spl20.tex

Lecture 20. 22.11.2010

4. Other classes of process

Gaussian Processes

A vector $\mathbf{X} \in \mathbf{R}^n$ has the *multivariate normal distribution* in n dimensions if all linear combinations $\mathbf{a}'\mathbf{X} = \sum_{i=1}^n a_i X_i$ of its components are normally distributed (in one dimension). Such a distribution is determined by a vector μ of means and a non-negative definite $n \times n$ matrix $\mathbf{\Sigma}$ of covariances, and is written $N(\mu, \mathbf{\Sigma})$. Then $\mathbf{X}$ has distribution $N(\mu, \mathbf{\Sigma})$ if and only if it has characteristic function

$$\phi_{\mathbf{X}}(\mathbf{t}) := E[\exp\{i\mathbf{t}' \cdot \mathbf{X}\}] = \exp\{i\mathbf{t}' \cdot \mu - \frac{1}{2}\mathbf{t}'\mathbf{\Sigma}\mathbf{t}\} \qquad (\mathbf{t} \in \mathbf{R}^n).$$

Further, if $\mathbf{\Sigma}$ is positive definite (so non-singular), $\mathbf{X}$ has density

$$f_{\mathbf{X}}(\mathbf{x}) = \frac{1}{(2\pi)^{n/2}|\mathbf{\Sigma}|^{1/2}} \exp\left\{-\frac{1}{2}(\mathbf{x}-\mu)'\mathbf{\Sigma}^{-1}(\mathbf{x}-\mu)\right\},$$

(Edgeworth's formula: F.Y. EDGEWORTH (1845-1926) in 1892). A process $X = (X(t))_{t \geq 0}$ is *Gaussian* if all its finite-dimensional distributions are Gaussian. Such a process can be specified by:
(i) a measurable function $\mu = \mu(t)$ with $E(X(t)) = \mu(t)$, the *mean function*;
(ii) a non-negative definite function $\sigma(s,t)$ with $\sigma(s,t) = cov(X(s), X(t))$, the *covariance function*.

Gaussian processes have many interesting properties. Among these, we quote *Belyaev's dichotomy*: with probability one, the paths of a Gaussian process are either continuous, or extremely pathological: for example, unbounded above and below on any time interval, however short. Naturally, we shall confine attention in this book to continuous Gaussian processes.

Markov Processes

X is *Markov* if for each t, each $A \in \sigma(X(s) : s > t)$ (the 'future') and $B \in \sigma(X(s) : s < t)$ (the 'past'),

$$P(A|X(t), B) = P(A|X(t)).$$

That is, if you know where you are (at time t), how you got there doesn't matter so far as predicting the future is concerned – equivalently, past and future are conditionally independent given the present. X is said to be *strong Markov* if the above holds with the fixed time t replaced by a stopping time τ (a random variable). This is a real restriction of the Markov property in the continuous-time case (though not in discrete time). Perhaps the simplest example of a Markov process that is not strong Markov is given by

$$X(t) := 0 \quad (t \leq \tau), \quad t - \tau \quad (t \geq \tau),$$

where τ is an exponentially distributed random variable. Then X is Markov (from the lack of memory property of the exponential distribution), but not strong Markov (the Markov property fails at the stopping time τ). One must expect the strong Markov property to fail in cases, as here, when 'all the action is at random times'. Another example of a Markov but not strong Markov process is a left-continuous Poisson process – obtained by taking a Poisson process (see below) and modifying its paths to be left-continuous rather than right-continuous.

Diffusions

A *diffusion* is a path-continuous strong Markov process such that for each time t and state x the following limits exist:

$$\mu(t,x) := \lim_{h\downarrow 0} \frac{1}{h} E[(X(t+h) - X(t))|X(t) = x],$$

$$\sigma^2(t,x) := \lim_{h\downarrow 0} \frac{1}{h} E[(X(t+h) - X(t))^2|X(t) = x].$$

Then $\mu(t,x)$ is called the *drift*, $\sigma^2(t,x)$ the *diffusion coefficient*.

The term diffusion derives from physical situations involving Brownian motion (below). The mathematics of heat diffusing through a conducting medium (which goes back to Fourier in the early 19th century) is intimately linked with Brownian motion (the mathematics of which is 20th century).

