Multivariate Normal Exercise

February 25, 2014

In this exercise, we will show, for a simple case, that the conditional and marginal distributions of a multivariate normal distribution are themselves normal. We will consider the simple case when $\mu = 0$ and when we wish to condition or marginalize solely over the final dimension.

The inverse of the covariance matrix Σ is called the precision matrix, which we here give the name D:

$$D \equiv \Sigma^{-1}$$

First, we break the variable x and the precision matrix D into their block representations, where we split off only the last dimension:

$$x = \begin{bmatrix} x_a \\ \chi_b \end{bmatrix} \text{ and } D = \begin{bmatrix} D_{aa} & d \\ d^T & \delta \end{bmatrix}$$

Note that we are using the convention that capital letters are matrices, lower case latin letters are vectors, and lower case greek letters are scalars.

Conditioning

1. Show that
$$x^T \Sigma^{-1} x = x_a^T D_{aa} x_a - 2\chi_b x_a^T d + \delta \chi_b^2.$$

2. To condition over the final dimension, we simply hold χ_b constant. (Why?) With χ_b held constant, find constants μ and K such that
$$x^T \Sigma^{-1} x = (x_a - \mu)^T D_{aa} (x_a - \mu) + K.$$

3. Conclude that a multivariate gaussian with mean zero conditioned on the last dimension is itself an unnormalized multivariate gaussian with potentially non-zero mean.

Marginalizing

1. To marginalize out the final dimension, we integrate over χ_b:
$$p(x_a) = \int_{-\infty}^{\infty} p(x)\, d\chi_b.$$
Perform this integral and conclude that the marginalization of a multivariate gaussian is itself a multivariate gaussian with the same mean.

2. Argue that the marginalized gaussian is already normalized.