CS195: Computer Vision

Continuous Probability Distribution

Gaussian/Normal Distribution

Multivariate Gaussian/Normal Distribution

Md Alimoor Reza

Assistant Professor of Computer Science

Recap: Probability Basics

- Discrete Probability Distribution
- Joint Probability Distribution
- Marginal Probability Distribution
- Conditional Probability Distribution

Today

- **Continuous Probability Distribution**
- **Joint Probability Distribution**
- **Marginal Probability Distribution**
- **Conditional Probability Distribution**
- **Bayes Rule**

Continuous Probability Distribution

- **Probability Distribution**
  - Discrete Probabilities
  - Continuous Probabilities

- **Continuous Probability Distribution and Random Variable**
  - Probability distribution is a function which will depend on a random variable eg, $\boldsymbol{x}$
  - If the random variable $\boldsymbol{x}$ takes continuous values then you get continuous probabilities

Continuous Probability Distribution

- **Probability Density Function (pdf)**
  - The random variable ***x*** takes on continuous values or real numbers
  - The probability is now expressed in terms of *probability density function (pdf)*

$$q(x) \geq 0$$

$$\int_{-\infty}^{+\infty} q(x)dx = 1$$

  - It specifies the probability of the continuous random variable falling *within a particular range of values (eg, between **a** and **b**)*, as opposed to taking on any one value

$$P(a \leq x \leq b) = \int_a^b q(x)dx$$

The continuous random variable to take on any particular value is 0

Gaussian Distribution

A Gaussian distribution has two parameters:

- mean μ
- standard deviation σ

A one-dimensional Gaussian distribution can be expressed using the following probability density function (pdf)

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

A one-dimensional Gaussian distribution with $\mu = 0, \sigma = 1$ can be expressed using the following probability density function (pdf)

$$P(x) = \frac{1}{\sqrt{(2\pi)}} exp^{\frac{-(x)^2}{2}}$$

Exercise: Gaussian Distribution (1D)

Use the given notebook and generate a 1-D Gaussian distribution by varying the mean and standard deviation. Observe how the shape of the Gaussian distribution changes according to these parameter adjustments.

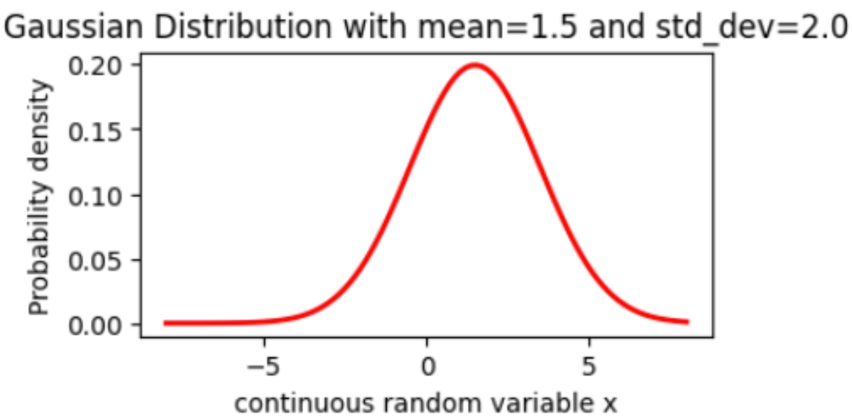

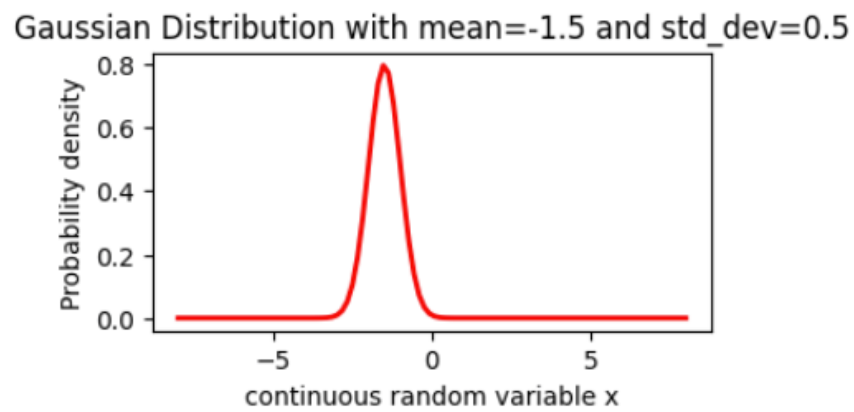

Drake UNIVERSITY

Exercise: Gaussian Distribution (1D)

```
plt.figure(figsize=(4,2))

mean     = -1.5
std_dev  = 0.5
xmin     = -8
xmax     = 8
x        = np.linspace(xmin, xmax, 100)
p        = np.exp(-0.5*((x-mean)/std_dev)**2) / (std_dev * np.sqrt(2*np.pi))
plt.plot(x, p, 'r', linewidth=2)

plt.title('Gaussian distribution with mean='+str(mean)+', std_dev='+str(std_dev))
plt.xlabel('continuous random variable x' )
plt.ylabel('Probability density')
plt.show()
```

