2-Month AI Course Syllabus

Course Overview

- **Duration:** 8 weeks
- **Level:** Beginner → Intermediate
- **Prerequisites:**
  - Basic computer knowledge
  - Basic Python skills are helpful but not mandatory
- **Tools Used:**
  - Python
  - Jupyter Notebook / Google Colab
  - scikit-learn
  - Pandas, NumPy
  - TensorFlow/Keras or PyTorch (basic level)
  - Matplotlib, Seaborn

Week 1: Foundations of AI and Python

Introduction to Artificial Intelligence

- What is AI?
- Brief history and evolution
- Types of AI: Narrow, General, Super AI
- Different areas of AI
- Real-life applications of AI

Python Refresher (Part 1)

- Basic Python syntax
- Data types (int, float, string, boolean)
- Variables
- Conditional statements (if/else)

Python Refresher (Part 2)

- Loops (for, while)
- Lists, Tuples, Dictionaries
- Functions
- Introduction to Jupyter/Colab

Working with Data in Python

- Introduction to NumPy
- Introduction to Pandas
- Loading and cleaning data

Data Visualization Basics

- Basics of Matplotlib
- Basics of Seaborn
- Creating graphs and charts

Week 2: Introduction to Machine Learning

Fundamentals of Machine Learning

- What is Machine Learning?
- Supervised vs Unsupervised Learning
- Regression vs Classification

Machine Learning Workflow

- Data preprocessing
- Feature scaling
- Model training and evaluation

Linear Regression

- Simple Linear Regression
- Multiple Linear Regression
- Practical example using scikit-learn

Model Evaluation Metrics

- MSE, RMSE
- R^2 Score
- Visualization of predictions

Logistic Regression

- Introduction to classification problems
- Logistic function
- Binary vs multi-class classification
- Practical example

Week 3: More Supervised Learning Models

Decision Trees

- How decision trees work
- Entropy and Gini Index
- Visualizing decision trees

K-Nearest Neighbors (KNN)

- How KNN works
- Distance metrics
- Classification example

Naive Bayes Classifier

- Bayes Theorem
- Application in text classification
- Practical example

Model Evaluation (Classification)

- Confusion Matrix
- Accuracy, Precision, Recall, F1-Score

Cross-Validation and Hyperparameter Tuning

- What is cross-validation?
- Grid Search
- Randomized Search

Week 4: Unsupervised Learning

Introduction to Unsupervised Learning

- Clustering and Dimensionality Reduction
- Use cases

K-Means Clustering

- How K-Means works
- Choosing K using the Elbow Method
- Practical example in scikit-learn

Hierarchical Clustering

- Agglomerative and Divisive Clustering
- Dendrograms

PCA (Principal Component Analysis)

- Importance of dimensionality reduction
- Variance and principal components
- Visualization

Assignment / Practice

- Practice with K-Means and PCA
- Prepare your own data analysis report

Week 5: Introduction to Deep Learning

Introduction to Deep Learning

- What is Deep Learning?
- Traditional ML vs Deep Learning
- Neural network architecture basics

Components of Artificial Neural Networks

- Layers and neurons
- Activation Functions (ReLU, Sigmoid, Softmax)
- Backpropagation basics

Using TensorFlow/Keras

- Basic Keras code
- Building a Sequential Model

Building a Simple Neural Network Project

- Classification problem with simple ANN
- Checking accuracy and results

Overfitting and Regularization

- What is overfitting?
- Dropout layers
- Early stopping

Week 6: Introduction to Computer Vision

Day 26 — What is Computer Vision?

- How images are stored digitally

- Use cases of computer vision

Basics of Image Processing

- Image resizing
- Normalization
- Data augmentation

Convolutional Neural Networks (CNNs)

- Convolution layers
- Pooling layers
- Filters and feature maps

Building a CNN Model

- CNN on MNIST data
- Visualizing accuracy and results

Introduction to Pre-trained Models

- What is transfer learning?
- Using MobileNet, VGGNet

Week 7: Natural Language Processing (NLP)

Introduction to NLP

- What is NLP?
- Challenges with text data

Text Processing Basics

- Tokenization
- Stop words removal
- Lemmatization and stemming

Feature Extraction in NLP

- Bag of Words
- TF-IDF
- Practical example in scikit-learn

Text Classification Project

- Sentiment Analysis
- Working on a movie reviews dataset

Word Embeddings

- Introduction to Word2Vec
- Embedding layers in neural networks

Week 8: Projects and Advanced Topics

Ethics and Responsible AI

- Bias and fairness in AI
- Privacy concerns
- Explainable AI (XAI)

Introduction to Reinforcement Learning

- What is Reinforcement Learning?
- Agents, environments, and rewards

Introduction to Time Series Data

- What is Time Series data?
- Basic forecasting techniques

Mini Project Work

- Choose a project (see suggestions below)
- Implement end-to-end solution

Project Presentations and Recap

- Present projects
- Q&A session
- Career pathways in AI
- Resources for further learning

Skills You'll Gain After This Course

- Good understanding of ML concepts
- Ability to preprocess and visualize data

- Building basic to intermediate ML models
- Understanding neural networks basics
- Capable of building simple projects in computer vision and NLP
- Experience building end-to-end AI projects

Project Ideas

- House price prediction (regression)
- Spam vs ham email classification
- Customer segmentation
- Sentiment analysis on tweets
- Image classification (MNIST, CIFAR-10)
- Movie recommendation system
- Time series forecasting (e.g. stock prices)

Resources

- scikit-learn documentation
- TensorFlow/Keras and PyTorch tutorials
- Kaggle datasets
- Google Colab
- Recommended books:
  - "Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow" – Aurélien Géron
  - "Python Machine Learning" – Sebastian Raschka
