A stochastic differential equation is a particular type of noisy dynamical system. Applications call for generating SDE sample paths subject to a variety of constraints.

1 Multivariate normal

Let $X \in R^n$ be an n component normal with covariance matrix C. Choose an orthonormal family of vectors v_k and eigenvalues λ_k so that $Cv_k = \lambda_k v_k$. The λ_k are the *principle components* of X (or of C) and the v_k the corresponding *principle components* (the terminology differs slightly from place to place). The random numbers $Y_k = v_k^t X$ are independent mean zero Gaussians with mean zero and variance λ_j. Said another way, if V is the orthogonal matrix whose columns are the v_k, then $Y = V^t X$ is a multivariate normal with mean zero and covariance Λ, the diagonal matrix with λ_k in the (k,k) position. The Y_k also may be called principle components, or possibly components in the principle directions. To summarize, the eigenvalue/eigenvector decomposition of C is the *principle component analysis* of X. When using PCE, we usually order the principle values from largest to smallest: $\lambda_k \geq \lambda_{k+1}$ for all k.

On the other hand, suppose we know the v_k and λ_k. We can construct a sample Y_k as $Y_k = \sigma_k Z_k$, where $\sigma_k = \sqrt{\lambda_k}$ and the Z_k are independent standard normals. Therefore, the formula $X = \sum_{k=1}^n \sigma_k Z_k$ provides a sample of X from independent standard normals. This kind of PCA is helpful in many ways in Monte Carlo practice. If we want to know $A = E[f(X)]$, it is likely that the larger principle components are more important than the smaller ones. For example, the quantity $B = \int f(\sigma_1 v_1 z_1) e^{-z_1^2/2} dz/\sqrt{2\pi}$ may be a helpful control variate for evaluating A.

2 Brownian motion

A standard univariate Brownian motion path is a continuous random function $X(t)$ with $X(0) = 0$ and *increments*, $\Delta X_k = X(t_{k+1}) - X(t_k)$ independent Gaussian random variables with mean zero and variance $t_{k+1} - t_k$. Here t_k is any sequence of positive numbers with $t_{k+1} > t_k$. Many problems involve expected values of functions of Brownian motion and the intuitions gained by studying Brownian motion help design methods for related problems.

We start with the PCA of standard Brownian motion. Since Brownian motion is infinite dimensional, we start with finite dimensional approximations then get the PCA of Brownian motion itself as a limit. Suppose we fix a $T > 0$. For any integer $n > 0$, we divide the interval $[0, T]$ into n equal subintervals using the notation $h = T/n$ and $t_k = kh$. The multivariate normal $X^h \in R^n$ has components $X_k^h = X(t_k)$, for $1 \leq k \leq n$. We showed before that the covariance matrix for X^h has $H = C^{-1}$ that is tri-diagonal of the form .... . Clearly the eigenvalues of C are the inverses of the eigenvalues of H while the eigenvectors are the same. It happens (we will get more insight into this later)

that the eigenvectors of H have the form $v^h_{k,j} = \sin(\omega^h_k t^h_j)$, where ... . If we take the limit $h \to 0$ and $n \to \infty$ with $T = nh$ fixed, the sequence of eigenvalues of H converges to $\mu_k = (k+1)^2\pi^2/T^2$, as we verify by direct computation.

The PCA representation of a Brownian motion path is

$$X(t) = \frac{T}{\pi} \sum_{k=1}^{\infty} \frac{Z_k}{k+\frac{1}{2}} \sin\left(\frac{k+\frac{1}{2})\pi}{T}\, t\right) \quad . \tag{1}$$

This representation is valid only for t in the range $0 \le t \le T$. Moreover, it presents serious mathematical challenges. It is not straightforward to see that the sum converges. Indeed, the Z_k and the sine factors are of the order of unity, so the terms in the sum seem to have the order of magnitude $1/k$, which would not by itself lead to a convergent sum. The series does converge in the L^2 sense because

$$E\left[\sum_{k=1}^{\infty} \frac{Z_k^2}{\left(k+\frac{1}{2}\right)^2}\right] \;=\; \sum_{k=1}^{\infty} \frac{1}{\left(k+\frac{1}{2}\right)^2} \;<\; \infty\,.$$

However, this kind of convergence does not even imply that $X(t)$ is continuous. Many discontinuous are represented by L^2 convergence Fourier series. If we wish to generate a Brownian motion path using (1), we can use the FFT to calculate the sums on the right efficiently.

The Brownian Bridge construction is another good way to generate Brownian motion paths. For this, we consider a *diadic* collection of times. On level k, the interval $[0,T]$ is divided into 2^k equal subintervals. At the top is level zero with a single interval $[0,T]$ whose endpoints may be written $t_{0,0} = 0$ and $t_{0,1} = T$. At the next level $[0,T]$ is broken into two halves bounded by the points $t_{1,0} = 0$, $t_{1,1} = T/2$, and $t_{1,2} = T$. The diadic points at level k are $t_{k,j} = j2^{-k}T$, for integers j in the range $0 \le j \le 2^k$.

The construction itself is an algorithm that generates values for $X_{k,j} = X(t_{k,j})$ recursively proceeding from k to $k+1$. At the top level we already have $X_{0,0} = X(0) = 0$ and we need $X_{0,1} = X(T)$. But $X(T)$ is a mean zero Gaussian with variance T, so we can take $X_{0,1} = \sqrt{T}Z_{0,1}$, where $Z_{0l,1}$ is a standard normal. Things get interesting at the next level. From the properties of Brownian motion, we know that the values $(X(t_{1,1}), X(t_{1,2})) = (X_{1,1}, X_{1,2})$ is a bivariate normal. Therefore, conditional on the value of $X_{1,2}$, which already is known from the previous level, the value of $X_{1,1}$ is a univariate normal. Simple calculations show that $E[X_{1,1} \mid X_{1,2}] = \frac{1}{2}X_{1,2}$, and $\mathrm{var}[X_{1,1} \mid X_{1,2}] = \frac{1}{4}T$. The first result is very natural, that the expected value of the midpoint value should be the average of the two end values. The second result may be surprising, since the variance of $X_{1,1}$ without information about $X_{1,2}$ is $T/2$ instead of $T/4$. Specifying both endpoint values constrains $X_{1,1}$ to have a smaller variance. With this information the sampling can be done using $X_{1,1} = \frac{1}{2}X_{0,1} + \frac{\sqrt{t}}{2}Z_{1,1}$. As usual, $Z_{1,1}$ is a standard normal independent of all others.

The step of going from level $k-1$ to level k is clear. We already have values $X_{k,2j} = X_{k-1,j}$ from the previous levels. We need to generate values for the new points $t_{k,2j+1}$. The Markov property (which is the main thing that makes this

construction work) implies that the conditional distribution of $X_{k,2j+1}$ given all the values on level $k-1$ is the same as the distribution given the level $k-1$ endpoints $X_{k,2j} = X_{k-1,j}$, and $X_{k,2j+2} = X_{k-1,j+1}$. As above, the conditional expectation of $X_{k,2j+1}$ is $\frac{1}{2}(X_{k-1,j} + X_{k-1,j+1})$. Let $\Delta t_k = 2^{-k}T$ be the length of the diadic intervals at level k. The conditional variance of $X_{k,2j+1}$ is $\Delta t_k/2$ (which again is half the conditional variance if only the left value $X_{k,2j}$ were specified). Therefore, the sampling may be done by

$$X_{k-2j+1} = \frac{1}{2}(X_{k-1,j} + X_{k-1,j+1}) \; + \; \sqrt{\frac{\Delta t_k}{2}} Z_{k,2j+1} \; .$$

One use of this Brownian bridge construction is similar to the PCA construction, to have a way to separate random variables that have a large impact on X from others that have a smaller impact.