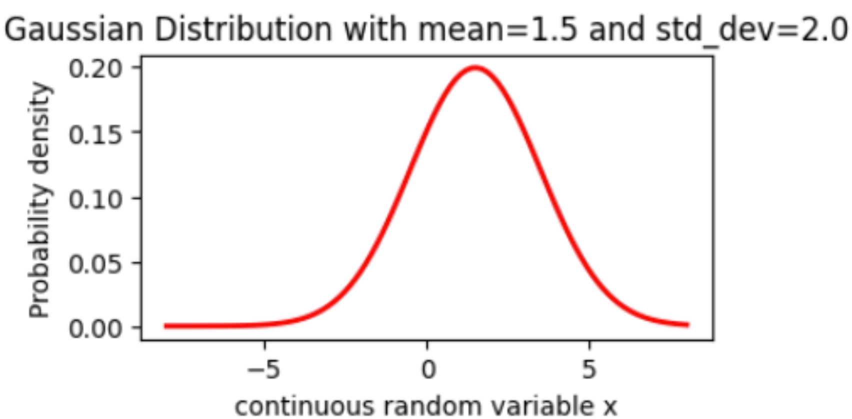

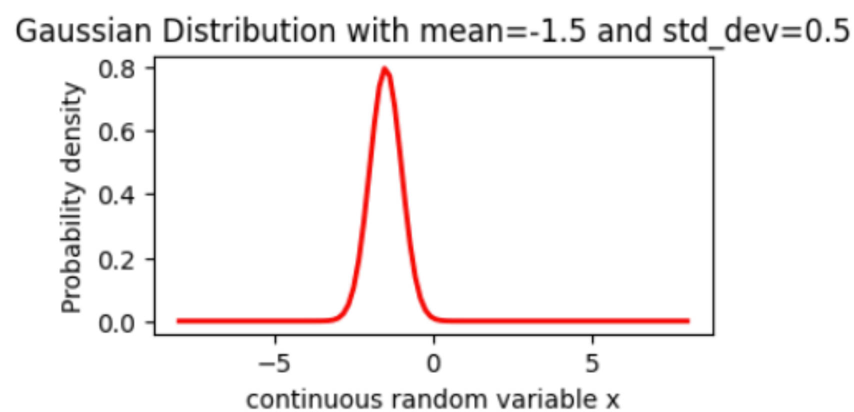

Today

- **Continuous Probability Distribution**
- **Joint Probability Distribution**
- **Marginal Probability Distribution**
- **Conditional Probability Distribution**

2D (Multivariate) Normal Distribution

A multivariate normal distribution is a Gaussian distribution for two or more continuous random variables. For example, a 2-dimensional Gaussian distribution has a probability density function (pdf) expressed with P(x, y)

For example, for 2D normal distribution, the two random variable x, y can be packed into a vector $\mathbf{z}$

A multivariate normal distribution has two parameters:

- mean vector $\boldsymbol{\mu}$
- covariance matrix $\boldsymbol{\Sigma}$

For example, the mean $\boldsymbol{\mu}$ could be a 2D vector defining the position of the distribution
The covariance $\boldsymbol{\Sigma}$ is a symmetric 2x2 matrix that describes the shape of the distribution

$$\boldsymbol{\Sigma} = \begin{bmatrix} \sigma_{11}^2 & \sigma_{12}^2 \\ \sigma_{21}^2 & \sigma_{22}^2 \end{bmatrix}$$

2D (Multivariate) Normal Distribution

A 2D multivariate normal distribution can be expressed using the following probability density function:

$$P(x, y) = P(\mathbf{z}) = \frac{\mathbf{1}}{\mathbf{2}\pi\sqrt{(|\boldsymbol{\Sigma}|)}}\mathbf{exp}^{\frac{-(\mathbf{z}-\mu)^{\mathbf{T}}\boldsymbol{\Sigma}^{-\mathbf{1}}(\mathbf{z}-\mu)}{\mathbf{2}}}$$

Determinant of a matrix
(here covariant matrix)

The two random variables' vector $\mathbf{z} = \begin{bmatrix} x \\ y \end{bmatrix}$

Mean vector $\mu = \begin{bmatrix} \mu_1 \\ \mu_1 \end{bmatrix}$

Covariance matrix $\Sigma = \begin{bmatrix} \sigma_{11}^2 & \sigma_{12}^2 \\ \sigma_{21}^2 & \sigma_{22}^2 \end{bmatrix}$

For example $\mu = \begin{bmatrix} 2.0 \\ 1.0 \end{bmatrix}$

For example $\Sigma = \begin{bmatrix} 0.5 & 0.0 \\ 0.0 & 0.5 \end{bmatrix}$

2D (Multivariate) Normal Distribution

Generate a 2D Gaussian distribution and plot it to visualize the probability density at different locations on the 2D grid.

```
# let's generate a spherical 2D Gaussian distribution
mean              =  [2.0, \
                      1.0]
covariance_matrix = [[0.5, 0.0], \
                     [0.0, 0.5]]

generate_2d_gaussian_distribution(mean, covariance_matrix)
```

