INDIA's Finest

JOB GUARANTEED

FullStack PYTHON

MACHINE LEARNING & AI

For Fresh Graduates

&

For Experienced Professionals

By Our Lead Faculty

Sri.Khaja

Sr. Cloud Architect

15+ Years Industry Experience

10,000 + Students & Professional Trained

Quality Thought®

/qttworld

Data Science, Artificial Intelligence, Machine Learning, NLP and Computer Vision

QualityThought®

/qttworld

TensorFlow

Amazon
Lex

Amazon SageMaker

Amazon Polly

QualityThought®

/qttworld

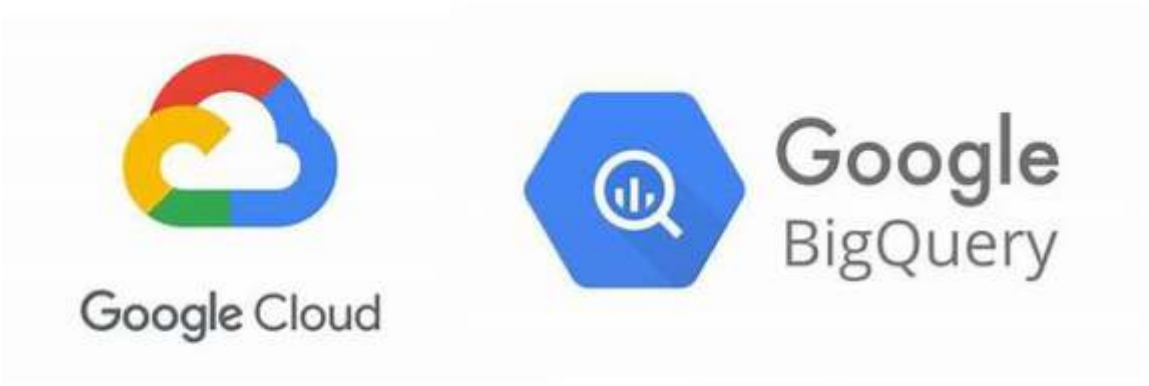

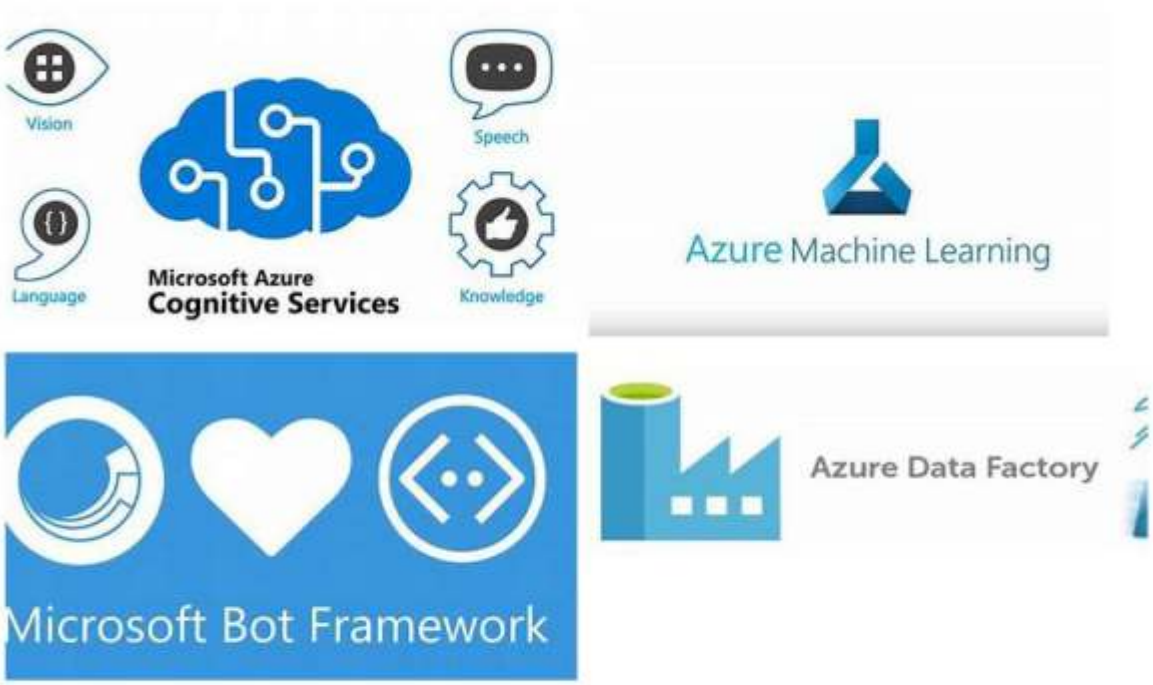

Introduction to Data Science

- Roles in Data Science Project
- Stages of data science project
  - Defining the goal
  - Data collection and management
  - Modelling
  - Model evaluation and critique
  - Presentation and documentation
  - Model Deployment and maintenance

Introduction to Big Data

Introduction to a Hadoop

Introduction to PySpark

Linear Algebra

- What is matrix?
- Scalars and Vectors
- Linear Algebra and geometry
- What is tensor
- Operations on Matrices
- Transpose
- Dot Product of vectors and matrices
- Why are matrices useful?

Statistics-Overview

Sample and population Data

- Fundamentals of descriptive statistics
- Measures of Central Tendency, asymmetry, and variability
- Distributions
- Estimator and estimates
- Confidence Intervals
- Inferential Statistics
- Hypothesis Testing
- Fundamentals of regression analysis
- Subtleties of regression analysis
- Assumptions of linear regression analysis
- Dealing with Categorical Data

Statistics-For Data Science and Machine Learning

Exploratory Data Analysis

- Elements of Structured Data
- Rectangular Data
- Estimates of Location
- Estimates of Variability
- Exploring the Data Distribution
- Exploring Binary and Categorical Data
- Correlation
- Exploring Two or More Variables

Data and Sampling Distributions

- Random Sampling and Sample Bias
- Selection Bias
- Sampling Distribution of a Statistic
- The Bootstrap

- Confidence Intervals
- Normal Distribution
- Long-Tailed Distribution
- Student's t-Distribution
- Binomial Distribution
- Chi-Square Distribution
- F-Distribution
- Poisson and Related Distributions

