

AI, ML & Deep Learning: Scikit-learn & Keras CNNs

Learn to build, evaluate and deploy machine learning and deep learning models with Python, Scikit-learn and Keras.

10 Days

60 HOURS

10

MODULES

45

TOPICS

35

HANDS-ON ACTIVITIES

COURSE OVERVIEW

This course begins with setting up a Python data science environment and refreshing statistics, probability and visualization concepts. It then covers core machine learning algorithms including regression, Bayesian methods, clustering, KNN and dimensionality reduction, along with model evaluation, validation and handling real-world data issues such as bias, outliers and normalization. Experimental design through A/B testing and statistical testing is introduced before moving into deep learning, covering neural network foundations, TensorFlow and Keras, and architectures such as CNNs, RNNs, LSTMs and GRUs. The course closes with a look at fairness, bias, privacy and accountability in AI systems.

WHO THIS COURSE IS FOR

Data Analysts

Machine Learning Engineers

Data Scientists

Python Developers

AI Researchers

LEARNING OBJECTIVES

- Set up a Python data science environment with Jupyter Notebooks
- Calculate statistical measures and probability distributions with Python
- Apply data visualization techniques using Matplotlib
- Build predictive models such as linear and polynomial regression
- Apply clustering, KNN and PCA to sample datasets
- Evaluate classification and regression models using key metrics
- Identify bias, variance and outliers in real-world data
- Build neural networks and CNNs using TensorFlow and Keras

10 Modules, 45 Topics

01

6 hours

4 Hands-on Labs

Getting Started

TOPICS

- Installing a Python Data Science Environment
- Using and understanding iPython (Jupyter) Notebooks
- Python basics (Part 1 & 2)
- Understanding Python code & importing modules
- Running Python scripts

02

6 hours

3 Hands-on Labs

Statistics and Probability Refresher with Python

TOPICS

- Types of data
- Mean, median, mode, standard deviation, variance
- Probability density & mass functions
- Types of distributions, percentiles, moments

03

3 hours

2 Hands-on Labs

Visualization & Advanced Probability

TOPICS

Crash course in Matplotlib · Covariance and correlation · Conditional probability · Bayes' theorem

04

14 hours

9 Hands-on Labs

Core Machine Learning

TOPICS

- | | |
|---|---|
| ■ Data preparation techniques | ■ Bayesian methods and Naïve Bayes classifier |
| ■ Linear & non-linear algorithms | ■ K-Means clustering & recommender systems |
| ■ Ensemble methods | ■ KNN concepts & applications |
| ■ Predictive models (linear, polynomial, multivariate regression) | ■ Dimensionality reduction (PCA) |
| | ■ Reinforcement learning basics |

05

4 hours

4 Hands-on Labs

Model Evaluation & Validation

TOPICS

- Train/Test split, Cross-validation
- Overfitting vs Underfitting
- **Metrics:** Accuracy, Precision, Recall, F1, ROC-AUC
- Evaluate classification & regression models

06

4 hours

4 Hands-on Labs

Dealing with Real-World Data

TOPICS

Bias/variance trade-off · K-fold cross-validation · Data cleaning & normalization · Outlier detection

07

3 hours

Testing & Experimental Design

TOPICS

A/B testing concepts · T-test and p-value · Experiment duration & common pitfalls

08

11 hours

4 Hands-on Labs

Deep Learning Foundations

TOPICS

- Introduction to Neural Networks & Perceptrons
- Backpropagation & activation functions
- Introduction to TensorFlow & Keras
- Building and training a feedforward neural network

09

7 hours

2 Hands-on Labs

Deep Learning Architectures

TOPICS

- Convolutional Neural Networks (CNNs) - image recognition
- Recurrent Neural Networks (RNNs) - sequence & time-series
- LSTMs & GRUs
- Transfer Learning basics

AI Ethics & Responsible AI

TOPICS

Fairness, bias & explainability in AI · Privacy & data governance issues · Ethical dilemmas in AI applications · Transparency & accountability in ML/DL models

NOTES

- Azure Machine Learning Complete end to end project (Optional if access is available)

HANDS-ON ACTIVITIES

Practise What You Learn

#	MODULE	ACTIVITY
01	Getting Started	■ Install Anaconda/Miniconda & set up Jupyter Notebooks ■ Write first Python program & import external modules (NumPy, Pandas, Matplotlib) ■ Explore Python data structures (lists, dictionaries, tuples, sets) ■ Run and execute Python scripts in terminal/VS Code
02	Statistics and Probability Refresher with Python	■ Use Pandas to calculate mean, median, mode, variance & standard deviation on a dataset ■ Plot different probability distributions (normal, binomial, Poisson) with Matplotlib ■ Generate synthetic data and calculate percentiles & quantiles
03	Visualization & Advanced Probability	■ Create histograms, bar plots, and scatter plots in Matplotlib ■ Calculate correlation & covariance using Pandas and Excel
04	Core Machine Learning	■ Data preprocessing (handling missing values, scaling, encoding) ■ Excel Lab: Build a simple predictive linear regression model in Excel using the Data Analysis ToolPak ■ Build a linear regression model in Python ■ Train a polynomial regression model & visualize predictions ■ Implement a Naïve Bayes classifier on a text dataset (spam classification) ■ Apply K-Means clustering to group customers by purchasing behavior ■ Build a simple user-item recommender system using cosine similarity ■ Implement a KNN classifier for image recognition (MNIST subset) ■ Apply PCA on a dataset to reduce dimensions & visualize clusters
05	Model Evaluation & Validation	■ Split dataset into train/test sets using Scikit-learn ■ Perform k-fold cross-validation ■ Plot ROC curves & calculate AUC scores ■ Compare performance metrics across multiple models

#	MODULE	ACTIVITY
06	Dealing with Real-World Data	■ Handle missing values (imputation techniques) ■ Normalize and scale features using MinMaxScaler & StandardScaler ■ Detect outliers using plots ■ Visualize impact of bias/variance trade-off on model performance
08	Deep Learning Foundations	■ Build a single-layer perceptron in NumPy ■ Visualize activation functions (sigmoid, tanh, ReLU) ■ Implement a feedforward neural network in Keras for MNIST classification ■ Train and evaluate network performance with TensorFlow/Keras
09	Deep Learning Architectures	■ Build a CNN in Keras for image classification (CIFAR-10 dataset) ■ Implement a simple RNN for text sequence prediction
10	AI Ethics & Responsible AI	■ Detect bias in a dataset (e.g., gender bias in hiring dataset) ■ Use LIME/SHAP for explainability of ML predictions ■ Discuss case studies of ethical dilemmas in AI

WHAT IS INCLUDED

Everything You Need to Succeed

- Koenig course material for every module
- Delivery by certified instructors and subject matter experts
- Hands-on exercises throughout the course
- Course completion certificate from Koenig

Who We Are

Established in 1993, Koenig Solutions is a globally recognised training provider with 35+ technology partnerships, 5,000+ training programmes and a network of certified instructors and subject matter experts, supporting individuals and enterprise teams worldwide.

Learn the Way That Suits You

Classroom Training

Live Online Classes (ILO)

Fly-Me-A-Trainer (FMAT)

Flexi (Self-Paced)

1-on-1 Training

Ready to Get Started?

Talk to a learning advisor about schedules, group training and customised delivery.

Visit koenig-solutions.com

✉ info@koenig-solutions.com