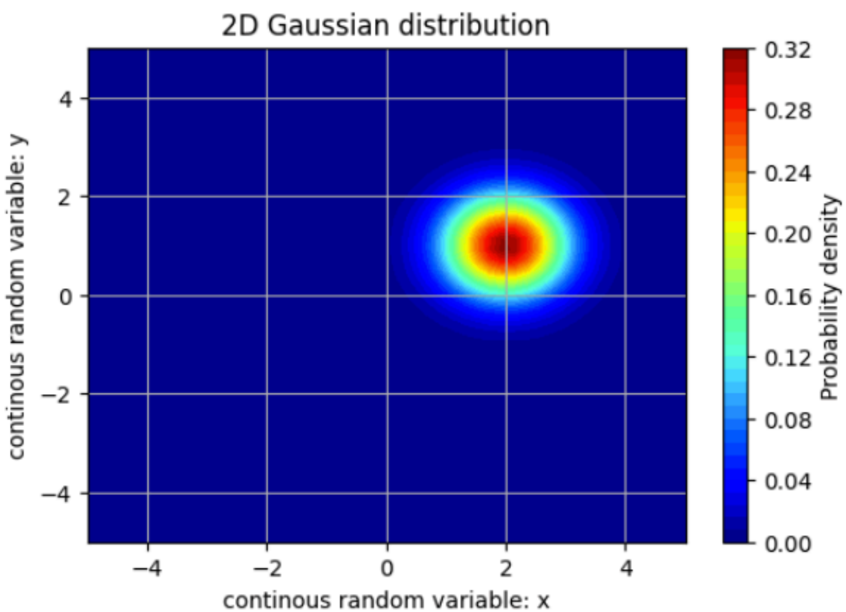


```
# let's generate a spherical 2D Gaussian distribution
mean              =  [2.0, \
                      -1.6]

# let's increase the magnitude of the covariance matrix entries
covariance_matrix = [[2.0, 0.0], \
                     [0.0, 2.0]]
generate_2d_gaussian_distribution(mean, covariance_matrix)
```

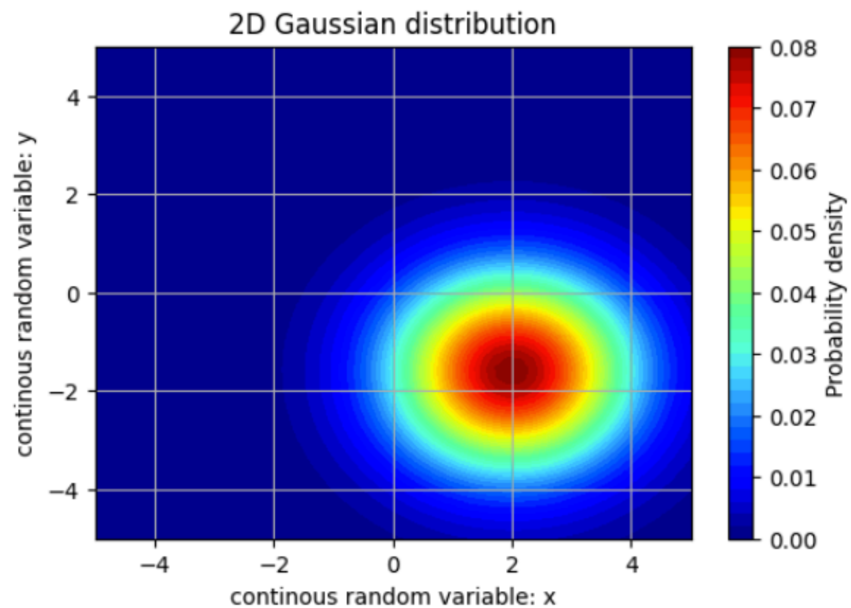

Exercise #2: 2D (Multivariate) Normal Distribution

Change the mean vector and covariance matrix to generate a 2D Gaussian distribution of different shapes and plot it to visualize the probability density at different locations on the 2D grid.

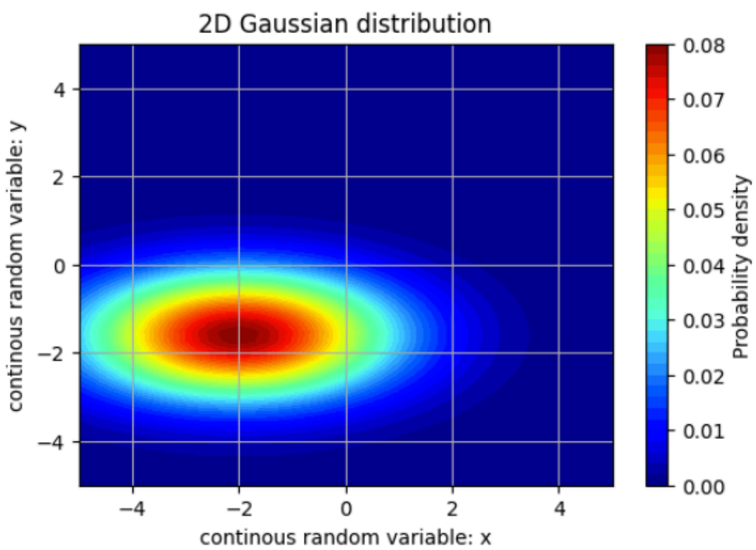

<u>Your turn:</u> Generate a diagonal covariance matrix (which will spread the probability density diagonally, more towards the x-coordinate).

