}}{\lambda_i}$$

Also for an absorbing state: $\tilde{P}_{ij} = \delta_{ij}$

Positive recurrent vs. null recurrent?

- As for discrete-time Markov chains, this question is resolved precisely by determining whether or not the class treated as its own Markov chain has a stationary distribution. If it does, then it's positive recurrent, otherwise not.
- The stationary distribution for the continuous-time Markov chain is not the same as the stationary distribution for the associated discrete-time Markov chain.

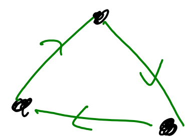

Computing the stationary distribution for a continuous-time Markov chain

$$\pi_j = P(X(t) = j)$$

should be time-independent solution of Kolmogorov forward equation for the probability distribution of states:

$$\frac{d\vec{\phi}}{dt} = \vec{\phi} A \qquad \phi_j(t) = P(X(t) = j)$$

$$\Rightarrow \boxed{\begin{array}{c} \vec{\pi} A = 0 \\ 0 \le \pi_j \le 1 \\ \sum_{j \in S} \pi_j = 1 \end{array}}$$

Note that periodicity is not an issue for continuous-time Markov chains because the random transition times randomize the state of the system even if the transitions have a periodic structure.

Detailed balance solutions still play an important role (when they exist):

$$\pi_i A_{ij} = \pi_j A_{ji} \quad \text{for all } i, j \in S$$

These are associated with time-reversible and some other systems, and when detailed balance solutions can be found, they are also guaranteed to be stationary distributions (sum over i).

Absorption probabilities can be computed through the associated discrete-time Markov chain.

But we have to revisit the formula for the accumulated reward when this occurs continuously in time because the accumulated reward will depend on how long one spends in a given state.

The formula for the continuous-time Markov chain is as follows:

$$w_i = \mathbb{E}\left[\int_0^{\tau} f(X(t))\, dt \,\middle|\, X(0) = i \right]$$

$$\tau = \inf\{t \geq 0 : t \notin T\}$$

transient states

$$-A\vec{w} = \vec{f} \qquad \text{with}$$

$$w_i = f_i = 0 \quad \text{when} \quad i \notin T$$

Proper derivation based on first-step analysis can be found in Lawler Sec. 3.3, Karlin and Taylor Ch. 4

Can also formally derive this by considering a discretization of time into time intervals of fixed width and then taking the limit of infinitely fine mesh for the discrete-time formula.