Jointly Gaussian Random Variables

- U and V are independent, standard Gaussian random variables if U is Gaussian(0,1), V is Gaussian(0,1), and U and V are independent.

→ Means: $\mathbb{E}[U] = 0$, $\mathbb{E}[V] = 0$

→ Variances: $\mathrm{Var}[U] = 1$, $\mathrm{Var}[V] = 1$

→ Covariance: $\mathrm{Cov}[U,V] = 0$

→ Marginal PDFs: $f_U(u) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{u^2}{2}\right)$ $f_V(v) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{v^2}{2}\right)$

→ Joint PDF: $f_{U,V}(u,v) \underset{\substack{\uparrow \\ \text{independence}}}{=} f_U(u) f_V(v) = \frac{1}{2\pi} \exp\left(-\frac{1}{2}(u^2+v^2)\right)$

→ Contour Plot:

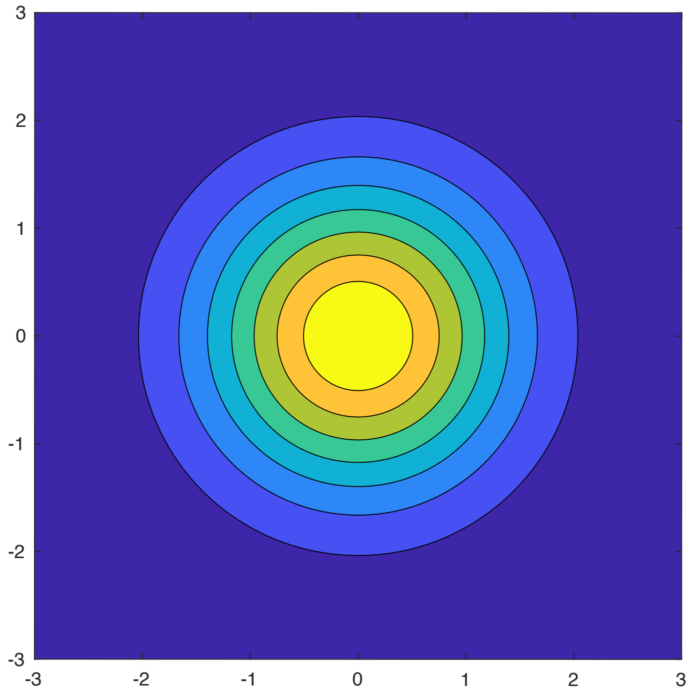

Circularly symmetric about the origin.

- X and Y are jointly Gaussian random variables if they can be expressed as linear functions of independent, standard Gaussian random variables:

$$X = aU + bV + c \qquad Y = dU + eV + f$$

- We usually specify the distribution via these 5 parameters
  → Means: $\mu_X = \mathbb{E}[X]$ and $\mu_Y = \mathbb{E}[Y]$
  → Variances: $\sigma_X^2 = \text{Var}[X]$ and $\sigma_Y^2 = \text{Var}[Y]$
  → Covariance $\text{Cov}[X,Y]$ or Correlation Coefficient $\rho_{X,Y} = \dfrac{\text{Cov}[X,Y]}{\sqrt{\text{Var}[X]\,\text{Var}[Y]}}$

  and leave the linear functions as implicit.

- Linear functions of jointly Gaussian random variables are themselves jointly Gaussian. Only need to update parameters.

