

Visualizing the bivariate normal distribution in R

Karl Gregory

11/28/2017

Bivariate Normal distribution

The `bivNormal`($\mu_X, \mu_Y, \sigma_X^2, \sigma_Y^2, \rho$) pdf is given by

$$f_{(X,Y)}(x,y) = \frac{1}{2\pi} \frac{1}{\sigma_X \sigma_Y \sqrt{1-\rho^2}} \exp\left(-\frac{1}{2} \left[\left(\frac{x-\mu_X}{\sigma_X}\right)^2 - 2\rho \left(\frac{x-\mu_X}{\sigma_X}\right) \left(\frac{y-\mu_Y}{\sigma_Y}\right) + \left(\frac{y-\mu_Y}{\sigma_Y}\right)^2 \right]\right).$$

The following R code plots contour lines of the bivariate Normal pdf under $\mu_X = \mu_Y = 0$ and several different values of σ_X , σ_Y , and ρ and overlays scatterplots of 50 realizations $(x_1, y_1), \dots, (x_{50}, y_{50})$ of the random variable pair $(X, Y) \sim \text{bivNormal}(\mu_X, \mu_Y, \sigma_X^2, \sigma_Y^2, \rho)$.

```
library(mvtnorm)
grid.size <- 100
x.grid <- y.grid <- seq(-3.25,3.25,length=grid.size)
z <- matrix(NA,grid.size,grid.size)
par(pty="s",mfrow=c(3,3),mar=c(2,0,2,0))

# consider 9 different combinations of the parameters:
mu.x <- c(0,0,0,0,0,0,0,0,0)
mu.y <- c(0,0,0,0,0,0,0,0,0)
sigma.x <- c(1,1,1,1/2,1/2,1/2,1,1,1)
sigma.y <- c(1,1,1,1,1,1,1/2,1/2,1/2)
rho <- c(0,.5,.9,0,.5,.9,0,-.5,-.9)

for(k in 1:9)
{
  mu <- c(mu.x[k],mu.y[k])
  sigma <- matrix( c( sigma.x[k]^2, rho[k]*sigma.x[k]*sigma.y[k],
                    rho[k]*sigma.x[k]*sigma.y[k], sigma.y[k]^2),nrow=2,byrow=TRUE )
  for(i in 1:grid.size)
    for(j in 1:grid.size)
    {
      z[i,j] <- dmvnorm( c(x.grid[i],y.grid[j]) , mean = mu ,sigma = sigma )
    }
  contour(x.grid,y.grid,z,drawlabels=FALSE,main=substitute(
    paste(sigma[X], "=", sigma.x, ", ", sigma[Y], "=", sigma.y, ", ", rho,"=",rho.xy),
    list(sigma.x = sigma.x[k], sigma.y = sigma.y[k], rho.xy = rho[k])))
  points(rmvnorm(50,mean = mu, sigma = sigma),col="red",cex=.5,pch=19)
}
```