Statistical Experiments and Significance Testing

- A/B Testing
- Hypothesis Tests
- Resampling
- Statistical Significance and p-Values
- t-Tests
- Multiple Testing
- Degrees of Freedom
- ANOVA
- Chi-Square Test
- Multi-Arm Bandit Algorithm
- Power and Sample Size

Regression and Prediction

- Simple Linear Regression
- Multiple Linear Regression
- Prediction using Regression.
- Factor Variables in Regression
- Interpreting the Regression Equation
- Regression Diagnostics
- Polynomial and Spline Regression

Classification

- Naïve Bayes
- Discriminant Analysis
- Logistic Regression
- Evaluation Classification Models
- Strategies for Imbalanced Data

Statistical Machine Learning

- K-Nearest Neighbours
- Tree Models
- Bagging and Random Forest

- Boosting

- Statistical Unsupervised Learning
  - Principal Components Analysis
  - K-Means Clustering
  - Hierarchical Clustering
  - Model Based Clustering
  - Scaling and Categorical Variables

R Programming for Statistics and Data Science

- Introduction
  - Downloading and Installing R & RStudio
  - Quick guide to the RStudio user interface
  - Installing packages in R and using the library
- Building Blocks of R
  - Creating an object in R
  - Data types in R
  - Coercion rules in R
  - Functions in R
  - Functions & Arguments
- Vector and Vector Operations
  - Introduction to Vectors
  - Vector recycling
  - Naming a vector in R
  - Slicing and indexing a vector in R
  - Changing the dimensions of an object in R
- Matrices
  - Creating a matrix in R
  - Indexing an element from matrix
  - Slicing a matrix in R
  - Matrix arithmetic
  - Matrix Operations
  - Categorical data
  - Creating a factor in R
  - Lists in R
- Fundamentals of Programming in R
  - Relational Operations in R
  - Logical Operations in R

- Vectors and logical operators
- Conditional statements and looping constructs
- Functions in R

Data frames

- Creating a data frame in R
- The Tidyverse package
- Data Import in R
- Importing a CSV in R
- Data export in R
- Getting a sense of your data frame
- Indexing and slicing a data frame in R
- Extending a data frame in R
- Dealing with missed data in R

