



Hewlett Packard Enterprise

Course Datasheet

Machine Learning & Deep Learning

Education Services course product number – HPE-MLDL-v1.0

Course length – 90 Hrs.

Delivery mode – Instructor Led Training (ILT)

Virtual Instructor Led Training (vILT)

Machine Learning and Deep Learning are so persistent in today's world that one may probably be using it more than 10 times a day without knowing it. Facebook tag suggestions is one of the very common example linked to it. One can think of Artificial Intelligence (AI), Machine Learning and Deep Learning as a set of items nested within each other. Deep Learning is a subset of Machine Learning and Machine Learning is a subset of AI which is simply the science and engineering of making intelligent machines. Machine Learning uses algorithms to analyse data, learn from the data, and then make informed decision accordingly. However, Deep learning structures algorithms in layers and high-level representations to create an 'Artificial Neural Network' that can learn, understand and make intelligent decisions. As the term specifies, Deep Learning is stimulated by the human brain and is based on 'Artificial Neural Network' which is developed on a similar architecture of human brain.

Course Objective

This comprehensive 90 hrs course is designed to help learner master the concepts of Machine Learning and Deep Learning in a structured manner. As part of this course, a learner will be made to go through concepts like Machine Learning Algorithms, Statistical Concepts, Modelling Techniques, Regression, Classification, Clustering, Reinforcement Learning and R Programming as part of Machine Learning section. To have a deeper understanding, Deep learning using TensorFlow will help learner master the deep learning techniques and build deep learning models using TensorFlow covering a range of topics like Basic Neural Network to Convolutional and Recurrent Neural Network.

Prerequisite

Participants should have understanding on the fundamentals of Artificial Intelligence, Basic knowledge of Statistical concepts, and Python Programming.

Course Modules

Chapter 01 - Fundamentals of Machine Learning

- What is Machine Learning
- History of Machine Learning
- Traditional Programming vs Machine Learning
- Where and Why Machine Learning is Used
- Challenges of Machine Learning
- Machine Learning Languages
- Elements of Machine Learning
- Machine Learning in Practice

Chapter 02 - Types of Machine Learning Algorithms

- Ensemble Learning Method
- Supervised Learning
- Unsupervised Learning
- Semi-supervised Learning and Reinforcement Learning

Chapter 03 - Planning for Machine Learning

- The Machine Learning Cycle
- One Solution Fits All?
- Defining the Process
- Data Preparation
- Datasets
- Data Processing
- Data Storage, Data Privacy, and Data Quality
- Thinking about Input and Output Data

Chapter 04 - Mathematical and Statistical Concepts – A Quick Recap

- Concepts of Linear Algebra
- Mean, Median and Mode
- Sampling Techniques
- Probability and Statistics
- Calculus
- ANOVA

Chapter 05 - Regression

- Regression and its types
- Linear Regression
 - Univariate Regression
 - Multi Variate Regression
 - Cost Function
 - Gradient Descent
 - Gradient Descent Intuition
 - Gradient Descent For Linear Regression
- Regression Modelling
 - Linear Regression Model
 - Logistic Regression Model
 - The Problem of Overfitting

Course Datasheet

Chapter 06 - Tree Modelling

- Introduction to Tree Modelling
- Decision Tree
 - What is a Decision Tree and How it works
 - Advantages and Limitations of Decision Trees
 - Types of Decision Tree Algorithms
 - Types of Decision Trees
 - Univariate Trees and Multivariate Trees
 - Pruning
 - Gini Index and Chi-Square
 - Learning Rules from Data
- Random Forest
 - Working of Random Forest
 - Advantages and Disadvantages of Random Forest
- Boosting Models

Chapter 07 - Classification

- Classification
- Logistic Regression
- K-nearest neighbours
- Naïve Bayes
- Support Vector Machines
- Time Series Models

Chapter 08 - Clustering

- What is clustering?
- k-means Clustering
- Expectation-Maximization Algorithm
- Hierarchical Clustering
- Spectral Clustering
- Linkage Based Clustering Algorithms

Chapter 09 - Reinforcement Learning

- Introduction
- Elements of Reinforcement Learning
- Model-Based Learning
- Temporal Difference Learning

Chapter 10 - Machine Learning with R

- Installing R and R-Studio
- The R Basics
 - Variables and Vectors
 - Lists and Matrices
 - Data Frames
 - Installing Packages
- Simple Statistics
- Simple Linear Regression
- Apriori Association Rules
- Plotting and Visualization Techniques

Chapter 11 - Deep Learning

- Meaning and Importance of Deep Learning
- Difference Between Machine Learning, Deep Learning and Artificial Intelligence
- Deep Learning Primitives
- Deep Learning Frameworks

Chapter 12 - Going Deep Using TensorFlow

- Introduction to TensorFlow
- Installing TensorFlow
- Understanding TensorFlow Basics
- Computation Graphs
 - What is a Computation Graph
 - Benefits of Computation Graph
- Graphs, Sessions and Fetches
 - Creating a Graph and running it in a Session
 - Managing Graphs
 - a. Fetches
- Flowing Tensors
 - Nodes and Edges
 - Data Types
 - Arrays and Shapes
 - Names
- Sharing Variables, Placeholders and Optimization
- TensorBoard

Chapter 13 - Artificial Neural Networks (ANN)

- What is a Neural Network and its uses & applications
- Advantages of ANN and associated risks
- Breaking down the Artificial Neural Network
- Non-linear Hypothesis
- Neurons and the Brain
- Examples and Intuitions
- Multiclass Classification
- Backpropagation Algorithm
- Backpropagation Intuition
- Gradient Checking
- Random Initialization

Chapter 14 - Convolutional Neural Networks (CNN)

- Introduction to CNN
- Kernel Filter
- Multiple Filters
- CNN Applications

Chapter 15 - Recurrent Neural Networks (RNN)

- Introduction to RNN
- The importance of sequence Data
- RNN for Text Sequences

Course Datasheet

- LSTM
- RNN Applications

Chapter 16 - Distributed TensorFlow

- Distributed Computing
- TensorFlow Elements
 - Clusters and Servers
 - Replicating a Computational Graph Across Devices

Chapter 17 - Deep Learning Applications

- Image Processing
- Natural Language Processing
- Speech Recognition