Poisson subordination (from last lecture)

- Given continuous-time Markov chain X_t with generator matrix Q. If $\lambda \geq \max(q_, \ldots, q_k)$ then
  - $R = \frac{1}{\lambda}Q + I$ is a stochastic matrix.
  - The transition matrix for X_t becomes

$$P(t) = e^{tQ} = e^{-t\lambda I}e^{t\lambda R} = e^{-t\lambda}\sum_{n=0}^{\infty}\frac{(t\lambda R)^n}{n!} = \sum_{n=0}^{\infty}R^n\frac{e^{-\lambda t}(\lambda t)^n}{n!}.$$

- Define $Z_t = Y_{N_t}$ where N_t is a Poisson process with parameter λ and Y_t is a discrete-time Markov chain with transition matrix R. The transition matrix for Z_t also becomes

$$P(t) = \sum_{n=0}^{\infty}R^n\frac{e^{-\lambda t}(\lambda t)^n}{n!}.$$

- The above can be used to simulate from a Markov chain, and also to compute $P(t) = e^{tQ}$ with better accuracy.
- NOTE: The Poisson subordinate discrete chain Y_t has the same stationary distribution as the continuous chain X_t.

MVE550 2022 Lecture 12
Dobrow Sections 8.1 - 8.4
Introduction to Brownian motion
Gaussian processes

Petter Mostad

Chalmers University

December 8, 2022

Continuous-time continuous state space processes

- Having looked at
  - Discrete-time discrete state space processes. (Discrete Markov chains and Branching processes).
  - Discrete-time continuous state space processes (not so much but we had some MCMC examples).
  - Continuous-time discrete state space processes (Poisson processes and more generally continuous-time Markov chains).
  - we now look at continuous-time continuous state space processes.
- We will look at two examples:
  - Brownian motion.
  - More generally, Gaussian processes.

Brownian motion

- In a gas, atoms bump into each other and change course. Over time, how does a single atom move, on average?
- If $f(x,t)$ represents the probability density for the position x of an atom at time t moving along a line, Albert Einstein showed that

$$\frac{\partial}{\partial t} f(x,t) = \frac{1}{2} \frac{\partial^2}{\partial x^2} f(x,t).$$

- The solution is

$$f(x,t) = \frac{1}{\sqrt{2\pi t}} e^{-x^2/2t}.$$

  So $x \sim \text{Normal}(0,t)$ at time t.
- It turns out a single atom will move as simulated below. These paths are sampled from a model called Brownian motion.

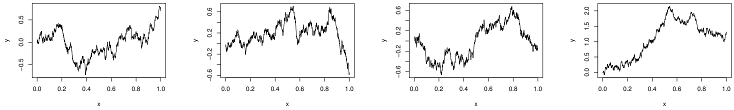

Definition of Brownian motion

Brownian motion is a continuous-time stochastic process $\{B_t\}_{t\geq 0}$ with the following properties:

1. $B_0 = 0$.
2. For $t > 0$, $B_t \sim \text{Normal}(0, t)$ (so the *variance* is t, not the standard deviation).
3. For $s, t > 0$, $B_{t+s} - B_s \sim \text{Normal}(0, t)$.
4. For $0 \leq q < r \leq s < t$, $B_t - B_s$ is independent from $B_r - B_q$.
5. The function $t \mapsto B_t$ is continuous with probability 1.

Simulation of Brownian motion

- Given time points $t_1 < t_2 < \cdots < t_n$, we can write

$$B_{t_i} = B_{t_{i-1}} + (B_{t_i} - B_{t_{i-1}}) = B_{t_{i-1}} + Z_i$$

  where $Z_i \sim \text{Normal}(0, t_i - t_{i-1})$.
- Writing $t_0 = 0$, we get for independent $Z_1, \ldots, Z_n$,

$$B_{t_n} = \sum_{i=1}^{n} Z_i.$$

- A good way to simulate the path $t \mapsto B_t$ on $t \in [0, a]$ is to set $t_i = ai/n$, simulate independently

$$Z_i \sim \text{Normal}(0, a/n)$$

  and compute

$$B_{t_i} = \sum_{j=1}^{i} Z_j.$$

- Note that we can also write $Z_i = \sqrt{a/n}\,Y_i$, where $Y_i \sim \text{Normal}(0, 1)$.