Manipulating data

- Data transformation with R – the Dplyr package
- Sampling data
- Using the pipe operator
- Tidying the data in R – gather() and separate()
- Tidying the data in R – unite() and spread()

Visualizing data

- Introduction to data visualization
- Intro to ggplot2
- Variables
- Building a histogram
- Building a bar chart
- Building a box and whiskers plot
- Building a scatterplot

Exploratory data analysis

- Population vs sample
- Mean, Median and mode
- Skewness
- Variance, standard deviation, and coefficient of variability
- Covariance and correlation

Hypothesis Testing

- Distributions
- Standard Error and Confidence Intervals
- Hypothesis testing
- Type I and Type II errors

- Test for the mean
- The P-value
- Comparing two means – Dependent samples
- Comparing two means – Independent samples

Linear Regression Analysis

- The linear regression model
- Correlation vs regression
- Geometrical representation
- First regression in R
- How to interpret the regression table
- Decomposition of variability: SST, SSR, SSE
- R-squared

Data Science in Python

Introduction

- Application of Data Science
- Quick Overview of Python

Python for Data Science

- The Pandas Package
- NumPy Package
- Scikit-Learn

Introduction to Data Wrangling with Python

- Importance of Data Wrangling
- Python for Data Wrangling
- List Functions
- Advanced Data Structures

Introduction to NumPy, Pandas and Matplotlib

- NumPy Arrays
- Advanced Mathematical Operations
- Statistics and Visualizations with NumPy and Pandas
- The Definition of Statistical Measures – Central Tendency and Spread
- Data Wrangling in Statistics and Visualization

A Deep Dive into Data Wrangling with Python

- Subsetting, Filtering and Grouping
- Detecting Outliers and Handling Missing Values
- Concatenating, Merging and Joining
- Useful Methods of Pandas

Getting Comfortable with Different Kinds of Data Sources

- Reading Data from Different Sources
- Introduction to Beautiful Soup 4 and Web Page Parsing

Learning the Hidden Secrets of Data Wrangling

- Advanced List Comprehension and zip function
- Data Formatting
- Identifying and Cleaning Outliers
- Levenshtein Distance

Advanced Web Scraping and Data Gathering

- The Requests and Beautiful Soup Libararies
- Reading Data from XML
- Reading Data from an API
- Fundamentals of Regular Expressions

Regression

Simple Linear Regression

- Multiple Linear Regression
- Conducting Regression Analysis using python
- Multiple Regression Analysis

Binary Classification

- Understanding the Business Context
- Feature Engineering
- Data-Driven Feature Engineering
- Correlation Matrix and Visualization

Multiclass Classification with Random Forest

- Training a Random Forest Classifier
- Evaluating a Model's Performance
- Maximum Depth
- Minimum Sample in Leaf
- Maximum Features

Performing Your First Cluster Analysis

- Clustering with k-means
- Interpreting k-means Results
- Choosing the Number of Clusters
- Initializing Clusters
- Calculating the Distance to the Centroid

- How to Assess Performance
  - Splitting Data
  - Assessing Model Performance for Regression Models
  - R2 Score
  - Mean Absolute Error
  - Assessing Model Performance for Classification Models
  - The Confusion Matrix
  - Receiver Operating Characteristic Curve
  - Area Under the ROC Curve
  - Saving and Loading Models
- Generalization of Machine Learning Models
  - Overfitting
  - Underfitting
  - Data
  - Random State
  - Cross Validation
  - Logistic Regression CV
  - Hyper parameter Tuning with Grid Search CV
  - Hyper parameter Tuning with Randomized Search VC
  - Model Regularization with Lasso Regression
  - Ridge Regression
- Hyperparameter Tuning
  - What are Hyper Parameters?
  - Finding the Best Hyper parametrization
  - Tuning using Grid Search.
  - Grid Search CV
  - Random Search

- Interpreting a Machine Learning Model
  - Linear Model Coefficients
  - Random Forest Variable Importance
  - Variable Importance via Permutation
  - Partial Dependence Plots
  - Analyzing Dataset
  - Analyzing the Content of Categorical Variable
  - Summarizing Numerical Variables
  - Visualizing Your Data
  - Boxplots
- Data Preparation

