§31. Multivariate Gaussian

The normal pdf is

$$p(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

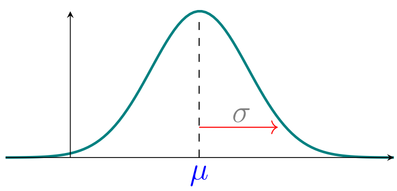

The bivariate normal pdf is

$$p(x,y) = \frac{1}{2\pi\sigma_x\sigma_y\sqrt{1-\rho^2}} \exp\left\{-\frac{1}{2\left(1-\rho^2\right)}\left[\left(\frac{x-\mu_x}{\sigma_x}\right)^2 - 2\rho\left(\frac{x-\mu_x}{\sigma_x}\right)\left(\frac{y-\mu_y}{\sigma_y}\right) + \left(\frac{y-\mu_y}{\sigma_y}\right)^2\right]\right\} \quad (*)$$

This is a 5–parametric pdf :

- means μ_x, μ_y
- variances σ_x^2, σ_y^2
- correlation coefficient ρ

As expected the marginals are: $X \sim N\left(\mu_x, \sigma_x\right); Y \sim N\left(\mu_Y, \sigma_Y\right)$.

Let's consider the centred, standardized bivariate normal with
$\mu_x = \mu_y = 0 \quad ; \quad \sigma_x = \sigma_y = 1$.

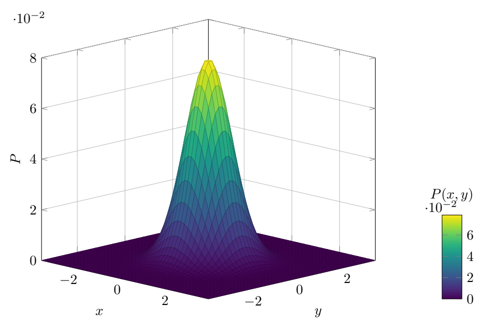

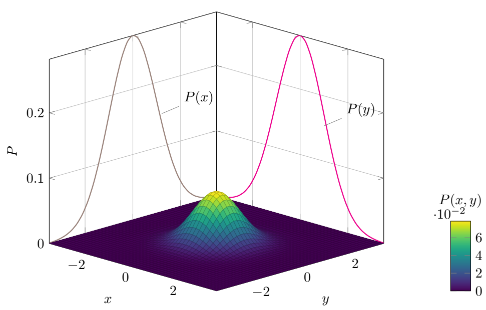

The density is

$$p(x,y) = \frac{1}{2\pi\sqrt{1-\rho^2}} \exp\left\{-\frac{\left(x^2 - 2\rho xy + y^2\right)}{2\left(1-\rho^2\right)}\right\}$$

What is the conditional density, $P_{X|Y=y}(x) = p(x \mid y)$?

$$p(x \mid y) = \frac{p(x,y)}{p(y)} = \frac{1}{\sqrt{2\pi\left(1-\rho^2\right)}} e^{-\frac{(x-\rho y)^2}{2(1-\rho^2)}} \sim N\left(\rho y, \sqrt{1-\rho^2}\right)$$

For the general bivariate Gaussian in $(*)$

$$p(x \mid y) \sim N\left(\mu_x + \rho\frac{\sigma_x}{\sigma_y}\left(y-\mu_y\right), \sigma_x\sqrt{1-\rho^2}\right)$$
$$p(y \mid x) \sim N\left(\mu_y + \rho\frac{\sigma_y}{\sigma_x}\left(x-\mu_x, \sigma_y\sqrt{1-\rho^2}\right)\right)$$

The conditional expected value is

$$E(Y \mid X = x) = \mu_y + \rho\frac{\sigma_y}{\sigma_x}\left(x-\mu_x\right) \quad - \quad \text{straight (regression) line.}$$

Remark
When $\rho = 0$, $p(x,y) = p(x)p(y)$ and X and Y are independent.

The **multivariate Gaussian** distribution of a k−dimensional random vector $\vec{x} = (x_1, \ldots, x_k)^\top$ has pdf

$$p(\vec{x}) = p\left(x_1, \ldots, x_u\right) = \frac{1}{\sqrt{(2\pi)^k|\Sigma|}} \exp\left\{-\frac{1}{2}(\vec{x}-\vec{\mu})^\top\Sigma^{-1}(\vec{x}-\vec{\mu})\right\}$$

where $\vec{\mu} = E(\vec{x})$ is the mean, and $\sum$ is the covariance matrix

$$\sum_{1j} = \operatorname{Cov}\left(x_i, x_j\right) = E\left[\left(x_i-\mu_i\right)\left(x_j-\mu_j\right)\right]$$

1. The marginal distributions are all Gaussian. In fact to obtain the marginal distribution over a subset of the variables, one only needs to drop the irrelevant variables from the mean vector and the covariance matrix.

2. All the conditionals are also Gaussian. If $\vec{x}$ is partitioned as $\vec{x} = \begin{pmatrix} \vec{x}_1 \\ \vec{x}_2 \end{pmatrix}$ and $\vec{\mu} = \begin{pmatrix} \vec{\mu}_1 \\ \vec{\mu}_2 \end{pmatrix}$ and $\Sigma = \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix}$ are the corresponding partitions of the mean vector and the covariance matrix then

$$P_{\vec{x}_1|\vec{x}_2} = \vec{v} \sim N(\tilde{\mu}, \tilde{\Sigma})$$

where

$$\tilde{\mu} = \mu_1 + \Sigma_{12}\Sigma_{22}^{-1}\left(\vec{v}-\mu_2\right)$$
$$\tilde{\Sigma} = \Sigma_{11} - \Sigma_{12}\Sigma_{22}^{-1}\Sigma_{21}$$

Visual View of the Covariance Matrix; Bivariate Gaussian