COURSE SYLLABUS

Python Machine Learning Bootcamp

Learn the fundamentals of machine learning, including regression analysis and classification algorithms, in this practical, hands-on course. Gain the skills needed to solve real-world problems using machine learning, with a focus on Python programming and data science libraries.

Group classes (NYC) and onsite training are available for this course. For more information, email corporate@nobledesktop.com or visit: https://www.nobledesktop.com/classes/python-machine-learning



hello@nobledesktop.com • (212) 226-4149

Course Outline

1. Course Kick-off & Python Refresher

- Data Science tool recap - Pandas and indexing
- Exploratory data analysis (EDA): standard deviations and uniform vs. normal distributions using NumPy/Pandas
- Hands-on: loading CSVs, basic plotting with Matplotlib

2. Data Visualization & Simple Linear Regression

- Crafting clear scatterplots: labels, grids, styling
- Single-variable linear regression (attendance → concessions)
- Train-test splitting and dealing with outliers
- Evaluating models with R^2; interpreting residuals
- Extended example: car-sales dataset, predicting price from one feature

3. Binary Classification & Logistic Regression

- From regression to classification: why logistic vs. linear
- Implementing logistic regression on an employee “stay/leave” dataset
- Classification metrics deep dive: accuracy, precision, recall, F1 score, ROC curve
- Understanding variability: train-test ratios, data shuffling, sample size effects
- Confusion matrix analysis

4. k-Nearest Neighbors & the Iris Dataset

- Introduction to k-NN: distance metrics, choosing k
- Dataset exploration: sepal/petal measurements, plotting clusters
- Preprocessing: label encoding categorical data, feature scaling
- Model training, hyperparameter tuning, evaluating with confusion matrix and classification report
- Brief intro to decision-tree logic (setting up for ensembles)

5. Ensemble Methods & Neural Networks

- Random forest classifiers on the Titanic dataset: feature engineering, importance scores
- Kaggle workflow: generating predictions, submitting to competition
- Neural network primer: perceptron to multilayer architectures
- Hands-on MNIST digit classification with Keras/TensorFlow in Colab