- Handling Row Duplication
- Converting Data Types
- Handling Incorrect Values
- Handling Missing Values

Feature Engineering

- Merging Datasets
- Binning Variables
- Manipulating Dates
- Performing Data Aggregation

mbalanced Datasets

- Understanding the Business Context
- Challenges of Imbalanced Datasets
- Strategies for Dealing with Imbalanced Datasets
- Generating Synthetic Samples

Dimensionality Reduction

- Creating a High-Dimensional Dataset
- Strategies for Addressing High-Dimensional Datasets
- Comparing Different Dimensionality Reduction Techniques

Ensemble Learning

- Ensemble Learning
- Simple Methods for Ensemble Learning
- Advanced Techniques for Ensemble Learning

Machine Learning

Introduction to Scikit-Learn

- Introduction to Machine Learning
- Scikit Learn
- Data Representation
- Data Pre-processing
- Scikit-Learn API
- Supervised and Unsupervised Learning

Introduction to Unsupervised Learning

- Clustering
- Exploring a Dataset
- Data Visualization
- Mean-Shift Algorithm
- DBSCAN Algorithm

Supervised Machine Learning

- Auto regression
  - Auto regression Models
- Classification Techniques
  - Ordinary Least Squares as a Classifier
  - Logistic Regression
  - Classification using K-Nearest Neighbours
  - Classification using Decision trees
  - Artificial Neural Networks
- Ensemble Modelling
  - One-Hot Encoding
  - Overfitting and Underfitting
  - Bagging
  - Bootstrapping
  - Boosting
  - Stacking
- Model Evaluation
  - Importing the Modules and Preparing our Dataset
  - Evaluation Metrics
  - Splitting a Dataset
  - Performance Improvement Tactics

Unsupervised Learning

- Introduction to clustering
  - Unsupervised vs Supervised Learning
  - Clustering
  - Introduction to k-means Clustering
- Hierarchical Clustering
  - Clustering Refresher
  - The Organization of Hierarchy
  - Introduction to Hierarchical Clustering
  - Linkage
  - Agglomerative versus Divisive Clustering
  - k-means vs Hierarchical Clustering
- Neighbourhood Approaches and DBSCAN
  - Clusters as Neighbourhoods
  - Introduction to DBSCAN
  - DBSCAN vs k-means and Hierarchical Clustering
- Dimensionality Reduction Techniques and PCA
  - What is Dimensionality Reduction?
  - Overview of Dimensionality Reduction Techniques
  - Principal Component Analysis

- Auto encoders
  - Fundamentals of Artificial Neural Networks
  - Autoencoders
- t-Distributed Stochastic Neighbour Embedding
  - Stochastic Neighbour Embedding (SNE)
  - t-Distributed SNE
  - Interpreting t-SNE Plots
- Topic Modelling
  - Topic Models
  - Cleaning Text Data
  - Latent Dirichlet Allocation
  - Non-Negative Matrix Factorization
- Market Basket Analysis
  - Market Basket Analysis
  - Characteristics of Transaction Data
  - The Apriori Algorithm
  - Association Rules
- Hotspot Analysis
  - Spatial Statistics
  - Kernel Density Estimation
  - Hotspot Analysis

Deep Learning

- Building Blocks of Deep Learning
  - Introduction
    - *AI, Machine Learning and Deep Learning*
    - *Machine Learning*
    - *Deep Learning*
    - *Using Deep Learning to Classify an Image*
    - *Deep Learning Models*
    - *Generative Adversarial Networks*
  - Introduction to TensorFlow
    - *Constants*
    - *Variables*
    - *Linear Algebra with TensorFlow*
    - *The reshape Function*
    - *The argmax Function*
    - *Optimizers*
  - Introduction to PyTorch
    - *GPUs in PyTorch*
    - *What are Tensors?*
    - *Advantages of Using PyTorch*