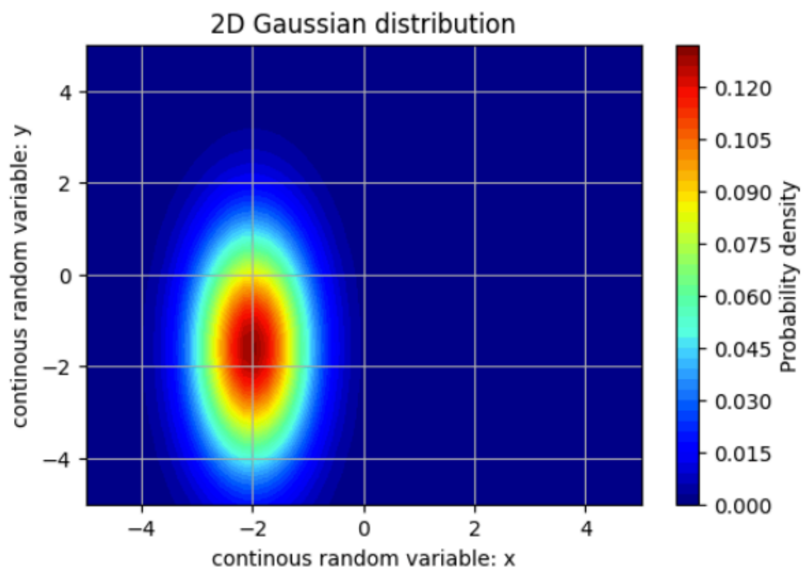

<u>Your turn:</u> Generate a diagonal covariance matrix (which will spread the probability density diagonally, more towards the y-coordinate).

Exercise #2: 2D (Multivariate) Normal Distribution

Change the mean vector and covariance matrix to generate a 2D Gaussian distribution of different shapes and plot it to visualize the probability density at different locations on the 2D grid.

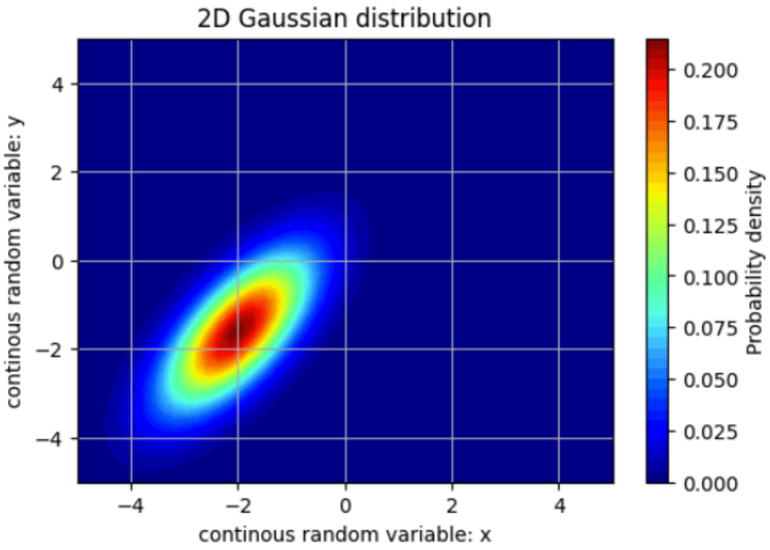

<u>Your turn:</u> Generate a full covariance matrix (which will spread the probability density diagonally, more towards the x-coordinate).

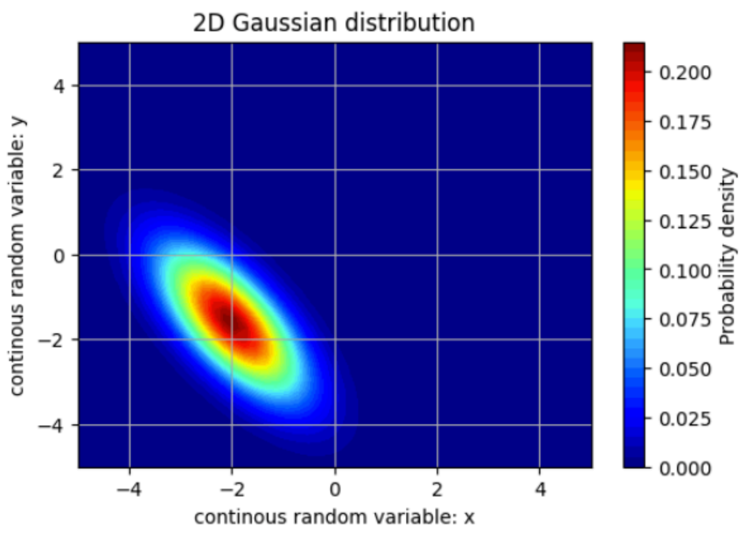

<u>Your turn:</u> Generate a full covariance matrix (which will spread the probability density diagonally, more towards the y-coordinate).

Other Common Probability Distributions

Data Type	Domain	Distribution
univariate, discrete, binary	$x \in \{0,1\}$	Bernoulli
univariate, discrete, multivalued	$x \in \{1,2,\ldots,K\}$	categorical
univariate, continuous, unbounded	$x \in \mathbb{R}$	univariate normal
univariate, continuous, bounded	$x \in [0,1]$	beta
multivariate, continuous, unbounded	$\mathbf{x} \in \mathbb{R}^K$	multivariate normal
multivariate, continuous, bounded, sums to one	$\mathbf{x} = [x_1, x_2, \ldots, x_K]^T$ $x_k \in [0,1], \sum_{k=1}^{K} x_k = 1$	Dirichlet
bivariate, continuous, x_1 unbounded, x_2 bounded below	$\mathbf{x} = [x_1, x_2]$ $x_1 \in \mathbb{R}$ $x_2 \in \mathbb{R}^+$	normal-scaled inverse gamma
vector $\mathbf{x}$ and matrix $\mathbf{X}$, $\mathbf{x}$ unbounded, $\mathbf{X}$ square, positive definite	$\mathbf{x} \in \mathbb{R}^K$ $\mathbf{X} \in \mathbb{R}^{K \times K}$ $\mathbf{z}^T\mathbf{X}\mathbf{z} > 0 \quad \forall\, \mathbf{z} \in \mathbb{R}^K$	normal inverse Wishart