3 The Ornstein Uhlenbeck Process

The *Ornstein Uhlenbeck* process is a simple modification of Brownian motion to add *mean reversion*. Starting here, we refer to the proces of interest as $X(t)$ while $W(t)$ will be Brownian motion. The general Ornstein Uhlenbeck process satisfies the SDE

$$dX = -\mu X dt + \sigma dW \; . \tag{2}$$

This means that we can write

$$X(T) = e^{-\mu T} X(0) + \sigma \int_0^T e^{-\mu(T-t)} dW(t) \; . \tag{3}$$

We see from this that X is a linear transformation of the Gaussian W so it is Gaussian. This is contingent on $X(0)$ being Gaussian if it is random. From this we can calculate that $X(T)$ is a gaussian with mean $e^{-\mu T}X(0)$, and variance

$$\sigma^2 \int_0^T e^{-2\mu(T-t)} dt \; = \; \frac{\sigma^2}{2\mu}\left(1 - e^{2\mu T}\right) \; \rightarrow \; \frac{\sigma^2}{2\mu} \text{ as } T \rightarrow \infty \; . \tag{4}$$

As $T \rightarrow \infty$, the distribution of $X(T)$ converges to a gaussian that is independent of the initial condition (if the mean and variance converge and if it's Gaussian, the distribution converges).

Much Monte Carlo depends on time dependent random processes like this so we analyze this in more detail, focusing on features that have analogues in many Monte Carlo applications. Suppose $u(x,t)$ is the probability density function for $X(t)$. This satisfies the Kolmogorov forward equation (other names include, but are not limited to, the forward equation, the Fokker Planck equation, and the Chapman Kolmogorov equation)

$$\partial_t u \; = \; \frac{\sigma^2}{2}\partial_x^2 u + \mu \partial_x (xu) \; . \tag{5}$$

We write this as

$$\partial_t u = L^* u \ , \quad \text{where } L^* u \ = \ \frac{\sigma^2}{2}\partial_x^2 u + \mu \partial_x \left(x u \right) \ . \tag{6}$$

It is traditional to call the operator in (6) L^* so that it's adjoint that appears in the backward equation can be called L.

If the distribution of $X(t)$ has a limit as $t \to \infty$, then we expect $u(x,t) \to u_0(x)$ as $t \to \infty$. From the equation (5), this suggests that u_0 is a solution with $L^* u_0 = 0$. This gives

$$\frac{\sigma^2}{2}\partial_x^2 u + \mu \partial_x \left(x u \right) = 0 \ ,$$

If $u_0 \to 0$ as $x \to \infty$ together with its derivatives, this may be integrated to

$$\frac{\sigma^2}{2}\partial_x u_0 + \mu x u_0 = 0 \ .$$

The solution takes the form $u_0(x) = Ce^{-x^2/2\rho}$, with $\rho = \sigma^2/2\mu$.

We are interested in the rate of convergence $u(\cdot,t) \to u_0$,which depends on the other eigenvalues of the operator L^*. That is, we seek functions

$$L^* u_k \ = \ \frac{\sigma^2}{2}\partial_x^2 u_k(x) + \mu \partial_x \left(x u_k \right) \ = \ -\lambda_k u_k \ . \tag{7}$$

It turns out that the solutions have the form $u_k(x) = \partial_x^k u_0$. We verify this by induction on k, the case $k = 0$ having been done above. The trick for doing the calculation is the "differentiation by parts" formula $f\partial_x g = \partial_x(fg) - (\partial_x f)g$. If $f = \partial_x^{k-1} u_0$ and $g = x$, the result is

$$x\partial_x^k u_0 \ = \ \partial_x(x\partial_x^{k-1} u_0) - \partial_x^{k-1} u_0 \ .$$

Therefore, using $L^* u_{k-1} = -\lambda_{k-1} u_{k-1}$,

$$\begin{aligned} L^* u_k \ = \ L^* \partial_x u_{k-1} \quad &= \quad \frac{\sigma^2}{2}\partial_x^2 \partial_x u_{k-1} + \mu \partial_x (x \partial_x u_{k-1}) \\ &= \quad \partial_x \left(\frac{\sigma^2}{2}\partial_x^2 u_{k-1} \right) \ + \ \partial_x \left(\ \mu \partial_x \left(x u_{k-1} \right) \ \right) \ - \ \mu \partial_x u_{k-1} \\ &= \quad \partial_x \left(L^* u_{k-1} \right) - \mu u_k \quad = \quad -\left(\lambda_{k-1} + \mu \right) u_k \ . \end{aligned}$$

This shows that $L^* u_k = -\lambda_k u_k$ with $\lambda_k = \lambda_{k-1} + \mu$. To conclude, the eigenvalues are 0 (the steady state), $-\mu$, -2μ, etc. The corresponding eigenfunctions are derivatives of Gaussians, which is to say, given by Hermite polynomials (more on this below).

We use the above information to describe the behavior of the probability density $u(x,t)$. The solution of the evolution equation (5) is a linear combination of eigenfunctions:

$$u(x,t) = \sum_{k=0}^{\infty} a_k e^{-k\mu t} u_k(x) \ , \tag{8}$$

where the coefficients are determined by the initial data $u(x,0)$. The $k=0$ term corresponds to the steady state solution $a_0 u_0(x)$. All the other terms decay with various exponential rates. Slowest to decay is the $k=1$ term, which will make up most of the difference between $u(x,t)$ and $a_0 u_0(x)$ for large t:

$$u(x,t) - a_0 u_0(x) \approx a_1 e^{-\mu t} u_1(x) \ , \quad \text{for large } t.$$

It may not be surprising that the exponential decay rate μ in (2) results in the same exponential decay rate for $u(x,t) - a_0 u_0(x)$. What is new in the PDE solution is the family of faster decay rates 2μ, 3μ, etc. We will find versions of this phenomenon in any linear Gaussian process.

There are many practical applications for this kind of analysis. One concerns dynamic sampling methods, also called Markov Chain Monte Carlo. A complicated probability distribution may not have an efficient simple sampler, but there may be a dynamic process, a Markov Chain, that preservers the distribution. Here, the dynamical process (2) preserves the Gaussian $Ce^{-x^2/2\rho}$. A major issue for dynamic samplers is the rate of convergence: How long do you have to run the process before the distribution of $X(t)$ closely approximates $u_0(x)$. In this simple problem, the convergence time is of the order of $1/\mu$, which also is the decay rate for the linear Gaussian process (2).

3.1 Hermite polynomials

Hermite polynomials have many applications in Monte Carlo and other parts of applied probability. The n^{th} *Hermite polynomial*, $H_n(x)$, is defined by

$$H_n(x) e^{-x^2/2} \ = \ \pm \partial_x^n e^{-x^2/2} \ . \tag{9}$$

The calculations $\partial_x e^{-x^2/2} = -x e^{-x^2/2}$ and $\partial_x^2 e^{-x^2/2} = (x^2-1) e^{-x^2/2}$, and $\partial_x^3 e^{-x^2/2} = (-x^3 - 3x) e^{-x^2/2}$ imply that the first few polynomials are $H_0(x) = 1$, $H_1(x) = x$, $H_2(x) = x^2 - 1$ and $H_3(x) = x^3 - 3x$. It is common to choose the sign so that the leading term is positive and also to choose the sign itself always to be positive (which would give $H_1(x) = -x$). It is easy to verify by induction that (9) defines polynomials of degree n with leading coefficient equal to ± 1.