Introduction to Keras
Advantages of Keras
Disadvantages of Keras

- Neural Networks
  - Introduction
  - Neural Networks and Structure of Perceptron's
  - Training a Perceptron
  - Kera as a High-Level API
  - Exploring the Optimizers and Hyper parameters of Neural Networks
- Image Classification and Convolutional Neural Networks (CNNs)
  - Digital Images
  - Image Processing
  - Convolutional Neural Networks
  - Data Augmentation
  - Saving and Restoring Models
  - Transfer Learning
  - Fine-Tuning
- Deep Learning for Text-Embedding
  - Getting Started with Text Data Handling
  - Text Pre-processing
  - Stemming and Lemmatization
  - Text Representation Considerations
  - Classical Approaches to Text Representation
  - Distributed Representation for Text
- Deep Learning for Sequences
  - Working with Sequences
  - Recurrent Neural Networks
- LSTMs, GRUs and Advanced RNNs
  - Long-Range Dependence/Influence
  - The Vanishing Gradient Problem
  - Sequence Models for Text Classification
  - The Embedding Layer
  - Building the Plain RNN Model
  - Making Predictions on Unseen Data
  - LSTMs, GRUs, and Other Variants
  - Parameters in LSTM
  - LSTM vs Plain RNNs
  - Gated Recurrence Units
  - Bidirectional RNNs
  - Stacked RNNs
  - Summarizing all the Models
  - Attention Models
- Generative Adversarial Networks
  - Introduction

- Deep Convolutional GANs

Reinforcement Learning

- Introduction to Reinforcement Learning
  - Learning Paradigms
  - Fundamentals of Reinforcement Learning
    - *Elements of RL*
    - *The Agent-Environment Interface*
    - *Environment Types*
    - *An Action and Its Types*
    - *Policy*
    - *Goals and Rewards*
  - Reinforcement Learning Frameworks
    - *OpenAI Gym*
    - *OpenAI Baselines*
  - Applications of Reinforcement Learning
- Markov Decision Process and Bellman Equations
  - Introduction
  - Markov Process
  - Markov Property
  - Markov Chains
  - Markov Reward Process
  - Markov Decision Process
  - Solving MDPs
  - Gridworld
- Getting Started with OpenAI and TensorFlow for Reinforcement Learning
  - Introduction
  - OpenAI Gym
  - OpenAI Universe – Complex Environment
  - TensorFlow for Reinforcement Learning
  - Open AI Baselines
  - Training an RL Agent to Solve a Classic Control Problem
- Dynamic Programming
  - Introduction
  - Solving Dynamic Programming Problems
  - Identifying Dynamic Programming Problems
  - Dynamic Programming in RL
- Monte Carlo Methods
  - The Working of Monte Carlo Methods
  - Understanding Monte Carlo with Blackjack
  - Types of Monte Carlo Methods
  - Exploration versus Exploitation Trade-Off

- Importance Sampling
- Solving Frozen Lake Using Monte Carlo

Temporal Difference Learning

- Introduction to TD Learning
- TD(0) – SARSA and Q-learning
- N-Step TD and TD(λ) Algorithms
- The Relationship between DP, Monte -Carlo and TD Learning

The Multi-Armed Bandit Problem

What is Deep Q-Learning

Policy-Based Methods for Reinforcement Learning

Evolutionary Strategies for RL

Computer Vision

Basics of Image Processing

- NumPy Arrays
- Pixels in Images
- Introduction to OpenCV

Common Operations When Working with Images

- Geometric Transformations
- Image Arithmetic
- Binary Images

Working with Histograms

- Introduction to Matplotlib
- Histograms with OpenCV

Working with contours

- Contours – Basic Detection and Plotting
- Hierarchy
- Contour Matching

Face Processing in Image and Video

- Introduction to Haar Cascades
- GrabCut Technique

Object Tracking

- Naïve Tracker
- Kalman Filter
- Mean shift
- CAM shift
- The OpenCV Object Tracking API
- Object Tracking with Dlib

Object Detection and Face Recognition

- Face Recognition

- Object Detection
- Open VINO with OpenCV
  - Exploring Open VINO Toolkit
  - Model Conversion Using Model Optimizer

