Math 526: Brownian Motion Notes

Mike Ludkovski, 2007, all rights reserved.

Definition.

A stochastic process (X_t) is called *Brownian motion* if:

1. The map $t \mapsto X_t(\omega)$ is continuous for every ω.
2. $(X_{t_1} - X_{t_0}), (X_{t_2} - X_{t_1}), \dots, (X_{t_n} - X_{t_{n-1}})$ are independent for any collection of times $t_0 \le t_1 \le \dots \le t_n$.
3. The distribution of $X_t - X_s$ depends only on $(t - s)$.

Property 1 is called continuity of sample paths. Property 2 is called independent increments. Property 3 is called stationarity of increments.

Here are some immediate conclusions and comments:

- Property 2 and 3 imply that (X_t) is time-homogeneous (one may shift the time-axis) and space-homogeneous (since distribution of increments does not depend on current X_s).
- Property 2 and 3 also imply that (X_t) is Markov, since the future increments are conditionally independent of the past given the present.
- Properties 2 and 3 are also satisfied by the Poisson process. However the Poisson process certainly does not have continuous sample paths.
- It is not at all clear that such a process (X_t) exists, since the requirement of continuous paths is very stringent.
- The underlying probability space here is $\Omega = \{\text{continuous functions on } [0, \infty)\}$.

Brownian Motion has Gaussian Marginals.

We have the following

Theorem 1. *If (X_t) is a Brownian motion, then there exist constants μ, σ^2 such that $X_t - X_s \sim \mathcal{N}((t-s)/\mu, (t-s)\sigma^2)$.*

Proof. Without loss of generality take $s = 0$ and pick some $n \in \mathbb{Z}$ and write

$$X_t - X_0 = (X_{t/n} - X_0) + (X_{2t/n} - X_{t/n}) + \dots + (X_t - X_{(n-1)t/n}).$$

Each term on the right hand side is independent and identically distributed. Letting $n \to \infty$, by the central limit theorem, the right hand side must converge to some normally distributed random variable (the continuity of (X_t) makes sure that one can apply the CLT, as the increments must become smaller and smaller as n increases).

We conclude that $X_t - X_0 \sim \mathcal{N}(a(t), b(t))$ for some functions a, b. However, since $X_{t+s} - X_0 = (X_t - X_0) + (X_{t+s} - X_t)$ and the two are independent, we have that

$$X_{t+s} - X_0 \sim \mathcal{N}(a(t), b(t)) + \mathcal{N}(a(s), b(s)) \stackrel{d}{=} \mathcal{N}(a(t) + a(s), b(t) + b(s)).$$

On the other hand, we directly have $X_{t+s} - X_0 \sim \mathcal{N}(a(t+s), b(t+s))$. It follows that $a(t+s) = a(t) + a(s)$ and similarly for b. But this means that a, b are *linear* functions, i.e. $a(t) = \mu t$, $b(t) = \sigma^2 t$ (clearly b must have positive slope), for some μ, σ^2. □

The parameter μ is called drift of the Brownian motion and σ is called the volatility. To summarize, if (X_t) is a BM then the marginal distribution is

$$X_t - X_0 \sim \mathcal{N}(\mu t, \sigma^2 t).$$

Wiener Process

The special case $\mu = 0, \sigma^2 = 1, X_0 = 0$ is called the Wiener process. We write (W_t) in that case. Here are some computations for the Wiener process:

$$\begin{aligned}
\mathbb{E}[W_t] &= 0. \\
Var(W_t) &= t. \\
Cov(W_s, W_t) &= \mathbb{E}[W_s W_t] - \mathbb{E}[W_s]\mathbb{E}[W_t] \qquad s < t \\
&= \mathbb{E}\left[W_s(W_s + W_t - W_s)\right] \\
&= \mathbb{E}[W_s^2] + \mathbb{E}[W_s]\mathbb{E}[W_t - W_s] \\
&= Var(W_s) + 0 = s.
\end{aligned}$$

The last statement is written as $Cov(W_s, W_t) = \min(s, t)$.

