Lisbon School of Economics & Management
Universidade de Lisboa

LISBOA
UNIVERSIDADE DE LISBOA

Ano Letivo 2026/2027

MACHINE LEARNING

1. What is Machine Learning?

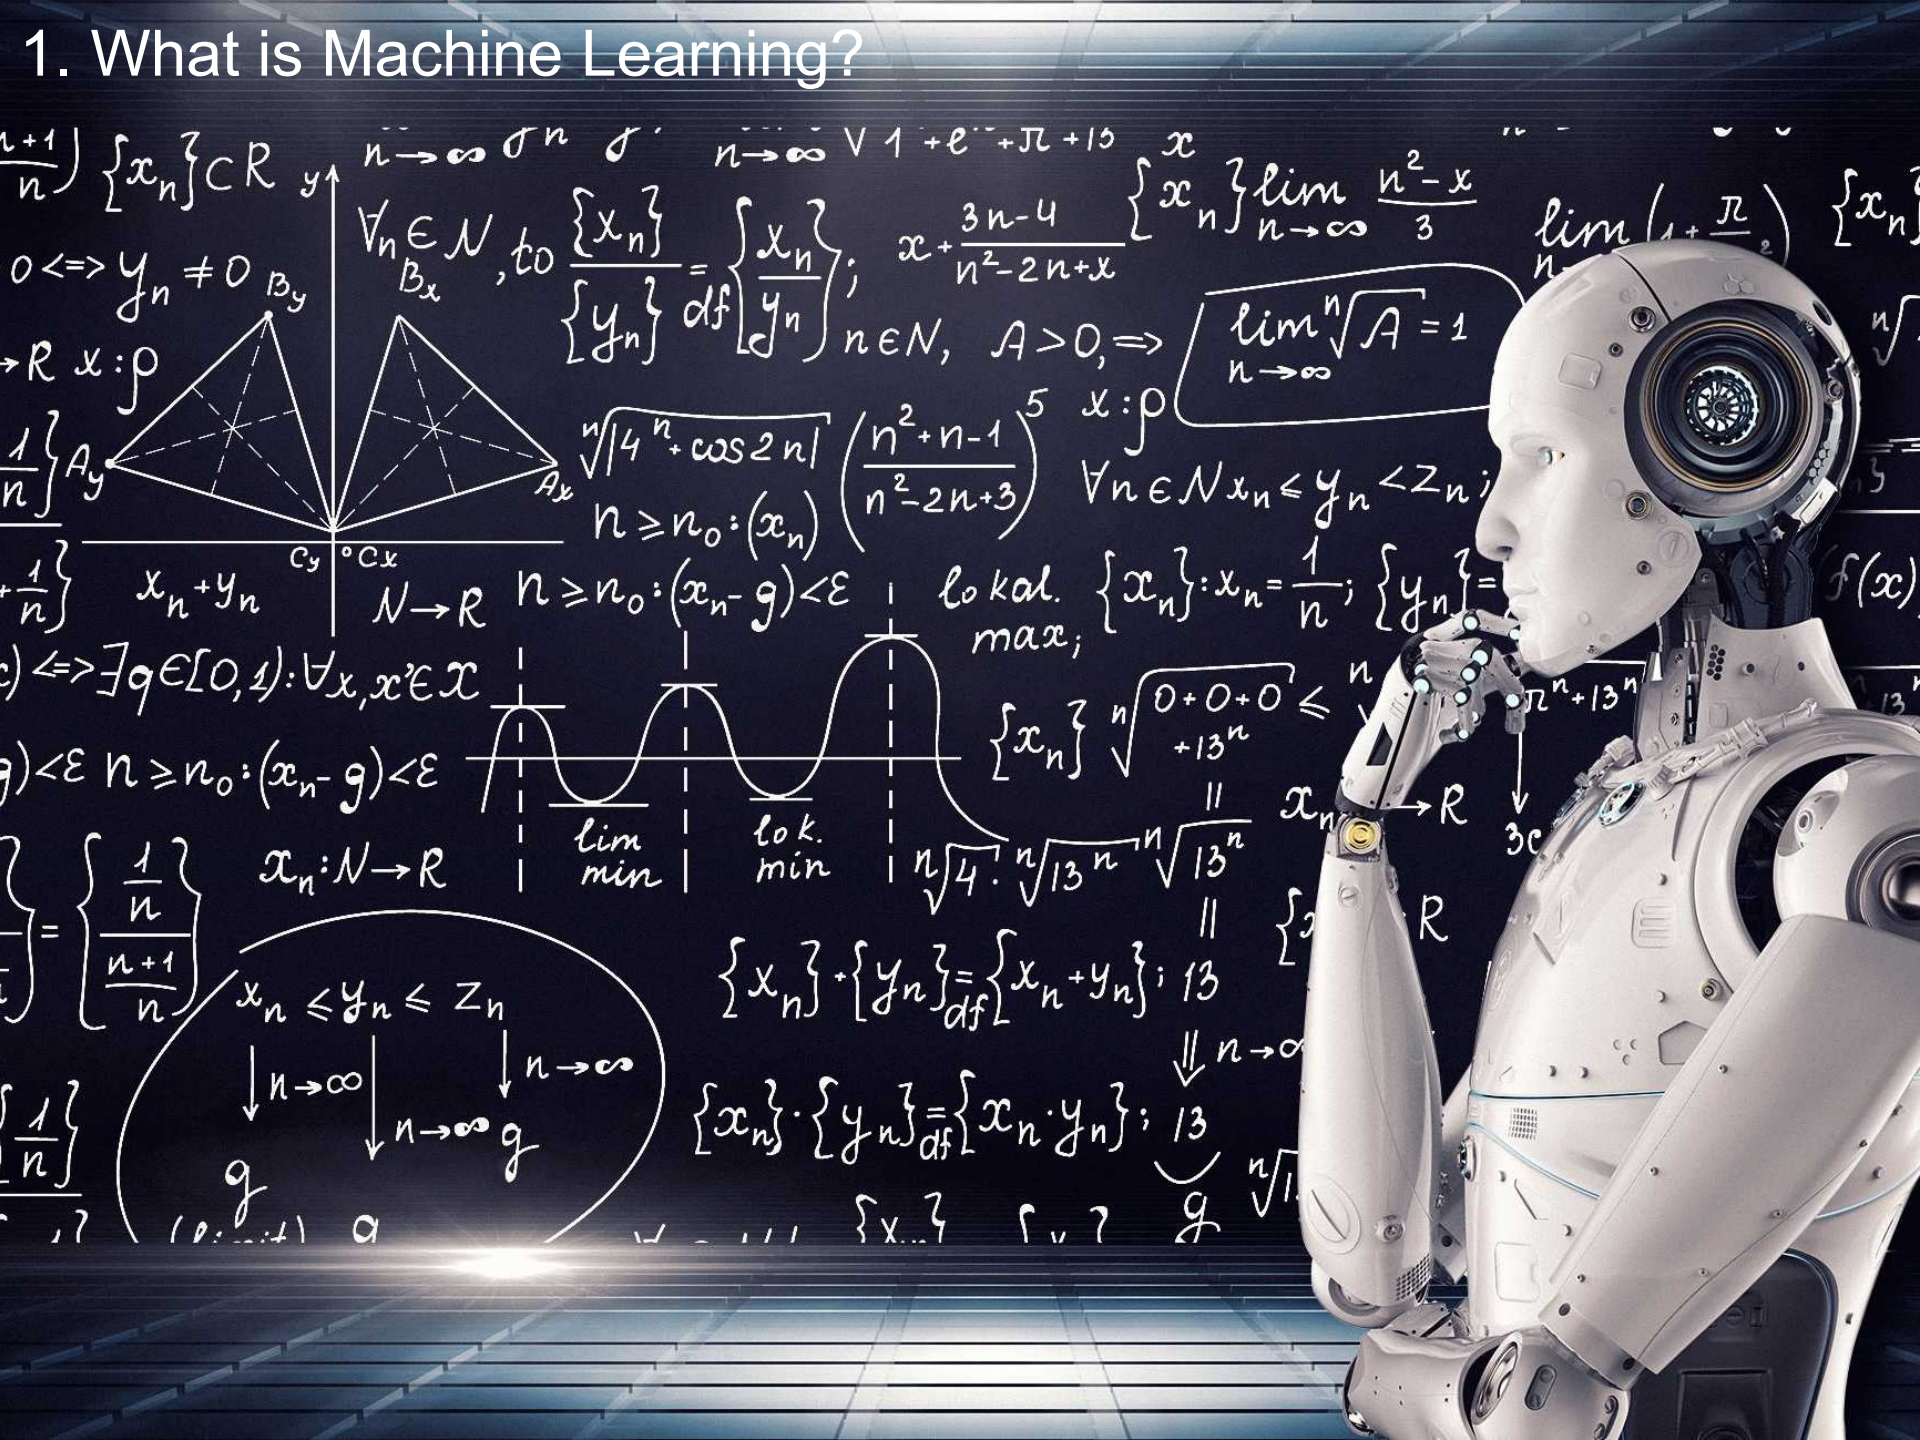

1. What is Machine Learning?

- Branch of Artificial Intelligence
- focuses on creating systems that can **learn from data** rather than relying on explicit programming.
- Instead of writing step-by-step instructions we provide data and algorithms that **detect patterns**, make predictions, or group observations.

