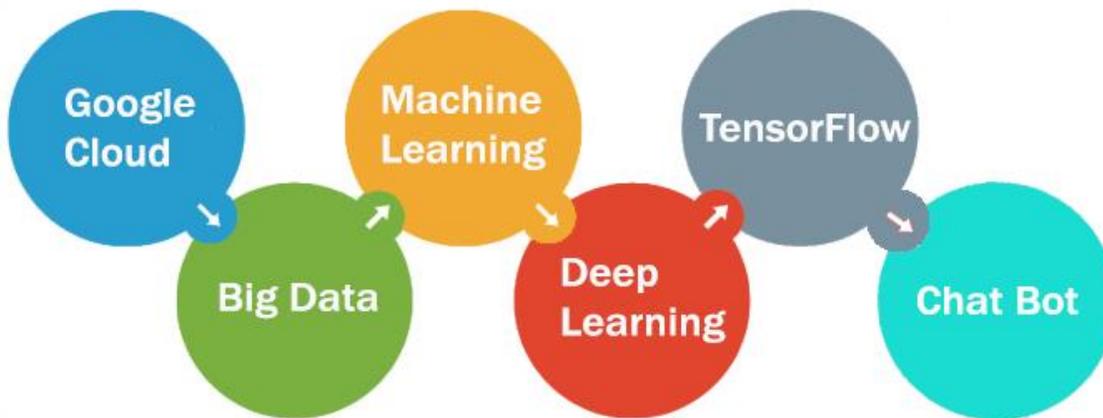


**Module Title** : Machine Learning and Artificial Intelligence

**Duration** : 5 days

## Course Structure

This course focuses on the practical aspects of Machine Learning, Deep Learning and Artificial Intelligence. The objective is to make use of TensorFlow for various types of neural networks. The participants will build and train deep learning models.



## Intended Audience

The intended audience for this course:

- Data Engineers
- Data Scientists
- Machine Learning Engineers
- Integration Engineers
- Architects

## Prerequisites and Lab details

Participants should preferably have some hands-on experience on Python

## Course Objectives

### Understanding the Big Picture

- Artificial Intelligence (AI) Overview

- What is Machine Learning (ML)?
- AI vs ML vs Data Science
- Relationship between Deep Learning (DL) and Machine Learning
- Practical Use cases
- Concepts and Terms
- Tools/Platforms for ML, DL and AI
- Machine Learning Project End to End Pipeline
- Scalable ML/AI: Big Data and Cloud fits into the Ecosystem

#### **GCP Introduction**

- GCP Introduction.
- Why Google Cloud Platform (GCP)?
- How Innovations at Google driving Data Engineering and Science globally?
- Key Google Products related to Data and Machine Learning
- Come on same page w.r.t. terms and concepts

#### **Getting holistic view: Architectures and Pipelines**

- Introduction to Serverless Architecture
- Current Challenges with On-Premise Architectures
- How Google enables higher productivity?
- How Key Google Products fit in Enterprise Architecture?
- How to design modern Data Analytics Pipeline on GCP?

#### **Machine Learning APIs**

- Introduction to Machine Learning APIs
- Key ML Use Cases
- Vision API
- Natural Language API
- Translate API
- Speech API

#### **Environment for Experiments**

- Installing Anaconda
- Setting up Jupyter Notebook
- Experiencing Notebooks
- Key Python Syntax Recap
- Hands-on Exercises

#### **Key Statistics**

- Summary Statistics
- Exploratory Data Analysis
- Numerical Computation using Python
- Hands-on Exercises

#### **Data Visualization**

- Overview
- Using MatPlot Lib
- Working with Seaborn
- Key types of plots
- Exploratory Analysis using Seaborn
- Hands-on Exercises

#### **Acquiring & Preparing Data**

- Content Acquisition Overview
- Working with Beautiful Soup
- Acquiring data using Rest Based APIs
- Data Cleaning & Wrangling using Pandas
- Missing Values and Outlier
- Cleansing Twitter Data
- Performing Twitter Sentiment Analysis
- Hands-on Exercises

#### **Feature Engineering**

- What is Feature Engineering?
- Why Feature Engineering?
- How to apply Feature Engineering?
- Discussions on various scenarios
- Hands-on Exercises

#### **Machine Learning using Scikit Learn**

- Types of Machine Learning
- Key Algorithms in Machine Learning
- Practical Applications of Machine Learning
- Various frameworks/Libraries popular for ML
- Concepts and Terms
- Why Scikit Learn?
- Code Walkthrough

- Hands-on Exercises

### **Supervised Machine Learning**

- Key Classification Algorithms
- Naïve Bayes Classifier
- Confusion Matrix
- Accuracy
- Key Regression Algorithms
- Linear, Logistic Regressions
- Gradient Descent
- Loss function
- Bias vs Variance Tradeoff
- Evaluating Models
- Hands-on Exercises

### **Un-Supervised Machine Learning**

- Why Un-supervised learning is important?
- Where it can be applied?
- Principal Component Analysis
- Performing Clustering of data
- Hands-on Exercises

### **Spark Introduction**

- Overview
- Understand What is Spark?
- Why Spark?
- Languages used in Spark Programming
- Logical Architecture
- Physical Architecture
- Pros & Cons of using Python, Java & Scala
- Key Terms & Concepts
- Ways to create RDDs
- Operations on RDD
- Pair RDD
- Key Spark Components
- Working with various types of data
- Hands-on exercises

### **Spark SQL**

- Introduction
- Using the SQL API – sqlContext.sql
- Dataframe API
- Working with various Datasources
- Inferred Schemas
- Querying DataFrames Using Column Expressions
- Data formats
- Hands-on Exercises

### **Spark Machine Learning**

- Quick introduction to Machine Learning
- Introduction to ML and MLLib Spark
- Data types - Vectors, Matrices, LabeledPoint
- Summary Statistics
- Calculating Correlations
- Transformers, Estimators and Pipelines
- Evaluation of Model
- Walkthrough of a Regression model
- Walkthrough of a K Means Clustering model
- Hands-on on Working with Machine Learning Pipelines

### **Introduction to Deep Learning**

- What is Deep Learning?
- Relationship between Deep Learning and Machine Learning
- Deep Learning Use cases
- Concepts and Terms
- How to implement Deep Learning?
- Various Libraries, Pros & Cons
- Hands-on: Recap on Machine Learning

### **Artificial Neural Network**

- Introduction to Neural Networks
- Introduction to Perceptron
- Neural Network Activation Functions
- Basic Neural Nets
- Concepts

### **TensorFlow API**

- What is TensorFlow?
- Installing TensorFlow
- TensorFlow Graphs
- Variables and Placeholders
- Activation Functions
- Hands-on exercises

### **Convolutional Neural Networks (CNN)**

- CNN History
- Understanding CNNs
- CNN Application
- Hands-on exercises

### **Interactive Data exploration using DataLab**

- Spinning up Cloud DataLab
- Experiencing Datalab notebooks
- Exploring Google Cloud Storage
- Leveraging relational data with Google Cloud SQL
- Reading and writing streaming Data with Google BigTable
- Querying Data from Google BigQuery
- Making Google API Calls from notebooks

### **CloudML: Scalable Models on GCP**

- Why Cloud ML?
- Running TensorFlow model in Local mode
- Porting TensorFlow models to GCP
- Deploying Models in Production
- Model Predictions
- Hands-on exercise(s)

### **Conversational AI**

- Intro to ChatBots
- Key options available
- Building ChatBot
- Hands-on Exercise using DialogFlow API