The H_n have the *orthogonality* property that

$$\int_{-\infty}^{\infty} H_n(x) H_m(x) e^{-x^2/2} dx \ = \ 0 \ , \quad \text{for } n \neq m. \tag{10}$$

We briefly indicate a proof of this by induction on m, assuming $m < n$, since a better proof will come shortly. Actually, we prove (10) holds for a given m and all $n > m$ by induction on m. It clearly is true for $m = 0$, because then it just is $\int \partial_x^n e^{-x^2/2} dx = 0$. Now assume it is true for $m-1$ and try to prove it for m. The rewrite of (9) as $H_m(x) = e^{x^2/2} \partial_x^m e^{-x^2/2}$ puts (10) in the form

$$\int_{-\infty}^{\infty} e^{x^2/2} \left(\partial_x^n e^{-x^2/2} \right) \left(\partial_x^m e^{-x^2/2} \right) dx \ = \ 0$$

Write $\partial_x^m = \partial_x \partial_x^{m-1}$ and integrate by parts, and we get the two terms

$$-\int_{-\infty}^{\infty} e^{x^2/2} \left(\partial_x^{n+1} e^{-x^2/2}\right) \left(\partial_x^{m-1} e^{-x^2/2}\right) dx$$
$$- \int_{-\infty}^{\infty} e^{x^2/2} \left(x\partial_x^n e^{-x^2/2}\right) \left(\partial_x^{m-1} e^{-x^2/2}\right) dx$$

The top term is zero by the induction hypothesis, which is why we used a more complex induction hypothesis. We study the bottom term using something similar to the differentiation by parts idea above, namely the formula $x\partial_x^n f = \partial_x^n (xf) - n\partial_x^{n-1} f$. This implies that the bottom term is

$$\int_{-\infty}^{\infty} e^{x^2/2} \partial_x^n \left(xe^{-x^2/2}\right) \left(\partial_x^{m-1} e^{-x^2/2}\right) dx$$
$$- n\int_{-\infty}^{\infty} e^{x^2/2} \left(\partial_x^{n-1} e^{-x^2/2}\right) \left(\partial_x^{m-1} e^{-x^2/2}\right) dx$$

The bottom term is zero because $m - 1 < n - 1$. The top term is zero because $xe^{-x^2/2} = -\partial_x e^{-x^2/2}$, so it is equal to

$$-\int_{-\infty}^{\infty} e^{x^2/2} \left(\partial_x^{n+1} e^{-x^2/2}\right) \left(\partial_x^{m-1} e^{-x^2/2}\right) dx \;=\; 0\,,$$

again by induction.

The orthogonality property of Hermite polynomials gives a different way to find them. This is because of the fact that if $p(x)$ is a polynomial of degree n if and only if it may be written as a sum of Hermite polynomials of order not more than n:

$$p(x) = \sum_{m=0}^{n} a_m H_m(x)\,.$$

Therefore, (10) is true for a given n and all $m < n$ if and only if

$$\int H_n(x) x^m e^{-x^2/2} dx \;=\; 0$$

for all $m < n$. This is the same as saying that $E[H_n(X)X^m] = 0$ for all $m < n$, where X is a standard normal random variable. We compute the first few that way. Of course, $H_0(x) = 1$. $H_1(x)$ is the polynomial of degree 1 that is orthogonal to $X^0 = 1$, which is to say $H_1(x) = x$. Next, $H_2(x)$ is the quadratic that is orthogonal to x and 1. We achieve $E[H_2(X)X] = 0$ if H has the form $x^2 - a$. We achieve $E[H_2(X)X^0] = E[H_2(X)] = 0$ if $a = E[X^2] = 1$, i.e. if $H_2(x) = x^2 - 1$. Finally, we want $H_3(x) = x^3 - ax^2 - bx - c$ to make H_3 orthogonal to 1, x, and x^2. The even powers are automatic if $H_3(x)$ is odd, i.e. if $H_3(x) = x^3 - bx$. The remaining orthogonality condition $E[H_3(X)X] = 0$ is $E[X^4] - bE[X^2] = 0$, which gives $b = 3$, as claimed above. It is the orthogonality properties of Hermite polynomials that makes them handy, for

example, in finding kernel functions for density estimation that satisfy several moment conditions.

In the slightly more general situation of $e^{-x^2/2\rho}$, we can define $H_n(x,\rho) = e^{x^2/2\rho}\partial_x^n e^{-x^2/2\rho}$. These polynomials are orthogonal in the natural sense

$$\int_{-\infty}^{\infty} H_n(x,\rho)H_m(x,\rho)e^{-x^2/2\rho}dx = 0$$

for $m \neq n$. The proof is the same as above.

3.2 Adjoint eigenfunctions

The operator L^* has an *adjoint* operator, L. The eigenvalues of L^* and L are the same but the eigenfunctions are different. The eigenfunctions of L are used to determine the expansion coefficients a_k in (8). Later we will find a way to derive the adjoint eigenfunctions directly from the u_n using the principle of *detailed balance*.

We explain the ideas in a finite dimensional setting. Suppose A is an $n \times n$ matrix. The right eigenvectors of A are v_k with $Av_k = \lambda_k v_k$. The left eigenvectors are row vectors u_k with $u_k A = \lambda_k u_k$. This is the same as $A^* u_k^* = \lambda_k u_k^*$. That is, the left eigenvectors of A (when written as column vectors) are the right eigenvectors of the adjoint (transpose) of A. The eigenvalues λ_k are the same. The right and left eigenvectors satisfy *biorthogonality relations* $u_k v_l = 0$ if $\lambda_k \neq \lambda_l$.

For column vectors f and g, let (f,g) be the l^2 inner product $(f,g) = f^*g = \sum_{k=1}^n f_k g_k$. The adjoint of the $n \times n$ matrix A is defined by the property that it should satisfy $(A^*f, g) = (f, Ag)$ for all vectors f and g (check this). It is possible to define other inner products. The one that will be useful for us is the *weighted* l^2 inner product $\langle f,g\rangle_w = \sum_{k=1}^n f_k g_k w_k$, where the w_k are n positive *weights*. Each weighted norm determines a different A^* through the property $\langle A^*f, g\rangle_w = \langle f, Ag\rangle_w$ for all f and g. In this case, the elements of A^* are given by $A^*_{jk} = A_{kj}w_k/w_j$ (check this). The eigenvalues of A^* are the same as the eigenvalues of A (for example, because A^* is similar to the transpose of A). The eigenvectors satisfy the natural biorthogonality relation: If $Av_j = \lambda_j v_j$ and $A^*u_k = \lambda_k u_k$ with $\lambda_j \neq \lambda_k$, then $\langle u_k, v_j\rangle_w = 0$ (check this).