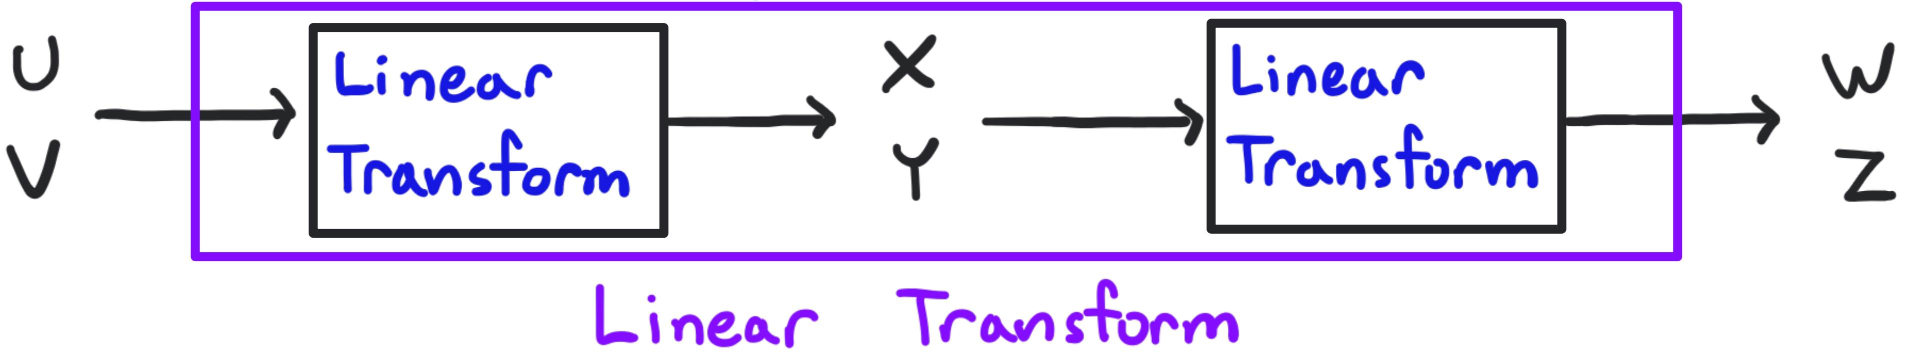

- Joint PDF for Jointly Gaussian X and Y:

$$f_{X,Y}(x,y) = \frac{1}{2\pi\sigma_X\sigma_Y\sqrt{1-\rho_{X,Y}^2}} \exp\left(-\frac{1}{2(1-\rho_{X,Y}^2)}\underbrace{\left(\frac{(x-\mu_X)^2}{\sigma_X^2} - 2\rho_{X,Y}\frac{(x-\mu_X)(y-\mu_Y)}{\sigma_X\sigma_Y} + \frac{(y-\mu_Y)^2}{\sigma_Y^2}\right)}_{\text{specifies an ellipse}}\right)$$

where μ_X is the mean of X
μ_Y is the mean of Y
σ_X^2 is the variance of X
σ_Y^2 is the variance of Y
$\rho_{X,Y}$ is the correlation coefficient of X and Y

→ Not important to memorize the formula for the joint PDF.

→ Instead, we should remember that the joint PDF is fully specified by the parameters μ_X, μ_Y, σ_X^2, σ_Y^2, $\rho_{X,Y}$. How does each parameter influence the contour plot?

- Changing the mean $\mu_x = \mathbb{E}[X]$ shifts the distribution along the x-axis.
- Changing the mean $\mu_Y = \mathbb{E}[Y]$ shifts the distribution along the y-axis.
- **Example:** Fix $\sigma_x^2 = 1$, $\sigma_Y^2 = 1$, $\mathrm{Cov}[X,Y] = 0$.