Here is another important computation: we find the conditional distribution of W_s given $W_t = z$, $s < t$:

$$\begin{aligned}
\mathbb{P}(W_s \in dx | W_t = z) &= \frac{\mathbb{P}(W_s \in dx, W_t \in dz)}{\mathbb{P}(W_t \in dz)} \\
&= \frac{\mathbb{P}(W_s \in dx, (W_t - W_s) \in (dz - x))}{\mathbb{P}(W_t \in dz)}
\end{aligned}$$

Using the Gaussian density and independent increments:

$$\begin{aligned}
&= \frac{\left(\frac{1}{\sqrt{2\pi s}} \mathrm{e}^{-x^2/(2s)} dx\right)\left(\frac{1}{\sqrt{2\pi(t-s)}} \mathrm{e}^{-(z-x)^2/(2(t-s))} dz\right)}{\frac{1}{\sqrt{2\pi t}} \mathrm{e}^{-z^2/(2t)} dz} \\
&= \exp\left(-\frac{1}{2}\left(x^2/s + (z-x)^2/(t-s) - z^2/t\right)\right) \cdot \frac{1}{\sqrt{2\pi s(t-s)/t}} \\
&= \exp(-(x - \frac{s}{t} z)^2/(2s(t-s)/t)) \cdot \frac{1}{\sqrt{2\pi s(t-s)/t}}.
\end{aligned}$$

The last line is the density of a Gaussian random variable and we conclude that

$$W_s|W_t = z \sim \mathcal{N}\left(\frac{s}{t}z, s(t-s)/t\right).$$

This is quite intuitive: $\mathbb{E}[W_s|W_t = z] = (s/t)z$ and $Var(W_s|W_t = z) = (s/t)^2(t-s)^2$.

Constructing Brownian Motion.

Here is the "physicist's" construction of BM. This idea was first advanced by Einstein in 1906 in connection with molecular/atomic forces.

Consider a particle of pollen suspended in liquid. The particle is rather large but light and is subject to being hit at random by the liquid molecules (which are small but very many). Let us focus on the movement of the particle along the $x-$axis, and record its position at time t as X_t^ϵ. Each time a particle is hit on the left, it moves to the right by a distance ϵ, and each time it is hit on the right, it moves left by an ϵ. Let L_t, M_t be the number of hits by time t on the left and right respectively. Then $X_t^\epsilon = \epsilon L_t - \epsilon M_t$ (for simplicity we took $X_0 = 0$). It is natural to assume that L_t, M_t are two independent stationary Markov processes with iid increments. That is, L_t, M_t are two Poisson processes with same intensity λ_ϵ. We immediately get $\mathbb{E}[X_t^\epsilon] = 0$ and

$$Var(X_t^\epsilon) = \epsilon^2\mathbb{E}[(L_t - M_t)^2] = \epsilon^2(2(\lambda_\epsilon t + (\lambda_\epsilon t)^2) - 2\mathbb{E}[L_t M_t] = 2\epsilon^2\lambda_\epsilon t.$$

So far ϵ was a dummy parameter and of course we are planning on taking $\epsilon \to 0$. To obtain an interesting (X_t) in the limit, we should make sure that the variance above converges to some non-zero finite limit. This would happen if we take $\lambda_\epsilon = C/(2\epsilon^2)$. Thus, as the displacement of the pollen particle from each hit is diminished we naturally increase the frequency of hits (and the correct scaling is quadratic).

To figure out what is the limiting (X_t) we compute the mgf (recall the mgf of $Poisson(\lambda)$ is $\mathrm{e}^{\lambda(\mathrm{e}^{s-1})}$):

$$\begin{aligned}\mathbb{E}[\mathrm{e}^{sX_t^\epsilon}] &= \exp\left(\lambda_\epsilon t(\mathrm{e}^{s\epsilon} - 1)\right) \cdot \exp\left(\lambda_\epsilon t(\mathrm{e}^{-s\epsilon} - 1)\right) \\ &= \mathrm{e}^{tC/(2\epsilon^2)(\mathrm{e}^{s\epsilon} + \mathrm{e}^{-s\epsilon} - 2)}.\end{aligned}$$