Natural Language Processing

Introduction to Natural Language Processing

- Introduction
- Sentiment Analysis
  - Introduction
  - Why is Sentiment Analysis Required?
  - The Growth of Sentiment Analysis
  - The Monetization of Emotion
  - Types of Sentiment
  - Key Ideas and Terms
  - Applications of Sentiment Analysis
  - Tools used for Sentiment Analysis.
  - The textblob library
  - Understanding Data for Sentiment Analysis
  - Training Sentiment Models
    - History of NLP
  - Steps in NLP
  - Word Sense Disambiguation
  - Sentence Boundary Detection
  - Kick Starting NLP Project
- Feature Extraction Methods
  - Introduction
  - Types of Data
  - Cleaning Text Data
  - Feature Extraction from Texts
  - Finding Text Similarity
- Developing a Text Classifier
  - Machine Learning
  - Supervised Learning
  - Developing a Text Classifier
  - Building Pipelines for NLP Projects
  - Saving and Loading Models
- Collecting Text Data with Web Scraping and APIs
  - Introduction
  - Collecting Data by Scraping Web Pages
  - Dealing with Semi Structured Data

- Topic Modeling
  - Introduction
  - Topic Discovery
  - Topic Modeling Algorithms
  - Key Input Parameters for LSA Topic Modeling
  - Hierarchical Dirichlet Processing (HDP)
- Vector Representation
  - Introduction
  - What is Vector?
  - Frequency Based Embedding
  - Learned Word Embedding
  - Using Pre-Trained Word Vectors
  - Document Vectors
- Text Generation and Summarization
  - Introduction
  - Generating Text with Markov Chains
  - Text Summarization
  - Key Input Parameters for TextRank
  - Recent Developments in Text Generation and Summarization
  - Practical Challenges in Extractive Summarization

Machine Learning with Google Cloud Platform (GCP)

- Introducing the GCP
  - ML and the cloud
  - Introducing the GCP
  - Getting Started with GCP
- Querying your Data with BigQuery
  - Approaching big data
  - Data Structuring
  - Querying the database
  - Google BigQuery
  - Visualizing data with Google Data Studio
- Transforming Your Data
  - How to clean and prepare the data using Google Cloud Dataprep
  - Finding outliers in the data
  - Run Job
  - Scale of features
  - Google Cloud Data flow
- Google Machine Learning APIs
  - Vision API

- Cloud Translation API
- Natural Language API
- Speech-to-text API
- Video Intelligence API

Creating ML Applications with Firebase

- Features of Firebase

Neural Networks with TensorFlow and Keras

- Overview of neural network
- Setting up Google Cloud Datalab
- Working details of simple neural network
- Implementing a simple neural network in Keras
- Implementing neural networks in TensorFlow

Evaluating Results with Ten sor Board

- Setting up Tensor Board
- Overview of Summary operations

Optimizing the Model through Hyperparameter Tuning

- Hyperparameter tuning in Google Cloud

Chatbots

- Chatbots fundamentals
- Building a Bot
- Chatbots
- Chatbots architecture
- Natural Language Processing
- Natural language Understanding
- Google Cloud Dialogflow
- Basic Dialogflow elements
- building a chatbot with Dialogflow

Machine Learning with Amazon Web Services (AWS)

AWS Application Services for AI/ML

- Analyzing images and videos with Amazon Rekognition
- Text to speech with Amazon Polly
- Speech to text with Amazon Transcribe
- Implementing natural language processing with Amazon Comprehend
- Translating documents with Amazon Translate
- Extracting text from documents with Amazon Textract
- Creating Chatbots on Amazon Lex

Amazon SageMaker Modeling

- Creating notebooks in SageMaker
- Model tuning

- Choosing instance types in Amazon SageMaker
- Securing SageMaker notebooks
- Creating alternative pipelines with Lambda Functions

Working with Step Functions

Machine Learning with Azure

Cognitive Services

- Cognitive Services for Vision APIs
- The Computer Vision API
- Face API
- Cognitive Services for Language APIs
- Cognitive Services for Speech APIs
- Cognitive Services for Knowledge APIs
- Cognitive Services for Search APIS

Bot Framework

- Bot Builder SDK
- Bot Framework
- QnA Maker
- Bot Service

Azure Machine Learning Studio

- Deploying an Azure AI Gallery template
- Building an experiment
- Deploying a model as a web service

Machine Learning Server

- What is Microsoft ML Server?
- Machine Learning with Python

Building Deep Learning Solutions

- Overview of the Azure Notebook service
- Overview of Azure Deep Learning Virtual Machine toolkits
- An overview of Tensor Flow on Azure

Integration with Other Azure Services

- Logic Apps
- Azure Functions
- Azure Data Lake Analytics
- Azure Data Factory