$\mu_x = 0,\ \mu_Y = 0$

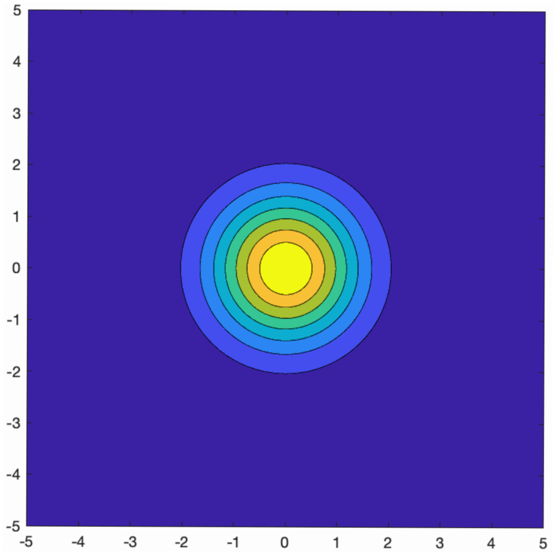

$\mu_x = 2,\ \mu_Y = 0$

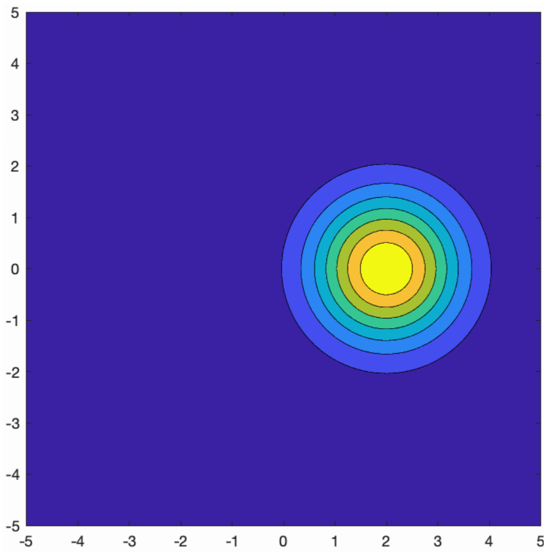

$\mu_x = 0,\ \mu_Y = -3$

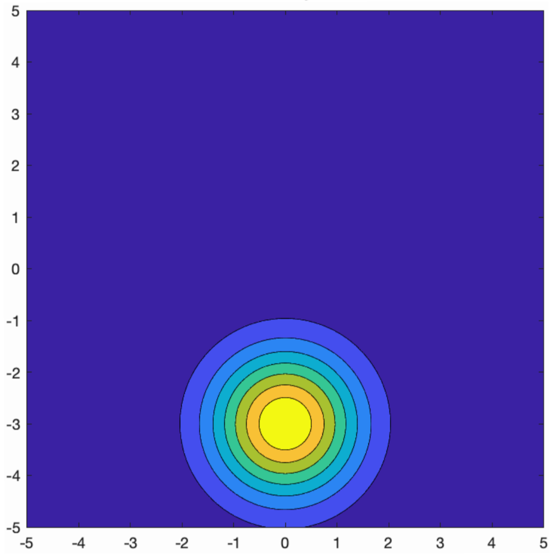

$\mu_x = 2,\ \mu_Y = -3$

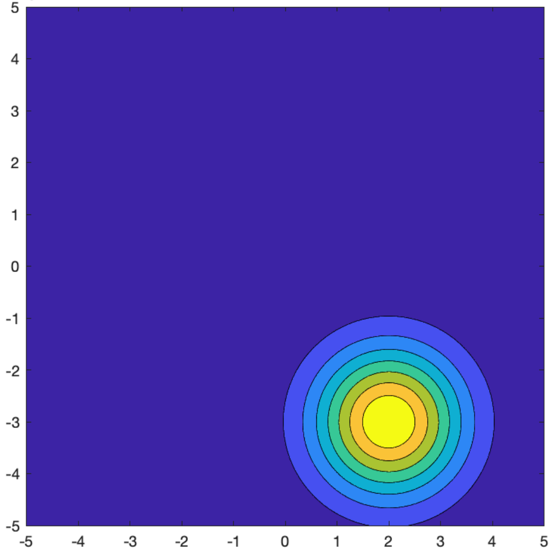

- Increasing the variance $\sigma_X^2 = \text{Var}[X]$ stretches the distribution along the x-axis.
- Increasing the variance $\sigma_Y^2 = \text{Var}[Y]$ stretches the distribution along the y-axis.