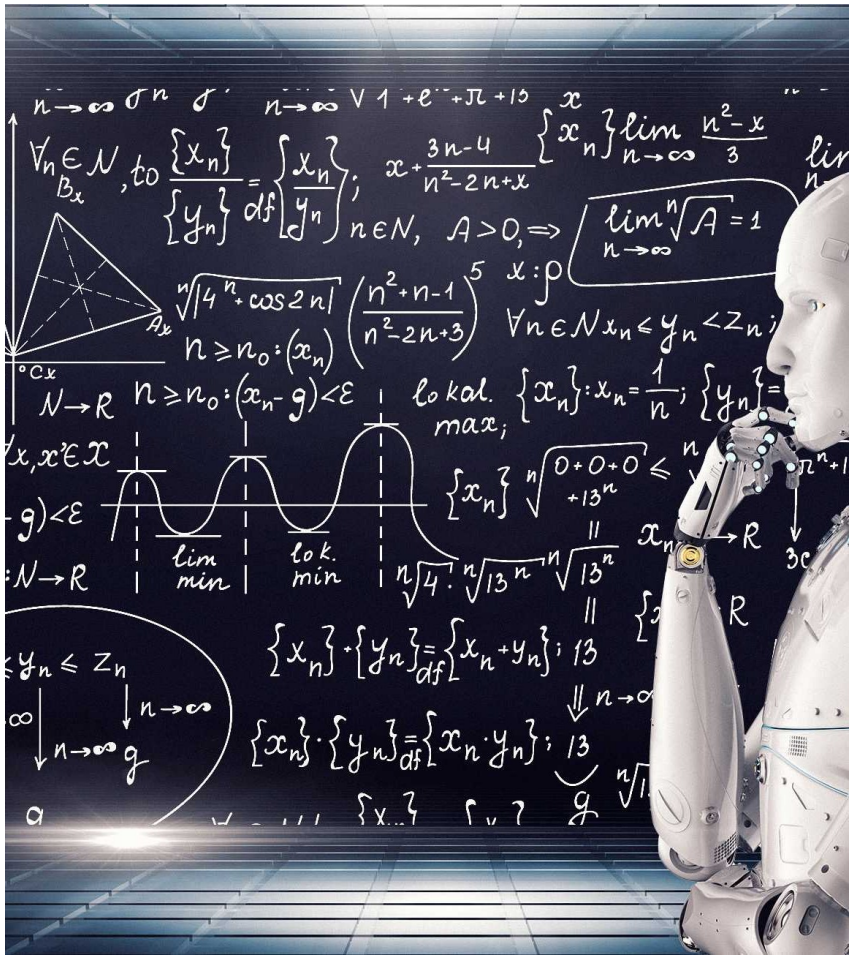

1. What is Machine Learning?

- **Key idea:** Algorithms approximate relationships in data.
- **Common applications:**
  - forecasting demand,
  - fraud detection,
  - customer segmentation,
  - recommendation systems,
  - sentiment analysis.

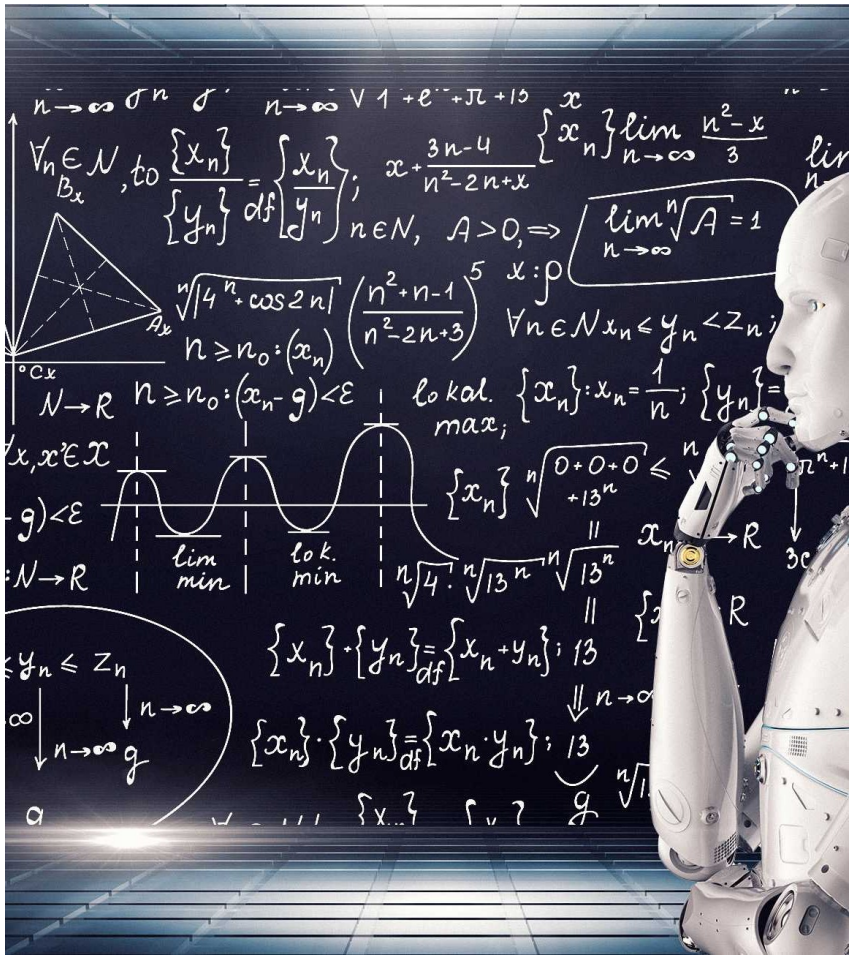

2. Categories of Machine Learning

- Supervised Learning
- Unsupervised Learning
- Reinforcement Learning

2.1 Supervised Learning

- **Definition:** Learning from labeled data, i.e., input-output pairs (X, y).
- **Goal:** Train a model that maps inputs $X$to outputs y.

2.1 Supervised Learning

- **Examples:**
  - Predicting house prices (regression)
  - Classifying emails as spam/not spam (classification)
- **Main algorithms:** Linear Regression, Logistic Regression, Decision Trees, Random Forests, Support Vector Machines, Neural Networks.

2.2 Unsupervised Learning

- **Definition:** Learning from unlabeled data, i.e., only inputs X.
- **Goal:** Discover hidden structures or patterns.

