



MACHINE LEARNING

REGRESSION . CLASSIFICATION
FEATURE ENGINEERING . MODEL
KIT-LEARN . EVALUATION . DEPLOYMENT

DURATION

70 Hours

LEVEL

**Beginner to
Advanced**

MODULES

9

MODE

**Classroom ·
Online · Hybrid**

01 COURSE OVERVIEW

A practical machine learning course built on Scikit-Learn and real business datasets.

The course covers the supervised and unsupervised toolkit, the mathematics needed to reason about it, and the discipline around it: data splitting, leakage, cross-validation, metric choice and hyperparameter search.

Learners complete three predictive projects — churn, price and spam — and deploy one as an API.

WHO IT IS FOR

Graduates entering data roles

Analysts moving from reporting to prediction

Developers adding ML to products

Students in engineering and statistics

PREREQUISITES

Basic Python and school-level mathematics. NumPy and Pandas are revised in Module 2.

02 LEARNING OUTCOMES

01 Frame a business question as a learning problem.

03 Train regression, classification and clustering models.

05 Tune models with grid and randomised search.

07 Serve a trained model behind an API.

02 Prepare, encode and scale features without leakage.

04 Select metrics that match the business cost of error.

06 Explain model behaviour with feature importance and SHAP.

03 MODULE STRUCTURE

9 MODULES

01 FOUNDATIONS

What machine learning is

Reinforcement learning overview

Bias and variance

Supervised vs unsupervised

Train, validation and test

Overfitting

02 PYTHON FOR ML

NumPy

Matplotlib

Data loading

Pandas

Seaborn

Vectorised operations

LAB

DATASET EXPLORATION

03 DATA PREPARATION

Missing values

Encoding

Feature engineering

Data leakage

Outliers

Scaling

Pipelines

LAB

PREPROCESSING PIPELINE

04 REGRESSION

Linear regression

Regularisation

Residual analysis

Polynomial regression

Loss functions

Metrics

LAB

HOUSE PRICE PREDICTION

05 CLASSIFICATION

Logistic regression

Naive Bayes

Support vector machines

K-nearest neighbours

Decision trees

Class imbalance

LAB

SPAM DETECTION

06 ENSEMBLES

Bagging

Boosting

LightGBM

Random forest

XGBoost

Stacking

LAB

CUSTOMER CHURN PREDICTION

07 UNSUPERVISED LEARNING

K-means

DBSCAN

Anomaly detection

Hierarchical clustering

Principal component analysis

LAB

CUSTOMER SEGMENTATION

08 EVALUATION AND TUNING

Cross-validation

ROC and AUC

Grid and random search

Confusion matrix

Precision-recall

Model selection

09 INTERPRETABILITY AND DEPLOYMENT

Feature importance

Model cards

FastAPI serving

SHAP

Pickle and joblib

Monitoring drift

PROJECT

DEPLOYED PREDICTION SERVICE

04 ASSESSMENT & CERTIFICATION

Module assignments and labs

25%

Mini projects

20%

Internal assessments

15%

Final project and review

40%

CERTIFICATION

Learners who complete all modules, submit the final project and clear the review receive a course completion certificate from Kaizen Infinities Private Limited. Project work is documented for the learner's portfolio, and interview preparation is included in the closing sessions.

05 CAREER OUTCOMES

ROLES THIS PROGRAMME PREPARES FOR

MACHINE LEARNING ENGINEER

DATA SCIENTIST

PREDICTIVE ANALYTICS ANALYST

PYTHON DEVELOPER (ML)

AI ENGINEER

CALL US

+91 73051 33996

EMAIL US

apply@kaizeninfinities.com

VISIT OUR WEBSITE

www.kaizeninfinities.com