Since by L'Hopital rule

$$\begin{aligned}\lim_{\epsilon \to 0} \frac{\mathrm{e}^{s\epsilon} + \mathrm{e}^{-s\epsilon} - 2}{\epsilon^2} &= \lim_{\epsilon \to 0} \frac{s\mathrm{e}^{s\epsilon} - s\mathrm{e}^{-s\epsilon}}{2\epsilon} \\ &= \lim_{\epsilon \to 0} \frac{s^2\mathrm{e}^{s\epsilon} + s^2\mathrm{e}^{-s\epsilon}}{2} \\ &= s^2\end{aligned}$$

we have

$$\lim_{\epsilon \to 0} \mathbb{E}[\mathrm{e}^{sX_t^\epsilon}] = \mathrm{e}^{tCs^2/2}.$$

The latter is the moment generating function of a $\mathcal{N}(0, tC)$ random variable. Since for every $\epsilon > 0$, (X_t^ϵ) had stationary and independent increments (because L_t, M_t do), so does the limiting (X_t). Moreover, since the displacement $\epsilon \to 0$, (X_t) should be continuous. Putting it all together we conclude that (X_t) is a Brownian motion with zero drift and volatility C. If $C = 1$ then we get the Wiener process.

The name Brownian motion comes from the botanist Robert Brown who first observed the irregular motion of pollen particles suspended in water in 1828. As you can see it took 80 years until Einstein could provide a good mathematical model for these observations. Einstein's model was a crucial step in convincing physicists that molecules really exist and move around randomly when in liquid or gaseous form. Since then, this model has indicated to physicists that any random movement in physics should be based on Brownian motion, at least on the atomic/molecular level where external forces, such as gravity, are minimal.

Remark: note that (X^ϵ_t) above is a scaled simple birth-and-death process.

Recognizing Brownian Motion.

We first observe that if (X_t) is a Brownian motion with $\mu = 0$ then (X_t) is a martingale with respect to itself. Indeed, let $\mathcal{F}_t = \{X_u : 0 \le u \le t\}$ be the history of X up to time t. Then for $s > t$

$$\mathbb{E}[X_s|\mathcal{F}_t] = \mathbb{E}[(X_s - X_t) + X_t|\mathcal{F}_t] = \mathbb{E}[X_s - X_t] + X_t = X_t,$$

since the expected value of any increment is zero. More generally, $Y_t = X_t - \mu t$ is a martingale for any Brownian motion with drift μ.

The above property has far-reaching consequences. In fact, it turns out that the Wiener process is the canonical continuous martingale. This fact forms the basis for stochastic calculus and underlines the importance of understanding the behavior of BM. One basic application is the following Levy characterization of Wiener process:

Theorem 2. *Suppose that (Z_t) is a continuous-time stochastic process such that:*

- *The paths of Z are continuous.*
- *(Z_t) is a martingale with respect to its own history.*
- *$Var(Z_t - Z_s) = (t - s)$ for any $t > s > 0$.*

Then $(Z_t - Z_0)$ is a Wiener process.

As a corollary, if (Z_t) is a continuous process such that $Z_t - \mu t$ is a martingale and $Var(Z_t) = \sigma^2 t$ then (Z_t) is a Brownian motion with drift μ and volatility σ.

From Random Walk to Brownian Motion.

Here is another construction of Brownian motion. Let (S^δ_t) be a simple symmetric random walk that makes steps of size $\pm\delta$ at times $t = 1/n, 2/n, \ldots$. We know that $S(^\delta_t)$ is a time- and space-stationary discrete-time martingale. In particular, $\mathbb{E}[S^\delta_t] = 0$, and $Var(S^\delta_t) = \delta^2 \lfloor nt \rfloor$. To pass to the limit $\delta \to 0$ we therefore select $n = 1/\delta^2$. As $\delta \to 0$, we obtain a continuous-time, continuous-space stochastic process that has stationary and independent increments, continuous-paths and the martingale property. Moreover, $Var(S_t^0) = t$, so (S_t^0) must be the Wiener process.

One can also pick other random walks to make the convergence faster. For example a good choice is to take the step distribution as

$$X_n = \begin{cases} +\delta & \text{with prob. } 1/6 \\ 0 & \text{with prob. } 2/3 \\ -\delta & \text{with prob. } 1/6 \end{cases}$$

and $n = 3/\delta^2$. The corresponding random walk $S_n = X_1 + \ldots + X_n$ can go up, go down, or stay at the same level.