Today

- Continuous Probability Distribution
- Joint Probability Distribution
- **Marginal Probability Distribution**
- Conditional Probability Distribution

Marginal Probability Distribution

- The marginal probability distribution expresses probability distribution of any single random variable from a joint probability distribution by integrating over all other variables

Joint distribution

$$P(x, y) = P(\mathbf{z}) = \frac{1}{2\pi\sqrt{(|\boldsymbol{\Sigma}|)}}\mathbf{exp}^{\frac{-(\mathbf{z}-\mu)^{\mathbf{T}}\boldsymbol{\Sigma}^{-1}(\mathbf{z}-\mu)}{2}}$$

Marginal distribution

$$P(x) = \int P(x, y)\, dy$$

Marginal distribution

$$P(y) = \int P(x, y)\, dx$$

$$P(X) = \sum_y P(X, Y = y)$$

Notice the similarity, the discrete marginal distribution has a summation term instead of integration

Marginal Probability Distribution

- Let's consider a joint probability distribution of a 2D Gaussian distribution as follows:

```
# Let's generate a full covariance matrix
# 2D Gaussian is centered at (0, 0) coordinate
mean              =  [0.0, \
                      0.0]

# covariance matrix spreads the probability density in all directions.
covariance_matrix = [[0.8, 0.7], \
                     [0.7, 1.3]]
generate_2d_gaussian_distribution(mean, covariance_matrix)
```

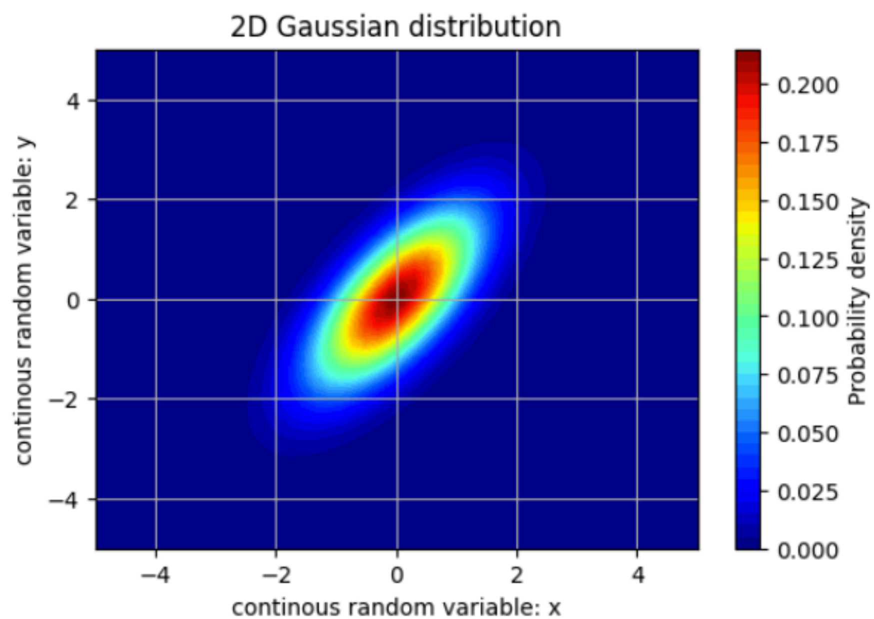

Marginal Probability Distribution

- The marginal probability distribution expresses probability distribution of any single random variable from a joint probability distribution by integrating over all other variables

