Lecture 21

David Aldous

16 October 2015

In continuous time $0 \le t < \infty$ we specify transition **rates**

$$q_{ij} = \lim_{\delta \downarrow 0} \frac{\mathbb{P}(X(t+\delta)=j | X(t)=i, \text{ past })}{\delta}$$

or informally

$$\mathbb{P}(X(t+dt) = j | X(t) = i) = q_{ij} dt$$

but note these are defined only for $j \neq i$. The time-t distribution $\pi(t)$ evolves as

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

where $\mathbf{Q}$ is the matrix with off-diagonal entries (q_{ij}) and with diagonal entries defined by

$$q_{ii} = -q_i = -\sum_{j \neq i} q_{ij}.$$

There is an alternative "jump and hold" description of a continuous-time Markov chain.

- After jumping into a state i, the process remains in state i for a random time with Exponential(q_i) distribution.
- Then it jumps to some other state, to state $j \neq i$ with probability $\widehat{p}_{ij} = q_{ij}/q_i$.

So the matrix

$$\widehat{\mathbf{P}} = (\widehat{p}_{ij}), \quad \text{where } \widehat{p}_{ii} = 0$$

is the transition matrix for the discrete-time **jump chain** $\hat{X}(0), \hat{X}(1), \ldots$ that shows the successive states visited.

Example: Yule process

- parameter $\beta > 0$
- states $1, 2, 3, \ldots$
- transition rates $q_{i,i+1} = \beta i$
- $X(0) = 1$.

The differential equations are

$$\frac{d}{dt}\pi_j(t) = \beta[(j-1)\pi_{j-1}(t) - j\pi_j(t)].$$

One can solve these equations – see [PK] section 6.1.3

$$\pi_j(t) = \mathbb{P}(X(t) = j) = e^{-\beta t}(1 - e^{-\beta t})^{j-1},\ j = 1, 2, \ldots$$

In other words $X(t)$ has Geometric $e^{-\beta t}$ distribution, so $\mathbb{E}X(t) = e^{\beta t}$.

The Yule process is a basic example of a continuous-time branching process [picture on board]

The Yule process is also an example of a "pure birth" process, meaning the only transitions are $i \to i+1$. For such processes the distribution of $X(t)$ can be related to the sum of independent Exponentials RVs – see [PK] section 6.1.2.

Example: Linear pure death process [PK] section 6.2.1.

- parameter $\mu > 0$
- states $0, 1, 2, 3, \ldots, N$
- transition rates $q_{i,i-1} = \mu i$
- $X(0) = N$.

The differential equations are

$$\frac{d}{dt}\pi_j(t) = \mu[(j+1)\pi_{j+1}(t) - j\pi_j(t)].$$

But one can find the time-t distribution easily via an alternative description of the process.

- N individuals; initially alive, each dies at rate μ.
- $X(t) =$ number alive at time t,

Clearly $X(t)$ has Binomial$(N, e^{-\mu t})$ distribution.

$$\pi_j(t) = \binom{N}{j} e^{-\mu t j}(1 - e^{-\mu t})^{N-j}.$$

Note the **general** pure death process (only transitions are $i \to i-1$) is mathematically the same as the general pure birth process.

Some theory – similar to discrete-time setting.

If the chain is irreducible, and either finite-state or infinite state and positive-recurrent, then a unique stationary distribution π exists, and is the solution of $\pi \mathbf{Q} = 0$, that is

$$\sum_{i \neq j} \pi_i q_{ij} = \pi_j q_j \quad \text{for each } j.$$

If you can find weights $w_i > 0$ such that

$$w_i q_{ij} = w_j q_{ji} \quad \text{for each } i, j \qquad \text{(detailed balance)}$$

then the stationary distribution is

$$\pi_i = w_i / w, \quad w = \sum_j w_j$$

provided (in the infinite-state case) $w < \infty$.

Birth-and-death chains.

These have states $\{0, 1, 2, \ldots, N\}$ or $\{0, 1, 2, \ldots \ldots\}$ and the only transitions are $i \to i \pm 1$. Write

$$\lambda_i = q_{i,i+1} \ \text{(birth rate)}; \qquad \mu_i = q_{i,i-1} \ \text{(death rate)}.$$

For these chains we can solve the detailed balance equations: [board]

$$w_i = \prod_{j=1}^{i} \frac{\lambda_{j-1}}{\mu_j}; \quad w = \sum_{i \geq 0} w_i.$$

So the stationary distribution is

$$\pi_i = w_i / w$$

provided (in the infinite-state case) $w < \infty$.

Example. Take $\lambda_i = \lambda$, $\mu_i = \mu i$. Then [board] π is the Poisson(λ/μ) distribution.

Note that if the stationary distribution π exists for an infinite-state birth-and-death process, then for the same process on states $\{0, 1, 2, \ldots, N\}$ the stationary distribution is

$$\pi_i^{[N]} = \pi_i / s. \quad s = \sum_{j=0}^{N} \pi_j.$$

In other words, taking π as the distribution of a RV Z, $\pi^{[N]}$ is the conditional distribution of Z given $\{Z \leq N\}$.

More theory – similar to discrete-time setting.

[Assume chain is irreducible, and either finite-state or infinite state and positive-recurrent, so a unique stationary distribution π exists.]

- For any initial distribution, $\mathbb{P}(X(t) = i) \to \pi_i$ as $t \to \infty$.
- Writing $N_i(t) =$ length of time chain spends in state i during $[0, t]$, we have $N_i(t)/t \to \pi_i$ as $t \to \infty$.
- $\mathbb{E}_i T_i^+ = 1/(\pi_i q_i)$, where T_i^+ is the first **return time** to i (after leaving i).

Note we don't need "aperiodic" in the first result. The third result can be seen by a general "cycle argument" [next slide and board].

(R,T) dependent $\quad 0 < T,\ ET < \infty,\ E|R| < \infty.$

$(R_i, T_i),\ i = 1, 2, 3, \quad$ IID copies of (R,T)

Get reward R_i after interval of length T_i,
that is at time $T_1 + T_2 + \ldots + T_i$.

$Y(t)$ = total reward up to time t

$$= \sum_{i \geq 1} R_i \mathbb{1}_{(T_1 + \ldots + T_i \leq t)}$$

Theorem "cycle trick" or "reward renewal theorem"

$$\frac{Y(t)}{t} \to \frac{ER}{ET} \quad \text{as } t \to \infty.$$