The theory of diffusions can be split according to dimension. For one-dimensional diffusions, there are a number of ways of treating the theory. For higher-dimensional diffusions, there is basically one way: via the stochastic differential equation methodology (or its reformulation in terms of a martingale problem). This shows the best way to treat the one-dimensional case: the best method is the one that generalizes. It also shows that Markov processes and martingales, as well as being the two general classes of stochastic process with which one can get anywhere mathematically, are also intimately linked technically. We will encounter diffusions largely as solutions of stochastic differential equations.

5. Brownian motion

Brownian motion originates in work of the botanist Robert Brown in 1828. It was introduced into finance by Louis Bachelier in 1900, and developed in physics by Albert Einstein in 1905 (see the handout for background and references).

The fact that Brownian motion *exists* is quite deep, and was first proved by Norbert WIENER (1894–1964) in 1923. In honour of this, Brownian motion is also known as the *Wiener process*, and the probability measure generating it – the measure P^* on $C[0,1]$ (one can extend to $C[0,\infty)$) by

$$P^*(A) = P(W_. \in A) = P(\{t \to W_t(\omega)\} \in A)$$

for all Borel sets $A \in C[0,1]$ – is called *Wiener measure.*

Definition and Existence

Definition. A stochastic process $X = (X(t))_{t \geq 0}$ is a standard (one-dimensional) Brownian motion, BM or $BM(\mathbf{R})$, on some probability space $(\Omega, \mathcal{F}, \mathcal{P})$, if
(i) $X(0) = 0$ a.s.,
(ii) X has *independent increments*: $X(t+u) - X(t)$ is independent of $\sigma(X(s) : s \leq t)$ for $u \geq 0$,
(iii) X has *stationary increments*: the law of $X(t+u) - X(t)$ depends only on u,
(iv) X has *Gaussian increments*: $X(t+u) - X(t)$ is normally distributed with mean 0 and variance u, $X(t+u) - X(t) \sim N(0,u)$,
(v) X has *continuous paths*: $X(t)$ is a continuous function of t, i.e. $t \to X(t,\omega)$ is continuous in t for all $\omega \in \Omega$.

The path continuity in (v) can be relaxed by assuming it only a.s.; we can then get continuity by excluding a suitable null-set from our probability space.

We shall henceforth denote standard Brownian motion $BM(\mathbf{R})$ by $W = (W(t))$ (W for Wiener), though $B = (B(t))$ (B for Brown) is also common. Standard Brownian motion $BM(\mathbf{R}^d)$ in d dimensions is defined by

$W(t) := (W_1(t), \ldots, W_d(t))$, where $W_1, \ldots, W_d$ are independent standard Brownian motions in one dimension (independent copies of $BM(\mathbf{R})$).

We turn next to Wiener's theorem, on existence of Brownian motion. The proof (cf. [BK], 5.3.1) is a streamlined version of the classical one due to Lévy in his book of 1948 and Cieselski in 1961 (see below for references).

Theorem (Wiener, 1923). Brownian motion exists.

Covariance. Before addressing existence, we first find the covariance function. For $s \le t$, $W_t = W_s + (W_t - W_s)$, so as $E(W_t) = 0$,

$$cov(W_s, W_t) = E(W_s W_t) = E(W_s^2) + E[W_s(W_t - W_s)].$$

The last term is $E(W_s)E(W_t - W_s)$ by independent increments, and this is zero, so

$$cov(W_s, W_t) = E(W_s^2) = s \qquad (s \le t): \qquad cov(W_s, W_t) = \min(s, t).$$

A Gaussian process (one whose finite-dimensional distributions are Gaussian) is specified by its mean function and its covariance function, so among centered (zero-mean) Gaussian processes, the covariance function $\min(s, t)$ serves as the signature of Brownian motion.

Finite-dimensional Distributions. For $0 \le t_1 < \ldots < t_n$, the joint law of $X(t_1), X(t_2), \ldots, X(t_n)$ can be obtained from that of $X(t_1), X(t_2) - X(t_1), \ldots, X(t_n) - X(t_{n-1})$. These are jointly Gaussian, hence so are $X(t_1), \ldots, X(t_n)$: the finite-dimensional distributions are *multivariate normal*. Recall that the multivariate normal law in n dimensions, $N_n(\mu, \Sigma)$ is specified by the mean vector μ and the covariance matrix Σ (non-negative definite). So to check the finite-dimensional distributions of BM – stationary independent increments with $W_t \sim N(0, t)$ – it suffices to show that they are multivariate normal with mean zero and covariance $cov(W_s, W_t) = \min(s, t)$ as above.