It is often helpful to think of Brownian motion as the limit of symmetric random walks. As we will see, most of the general results about random walks (including recurrence, reflection principle, ballot theorem, ruin probabilities) carry over nicely to Brownian motion.

Hitting Time Distribution.

Let (W_t) be the Wiener process and

$$T_b(\omega) = \min\{t \geq 0 : W_t(\omega) = b\}$$

be the first time (W_t) hits level b. We are interested in computing the distribution of T_b. Since the behavior of the Wiener process is symmetric about the x-axis, we take $b > 0$.

Define $\hat{W}_t = W_{T_b+t} - W_{T_b}$ to be the future of W after time T_b. Note that T_b is random, so it is not clear how $\hat{W}$ looks like. However, it is possible to verify all the conditions of Levy's Theorem and conclude that $\hat{W}$ is again a Wiener process. This is because (W_t) is a *strong Markov* process, meaning that the Markov property of (W_t) continues to hold when applied at random (stopping!) times, such as T_b. Thus, the future of (W_t) after T_b is independent of its history up to T_b.

From this fact we obtain

$$\mathbb{P}(W_t > b | T_b < t) = \mathbb{P}(\hat{W}_{t-T_b} > 0) = 1/2,$$

since $\mathbb{P}(\hat{W}_s > 0) = 1/2$ for any time s by symmetry. However, the LHS above can also be written as

$$\mathbb{P}(W_t > b | T_b < t) = \frac{\mathbb{P}(W_t > b, T_b < t)}{\mathbb{P}(T_b < t)} = \frac{\mathbb{P}(W_t > b)}{\mathbb{P}(T_b < t)},$$

since the only way that W is above b at t is if the hitting time of b has already occurred. Comparing we conclude that

$$\mathbb{P}(T_b < t) = 2\mathbb{P}(W_t > b) = 2\int_{b/\sqrt{t}}^{\infty} \frac{1}{\sqrt{2\pi}} \mathrm{e}^{-x^2/2} dx.$$

An immediate corollary is that if we take $t \to \infty$ then the RHS goes to $2\int_0^\infty \frac{1}{\sqrt{2\pi}} \mathrm{e}^{-x^2/2} dx = 1$, which means that $\mathbb{P}(T_b < \infty) = 1$, irrespective of value of b. Therefore, the Wiener process hits any level b with probability 1.

Furthermore, differentiating the above with respect to t we find that the density of T_b is given by:

$$\mathbb{P}(T_b \in dt) = \frac{|b| e^{-b^2/(2t)}}{\sqrt{2\pi t^3}}, \qquad b \in \mathbb{R}. \tag{1}$$

The latter belongs to the family of *stable* (or inverse Gamma) distributions (we say that T_b has *stable*(1/2)-distribution). In fact, here is a curious connection: let $Z \sim \mathcal{N}(0,1)$ and $Y = b^2/Z^2$. Then $T_b \sim Y$. Indeed,

$$\mathbb{P}(Y \leq t) = \mathbb{P}(b^2/Z^2 \leq t) = \mathbb{P}(|Z| \geq |b|/\sqrt{t}) = 2\mathbb{P}(Z \geq |b|/\sqrt{t}).$$

Differentiating,

$$\mathbb{P}(Y \in dt) = 2\frac{|b|}{2t^{3/2}} \frac{1}{\sqrt{2\pi}} e^{-(b/\sqrt{t})^2/2}$$

which matches with (1).

Using (1) one also shows that $\mathbb{E}[T_b] = +\infty$ for any b. So the conclusion is:

No matter how large b is, $\mathbb{P}(T_b < \infty) = 1$.

No matter how small $|b|$ is, $\mathbb{E}[T_b] = \infty$.

Maximum Process.