Zooming in on a Brownian motion realization

- What if we have a Brownian motion path simulated above, and want to plot it at twice the detail?
- The difference Z_i between the value at t_i and t_{i+1} can be written as a sum
$$Z_i = Z_{i0} + Z_{i1}$$
where $Z_{i0}, Z_{i1} \sim \text{Normal}(0, a/2n)$ independently.
- Reformulation: If we know Z_i, then $Z_{i0} \sim \text{Normal}(0, a/2n)$ as prior, with likelihood $Z_i \sim \text{Normal}(Z_{i0}, a/2n)$. Using conjugacy we get the posterior
$$Z_{i0} \mid Z_i \sim \text{Normal}\left(\frac{1}{2}Z_i, \frac{a}{4n}\right).$$
- We get the value at midpoint between t_i and t_{i+1} by simulating Z_{i0} and adding it to the value at t_i.
- NOTE: The resulting plot could be generated from scratch simply as Brownian motion using $a/2n$ instead of a/n: Scaling the x axis with factor $1/2$ scales the y axis with factor $\sqrt{1/2}$.

Computing with Brownian motion

- To compute probabilities for Brownian motion, we generally use the properties in the definition, e.g.,
  - $B_{t+s} - B_s \sim \text{Normal}(0, t)$
  - For $0 \leq q < r \leq s < t$, $B_t - B_s$ is independent from $B_r - B_q$.
- Example: Show that $B_1 + B_3 + 2B_7 \sim \text{Normal}(0, 50)$.
- Example: Show that $P(B_2 > 0 \mid B_1 = 1) = 0.8413$.
- Example: Show that $\text{Cov}(B_s, B_t) = \min\{s, t\}$.

Random walks: What happens when $n \to \infty$?

- Consider a symmetric random walk: A discrete time Markov chain $S_0, S_1, S_2, \ldots$ where

$$S_n = X_1 + X_2 + \cdots + X_n$$

  where $X_1, X_2, \ldots$ are independent random variables with expectation zero.
- If we assume $\mathrm{Var}(X_i) = 1$ we get $\mathrm{Var}(S_n) = n$.
- Interpolating between the values S_n we can make this into a continuous time process S_t (see Dobrow). $\mathrm{Var}(S_t) \approx t$.
- We may scale with an $s > 0$ to get processes $S_t^{(s)} = S_{st}/\sqrt{s}$ where we get $\lim_{s\to\infty} \mathrm{Var}(S_t^{(s)}) = t$.
- It turns out that the processes $S_t^{(s)}$ when $s \to \infty$ are *exactly* Brownian motion, no matter what type of X_i we start with!
- This is the Donsker invariance principle.
- We can see this effect in simulations.
- We can use this to find approximate properties of random walks.

Nowhere differentiable paths

- We have seen in our simulations that paths of Brownian motion are "jagged".
- We have also seen that this quality is unchanged if we change the scale, i.e., look at smaller intervals.
- Formally note that $B_{t+h} - B_t \sim \mathrm{Normal}(0, h)$ so that

$$\frac{B_{t+h} - B_t}{h} \sim \mathrm{Normal}(0, 1/h)$$

- Using these observations as starting points, one may show that the path (i.e., the function $t \mapsto B_t$) of a Brownian motion is nowehere differentiable, even though it is everywhere continuous.

The multivariate normal distribution (review)

- Definition (one of many): A set of random variables $X_1, \dots, X_k$ has a *multivariate normal distribution* if, for all real $a_1, \dots, a_k$, $a_1 X_1 + \cdots + a_k X_k$ is normally distributed.
- It is completely determined by the expectation vector $\mu = (\mathrm{E}(X_1), \dots, \mathrm{E}(X_k))$ and the $(k \times k)$ covariance matrix Σ, where $\Sigma_{ij} = \mathrm{Cov}(X_i, X_j)$.
- The joint density function on the vector $x = (x_1, \dots, x_k)$ is

$$\pi(x) = \frac{1}{|2\pi\Sigma|^{1/2}} \exp\left(-\frac{1}{2}(x-\mu)^T \Sigma^{-1} (x-\mu)\right).$$

  where $|2\pi\Sigma|$ is the determinant of the matrix $2\pi\Sigma$.
- All marginal distributions and all conditional distributions are also multivariate normal.

