Stochastic Optimisation

Problem Sheet 3

---

1. Let U be a random variable distributed uniformly on $[0,1]$, and let $\lambda > 0$ be a given constant. Show that the random variable $X = (-\log U)/\lambda$ has an exponential distribution with parameter λ. As usual, logarithms are natural, i.e., to base e.

2. Let X be a random variable with a Beta(α, β) distribution, and let $Y = 1 - X$. Show that Y has a Beta(β, α) distribution.

   Hint. You have a choice of two different ways to show this, both of which are pretty straightforward.

3. Let X and Θ be independent random variables, where X has an Exp($\frac{1}{2}$) distribution, and Θ is uniformly distributed on $[0, 2\pi)$. Let $V = \sqrt{X}\sin\Theta$ and $W = \sqrt{X}\cos\Theta$. Show that V and W are independent standard normal random variables, i.e., with zero mean and unit variance.

   Hint. Compute the joint density of V and W and show that it factorises as the product of the densities of two independent standard normal random variables.

4. Consider a Bayesian approach to the problem of inferring the mean of a normal distribution with unknown mean and known variance. More precisely, suppose $X \sim N(\theta, 1)$, where θ is unknown. Fix $\mu_0 \in \mathbb{R}$ and $\sigma_0^2 > 0$, and let π_0 denote the density of a $N(\mu_0, \sigma_0^2)$ distribution. Suppose π_0 denotes the prior distribution of θ. Let $\pi_1(\cdot|x)$ denote the posterior distribution, conditional on observing a sample of X which takes the value x. In other words,

   $$\pi_1(\theta|x) \propto \pi_0(\theta) f_\theta(x),$$

   where $\propto$ stands for 'proportional to' and f_θ denotes the density of a $N(\theta, 1)$ random variable. The constant of proportionality will be determined by the requirement that $\pi_1(\cdot|x)$ be a probability density, i.e., that it integrate to 1.

   Show that $\pi_1(\cdot|x)$ is the density of a $N(\mu_1, \sigma_1^2)$ random variable, where

   $$\mu_1 = \frac{\mu_0 + x\sigma_0^2}{1 + \sigma_0^2}, \quad \sigma_1^2 = \frac{\sigma_0^2}{1 + \sigma_0^2}.$$

   Remark. This shows that the family of normal distributions is a family of conjugate priors for the mean of a normal random variable with known variance.

   Hint. Unfortunately, this is a calculation-heavy exercise. You can lighten the burden somewhat by ignoring constants wherever possible.

5. Use the answer to the last question to formulate a version of the Thompson sampling algorithm for a two-armed bandit, where the rewards from arm i are iid, normally distributed, with unknown mean μ_i and unit variance. Explain your algorithm in sufficient detail to enable a non-expert to implement it. You may assume that the non-expert has access to a software package that will generate independent random variables with specified parameters, from any of a family of commonly used probability distributions.