- **Example:** Fix $\mu_X = 0, \ \mu_Y = 0, \ \text{Cov}[X,Y] = 0.$

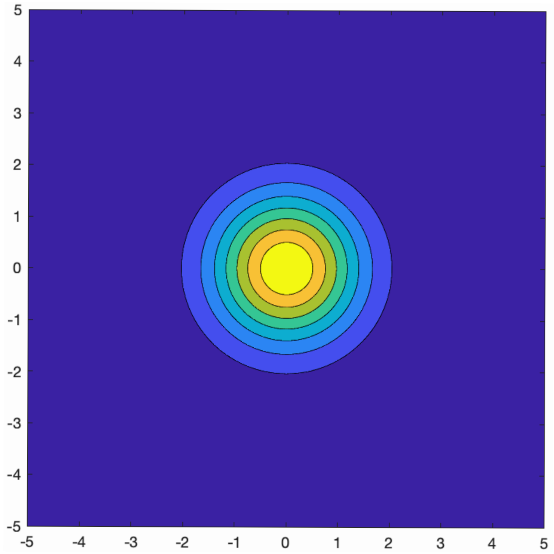
$\sigma_X^2 = 1, \ \sigma_Y^2 = 1$

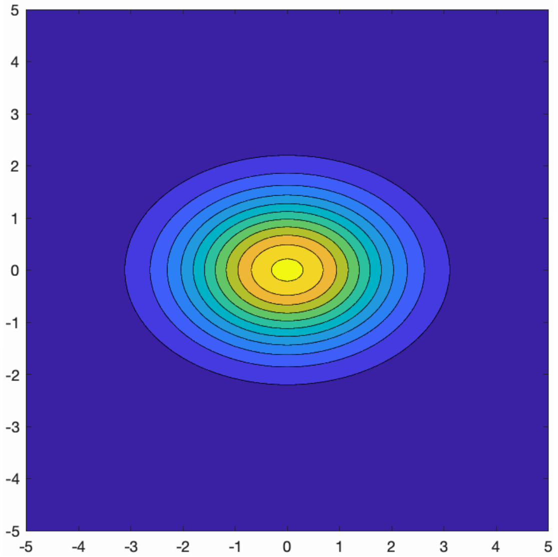
$\sigma_X^2 = 2, \ \sigma_Y^2 = 1$

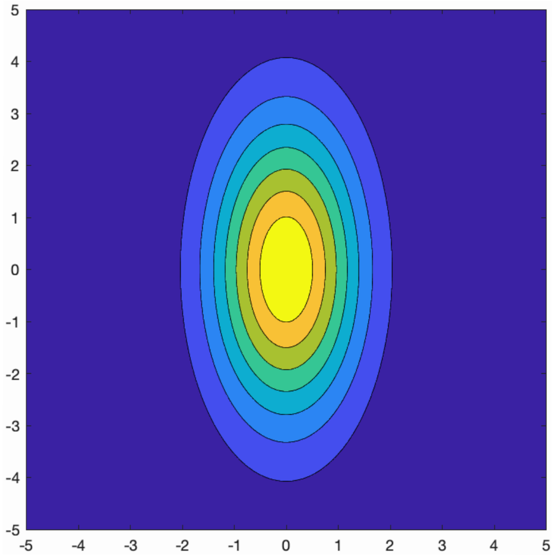
$\sigma_X^2 = 1, \ \sigma_Y^2 = 4$

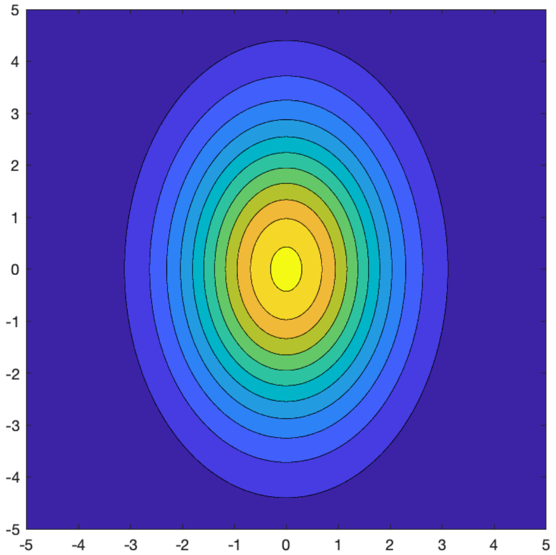
$\sigma_X^2 = 2, \ \sigma_Y^2 = 4$

- For $\rho_{X,Y} > 0$, increasing $\rho_{X,Y}$ towards 1 squeezes the distribution along the line $y = \frac{\sigma_Y}{\sigma_X}(x - \mu_X) + \mu_Y$.

- For $\rho_{X,Y} < 0$, decreasing $\rho_{X,Y}$ towards -1 squeezes the distribution along the line $y = -\frac{\sigma_Y}{\sigma_X}(x - \mu_X) + \mu_Y$.

- Example: ① Fix $\mu_X = 0$, $\mu_Y = 0$, $\sigma_X^2 = 1$, $\sigma_Y^2 = 1$.