Back to L^*, we compute the adjoint in the L^2 inner product $(f,g) = \int_{-\infty}^{\infty} f(x)g(x)dx$. This adjoint is defined by the requirement that for all functions for which the computation makes sense: $(L^*f, g) = (f, Lg)$. This is done by integration by parts:

$$\begin{aligned}
(L^*f, g) &= \int_{-\infty}^{\infty} \left(L^*f(x)\right) g(x)dx \\
&= \int_{-\infty}^{\infty} \left(\frac{\sigma^2}{2}\partial_x^2 f(x) + \mu\partial_x\left(xf(x)\right)\right) g(x)dx \\
&= \int_{-\infty}^{\infty} f(x)\left(\frac{\sigma^2}{2}\partial_x^2 g(x) - \mu x\partial_x g(x)\right) dx
\end{aligned}$$

$$= \int_{-\infty}^{\infty} f(x)\,(Lg(x))\,dx\ ,$$

with

$$Lg(x) = \frac{\sigma^2}{2}\partial_x^2 g(x) - \mu x \partial_x g(x)\ . \tag{11}$$

Note that the second derivative terms in L^* and L are the same while the first derivative terms change sign. Also, the x is inside the differentiation in L^* but outside in L.

It turns out that the eigenfunctions of L are Hermite polynomials $v_k(x) = H_k(x,\rho) = e^{x^2/2\rho}\partial_x^k e^{-x^2/2\rho}$. The biorthogonality relation $(u_j, v_k) = 0$ for $j \neq k$ is the orthogonality relation for Hermite polynomials, given that $u_j = \partial_x^j e^{-x^2/2\rho} = H_j(x,\rho)e^{-x^2/2\rho}$. The eigenvalue relation $Lv_k = -k\mu v_k$ can be verified by induction on k using the *recurrence relation* (one of many) $v_{k+1}(x) = \frac{-x}{\rho}v_x(x) + \partial_x v_k(x)$.

3.3 Detailed balance

So, the eigenfunctions of L^* and those of L differ only by the exponential factor $e^{-x^2/2\rho}$. The L^* eigenfunctions have it and the L eigenfunctions do not. This suggests that is a relationship between L and L^* that goes beyond one being the adjoint of the other. This relation is that in a suitable inner product, the adjoint of L is equal to L itself. This implies that L is similar to L^* with the similarity given by the weight used for the inner product. All of this is related to the[1] "principle" *detailed balance*.

There are many formulations of the detailed balance condition. One of them (which is not the origin of the name) is that L should be self adjoint in the inner product given by the steady state u_0. If $w(x)$ is a non negative *weight function*, the weighted L^2 inner product is

$$\langle f, g\rangle_w = \int_{-\infty}^{\infty} f(x)g(x)w(x)dx\ .$$

L is self adjoint in this inner product if $\langle Lf, g\rangle_w = \langle f, Lg\rangle_w$, for all sufficiently regular functions f and g. A direct calculation shows that the L in (11) is self-adjoint in the weighted norm with $w(x) = u_0(x) = e^{-x^2/2\rho}$.

Detailed balance implies a relation between the eigenfunctions of L and those of L^*. We can see this from a formula relating L to L^* derived as follows:

$$\begin{aligned}
\langle f, Lg\rangle_{u_0} &= \int_{-\infty}^{\infty} f(x)\Big(Lg(x)\Big)u_0(x)dx \\
&= \int_{-\infty}^{\infty} \Big(L^*\big(u_0 f\big)(x)\Big)g(x)dx \\
&= \int_{-\infty}^{\infty} \left(\frac{1}{u_0(x)}L^*\big(u_0 f\big)(x)\right) g(x)u_0(x)dx
\end{aligned}$$

[1] The *principle* of detailed balance is that certain physical processes satisfy detailed balance.

$$\langle f, Lg\rangle_{u_0} \quad = \quad \langle\ \frac{1}{u_0} L^*\big(u_0 f\big)\ ,\ g\ \rangle_{u_0}\ .$$

This, together with the detailed balance condition $\langle f, Lg\rangle_{u_0} = \langle Lf, g\rangle_{u_0}$, implies (because it holds for every g) that

$$u_0 Lf = L^*\big(u_0 f\big)\ . \tag{12}$$

Applying this with $f = v_j$ and $Lv_j = \lambda_j v_j$ shows that $u_j = u_0 v_j$ satisfies

$$L^* u_j = L^* u_0 v_j = u_0 L v_j = u_0 \lambda_j v_j = \lambda_j u_j\ ,$$

which is to say that u_j is an eigenfunction of L^* with eigenvalue λ_j. This is the eigenfunction relation we found above.

The detailed balance condition has a reformulation in terms of transition probabilities. The *transition probability density* is the probability density for $X(t)$ given that $X(0) = x_0$. We write this as $u(x_0, x, t)$, the probability (density) of going from x_0 to x in time t. From (3) and (4), we see that the conditional distribution of $X(t)$ is that of a Gaussian with mean $x_0 e^{-\mu t}$ and variance $\rho(t) = \frac{\sigma^2}{2\mu}\left(1 - e^{-2\mu t}\right)$. This implies that

$$u(x_0, x, t) = \frac{1}{\sqrt{2\pi\rho(t)}} \exp\left(-\left(x - x_0 e^{-\mu t}\right)^2 / 2\rho(t)\right)$$

Now suppose we sample x_0 from the steady state density $u_0(x_0) = Ce^{-x_0^2/2\rho}$. The probability density to observe a transition from x_0 to x is the probability of first choosing $X(0) = x_0$, then making a transition from that x_0 to x:

$$\begin{aligned} r(x_0, x, t) &= u_0(x_0) u(x_0, x, t) \\ &= Ce^{-x_0^2/2\rho} \frac{1}{\sqrt{2\pi\rho(t)}} \exp\left(-\left(x - x_0 e^{-\mu t}\right)^2 / 2\rho(t)\right)\ . \end{aligned}$$

Calculating the exponents shows that

$$r(x_0, x, t) \ = \ Ce^{-x_0^2/2\rho} \frac{1}{\sqrt{2\pi\rho(t)}} \exp\left\{\frac{-1}{2\rho}\left(\frac{x_0^2 + x^2}{1 - e^{-2\mu t}} - \frac{2x_0 x e^{-\mu t}}{1 - e^{-2\mu t}}\right)\right\}\ .$$

The interesting thing about this formula is that it symmetric in x and x_0. In other words, the probability of observing an $x_0 \to x$ transition is the same as the probability of observing an $x \to x_0$ transition.

This is the principle of detailed balance: In steady state, the probability of observing the transition $y \to x$ is the same as the probability of observing $x \to y$. It is called "detailed" balance because it is a more restrictive than the simple balance condition required for u_0 to be the steady state. Suppose $u(x_0, x, t)$ is any transition density, not necessarily related to the Ornstein Uhlenbeck process. A probability density $u_0(x)$ is a steady state for u if it satisfies the balance condition

$$u_0(x) = \int u_0(x_0) u(x_0, x, t) dx_0$$

for all x. If $u_0(x_0)u(x_0,x,t) = u_0(x)u(x,x_0,t)$ for all x_0 and x, then we can integrate with respect to x_0 using the fact that $\int u(x,x_0,t)dx_0 = 1$ for all x (because u is a transition density) and see that u_0 satisfies the balance condition to be a steady state. It will turn out that detailed balance is a particularly useful way to satisfy the balance condition, even though it is not necessary for balance.

These two forms of the detailed balance condition, $\langle f, Lg\rangle_{u_0} = \langle Lf, g\rangle_{u_0}$ and the symmetry of r, are closely related. In finite dimensional space and without weights, this relation is that if A is a symmetric matrix, then $S(t) = e^{tA}$ also is symmetric.

