



Digital Transformation and Innovation

Introduction to Machine Learning and Artificial Intelligence

Course Introduction

Artificial Intelligence (AI) and Machine Learning (ML) are currently changing the game for the industries. Many businesses take up artificial intelligence (AI) technology to try to reduce operational costs, increase efficiency, grow revenue and improve customer experience. For greatest benefits, businesses should look at putting the full range of smart technologies - including machine learning, natural language processing and more - into their processes and products. However, even businesses that are new to AI can reap major rewards.

This training course is designed to provide participants with the fundamental concepts of AI and machine learning and allow them to have an in-depth understanding of its benefits and uses in the organisation.

Target Audience

- Business and technology leaders
- Business Unit Managers
- Business Development Consultants
- General Managers / Regional Managers
- Senior and mid-level leaders
- individual leaders of all levels in the organization
- Art Director
- Marketing Consultants
- Marketing Development Manager

Learning Objectives

- The basic concept and ideas of machine learning and artificial intelligence.

- Python and Jupyter
- Statistics and Probability Refresher and Python Practice
- Matplotlib and Advanced Probability Concepts
- Algorithm Overview
- Predictive Models
- Applied Machine Learning
- Recommender Systems
- Dealing with Data in the Real World
- Machine Learning on Big Data (with Apache Spark)
- Testing and Experimental Design
- GUIs and REST: Build a UI and REST API for your Models

Course Outline

- **Day 01**

The Business of AI Adoption

- Defining artificial intelligence
- The uncertainties of new technology
- AI in the field
- The problem of trust
- Work is evolving
- Driverless transportation
- Trust and the machine
- The human-smart machine trust gap
- Trusting a smart machine
- Trusting the smart machine developer

Artificial Intelligence (AI)

- Artificial Intelligence and its application Areas
- AI and Machine learning
- Pattern recognition
- Supervised and unsupervised learning,
- Structured and unstructured data
- Discussion on Industry, industrial analytics

- **Day 02**

Machine Learning

- Basic ideas of machine learning
- Bias-variance complexity trade-off
- Model types
- Deep neural network
- Recurrent neural network or long short-term memory network
- Support vector machines
- Random forest or decision trees
- Self-organizing maps (SOM)
- Bayesian network and ontology
- Training and assessing a model
- How good is my model?
- Role of domain knowledge
- Optimization using a model

Getting Started with Python

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

Statistics and Probability Refresher and Python Practice

- Types of data
- Mean, median, and mode
- Using mean, median, and mode in Python
- Standard deviation and variance
- Probability density function and probability mass function
- Types of data distributions
- Percentiles and moments

• Day 03

Matplotlib and Advanced Probability Concepts

- A crash course in Matplotlib
- Covariance and correlation
- Conditional probability
- Bayes' theorem

Algorithm Overview

- Data Prep
- Linear Algorithms
- Non-Linear Algorithms
- Ensembles

Predictive Models

- Linear regression
- Polynomial regression
- Multivariate regression and predicting car prices
- Multi-level models

Applied Machine Learning with Python

- Machine learning and train/test
- Using train/test to prevent overfitting of a polynomial regression
- Bayesian methods: Concepts
- Implementing a spam classifier with Naïve Bayes
- K-Means clustering

• Day 04

Recommender Systems

- What are recommender systems?
- Item-based collaborative filtering
- How item-based collaborative filtering works?
- Finding movie similarities
- Improving the results of movie similarities
- Making movie recommendations to people
- Improving the recommendation results

More Applied Machine Learning Techniques

- K-nearest neighbors - concepts
- Using KNN to predict a rating for a movie
- Dimensionality reduction and principal component analysis
- A PCA example with the Iris dataset
- Data warehousing overview
- Reinforcement learning

Dealing with Data in the Real World

- Bias/variance trade-off
- K-fold cross-validation to avoid overfitting
- Data cleaning and normalization
- Cleaning web log data
- Normalizing numerical data
- Detecting outliers

• Day 05

Apache Spark: Machine Learning on Big Data

- Installing Spark
- Spark introduction
- Spark and Resilient Distributed Datasets (RDD)
- Introducing MLlib
- Decision Trees in Spark with MLlib
- K-Means Clustering in Spark
- TF-IDF
- Searching wikipedia with Spark MLlib
- Using the Spark 2.0 DataFrame API for MLlib

Testing and Experimental Design

- A/B testing concepts
- T-test and p-value
- Measuring t-statistics and p-values using Python
- Determining how long to run an experiment for
- A/B test gotchas

GUIs and REST

- Build a UI for your Models
- Build a REST API for your Models

Confirmed Sessions

FROM	TO	DURATION	FEES	LOCATION
April 7, 2025	April 11, 2025	5 days	4950.00 \$	England - London
June 23, 2025	June 27, 2025	5 days	2150.00 \$	Virtual - Online
Sept. 1, 2025	Sept. 5, 2025	5 days	4250.00 \$	UAE - Abu Dhabi
Dec. 8, 2025	Dec. 12, 2025	5 days	4950.00 \$	Austria - Vienna