1.9 : The Bivariate Normal Distribution

Let X and Y be RVs with the following joint PDF

$$p(x,y) \quad = \quad \frac{1}{2\pi\sigma_{\rm x}\sigma_{\rm y}\sqrt{1-\rho^2}} \exp\left[-\frac{1}{2(1-\rho^2)}Q(x,y)\right]$$

where the quadratic form, $Q(x,y)$ is given by

$$Q(x,y) \quad = \quad (\frac{x-\mu_{\rm x}}{\sigma_{\rm x}})^2 - 2\rho(\frac{x-\mu_{\rm x}}{\sigma_{\rm x}})(\frac{y-\mu_{\rm y}}{\sigma_{\rm y}}) + (\frac{y-\mu_{\rm y}}{\sigma_{\rm y}})^2$$

Then $p(x,y)$ is known as the **bivariate normal PDF** and is specified by the 5 parameters $\mu_{\rm x}$, $\mu_{\rm y}$, $\sigma_{\rm x}$, $\sigma_{\rm y}$ and ρ. This PDF is used often in astronomy to model the joint PDF of two random variables. Figure 8, for example shows the joint distribution of apparent magnitude and the logarithm of the 21cm line width (denoted by P), which is often modelled as a bivariate normal PDF in statistical studies of the Tully-Fisher distance relation for spiral galaxies.

The first 4 parameters of the bivariate normal PDF are, in fact, equal to the following expectation values:-

1. $E(X) = \mu_{\rm x}$
2. $E(Y) = \mu_{\rm y}$
3. $\text{var}(X) = \sigma_{\rm x}^2$
4. $\text{var}(Y) = \sigma_{\rm y}^2$

The parameter ρ is known as the **correlation coefficient** and satisfies

$$E[(X-\mu_{\rm x})(Y-\mu_{\rm y})] \quad = \quad \rho\sigma_{\rm x}\sigma_{\rm y}$$

Note that if $\rho = 0$ then X and Y are statistically independent.

$E[(X-\mu_{\rm x})(Y-\mu_{\rm y})]$ is known as the **covariance** of X and Y and is often denoted by $\text{cov}(X,Y)$.

The marginal PDFs of X and Y are just the univariate normal PDFs, i.e.

$$p_x(x) = N(\mu_{\rm x}, \sigma_{\rm x}) \qquad p_y(y) = N(\mu_{\rm y}, \sigma_{\rm y})$$

The conditional PDF of Y given x is also a univariate normal PDF, viz:-

$$p(y|x) \quad = \quad N(\mu_{\rm y} + \frac{\sigma_{\rm y}}{\sigma_{\rm x}}\rho(x - \mu_{\rm x}), \sigma_{\rm y}\sqrt{1-\rho^2})$$

with the corresponding expression for $p(x|y)$.

$\mu_{\rm y} + \frac{\sigma_{\rm y}}{\sigma_{\rm x}}\rho(x - \mu_{\rm x})$ is often referred to as the **conditional expectation** (value) of Y given x, and the equation

$$y = \mu_{\rm y} + \frac{\sigma_{\rm y}}{\sigma_{\rm x}}\rho(x - \mu_{\rm x})$$

is called the **regression line** of Y on X. We will say more about regression in Section 2.

Figure 8: The bivariate normal PDF

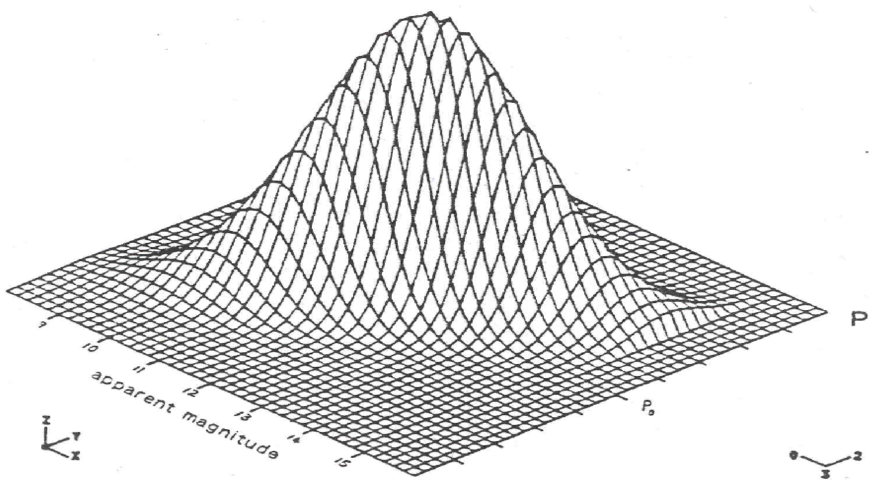

The bivariate normal PDF (and indeed any bivariate distribution function) can be represented as a **plot of isoprobability contours**. These contour curves, closely analogous to the contours on an OS map, denote those points on the (x, y) plane where the function $p(x, y)$ is constant. Examples of isoprobability contours for bivariate normal pdf's with different values of ρ, and the corresponding regression lines of Y on X for these pdf's, are shown on the handout provided on the website.