2.2 Unsupervised Learning

- **Examples:**
  - Customer segmentation (clustering)
  - Topic modeling in documents
  - Dimensionality reduction (e.g., PCA)
- **Main algorithms:** K-Means, Hierarchical Clustering, PCA, t-SNE, Autoencoders.

Example of Regression (0/8)

- This example illustrates Linear Regression to predict apartment prices based on size.
- This is a key tool for managers to forecast and make data-driven decisions.

Example of Regression (1/8)

1. Import Libraries

pandas for data handling, sklearn for machine learning.

```
import pandas as pd
from sklearn.model_selection import train_test_split
from sklearn.linear_model import LinearRegression
from sklearn.metrics import mean_squared_error, r2_score
import matplotlib.pyplot as plt
import numpy as np
```

Example of Regression (2/8)

2. Create Sample Data

- **Let's simulate data: 'Size' (X) and 'Price' (y).**
- **Imagine apartment sizes (e.g., 50-150 sq. meters) and their corresponding prices.**
- **For consistent results**
- **size = np.random.rand(100) * 100 + 50**
- **Apartment sizes from 50 to 150**

```
price = 200 * size + np.random.randn(100) * 10000 + 50000

df = pd.DataFrame({'Size_sqm': size, 'Price_EUR': price})

print("--- Sample Data (First 5 Rows) ---")
display(df.head())
.
```

Example of Regression (3/8)

3. Split Data for Training and Testing

- We separate our data so the model learns from one part (training) and is tested on another (test).
- This ensures the model can predict well on new, unseen data.

```
X = df[['Size_sqm']]
y = df['Price_EUR']

X_train, X_test, y_train, y_test = train_test_split(X, y,
test_size=0.2, random_state=42)

print(f"\nTraining data points: {len(X_train)}, Testing
data points: {len(X_test)}")
```

Example of Regression (4/8)

4. Train the Linear Regression Model

- The model finds the best-fit line to describe the relationship between Size and Price.

```
model = LinearRegression()
model.fit(X_train, y_train)

print("\n Model Training Complete ")
print("\n Learned relationship:",model.coef_[0],"* Size
+", model.intercept_)
print("\n Coefficient (Slope):, model.coef_[0])
print("\n Intercept (Base Price):, model.intercept_,"
(Estimated price when size is zero)")
```

Example of Regression (5/8)

5. Make Predictions

- Using the trained model to predict prices for our test apartments.

```
y_pred = model.predict(X_test)
```

Example of Regression (6/8)

6. Evaluate Model Performance

- **Mean Squared Error (MSE) measures how close our predictions are to actual prices.**
- **Lower MSE indicates a more accurate model.**

```
mse = mean_squared_error(y_test, y_pred)
print("Mean Squared Error (MSE): {mse} (Lower is better, indicates prediction accuracy)")
```

Example of Regression (7/8)

7. Visualize the Regression Line

- A plot helps us see how well the line fits the data.

```
plt.figure(figsize=(10, 6))
plt.scatter(X_test, y_test, color='blue', label='Actual Prices')
plt.plot(X_test, y_pred, color='red', linewidth=2, label='Predicted Prices (Regression Line)')
plt.title('Apartment Price Prediction: Actual vs. Predicted')
plt.xlabel('Apartment Size (sq. meters)')
plt.ylabel('Apartment Price (EUR)')
plt.legend()
plt.grid(True)
plt.show()
```

Example of Regression (8/8)