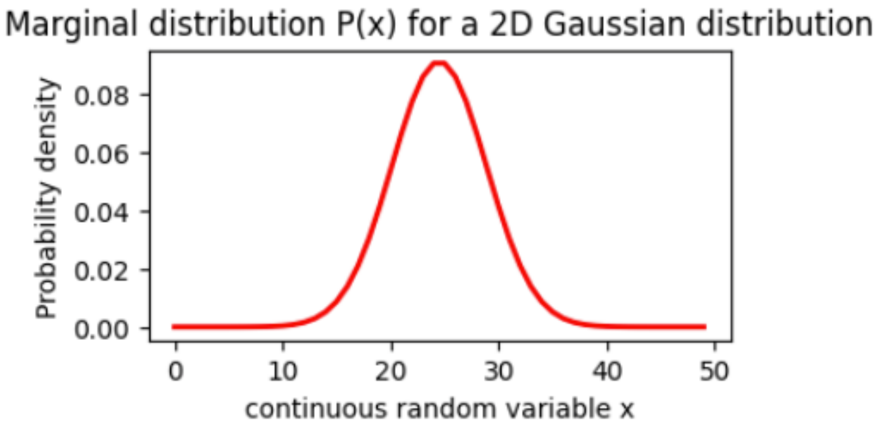

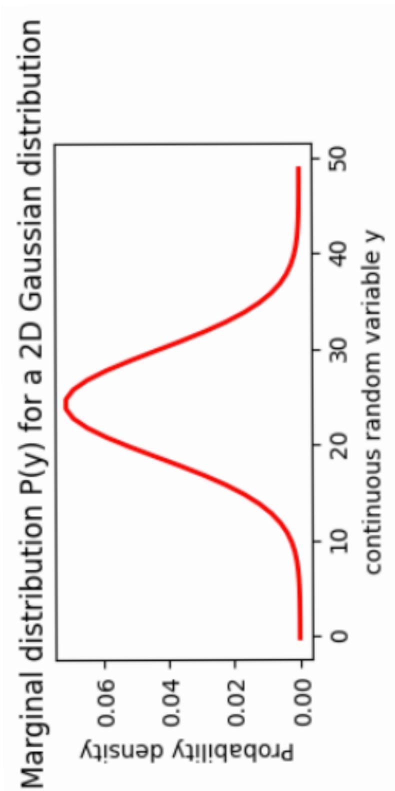

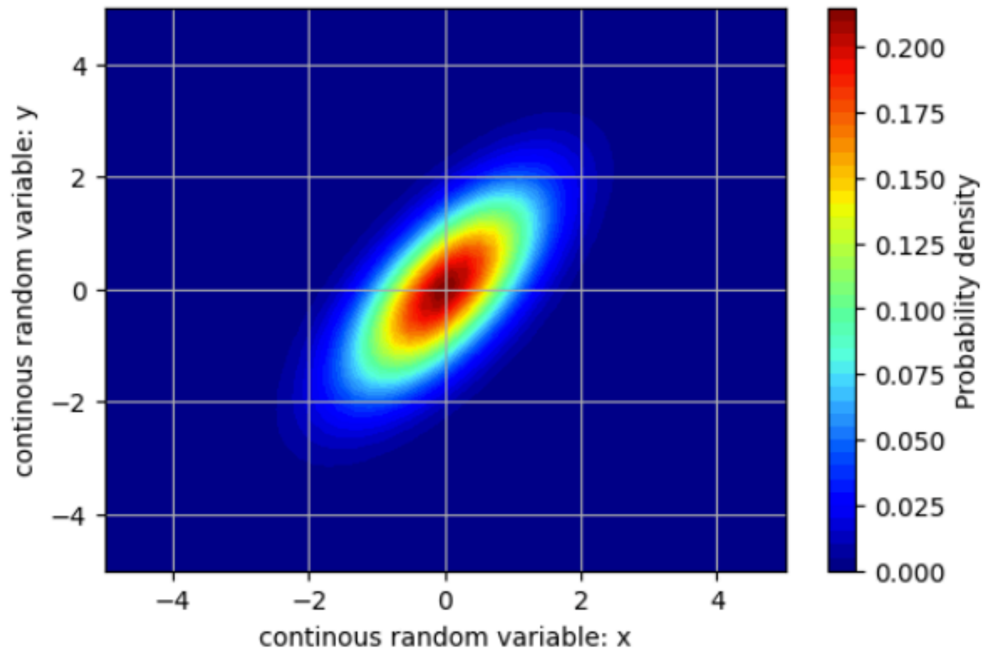

Today

- **Continuous Probability Distribution**
- **Joint Probability Distribution**
- **Marginal Probability Distribution**
- **Conditional Probability Distribution**

Conditional Probability Distribution

- The conditional probability distribution of any single random variable *(eg, x)* conditioned on other variable's *(eg, y)* value fixed to a particular value.

- **It tells us the relative propensity** of the random variable x to take different outcomes given that the random variable y is fixed to the value y^*

Conditional distribution

Joint distribution

$$P(x \mid y = y^*) = \frac{P(x, y)}{P(y = y^*)}$$

It's a marginal distribution for a fixed value of y*

$$= \frac{P(x, y)}{\int P(x, y = y^*)\ dx}$$

Notice the similarity, the discrete conditional distribution has a summation term instead of integration term.

$$P(X = a \mid Y = \alpha) = \frac{P(X = a, Y = \alpha)}{P(Y = \alpha)}$$

$$= \frac{P(X = a, Y = \alpha)}{P(X = a, Y = \alpha) + P(X = b, Y = \alpha) + P(X = c, Y = \alpha)}$$

If you don't recall, here is how it can be recomputed

Conditional Probability Distribution

- Let's consider a joint probability distribution of a 2D Gaussian distribution as follows:

```
# Let's generate a full covariance matrix
# 2D Gaussian is centered at (0, 0) coordinate
mean              =  [0.0, \
                      0.0]

# covariance matrix spreads the probability density in all directions.
covariance_matrix = [[0.8, 0.7], \
                     [0.7, 1.3]]
generate_2d_gaussian_distribution(mean, covariance_matrix)
```