Let $M_t = \max_{0 \leq s \leq t} W_s$ be the running maximum of Wiener process W. Then (M_t) is a non-decreasing process that goes to $+\infty$ since we know that (W_t) will eventually hit any positive level. Moreover,

$$\begin{aligned}\mathbb{P}(M_t > x) &= \mathbb{P}(T_x < t) \\ &= 2\mathbb{P}(W_t > x) = \mathbb{P}(\sqrt{t}|Z| > x)\end{aligned}$$

based on properties of normal r.v.'s. We conclude that the marginal distribution of M_t is the same as that of an absolute value of a Gaussian random variable. In particular,

$$\mathbb{E}[M_t] = \sqrt{2t/\pi}, \qquad Var(M_t) = (1 - 2/\pi)t.$$

Note that while (M_t) is increasing, it does so in an imperceptible creep: for any fixed t, (M_t) is flat in the neighborhood of t with probability 1. If you draw a typical path of (M_t) on an interval of say $[0,1]$, the length of flat areas will add up to 1 (even though $M_1 > M_0$ with probability 1). For those of you who know some analysis, the set of points where (M_t) icnreases is a Cantor set of measure zero.

By symmetry the minimum $m_t = \min_{0 \leq s \leq t} W_s$ has the same distribution as $-M_t$. Since both $m_t \to -\infty$, $M_t \to +\infty$ it must be that (W_t) recrosses zero an infinite number of times. Thus, the zero level (and by space-homogeneity, *any level*) is recurrent for (W_t).

The next computation shows that much more is true:

Probability of a Zero.

We are interested in computing the probability that the Wiener process W hits zero on the time interval $[s,t]$. Observe that if $W_s = a$, then the probability that W hits zero on $[s,t]$ is just

$$\mathbb{P}(T_a < (t-s)) = \int_0^{t-s} \frac{|a|}{\sqrt{2\pi y^3}} \mathrm{e}^{-a^2/(2y)}\, dy.$$

However, $W_s \sim \mathcal{N}(0,s)$, so putting it together (ie conditioning on W_s) we have:

$$\begin{aligned}
\int_{-\infty}^{\infty} \mathbb{P}(T_a < (t-s)) \frac{1}{\sqrt{2\pi s}} \mathrm{e}^{-a^2/(2s)}\, da &= \int_{-\infty}^{\infty} \int_0^{t-s} \frac{|a|}{\sqrt{2\pi y^3}} \mathrm{e}^{-a^2/(2y)}\, dy \frac{1}{\sqrt{2\pi s}} \mathrm{e}^{-a^2/(2s)}\, da \\
&= \int_0^{t-s} \frac{1}{2\pi s^{1/2} y^{3/2}} \left(\int_{-\infty}^{\infty} |a| \mathrm{e}^{-a^2/(2y) - a^2/(2s)} da \right) dy \\
&= \int_0^{t-s} \frac{1}{\pi s^{1/2} y^{3/2}} \left(\int_0^{\infty} |a| \mathrm{e}^{-a^2 * (y+s)/(2sy)} da \right) dy \\
&= \int_0^{t-s} \frac{1}{\pi s^{1/2} y^{3/2}} \frac{sy}{s+y}\, dy \\
&= \frac{\sqrt{s}}{\pi} \int_0^{t-s} \frac{dy}{(s+y)\sqrt{y}} \\
&= \frac{2}{\pi} \int_0^{\sqrt{(t-s)/s}} \frac{dx}{1+x^2} \quad \text{by substitution } y = sx^2 \\
&= \frac{2}{\pi} \arctan(\sqrt{(t-s)/s}) = \frac{2}{\pi} \arccos \sqrt{\frac{s}{t}}.
\end{aligned}$$

Corollary: if $s=0$ then probability of a zero on $[0,t]$ is $\frac{2}{\pi} \arccos 0 = 1$ for any t. Therefore, W returns to zero immediately after $t=0$. After some thinking and using the strong Markov property of W, it follows that W has infinitely many zeros on any interval $[0,t]$. This is quite counterintuitive given that $\mathbb{E}[T_\epsilon] = +\infty$ even for very small ϵ. Thus W really oscillates wildly around zero. This is one indication that W is nowhere differentiable: it oscillates around every point and the scale of oscillations over time interval of length h is on the order of $\mathcal{O}(\sqrt{h})$, which means that $\lim_{h\to 0} \frac{W_{t+h}-W_t}{h}$ cannot exist.

