COURSE SYLLABUS

Python Machine Learning Bootcamp

Take a step beyond normal programming, into using algorithms that can independently learn patterns and make decisions. Machine learning skills are in high demand, as these algorithms now run the majority of trading on Wall Street and the product recommendations at big companies like Amazon, Spotify, and Netflix.

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

nyc@careercenters.com • (212) 684-5151

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