3.4 Multidimensional processes

A multidimensional Ornstein Uhlenbeck process is a linear process with linear constant intensity random forcing. If μ is an $n \times n$ matrix and σ is an $n \times m$ matrix (neither needs to be symmetric), and $W(t) = (W_1(t), \ldots, W_m(t))$ is an m dimensional standard Brownian motion whose components $W_k(t)$ are independent, then we can interpret (2) as an equation for the evolution of $X(t) = (X_1(t), \ldots, X_n(t))$. The case $m < n$ is common. In this case, σ has more rows than columns. The case $m > n$ is pointless and we never consider it.

The solution formula (3) is almost correct, provided that we interpret $e^{\mu t}$ as the matrix exponential that satisfies $\partial_t e^{\mu t} = \mu e^{\mu t}$. The only modification is that $\sigma e^{\mu(T-t)}$) does not make sense because it is not compatible for matrix multiplication if $m < n$. The correct form is (check this)

$$X(T) = e^{-\mu t}X(0) + \int_0^T e^{\mu(T-t)}\sigma dW(t) \; . \tag{13}$$

This formula shows that if $X(0)$ is Gaussian or deterministic, then $X(T)$ also is Gaussian for all $T > 0$. Therefore, we can identify the limiting density of $X(T)$ as $T \to \infty$ by figuring out the limiting mean and covariance matrix. The limit will exist if the ODE $\dot{x} = -\mu x$ is strongly stable, which is the same as the eigenvalues of μ having positive real parts[2]. The mean $m(t) = E[X(t)]$ satisfies the simple differential equation

$$dm = dE[X(t)] = E[dX] = E[-\mu X dt] + \sigma E[dW] = -\mu \; m \; dt \; .$$

This shows that $m(t) \to 0$ as $t \to \infty$ as expected.

The covariance matrix calculation is similar. Let $\rho(x) = E[X(t)X^*(t)]$ be the covariance matrix at time t. Using the Ito calculus and the fact that $E[Q(t)dW(t)] = 0$ for any $Q(t)$ defined by time t, we have

$$\begin{aligned} d\rho &= dE\left[X(t)X^*(t)\right] \\ &= E\left[(dX)\, X^*\right] \; + \; E\left[X\,(dX^*)\right] \; + \; E\left[(dX)\,(dX^*)\right] \\ &= E\left[-\mu X dt X^*\right] \; + \; E\left[X\,(-X^*\mu^* dt)\right] \end{aligned}$$

[2] An eigenvalue with real part equal to zero would make the system neutrally stable or "weakly stable". The randomly forced neutrally stable system does not have a steady state.

$$
\begin{aligned}
&= -\mu E\left[XX^*\right] dt - E\left[XX^*\right] \mu^* dt \;+\; E\left[\sigma dW dW^* \sigma^*\right] \\
&= \left\{ -\left(\mu\rho + \rho\mu^*\right) + \sigma\sigma^* \right\} dt \; .
\end{aligned}
$$

The matrix differential equation

$$
\dot{\rho} = -\left(\mu\rho + \rho\mu^*\right) + \sigma\sigma^* \tag{14}
$$

is an example of a class of equations called *Liapounov* equations or *Ricatti* equations, though *Ricatti* equations usually are quadratically nonlinear instead of linear as (14). The limit $\rho = \lim_{t\to\infty} \rho(t)$ is found by setting $\dot{\rho} = 0$ and solving the resulting algebraic equations.

The correct multicomponent generalization of (4) is (check this)

$$
\rho(t) = \int_0^t e^{-(t-s)\mu} \sigma\sigma^* e^{-(t-s)\mu^*} ds \; .
$$

It is easy to check that this satisfies (14) and that the integral is positive (as a matrix) for all t. Even if σ does not have full rank, $\rho(t)$ will be positive definite if the matrix with many columns

$$
K = \left(\sigma \mid \mu\sigma \mid \mu^2\sigma \mid \cdots \mid \mu^n\right)
$$

has rank n. This is a theorem of R. Kalman (proof omitted but not that difficult). The limit value is

$$
\rho = \lim_{t\to\infty} \rho(t) = \int_0^\infty e^{-t\mu} \sigma\sigma^* e^{-t\mu^*} dt \; .
$$

Techniques of numerical linear algebra give a way to calculate, and show it is unique, ρ without doing the integral. As for any square matrix, μ may be put in *upper Schur form* using an orthogonal change of coordinates. That is, there is an $n \times n$ matrix, Q with $QQ^* = I$ so that $\mu = QLQ^*$, where L is upper triangular with the eigenvalues of μ on the diagonal. Writing $\sigma\sigma^* = D$, this puts $\mu\rho + \rho\mu^* = \sigma\sigma^*$ in the form

$$
QLQ^*\rho + \rho Q^* L^* Q = D \; .
$$

Note that the elements of L and Q may be complex, particularly if the eigenvalues of μ are not real. Now multiply from the left by Q^* and from the right by Q to get

$$
L\widetilde{\rho} + \widetilde{\rho} L^* = \widetilde{D} \; ,
$$

where $\widetilde{\rho} = Q^*\rho Q$ and $\widetilde{D} = Q^* D Q$. This allows us to calculate the elements of $\widetilde{\rho}$ one by one as is done in back substitution. First, $L_{11}\widetilde{\rho}_{11} + \widetilde{\rho}_{11}\overline{L}_{11} = \widetilde{D}_{11}$ (A_{jk} is the (j,k) entry of A, $\overline{z}$ is the complex conjugate of z). So, since $\mathrm{Re}(L_{11}) > 0 > 0$ (by hypothesis on μ), $\widetilde{\rho}_{11} = \widetilde{D}_{11}/2\mathrm{Re}(L_{11}) > 0$ is determined. The rest of the elements of $\widetilde{\rho}$ can be computed one by one in a similar way. Finally, we have $\rho = Q\widetilde{\rho}Q^*$.

Now that we know the multidimensional process has a steady state, we can look for eigenfunctions that describe the rate of convergence to steady state. The probability density for $X(t)$ satisfies the forward equation

$$\partial_t u = L^* u = \frac{1}{2} \sum_{ij} D_{ij} \partial_{x_i} \partial_{x_j} u + \sum_i \partial_{x_i} \left(\sum_j \mu_{ij} x_j u \right) ,$$

again with $D = \sigma\sigma^*$. A calculation shows that if $u_0(x) = Ce^{-x^* Hx/2}$ with $H = \rho^{-1}$ (so that $\mu\rho + \rho\mu^* = D$ is equivalent to $H\mu + \mu^* H = HDH$), then $L^* u_0 = 0$. This is the steady state density, which we already knew from the covariance calculations above.

