UCLA
Department of Engineering
Christoph Lippert

Los Angeles, January 8, 2014

Current Topics in Computational Biology

Winter 2014

Excercise 1

Hand-in date: Monday, January 13, 2014 at 2:00 pm

Task 1-1 Bayesian inference for the normal distribution

(a) You are given N independent observations of a normal distribution with known mean μ and but unknown variance. Derive the posterior distribution of the precision $\tau = 1/\sigma^2$ under a gamma distribution after N observations.

The gamma distribution is specified as

$$Gamma(\tau \mid \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \tau^{\alpha-1} e^{-\beta\tau},$$

where $\Gamma(a)$ is the gamma function of a.

(b) You are given N independent observations of a normal distribution with unknown mean and variance. The joint prior distribution $p(\mu, \tau)$ over mean μ and precision τ where the marginal distribution over the precision is a gamma distribution as above and the conditional distribution of the mean μ conditioned on the precision τ is a normal distribution as specified below. Derive the joint posterior distribution $p(\mu, \tau \,|\, \mathcal{D})$.

Joint normal-gamma prior distribution over μ and τ:

$$p(\mu, \tau) = \underbrace{\mathcal{N}\left(\mu \,|\, m_0\,,\, s_0^2/\tau\right)}_{p(\mu|\tau)} \cdot \underbrace{Gamma(\tau \mid \alpha, \beta)}_{p(\tau)}$$

Task 1-2 Bayesian inference for the binomial distribution

A coin is flipped N times and yielded k times heads. The conjugate prior for the probability of success θ in a binomial distribution $B(k \mid N, \theta)$, where N is the number of trials and k is the number of successes, is the beta distribution.

Derive the posterior distribution of the probability of success θ.

$$Binomial(k \mid N, \theta) = \left(\begin{array}{c} N \\ k \end{array}\right) \theta^k \left(1-\theta\right)^{N-k}.$$

$$Beta(\theta \mid \alpha, \beta) = \frac{\theta^{\alpha-1}(1-\theta)^{\beta-1}}{B(\alpha, \beta)},$$

where $B(a, b)$ denotes the Beta function of a and b.