Fundamentals of Artificial Intelligence (AI), Machine Learning & Deep Learning (60 Hours)

Duration: 10 Days × 6 Hours per day

---

1. Getting Started (6 Hours)

Topics:

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

Hands-on Labs:

- Install Anaconda/Miniconda & set up Jupyter Notebooks
- Write first Python program & import external modules (NumPy, Pandas, Matplotlib)
- Explore Python data structures (lists, dictionaries, tuples, sets)
- Run and execute Python scripts in terminal/VS Code

---

2. Statistics and Probability Refresher with Python (6 Hours)

Topics:

- Types of data
- Mean, median, mode, standard deviation, variance
- Probability density & mass functions
- Types of distributions, percentiles, moments

Hands-on Labs:

- Use Pandas to calculate mean, median, mode, variance & standard deviation on a dataset
- Plot different probability distributions (normal, binomial, Poisson) with Matplotlib
- Generate synthetic data and calculate percentiles & quantiles

---

3. Visualization & Advanced Probability (3 Hours)

Topics:

- Crash course in Matplotlib

- Covariance and correlation
- Conditional probability
- Bayes' theorem

Hands-on Labs:

- Create histograms, bar plots, and scatter plots in Matplotlib
- Calculate correlation & covariance using Pandas and Excel

---

4. Core Machine Learning (14 Hours)

Topics:

- Data preparation techniques
- Linear & non-linear algorithms
- Ensemble methods
- Predictive models (linear, polynomial, multivariate regression)
- Bayesian methods and Naïve Bayes classifier
- K-Means clustering & recommender systems
- KNN concepts & applications
- Dimensionality reduction (PCA)
- Reinforcement learning basics

Hands-on Labs:

- Data preprocessing (handling missing values, scaling, encoding)
- **Excel Lab:** Build a simple **predictive linear regression model in Excel** using the Data Analysis ToolPak
- Build a **linear regression** model in Python
- Train a **polynomial regression** model & visualize predictions
- Implement a **Naïve Bayes classifier** on a text dataset (spam classification)
- Apply **K-Means clustering** to group customers by purchasing behavior
- Build a simple **user-item recommender system** using cosine similarity
- Implement a **KNN classifier** for image recognition (MNIST subset)
- Apply **PCA** on a dataset to reduce dimensions & visualize clusters

---

5. Model Evaluation & Validation (4 Hours)

Topics:

- Train/Test split, Cross-validation
- Overfitting vs Underfitting
- Metrics: Accuracy, Precision, Recall, F1, ROC-AUC
- Evaluate classification & regression models

Hands-on Labs:

- Split dataset into train/test sets using Scikit-learn
- Perform **k-fold cross-validation**
- Plot ROC curves & calculate AUC scores
- Compare performance metrics across multiple models

---

6. Dealing with Real-World Data (4 Hours)

Topics:

- Bias/variance trade-off
- K-fold cross-validation
- Data cleaning & normalization
- Outlier detection

Hands-on Labs:

- Handle missing values (imputation techniques)
- Normalize and scale features using MinMaxScaler & StandardScaler
- Detect outliers using plots
- Visualize impact of bias/variance trade-off on model performance

---

7. Testing & Experimental Design (3 Hours)

Topics:

- A/B testing concepts
- T-test and p-value
- Experiment duration & common pitfalls

8. Deep Learning Foundations (11 Hours)

Topics:

- Introduction to Neural Networks & Perceptrons
- Backpropagation & activation functions
- Introduction to TensorFlow & Keras
- Building and training a feedforward neural network

Hands-on Labs:

- Build a **single-layer perceptron** in NumPy
- Visualize activation functions (sigmoid, tanh, ReLU)
- Implement a **feedforward neural network** in Keras for MNIST classification
- Train and evaluate network performance with TensorFlow/Keras

---

9. Deep Learning Architectures (7 Hours)

Topics:

- Convolutional Neural Networks (CNNs) – image recognition
- Recurrent Neural Networks (RNNs) – sequence & time-series
- LSTMs & GRUs
- Transfer Learning basics

Hands-on Labs:

- Build a **CNN in Keras** for image classification (CIFAR-10 dataset)
- Implement a **simple RNN** for text sequence prediction

---

10. AI Ethics & Responsible AI (2 Hours)

Topics:

- Fairness, bias & explainability in AI
- Privacy & data governance issues
- Ethical dilemmas in AI applications
- Transparency & accountability in ML/DL models

Hands-on Labs:

- Detect bias in a dataset (e.g., gender bias in hiring dataset)
- Use **LIME/SHAP** for explainability of ML predictions
- Discuss case studies of ethical dilemmas in AI

Note: Azure Machine Learning Complete end to end project (Optional if access is available)