Finding the rest of the eigenfunctions is simpler if we assume that μ has n real eigenvalues and linearly independent eigenvectors. The construction is an induction as in the one dimensional case. Suppose w is an eigenvector of μ and satisfies $\mu w = \alpha w$. Suppose that $L^* u = \lambda u$. Then $v = (w \cdot \nabla) u$ satisfies $L^* v = \lambda v - \alpha v$, as we now verify. These calculations are simpler using the summation convention, which allows formulas such as $(w \cdot \nabla) u = w_k \partial_{x_k} u$. The calculation concerns only the second term

$$\begin{aligned} \partial_{x_i} (\mu_{ij} x_j w_k \partial_{x_k} u) &= w_k \partial_{x_k} \{\partial_{x_i} (\mu_{ij} x_j u)\} - \partial_{x_i} (\mu_{ij} \delta_{jk} w_k u) \\ &= w_k \partial_{x_k} \{\partial_{x_i} (\mu_{ij} x_j u)\} - \partial_{x_i} (\mu_{ij} w_j u) \\ &= w_k \partial_{x_k} \{\partial_{x_i} (\mu_{ij} x_j u)\} - \alpha w_i \partial_{x_i} u \\ &= w_k \partial_{x_k} \{\partial_{x_i} (\mu_{ij} x_j u)\} - \alpha v \end{aligned}$$

Therefore,

$$L^* v = L^* \{(w \cdot \nabla) u\} = (w \cdot \nabla) L^* u - \alpha v = \lambda v - \alpha v ,$$

as claimed. Therefore, if $\alpha_1, \ldots, \alpha_n$ are the eigenvalues of μ with corresponding eigenvectors w_k, and if $p = (p_1, \ldots, p_n)$ is a sequence of non-negative integers, then

$$u_p = (w_1 \cdot \nabla)^{p_1} \cdots (w_n \cdot \nabla)^{p_n} u_u$$

satisfies

$$L^* u_p = - (p_1 \alpha_1 + \cdots + p_n \alpha_n) u_p .$$

The rest of the structure of the one dimensional problem also generalizes. Let $d = p_1 + \cdots + p_n$ be the *degree* of u_p, then $u_p = H_p(x, \rho) u_0$, where H_p is a polynomial of degree d. These are eigenfunctions of L: $LH_p = - (p_1 \alpha_1 + \cdots + p_n \alpha_n) H_p$. They satisfy the orthogonality relation $\langle H_p, H_q \rangle_{u_0} = 0$ if $p \neq q$. Let $\mathcal{S}_d$ be the space of all polynomials of degree at most d. Then the polynomials H_p with the degree of p not more than d is a basis for $\mathcal{S}_d$. Moreover, let $\mathcal{T}_d$ be the orthogonal complement of $\mathcal{S}_{d-1}$ in $\mathcal{S}_d$ in the u_0 weighted inner product. Then the H_p with the order of p equal to d form a basis of $\mathcal{T}_d$.

3.5 Example

We consider a discrete approximation to the heat equation in notation adapted to the above discussion. Let $x(r,t)$ be the temperature at location r at time t. Suppose this satisfies

$$\partial_t x = \frac{1}{2}\partial_r^2 x \ ,$$

with boundary conditions $x(0,t) = x(1,t) = 0$. Approximate the heat equation using n equally spaced points $r_j = j\Delta r$. Let $x_j(t)$ be the approximation to $x(r_j,t)$. The approximation to the heat equation is

$$\dot{x}_j = \frac{1}{2}\frac{x_{j+1} - 2x_j + x_{j-1}}{\Delta r^2} \ .$$

This can be written as $\dot{x} = -\mu x$, where μ (up to a constant multiple) is the $(n-1)\times(n-1)$ matrix we encountered earlier as the inverse of the covariance matrix for Brownian motion constrained to have $X(0) = X(1) = 0$.

If we add noise, the dynamics become

$$dX_j = \frac{1}{2\Delta r^2}\left(X_{j+1} - 2X_j + X_{j-1}\right) + \sigma dW_j \ . \tag{15}$$

The steady state covariance satisfies (μ being symmetric) $\mu\rho + \rho\mu = \sigma I$, so $\rho = \frac{1}{2\sigma}\mu^{-1}$. We have seen that the eigenvalues of μ are

$$\alpha_k = \frac{1 - \cos(k\pi\Delta r)}{\Delta r^2} \ .$$

For small k, these are approximately $\alpha_k = \pi^2 k^2/2$. Therefore, the smallest eigenvalues of L^* for (15) are $\lambda = -\alpha_1 \approx -\pi^2/2$, $\lambda = -\alpha_2 \approx -2\pi^2$, $\lambda = -2\alpha_1 \approx -\pi^2$, $\lambda = -\alpha_1 - \alpha_2 \approx \frac{5}{2}\pi^2$, etc. Note that the slowest mode is a linear mode, but the second slowest is quadratic and the third slowest is cubic ($-3\alpha_3$).

The fastest linear mode is $-\alpha_n \approx 2/\Delta r^2$. The matrix μ is *ill conditioned* for large n, which means that it has a wide range of eigenvalues: $\alpha_1/\alpha_n = O(n^2)$. This causes significant difficulties for Markov Chain Monte Carlo algorithms that use dynamics like (15).

4 Stochastic Differential Equations

A stochastic differential equation takes the form

$$dX = a(X(t),t)dt + b(X(t),t)dW(t) \ . \tag{16}$$

The solution, $X(t)$, is a continuous but random function of t. It in the multi-component case, $X(t) = (X_1(t),\ldots,X_n(t))$. In this case the *drift* "coefficient", $a(x,t)$, has n components and the *noise* coefficient, $b(x,t)$, is an $n\times m$ matrix with $m \le n$. The source of noise is m independent Brownian motion components $W(t) = (W_1(t),\ldots,W_m(t))$.

The interpretation of (16) is subtle in ways we will not explore completely here. However, one way to define the solution is as the limit of forward *Euler* approximations. Choose a time step h and times $t_k = kh$. The approximation solution, $X^h(t)$, is defined by

$$X^h(t_{n+1}) = X^h(t_n) + a(X^h(t_n), t_n)h + b(X^h(t_n), t)\,(W(t_{n+1}) - W(t_n)) \ . \quad (17)$$

According to this formula, to generate a forward Euler path, we must generate a Brownian motion path, then use (17) to generate the corresponding approximate sample path of (16) from it. More precisely, for fixed W and T, we imagine a sequence of h values converging to zero and define

$$X(W[0,T],T) \ = \ \lim_{h\to 0} X^h(W[0,T],T) \ .$$

Writing $f(W[0,T])$ means that f depends on the values of $W(t)$ for all t in the range $0 \leq t \leq T$. It is convenient for theoretical discussion to consider X^h to be defined for all t, not just the discrete times t_n. We do this using piecewise linear interpolation.

4.1 Strong accuracy

The *strong* accuracy of the Euler method (17) measures the expected difference between X^h and X. Here and everywhere below we do not indicate the dependence of X or X^h on W, though both depend on W. Strong accuracy of order p means that

$$E\left[\ \left|X^h(T) - X(T)\right|\ \right] \ = \ O\left(h^p\right) \ . \quad (18)$$

The Euler method (17) satisfies this with $p = 1/2$. We will see in examples that this is sharp in the sense that the expected error is of the order of $h^{1/2}$ and not smaller. The relation (18) is an L^1 estimate, which is hard to approach by direct computation. The standard way to approach it is with the L^2 estimate

$$E\left[\ \left|X^h(T) - X(T)\right|^2\ \right] \ = \ O\left(h^{2p}\right) \ . \quad (19)$$

This implies the L^1 estimate because of the Cauchy Schwarz inequality: If Y and Z are any random variables, then

$$E\left[\,|YZ|\,\right] \ \leq \ \left(\,E\left[Y^2\right]E\left[Z^2\right]\,\right)^{1/2} \ .$$

In particular, if $Z = 1$, then $E\left[\,|Y|\,\right] \ \leq \ E\left[Y^2\right]^{1/2}$, which is how (18) follows from (19). Some people prefer to cite *Jensen's inequality*, which states that if $\phi(y)$ is a a convex function of y, then

$$\phi\left(\,E\left[Y\right]\,\right) \ \leq \ E\left[\,\phi(Y)\,\right] \ .$$

Taking $\phi(y) = y^2$ then gives (18) from (19). Finally, versions of *Doob's inequality* from the theory of martingales allows us to strengthen (19) to a statement about the path up to time T:

$$E\left[\sup_{0\leq t\leq T}\left|X^h(t) - X(t)\right|^2\right] \ = \ O\left(h^{2p}\right) \ . \quad (20)$$