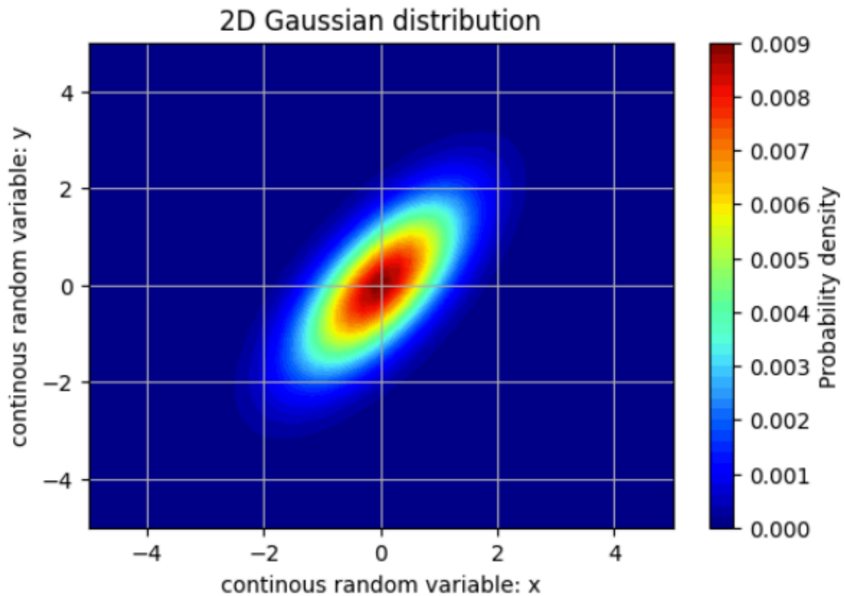

Conditional Probability Distribution

- The conditional probability distribution P(xly=y*) expresses the **relative propensity** of the random variable *x* to take different outcomes given that the random variable y is fixed to the value *y**

