

# Lab: Iris Flower Classification using Decision Theory

In this exercise, we will use decision theory on a simple multi-variable classification problem. In doing the exercise, you will learn to:

- Load a pre-installed dataset in the `sklearn` package.
- Estimate parameters of a multi-variable Gaussian from data
- Make multi-class predictions using linear and quadratic discriminants derived from the Gaussian parameters
- Evaluate the accuracy of the predictions on test data

For submission:

- Complete all sections labeled `#TODO`
- Run the notebook and print to PDF.
- Submit the PDF in CCLE. No other formats accepted. Do not submit the jupyter notebook.
- You may also do the lab in any other language (e.g. R) of your choice. But only submit a single PDF file. No source code.

We load the following packages.

```
In [1]: import numpy as np
import matplotlib.pyplot as plt
import matplotlib
matplotlib.rcParams.update({'font.size': 16})
```

## Loading and Visualizing the Data

In this lab, we will use the classic dataset [Iris flower dataset](https://en.wikipedia.org/wiki/Iris_flower_data_set) ([https://en.wikipedia.org/wiki/Iris\\_flower\\_data\\_set](https://en.wikipedia.org/wiki/Iris_flower_data_set)). The problem is to estimate the type of iris ('setosa' 'versicolor' 'virginica') from four features of the Iris flower. Since the data is widely-used in machine learning classes, it is included in the `sklearn` package in python. You can download the data with following command:

```
In [3]: from sklearn.datasets import load_iris
data = load_iris()
```

Set `X=data.data` and `y=data.target`. The matrix `X[i,j]` will then be the value of feature `j` in sample `i` and `y[i]` will be the index of the class for sample `i`. Also, print:

- Number of samples,
- Number of features per sample
- Number of classes

```
In [2]: # TODO
# X = ...
# y = ...
```

Print the feature and target names in `data.feature_names` and `data.target_names`

```
In [4]: # TODO
# print(...)
```

To visualize the data, create a *scatter* plot of two of the four features: (sepal length, petal length). In a scatter plot, each point is plotted as some marker, say a small circle. Different colors are used for different class labels. You can create a scatter plot using the `plt.scatter()` command. Make sure you label your axes using the `plt.xlabel()` and `plt.ylabel()` axes.

If you did the plot correctly, you should see that you can separate the points well even just using two features.

```
In [5]: # TODO
```

## Classifier Based on a Linear Discriminator

Before trying any classification methods, it is necessary to split the data into two components:

- Training samples: Used for fitting the classifier models
- Test samples: Used for testing the classifier

We will discuss splitting the training and test data in detail later in the class. But, the reason for the splitting is that we want to test the classifier on samples not used in training. This ensures we see how well it works on *new* samples that have not been seen.

For this purpose, divide the data into 75 samples ( $X_{tr}, y_{tr}$ ) for training and 75 samples ( $X_{ts}, y_{ts}$ ) for test. You must randomly shuffle the samples before splitting -- do not pick the first 75 for training. You can do the splitting manually or use the `sklearn.model_selection.train_test_split` function. If you use `train_test_split()`, set `shuffle=True`.

```
In [6]: from sklearn.model_selection import train_test_split

# TODO Split into training and test
# Xtr, Xts, ytr, yts = ...
```

We will first try using a linear discriminator. For linear discriminator we first need to estimate the sample mean of each class. Complete the following function which returns:

$\mu[:,j]$  = sample mean of  $X[i,:]$  for samples with  $y[i]=j$

```
In [7]: def fit_lin(X,y):  
        # TODO  
        # mu = ...  
        return mu
```

Run the function on the training data to get the sample means.

```
In [8]: # TODO  
        # mu= fit_lin(...)
```

Assuming a Gaussian model with equal and i.i.d. covariance matrices, the optimal estimator is given by the linear discriminator: Given a test sample  $x$ , we compute:

$$g[j] = \mu[:,j].\text{dot}(x) - 0.5*\text{sum}(\mu[:,j]**2)$$

for each class  $j$ . Then, we select  $\hat{y} = \text{argmax}_j g[j]$ . Complete the following code which takes a matrix of data samples  $X$  to compute a vector of class predictions  $\hat{y}$ .

```
In [10]: def predict_lin(X,mu):  
        # TODO  
        # g = ...  
        # yhat = ...  
        # return yhat  
        return yhat
```

Test the linear classifier on the test data. Specifically, estimate  $\hat{y}$  and measure the accuracy, which is the number of samples on which the classifier was correct. If you did everything correctly, you should get an accuracy of around 90%. But, it may be a little higher or lower depending on the random training / test split.

```
In [11]: # TODO  
        # yhat = predict_quad(...)  
        # acc = ...  
        # print(acc)
```

## Quadratic Discriminator

Now, we will try a more sophisticated classifier. In this case, we will estimate both the sample mean and covariance matrix. Complete the following code that computes:

- $\mu$  : Array where  $\mu[:,j]$  is the sample mean for the samples with  $y[i]=j$ . This is identical to the code you have above.
- $S$  : Array of covariance matrices where  $S[:, :, j]$  is the sample covariance matrix for the samples with  $y[i]=j$ . You can use the `np.cov()` method.

```
In [12]: def fit_quad(X,y):
          # TODO
          # mu = ...
          # S = ...
          # return mu, S
          return mu, S
```

Fit the quadratic model on the training data.

```
In [13]: # TODO
          # mu, S = fit_quad(...)
```

Given Gaussian models in each class the optimal decision rule is the following *quadratic* decision rule: Given a sample  $x$ , we compute the discriminant,

$$g[j] = 0.5*(x-\mu[:,j]).T.dot(Sinv[:, :, j]).dot(x-\mu[:,j]) + 0.5*\log(\det(S[:, :, j]))$$

where  $Sinv[:, :, j]$  is the matrix inverse of  $S[:, :, j]$ . Then, we take  $\hat{y} = \operatorname{argmin}_j g[j]$ . Complete the following code to compute the predictions based on the quadratic discriminants.

```
In [14]: def predict_quad(X,mu,S):
         """
         Quadratic discriminator using discriminants:

         
$$g = -\ln p(x|\mu,X) = 0.5*(x-\mu)*S^{-1}*(x-\mu) + 0.5*\log \det(S)$$


         """
         # TODO: For each class ic, compute
         #       Sinv[:,:,ic] = inverse of S
         #       Logdet[ic] = log det(S)

         # TODO: Compute
         # g[i,j] = discriminant for sample i and class j
         # yhat[i] = class estimate for sample i

         return yhat
```

Test the quadratic discriminator on the test data. You should get around 98% accuracy depending on the train / test split.

```
In [15]: # TODO
         # yhat = predict_quad(...)
         # acc = ...
```

```
In [ ]:
```

```
In [ ]:
```