Gaussian processes

- A *Gaussian process* is a continuous-time stochastic process $\{X_t\}_{t\geq 0}$ with the property that for all $n \geq 1$ and $0 \leq t_1 < t_2 < \cdots < t_n$, $X_{t_1}, \ldots, X_{t_n}$ has a multivariate normal distribution.
- Thus, a Gaussian process is completely determined by its mean function $\mathrm{E}(X_t)$ and its covariance function $\mathrm{Cov}(X_s, X_t)$.
- Gaussian processes are extremely versatile as models. One may generalize for example so that the index set (the t's) is $\mathbb{R}^n$.

Brownian motion and Gaussian processes

- Brownian motion is a Gaussian process, as we can show that any $a_1 B_{t_1} + \cdots + a_k B_{t_k}$ is normally distributed.
- A Gaussian process $\{X_t\}_{t \geq 0}$ is Brownian motion if and only if
  1. $X_0 = 0$.
  2. $\mathrm{E}(X_t) = 0$ for all t.
  3. $\mathrm{Cov}(X_s, X_t) = \min\{s, t\}$ for all s, t.
  4. The function $t \mapsto X_t$ is a continuous with probability 1.
- The proof is fairly straightforward (see Dobrow).
- One may use the above for example when proving that something is Brownian motion, if it is easier than using the definition directly.

Transformations of Brownian motion

- The following transformations of Brownian motion are again Brownian motion:
  - $\{-B_t\}_{t\geq 0}$.
  - $(B_{t+s} - B_s)_{t\geq 0}$ for any $s \geq 0$.
  - $\left\{\frac{1}{\sqrt{a}} B_{at}\right\}_{t\geq 0}$ for any $a > 0$.
  - The process $\{X_t\}_{t\geq 0}$ where $X_0 = 0$ and $X_t = tB_{1/t}$ for $t > 0$.
- The proofs are fairly straightforward.
- The process $X_t = x + B_t$ where B_t is Brownian motion and x is some real number is called "Brownian motion started at x".

First hitting time and stopping times

- For any fixed t, $(B_{t+s} - B_t)_{s \geq 0}$ is Brownian motion.
- Does this also happen if we start the chain anew from T when T is random? It depends!
- If T is the largest value less than 1 where $B_T = 0$, is $B_{T+s} - B_T$ Brownian motion?
- No!
- If T is the smallest value where $B_T = a$ for some constant a, is $B_{T+s} - B_T$ Brownian motion?
- Yes! The reason is that the event $T = t$ can be determined based on B_r where $0 \leq r \leq t$.
- Random T's that have this property are called *stopping times*. For these $B_{T+s} - B_T$ is Brownian motion.

The distribution of the first hitting time

- Given $a \neq 0$ what is the distribution of the first hitting time $T_a = \min\{t : B_t = a\}$?
- We prove below that

$$\frac{1}{T_a} \sim \text{Gamma}\left(\frac{1}{2}, \frac{a^2}{2}\right)$$

- Assuming that $a > 0$ and using that T_a is a stopping time we get for any $t > 0$ that $\Pr\left(B_{1/t} > a \mid T_a < 1/t\right) = \Pr\left(B_{1/t - T_a} > 0\right) = \frac{1}{2}$.
- We also have

$$\Pr\left(B_{1/t} > a \mid T_a < 1/t\right) = \frac{\Pr\left(B_{1/t} > a, T_a < 1/t\right)}{\Pr\left(T_a < 1/t\right)} = \frac{\Pr\left(B_{1/t} > a\right)}{\Pr\left(T_a < 1/t\right)}.$$

- It follows that $\Pr\left(T_a < 1/t\right) = 2\Pr\left(B_{1/t} > a\right)$ and so

$$\Pr\left(\frac{1}{T_a} < t\right) = 2\Pr\left(B_{1/t} < a\right) - 1 = 2\Pr\left(B_1 < at^{1/2}\right) - 1.$$

- Taking the derivative w.r.t. t we get the Gamma density

$$\pi_{1/T_a}(t) = 2\frac{1}{\sqrt{2\pi}}\exp\left(-\frac{1}{2}(at^{1/2})^2\right)\frac{a}{2}t^{-1/2}.$$

Maximum of Brownian motion

- Define $M_t = \max_{0 \leq s \leq t} B_s$.
- We may compute for $a > 0$ (using result from previous page)

$$\Pr(M_t > a) = \Pr(T_a < t) = 2\Pr(B_t > a) = \Pr(|B_t| > a)$$

- Thus M_t has the same distribution as $|B_t|$, the absolute value of B_t.
- Example: What is the probability that $M_3 > 5$?
- Example: Find t such that $\Pr(M_t \leq 4) = 0.9$.