$\rho_{X,Y} = 0$

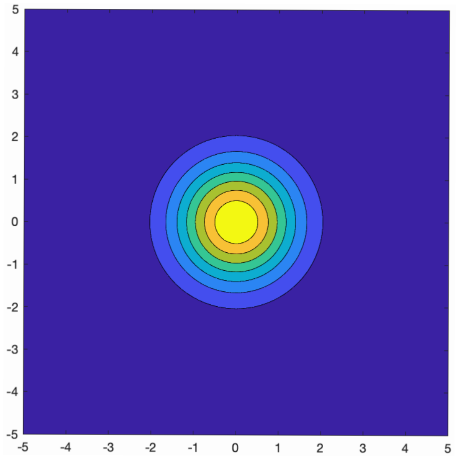

$\rho_{X,Y} = 0.5$

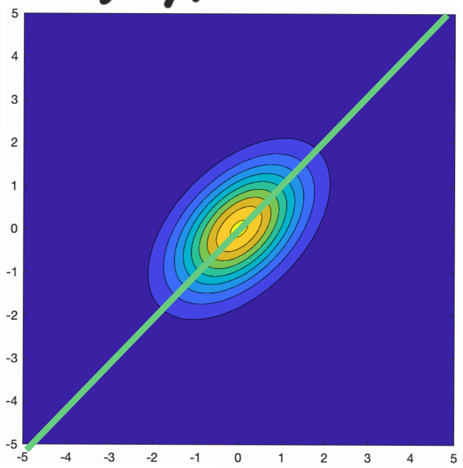

$\rho_{X,Y} = 0.95$

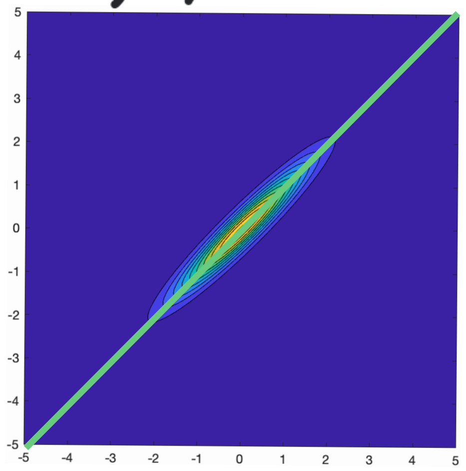

Line

$y = \frac{1}{1}(x - 0) + 0$

$y = x$

② Fix $\mu_X = -1$, $\mu_Y = -2$, $\sigma_X^2 = 4$, $\sigma_Y^2 = 1$.

$\rho_{X,Y} = 0$

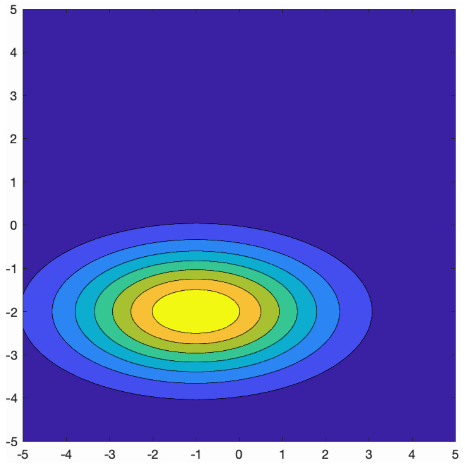

$\rho_{X,Y} = -0.5$

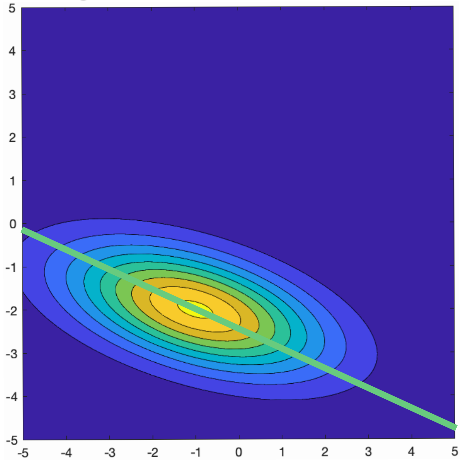

$\rho_{X,Y} = -0.95$

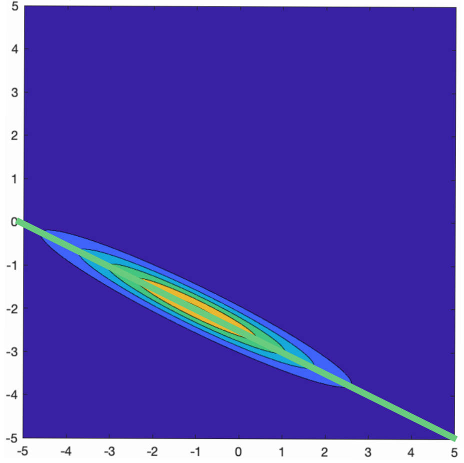

Line

$y = -\frac{1}{2}(x - (-1)) - 2$

$y = -\frac{1}{2}x - \frac{5}{2}$

• Properties of Jointly Gaussian X and Y:

→ If $W = aX + bY + c$ and $Z = dX + eY + f$ are linear functions of X and Y, then W and Z are jointly Gaussian with parameters $\mu_W, \mu_Z, \sigma_W^2, \sigma_Z^2, \text{Cov}[W, Z]$ that can be determined using the linearity of expectation and the variance and covariance of linear functions.

→ Marginal PDFs are Gaussian.

→ The conditional PDF of X given Y Gaussian$(\mathbb{E}[X|Y=y], \sigma_{X|Y}^2)$ where

$$\boxed{\begin{aligned} \mathbb{E}[X|Y=y] &= \mu_X + \frac{\text{Cov}[X,Y]}{\text{Var}[Y]}(y-\mu_Y) = \mu_X + \rho_{X,Y}\frac{\sigma_X}{\sigma_Y}(y-\mu_Y) \\ \sigma_{X|Y}^2 &= \text{Var}[X|Y=y] = \text{Var}[X] - \frac{(\text{Cov}[X,Y])^2}{\text{Var}[Y]} = \sigma_X^2(1-\rho_{X,Y}^2) \end{aligned}}$$

→ $\underbrace{\text{Uncorrelatedness}}_{\text{Cov}[X,Y]=0}$ implies independence and vice versa.

↑ Not true in general!

↑ Always true.