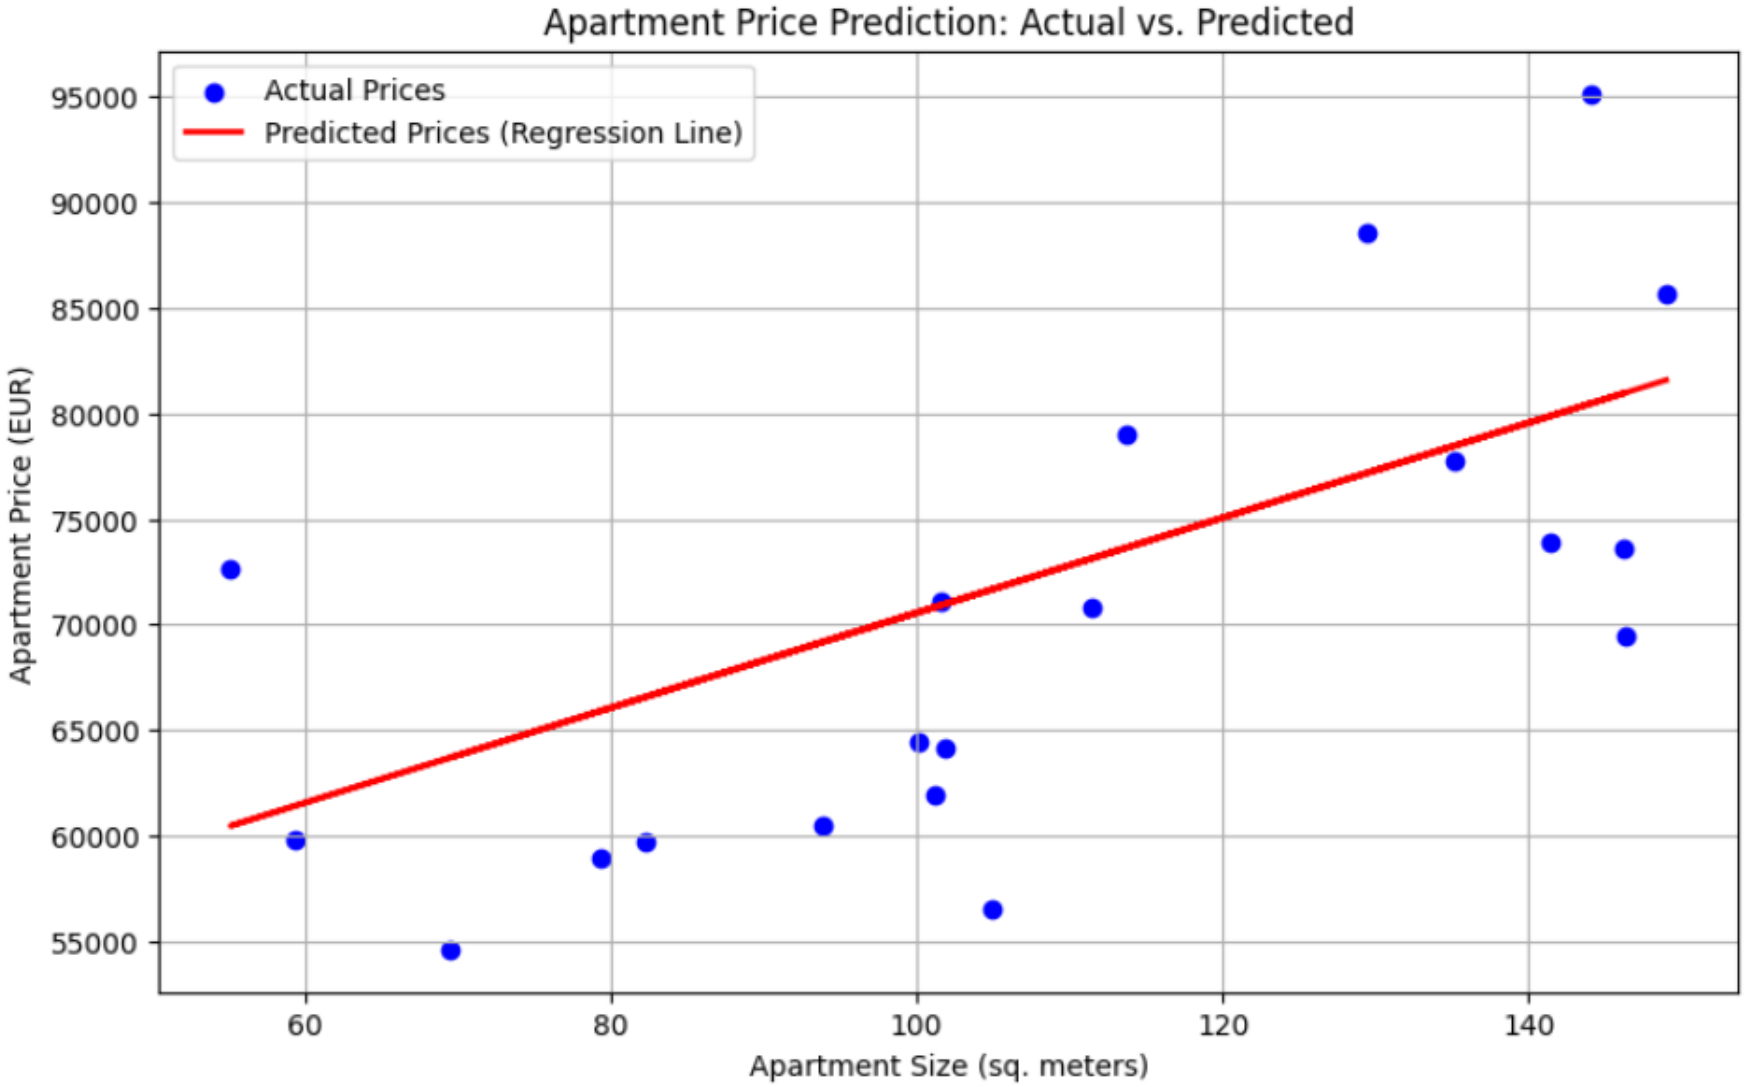

Example of Classification

- Predicting pass/fail based on study hours

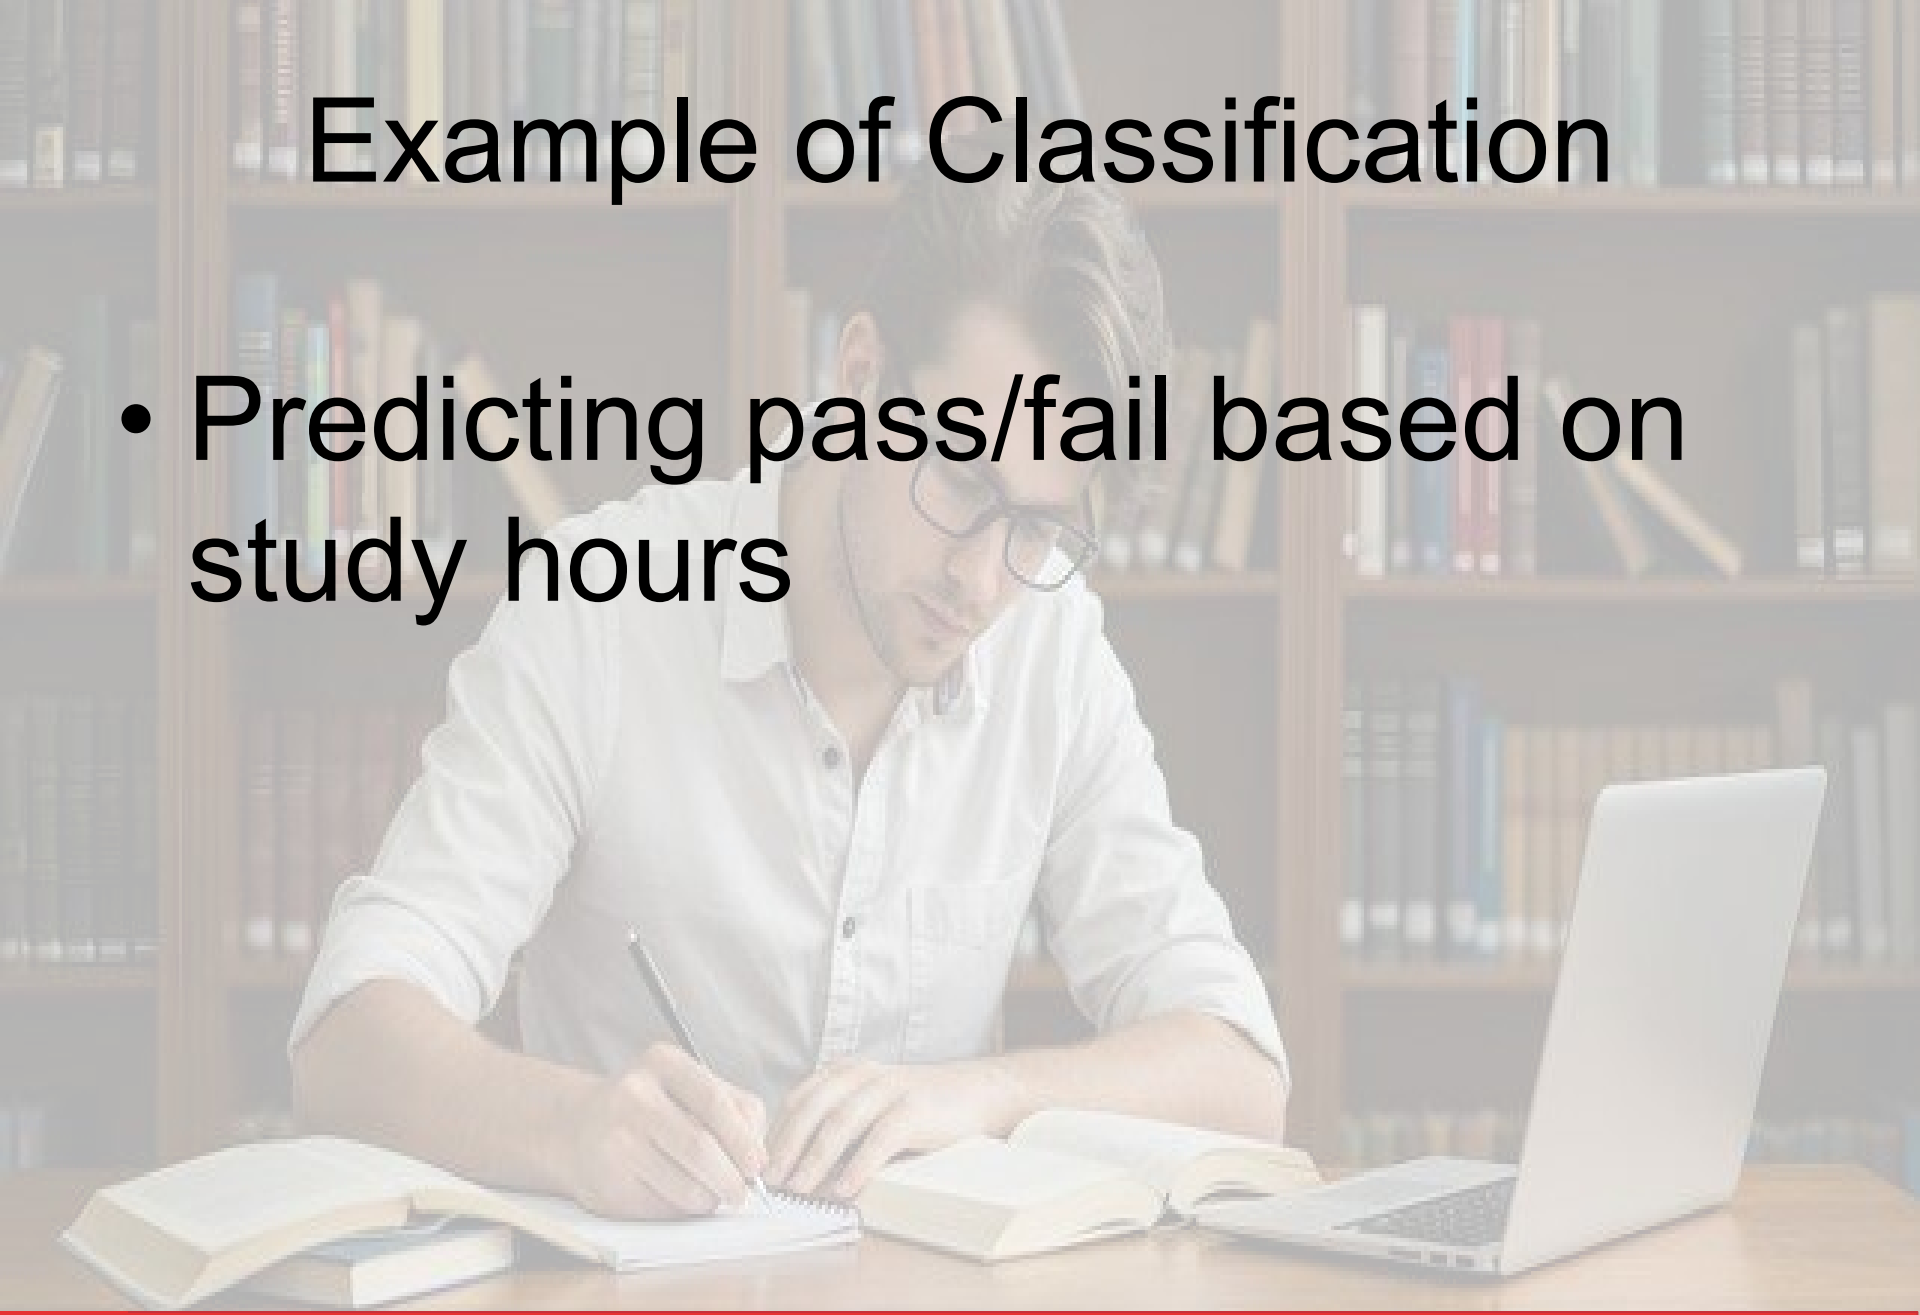

Example of Classification

```
import numpy as np
import pandas as pd
from sklearn.model_selection import train_test_split
from sklearn.linear_model import LogisticRegression
from sklearn.metrics import accuracy_score

hours = np.random.rand(100, 1) * 20
labels = (hours.squeeze() + np.random.randn(100)) > 10