Time of Last Zero.

Let L_t be the time of last zero of W before time t. Let D_t be the time of first zero of W after time t. We observe that $\{L_t < s\} = \{D_s > t\} = \{W \text{ has no zeros on } [s,t]\}$. From the last section, we therefore find that

$$\mathbb{P}(L_t < s) = 1 - \frac{2}{\pi} \arccos \sqrt{\frac{s}{t}} = \frac{2}{\pi} \arcsin \sqrt{\frac{s}{t}},$$

meaning that the density of L_t is:

$$\mathbb{P}(L_t \in ds) = \frac{1}{\pi\sqrt{s(t-s)}}, \qquad 0 \le s \le t.$$

These facts are known as the Arc-sine Law and show that the last zero of W is likely to be either close to zero or to t.

Here is another derivation of the same result using the Gaussian-Gamma connection. Note that conditional on $W_t = a$, the distribution of D_t is $D_t \sim t + T_{-a} \sim t + \frac{a^2}{\gamma_0}$ where $\gamma_0 \sim Gamma(1/2, 1/2)$. Therefore,

$$\begin{aligned}
\mathbb{P}(D_t \le (t+u)) &= \int_{-\infty}^{\infty} \mathbb{P}(D_t - t \le u | W_t = a) f_{W_t}(a)\, da \\
&= \int_{-\infty}^{\infty} \mathbb{P}(a^2/\gamma_0 \le u) f_{W_t}(a)\, da \\
&= \mathbb{P}(W_t^2/\gamma_0 \le u) \\
&= \mathbb{P}(t\gamma_1/\gamma_0 \le u) \\
&= \mathbb{P}\left(\frac{t\gamma_1}{\gamma_0} \le u\right),
\end{aligned}$$

where $\gamma_1 \sim Gamma(1/2, 1/2)$ is another Gamma rv, independent of γ_0 (since T_a is indep of W_0). So distribution of $D_t - t$ is same as $\frac{t\gamma_1}{\gamma_0}$ ie distribution of D_t is equal to that of $t(1 + \gamma_1/\gamma_0) = t/\frac{\gamma_0}{\gamma_0+\gamma_1}$. The latter expression $\beta = \frac{\gamma_0}{\gamma_0+\gamma_1}$ is known to have a $Beta(1/2, 1/2)$ distribution with density $f_\beta(u) = \frac{1}{\pi\sqrt{u(1-u)}}$, $0 \le u \le 1$. So $\mathbb{P}(D_t < (t+u)) = \mathbb{P}(t/\beta < (t+u)) = \mathbb{P}(\beta > t/(t+u))$.

However, as mentioned before,

$$\mathbb{P}(L_t < s) = \mathbb{P}(D_s > s + (t-s)) = \mathbb{P}(\beta < s/t) = \int_0^{s/t} \frac{1}{\pi\sqrt{u(1-u)}} du = \frac{2}{\pi} \arcsin\sqrt{\frac{s}{t}},$$

using the substitution $x = \sqrt{u}$. In particular note that the distribution of L_1 is $Beta(1/2, 1/2)$.

Planar BM

Let (W^1, W^2) be two independent Wiener processes, and take $X_t = X_0 + W_t^1$, $Y_t = Y_0 + W_t^2$. Then viewing (X_t, Y_t) as the (x, y)-coordinates of a moving particle, we obtain what is called a planar Brownian motion, starting at the initial location (X_0, Y_0). Here is one curious computation:

Suppose $(X_0, Y_0) = (0, b)$ so the particle starts somewhere on the positive y-axis. Let T be the first time that $Y_t = 0$ the particle hits the x-axis. We are interested in the distribution of X_T. Recall that from the remark after (1), $T \sim b^2/Z^2$. It follows that $X_T \sim \sqrt{T} Z_0$ where (Z_0, Z) are two independent standard normal random variables. In other words, $X_T \sim b\frac{Z_0}{Z}$ which we recall has a Cauchy distribution, namely

$$\mathbb{P}(X_T \in dx) = \frac{b}{(b^2 + x^2)\pi}.$$

Like the one-dimensional case, two-dimensional Brownian motion is recurrent and will hit every point in the plane with probability 1.