This is another advantage of L^2. Doob's inequality takes a weaker and more complicated form in L^1. Of course, (20) implies he L^1 bound

$$E\left[\sup_{0\leq t\leq T}\left|X^h(t)-X(t)\right|\right] \;=\; O\left(h^p\right)\;.$$

The proof of the L^2 bound has two main ingredients, *consistency* and *stability*. The consistency step is a bound for the error created in a single time step. The SDE (16) implies that[3]

$$X(t+h)-X(t) \;=\; \int_0^h a(X(t+s))ds \;+\; \int_0^h b(X(t+s))dW(s)\;. \tag{21}$$

We write this symbolically as

$$X(t+h) \;=\; X(t)+\Psi(X(t),W[t_n,t_{n+1}],h)\;.$$

In general, $x+\Psi(x,W[0,T],T)$ is the solution at time T to the SDE (16) with $X(0)=x$. If the coefficients do not depend on t, the starting time if irrelevant. For small h, we assume that $X(t+s)\approx X(t)$. Using this in (21) gives

$$X(t+h)-X(t) \;\approx\; a(X(t))h \;+\; b(X(t))\,(\,W(t+h)-W(t)\,)\;. \tag{22}$$

This is the basis of the Euler approximation (17). We write right hand side of (22) as

$$\Phi(x,W[0,T],T) \;=\; a(x)T \;+\; b(x)\big(W(T)-W(0)\big)\;,$$

so the Euler method takes the form

$$X^h(t_{n+1}) \;=\; X^h(t_n) \;+\; \Phi(X^h(t_n),W[t_n,t_{n+1}],h)\;.$$

The consistency bound controls the difference between Ψ and Φ. This depends on the difference between $X(t+s)$ and $X(t)$ for $s\leq h$. We assume that

$$|a(x)| \;\leq\; M\quad,\quad |b(x)| \;\leq\; M$$

for all x, and prove that

$$E\left[\,|X(t+s)-X(t)|^2\right] \;\leq\; Cs\;.$$

This is the same as showing that

$$E\left[\,|\Psi(x,W[0,T],T)|^2\right] \;\leq\; CT\;. \tag{23}$$

[3]We assume for notational simplicity that a and b do not depend on t. It is straightforward but more time consuming to treat the more general case.

Now, Ψ is the sum of two terms, which we bound separately using the fact that $(a+b)^2 \le 2a^2 + 2b^2$. The first term is handled using Cauchy Schwarz and the bound $|a| \le M$:

$$\begin{aligned} \left| \int_0^T a(X(t))\, dt \right|^2 &= \left| \int_0^T a(X(t)) \cdot 1\, dt \right|^2 \\ &\le \int_0^T |a(X(t))|^2\, dt \cdot \int_0^T 1\, dt \\ &\le M^2 T \cdot T \;=\; CT^2 \,, \end{aligned}$$

which is stronger than the needed bound for small T. Note that this bound is true for every path, and therefore for the expected value.

The second term satisfies

$$E\left[\left| \int_0^T b(X(t))\, dW(t) \right|^2 \right] \;\le\; CT \,. \tag{24}$$

This is a consequence of the *Ito isometry formula*, which states that if $f(t)$ is any adapted function of t, which we would write as $f(t) = g(W[0,t], t)$, then

$$E\left[\left| \int_0^T f(t)\, dW(t) \right|^2 \right] \;=\; \int_0^T E\left[|f(t)|^2 \right] dt \,. \tag{25}$$

This, together with $|b(x)| \le M$ gives (24).

We pause to indicate the proof of the Ito isometry formula (25) because it reveals the way lack of correlation in integrals with non-anticipating functions leads to L^2 bounds and the general utility of integrating squares to exploit cancellation due to lack of correlation. The Ito integral on the left of (25) is approximated by

$$\sum_{j=0}^{n-1} R_j \;\;, \quad \text{with} \quad R_j = f(t_j)\left(W(t_{j+1}) - W(t_j)\right) \,.$$

The main observation is that if $j \ne k$, then $E[R_j R_k] = 0$. Suppose, for example, that $k > j$. We have assumed that f is non anticipating, which implies that both $f(t_j)$ and $f(t_k)$ (and $W(t_{j+1}) - W(t_j)$) are determined by the values $W[0, t_k]$. But the independent increments property of Brownian motion then implies that $W(t_{k+1}) - W(t_k)$ is independent of all these values, so

$$E[R_j R_k] = E[(\text{det'd at time } t_k)(\text{indep of time } t_k)] = E[\cdots] \cdot 0 \,.$$

Now, the left side of (25) is approximately

$$E\left[\left(\sum_{j=0}^{n-1} R_j \right)^2 \right] \;=\; \sum_{j=0}^{n-1} \sum_{k=0}^{n-1} E\left[R_j R_k \right] \;=\; \sum_{j=0}^{n-1} E\left[R_j^2 \right] \,.$$

Moreover, again using the independent increments property,

$$E\left[R_j^2\right] \;=\; E\left[f(t_j)^2\right] E\left[\Delta W_j^2\right] \;=\; E\left[f(t_j)^2\right] \Delta t \;,$$

so $\sum_{j=0}^{n-1} E\left[R_j^2\right]$ is the Riemann sum approximation to the right side of (25). Taking $h \to 0$ finishes the proof.

The consistency estimate is

$$E\left[\, \left| \Psi(x, W[0,T], T) - \Phi(x, W[0,T], T) \right|^2 \right] \;\leq\; CT^2 \quad . \tag{26}$$

We give a proof under the hypothesis that the drift and noise coefficients are globally Lipschitz:

$$|a(x) - a(y)| \;\leq\; C\,|x-y| \quad , \quad |b(x) - b(y)| \;\leq\; C\,|x-y| \;. \tag{27}$$

Both Ψ and Φ are sums of a drift and a noise term. The bound (26) follows (as above) from bounds on these terms separately.

The drift term is (see above)

$$\begin{aligned} R_1 &= E\left[\left| \int_0^T \big(a(X(t)) - a(X(0)) \big)\, dt \right|^2 \right] \\ &\leq T \int_0^T \left| a(X(t)) - a(X(0)) \right|^2 dt \;. \end{aligned}$$

With this, and (23) and (27), we find

$$R_1 \leq hC \int_0^T E\left[\left(X(t) - X(0) \right)^2 \, dt \right] \;\leq\; CT \int_0^T t\, dt = CT^3 \;.$$

This is more, by one power of T, than is needed to bound the drift contribution to (26). Indeed, the largest errors are due to the noise.

The noise term may be handled using the Ito isometry formula, with (27) and (23):

$$\begin{aligned} R_2 &= E\left[\left| \int_0^T \left(b(X(t)) - b(X(0)) \right) dW(t) \right|^2 \right] \\ &= \int_0^T E\left[\, \left| \, b(X(t)) - b(X(0)) \, \right|^2 \right] \, dt \\ &\leq C \int_0^T E\left[\, \left| \, X(t) - X(0) \, \right|^2 \right] \, dt \\ &\leq C \int_0^T t\, dt = CT^2 \;. \end{aligned}$$

Convergence proofs of computational methods often take the form of *consistency* and *stability* arguments. Roughly speaking, consistency is the fact that

the computational method is consistent with the formulation of the limit problem. In the present case, the estimate (26) states that at least for one time step, the Euler method is consistent with the SDE. More technically, we might call $\Psi - \Phi$ *truncation error*. It is the source of the deviation between the paths $X(t)$ and $X^h(t)$. Stability is the statement that small deviations between X and X^h caused by truncation error are not amplified too much during the simulation process. Stability usually is the harder part, though here it is more complicated than hard. We do the stability argument only for the case with zero drift. Even though drift is the smaller source of error it is the larger source of technical complexity of a not very interesting kind.