X_train, X_test, y_train, y_test = train_test_split(hours, labels, test_size=0.2)

clf = LogisticRegression().fit(X_train, y_train)
y_pred = clf.predict(X_test)

print("Accuracy:", accuracy_score(y_test, y_pred))
```

Accuracy: 0.95

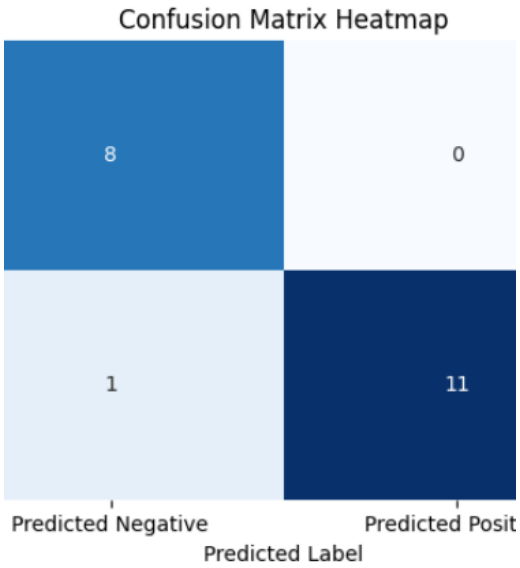

Example of Clustering

```
from sklearn.cluster import KMeans
import matplotlib.pyplot as plt

# Generate synthetic data: 3 clusters
np.random.seed(42)
X = np.vstack([
    np.random.normal([0,0], 0.5, (50,2)),
    np.random.normal([5,5], 0.5, (50,2)),
    np.random.normal([0,5], 0.5, (50,2))
])

# Apply KMeans
kmeans = KMeans(n_clusters=3, random_state=0).fit(X)
labels = kmeans.labels_

# Visualization
plt.scatter(X[:,0], X[:,1], c=labels, cmap="viridis")
plt.scatter(kmeans.cluster_centers_[:,0], kmeans.cluster_centers_[:,1],
            c="red", marker="X", s=200, label="Centers")
plt.title("KMeans Clustering")
plt.legend()
plt.show()
```