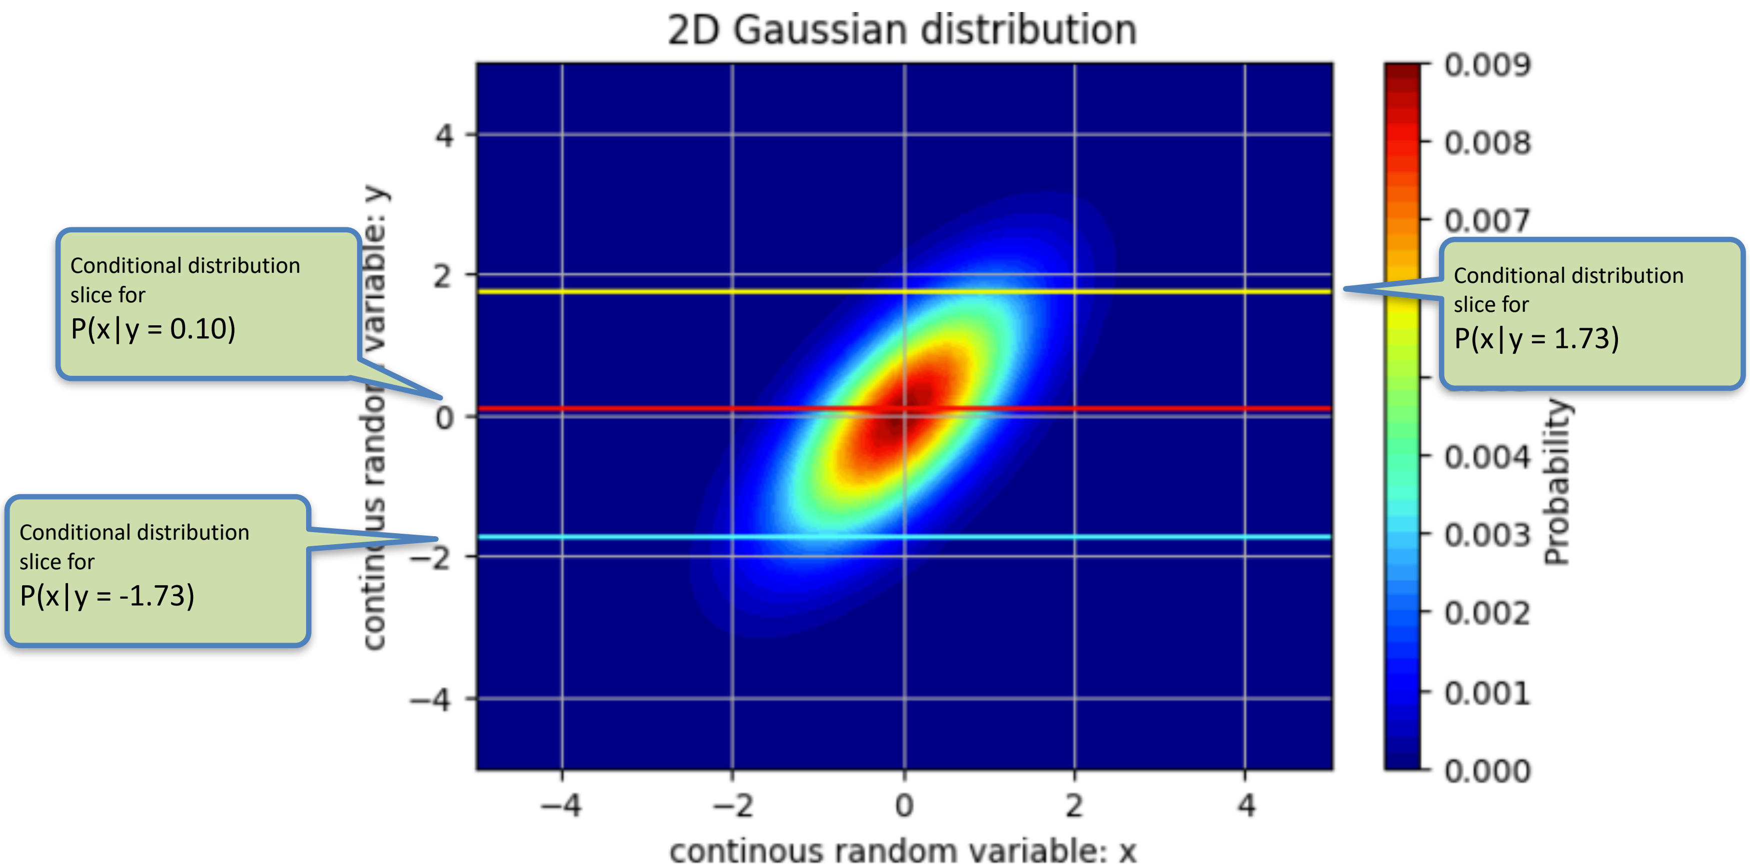

Conditional Probability Distribution

Conditional distribution
slice for
P(x|y = 1.73)

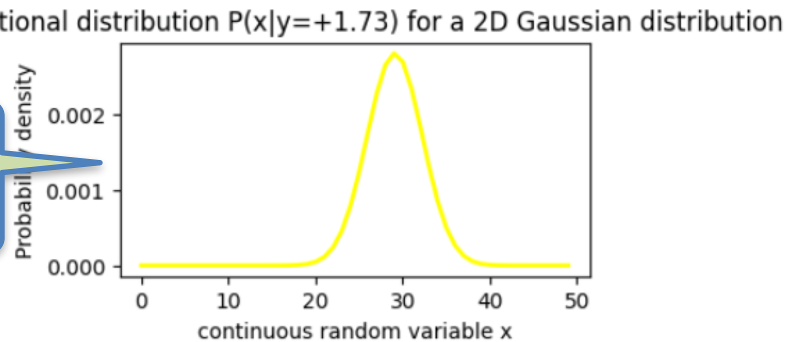

Conditional distribution
slice for
P(x|y = 0.10)

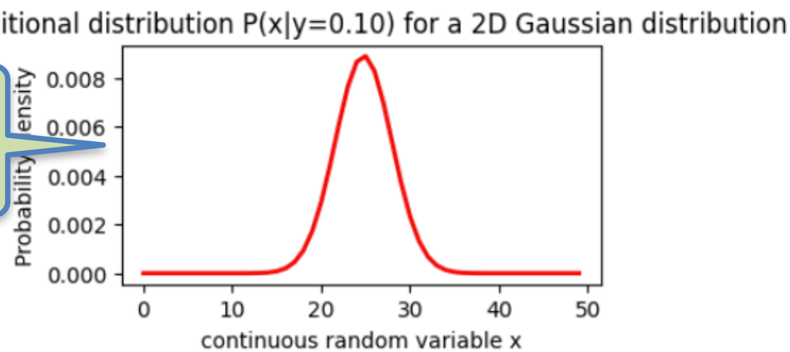

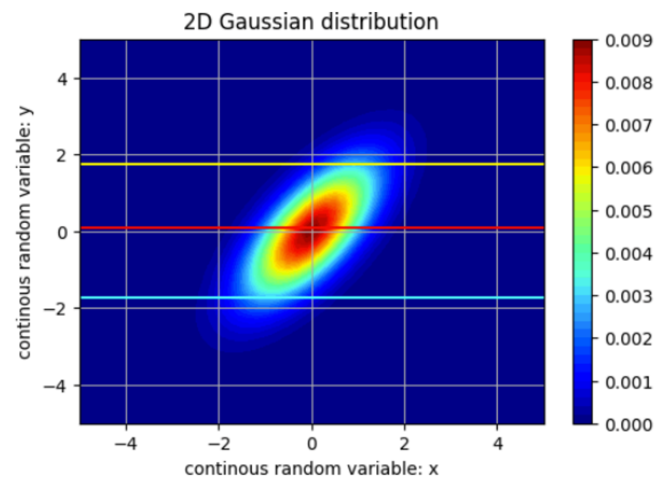

Conditional distribution
slice for
P(x|y = -1.73)

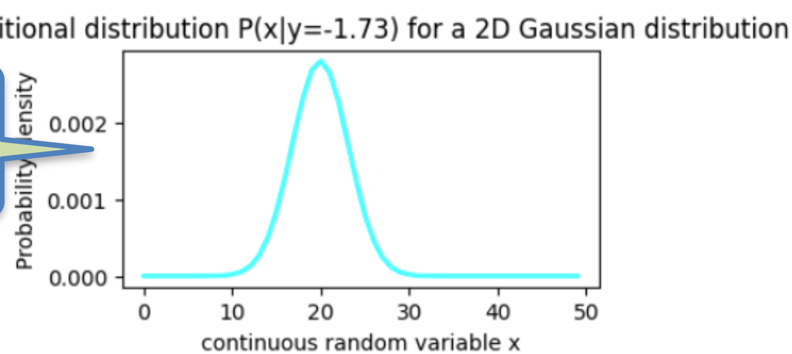

Exercise#4: 2D Normal Distribution

Change the mean vector and covariance matrix to generate a 2D Gaussian distribution of different shapes and recompute three conditional distributions
P(x| y= -1.73), P(x| y= 0.10) and P(x|y=+1.73).