We get a recurrence relation for the error by subtracting the X^h recurrence relation from that for X, and then subtracting and adding $\Phi(X(t_n), W[t_n, t_{n+1}], h)$:

$$\begin{aligned} X(t_{n+1}) - X^h(t_{n+1}) \;&=\; X(t_n) - X^h(t_n) \\ &+\quad \Psi(X(t_n), W[t_n, t_{n+1}], h) - \Phi(X(t_n), W[t_n, t_{n+1}], h) \\ &+\quad \Phi(X(t_n), W[t_n, t_{n+1}], h) - \Phi(X^h(t_n), W[t_n, t_{n+1}], h) \; . \end{aligned} \tag{28}$$

In the absence of drift,

$$E\left[\, \Psi(x, W[t_n, t_{n+1}], h) \;\mid\; \mathcal{F}_{t_n}\right] = 0 \; ,$$

and the same for the other three terms on the right side. Therefore we may calculate, again using $(a+b)^2 \le 2a^2 + 2b^2$,

$$E\left[\, \left| X(t_{n+1}) - X^h(t_{n+1}) \right|^2 \right] \;\le\; E\left[\, \left| X(t_n) - X^h(t_n) \right|^2 \right] \;+\; 2u_n + 2v_n \; ,$$

where

$$u_n = E\left[\, \left| \Psi(X(t_n), W[t_n, t_{n+1}], h) - \Phi(X(t_n), W[t_n, t_{n+1}], h) \right|^2 \right] \; ,$$

and

$$v_n = E\left[\, \left| \Phi(X(t_n), W[t_n, t_{n+1}], h) - \Phi(X^h(t_n), W[t_n, t_{n+1}], h) \right|^2 \right] \; .$$

The consistency bound (26) says that $u_n \le Ch^2$. For v_n, write $W(t_{n+1}) - W(t_n) = \Delta W_n$, so that

$$\begin{aligned} &\Phi(X(t_n), W[t_n, t_{n+1}], h) - \Phi(X^h(t_n), W[t_n, t_{n+1}], h) \\ &\qquad\qquad = \; \big(b(X(t_n)) - b(X^h(t_n)) \big) \Delta W \; . \end{aligned}$$

The crucial fact now is that ΔW is independent of everything up to time t_n, so the expected square is (again using (27))

$$\begin{aligned} v_n \;&\le\; E\left[\, \left| b(X(t_n)) - b(X^h(t_n)) \right|^2 \right] \cdot E\left[\, \left| \Delta W \right|^2 \right] \\ &\le\; ChE\left[\, \left| X(t_n) - X^h(t_n) \right|^2 \right] \; . \end{aligned}$$

This finally is enough estimates. Defining $e_n = E\left[\left|X(t_n) - X^h(t_n)\right|^2\right]$, we have shown that

$$e_{n+1} \leq e_n + Ch^2 + Che_n \ .$$

Of course, we start with $e_0 = 0$, so we get (using $t_n = nh$)

$$e_n \leq Ce^{Ct_n} nh^2 = Ce^{Ct_n} t_n h \ .$$

This implies that if we fix T and let n go to infinity while $h \to 0$ with $t_n \leq T$ always, then there is a $C(T)$ so that (19) is satisfied.

Milstein's method is a modification of the Euler method that improves the strong order of accuracy from $p = 1/2$ to $p = 1$. We explain the idea in the simplest case of a one component diffusion with zero drift. We want a $\Phi(x, W[0,T], T)$ that more closely resembles $\Psi(x, W[0,T], T)$ for small T. In the case of zero drift,

$$\Psi(x, W[0,T], T) \ = \ \int_0^T b(X(t))dW(t) \ ,$$

so we want a more accurate approximation to this integral than we get from taking $b(X(t)) \approx b(x)$. Since $X(t)$ is close to x, we expect $B(X(t))$ to be close to $b(x)$ and we make a Taylor approximation

$$b(X(t)) \ \approx \ b(x) \ + \ b'(x)\big(\, X(t) - x \,\big) \ .$$

The second term on the right is a small correction to the first term. Therefore, we may not need as accurate an approximation to it to improve the overall accuracy. Milstein used the Euler approximation

$$X(t) \ \approx \ x + b(x)\big(\, W(t) - W(0) \,\big) \ .$$

These two approximations lead to

$$\begin{aligned} \Psi(x, W[0,T], T) \ &\approx \ \Psi_M(x, W[0,T], T) \\ &\int_0^T \big(\, b(x) \ + \ b'(x)b(x)\big(\, W(t) - W(0) \,\big)\,\big)\, dW(t) \ , \end{aligned} \tag{29}$$

For one dimensional problems only, one can calculate the integral on the right to get the Milstein method

$$X^h(t_{n+1}) = X^h(t_n) + b(X^h(t_n))\Delta W_n + \frac{1}{2}b'(X^h(t_n))b(X^h(t_n))\Big(\Delta W_n^2 - \Delta t\Big) \ . \tag{30}$$

The accuracy of this method in the strong sense is determined by

$$\begin{aligned} &E\left[\left|\Psi(x, W[0,T], t) - \Phi_M(x, W[0,T], t)\right|^2\right] \\ &= \ E\left[\ \left|\ \int_0^T \Big(\, b(X(t)) \ - \ \big[\, b(x) + b'(x)b(x)W(t) \,\big]\Big)\, dW(t) \ \right|^2 \ \right] \\ &= \ \int_0^T E\left[\ \Big|\, b(X(t)) \ - \ \big[\, b(x) + b'(x)b(x)W(t) \,\big]\,\Big|^2 \ \right] dt \end{aligned}$$

The error in the last line may be decomposed into the sum of the error in the Taylor expansion of b and the error in the Euler approximation.

$$\begin{aligned} b(X(t)) \;-\; [\, b(x) + b'(x)b(x)W(t)\,] \\ \quad = \quad b(X(t)) \;-\; [\, b(x) + b'(x)\,(X(t)-x)\,] \\ \quad + \quad [\, b(x) + b'(x)\,(X(t)-x)\,] \;-\; [\, b(x) + b'(x)b(x)W(t)\,] \\ \quad = \quad b(X(t)) \;-\; [\, b(x) + b'(x)\,(X(t)-x)\,] \\ \quad + \quad b'(x)\,[\,(X(t)-x) \;-\; b(x)W(t)\,] \;. \end{aligned}$$

The expected square of the second term is bounded by the consistency estimate for the Euler method. The expected square of the first term depends on $M_4(t) = E\left[|X(t)-x|^4\right]$. We can bound this by computing $\partial_t M_4(t)$ and using Ito calculus. The result is

$$\begin{aligned} E\left[|\Psi(x,W[0,T],t) - \Phi_M(x,W[0,T],t)|^2\right] \\ \quad \leq \quad \int_0^T \left(Ct^2 + Ct^2\right) dt \\ \quad = \quad CT^3 \;. \end{aligned}$$

This is one power of T smaller than (26), the corresponding estimate for the Euler method.