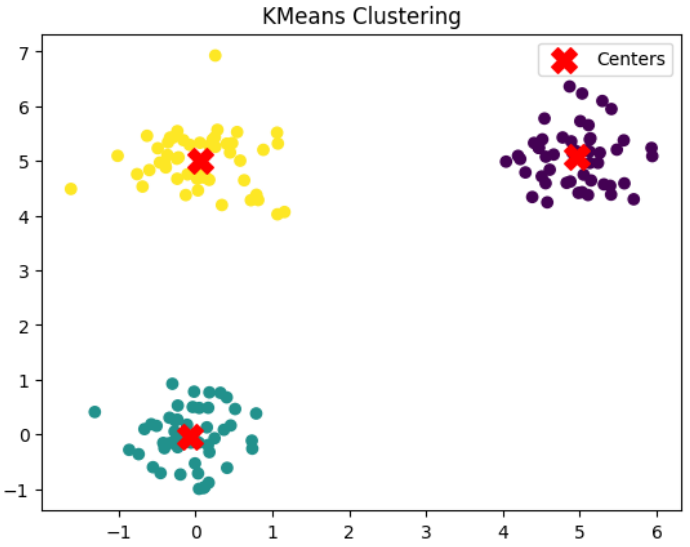

Sentiment Analysis

```
import nltk
from nltk.sentiment import SentimentIntensityAnalyzer

# Download VADER model (only needed once)
nltk.download('vader_lexicon')

# Initialize sentiment analyzer
sia = SentimentIntensityAnalyzer()

# Example texts
texts = [
    "I love this product, it’s fantastic!",
    "This is the worst experience I’ve ever had.",
    "The movie was okay, not great but not terrible either.",
    "Amazing customer support, very satisfied.",
    "Totally disappointed, waste of money."
]

# Analyze each text
for t in texts:
    score = sia.polarity_scores(t)
    sentiment = "Positive" if score['compound'] > 0.05 else ("Negative"
if score['compound'] < -0.05 else "Neutral")
    print(f"Text: {t}\n Score: {score} → Sentiment: {sentiment}\n")
```

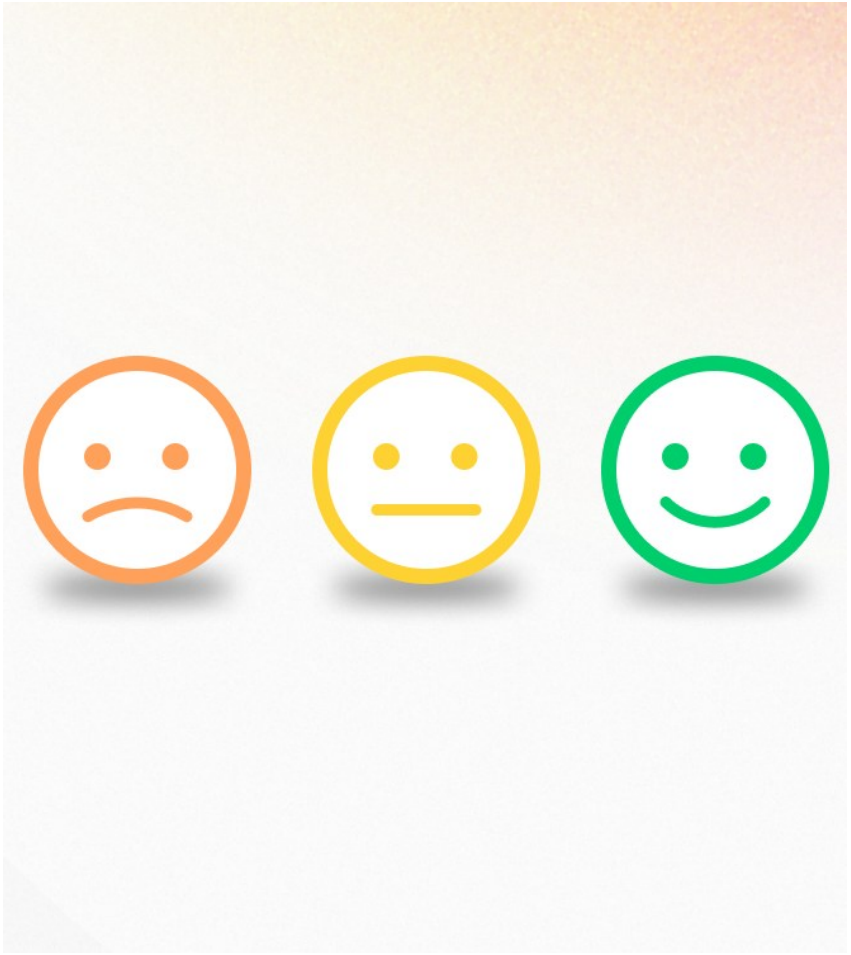

Sentiment Analysis

```
from sklearn.feature_extraction.text import CountVectorizer
from sklearn.model_selection import train_test_split
from sklearn.naive_bayes import MultinomialNB
from sklearn.metrics import accuracy_score

# Example labeled dataset
reviews = [
    "I love this phone", "This product is terrible",
    "Fantastic quality, very happy", "Worst service ever",
    "Not bad, but could be better", "Absolutely wonderful"
]
labels = [1, 0, 1, 0, 0, 1]  # 1 = Positive, 0 = Negative

# Convert text into numerical features (Bag of Words)
vectorizer = CountVectorizer()
X = vectorizer.fit_transform(reviews)

# Split train/test
X_train, X_test, y_train, y_test = train_test_split(X, labels, test_size=0.3, random_state=42)

# Train Naive Bayes classifier
clf = MultinomialNB().fit(X_train, y_train)
y_pred = clf.predict(X_test)

print("Accuracy:", accuracy_score(y_test, y_pred))
print("Predictions:", y_pred)
```

Sentiment Analysis

- How to use?

```
new_text = ["It is so good"]
X_new = vectorizer.transform(new_text)
prediction = clf.predict(X_new)

print("Prediction:", prediction)
```

4. Summary

- **Supervised Learning**: learns from labeled data (prediction, classification).
- **Unsupervised Learning**: finds hidden structures in unlabeled data (clustering, dimensionality reduction).
- Python provides powerful libraries like scikit-learn for easy implementation.