6-Month AI Course Syllabus

Course Overview

- **Duration:** 24 weeks (~120 sessions)
- **Level:** Beginner → Intermediate → Advanced
- **Prerequisites:**
  - Comfortable using computers
  - Basic familiarity with Python recommended
- **Tools & Libraries:**
  - Python, Jupyter Notebook / Google Colab
  - NumPy, Pandas, Matplotlib, Seaborn
  - Scikit-learn
  - TensorFlow, Keras, PyTorch
  - NLP libraries: NLTK, SpaCy, Hugging Face Transformers
  - OpenCV
  - Flask, FastAPI (for deployment)
  - MLflow, Streamlit, Docker basics
  - Cloud basics (AWS, GCP, Azure)

SYLLABUS DETAILS

Month 1 – Foundations of AI & Python Data Science

Week 1

- Introduction to AI & Machine Learning
  - AI definitions & history
  - Types of AI
  - AI in industries
- AI vs ML vs DL
- Overview of career paths in AI

Week 2

- Python refresher
  - Syntax, variables, data types
  - Lists, tuples, dictionaries
  - Functions, loops, file handling
- Working in Jupyter/Colab

Week 3

- NumPy deep dive
- Pandas for data manipulation
- Data loading (CSV, Excel, JSON)

Week 4

- Data cleaning
- Handling missing data
- Exploratory Data Analysis (EDA)
- Data visualization
  - Matplotlib
  - Seaborn

Month 2 – Core Machine Learning

Week 5

- Machine Learning overview
  - Types of learning
  - ML pipeline
- Regression basics
  - Simple Linear Regression
  - Multiple Linear Regression

Week 6

- Polynomial Regression
- Regularization: Ridge, Lasso
- Model evaluation metrics
  - MAE, MSE, RMSE, R^2

Week 7

- Classification basics
  - Logistic Regression
  - K-Nearest Neighbors (KNN)

Week 8

- Decision Trees
- Random Forests
- Naive Bayes
- Model evaluation for classification
  - Confusion matrix
  - Precision, Recall, F1-score
  - ROC, AUC

Month 3 – Advanced ML Techniques & Unsupervised Learning

Week 9

- Support Vector Machines (SVM)
- Ensemble methods
  - Gradient Boosting
  - XGBoost, LightGBM

Week 10

- Cross-validation techniques
- Hyperparameter tuning
  - Grid Search
  - Random Search

Week 11

- Unsupervised Learning
  - Clustering concepts
  - K-Means
  - Hierarchical Clustering

Week 12

- DBSCAN
- Dimensionality reduction
  - PCA
  - t-SNE
- Anomaly detection basics

Month 4 – Real-World ML & Deployment Basics

Week 13

- Working with imbalanced datasets
- Feature engineering
  - Encoding categorical data
  - Feature scaling
  - Binning, transformations

Week 14

- Pipelines in scikit-learn
- Model saving and loading (pickle, joblib)

Week 15

- Introduction to model deployment
  - Flask / FastAPI basics
  - Creating APIs for ML models

Week 16

- MLflow basics
- Introduction to cloud ML (AWS/GCP/Azure)
- Mini Project: Deploy an ML model

Month 5 – Deep Learning

Week 17

- Introduction to Neural Networks
  - Perceptron
  - Activation functions (ReLU, Sigmoid, Softmax)

Week 18

- Building Neural Networks with Keras/TensorFlow
- Loss functions & optimizers
- Backpropagation concept

Week 19

- CNNs (Convolutional Neural Networks)
  - Architecture
  - Pooling layers

Week 20

- Transfer learning

- Data augmentation
- CNNs for image classification (project)

Month 6 – NLP, Advanced Topics, and Capstone Project

Week 21

- Introduction to NLP
- Text preprocessing
  - Tokenization, stopwords, stemming, lemmatization

Week 22

- Feature extraction in NLP
  - Bag of Words
  - TF-IDF
  - Word Embeddings (Word2Vec, GloVe)

Week 23

- Introduction to Transformers
  - Hugging Face basics
- Building a sentiment analysis model

Week 24

- Time Series basics
- Explainable AI (XAI)
- AI Ethics
- Capstone Project week

Capstone Project

During the final weeks, students work on a capstone project end-to-end:

- Data collection & cleaning
- Model development
- Model evaluation
- Deployment (if possible)
- Presentation of results

Example Capstone Ideas:

- Sales forecasting
- Image classification for custom data
- Text summarization or chatbot
- Fraud detection
- Recommendation system
- Object detection for real-world use
- Sentiment analysis on social media data

Skills You'll Gain

- Strong Python data science skills

- ML algorithms (regression, classification, clustering, ensemble models)
- Advanced model tuning & evaluation
- Deep Learning with TensorFlow/Keras
- Basic deployment skills (Flask/FastAPI, MLflow)
- Computer vision & NLP understanding
- Exposure to real-world projects
- Foundation for cloud-based ML work

Recommended Learning Resources

- Books:
  - Hands-On Machine Learning – Aurélien Géron
  - Python Machine Learning – Sebastian Raschka
  - Deep Learning with Python – François Chollet
- Platforms:
  - Kaggle
  - Coursera
  - YouTube channels: StatQuest, Data School
- Official documentation:
  - scikit-learn
  - TensorFlow
  - PyTorch
  - Hugging Face

This curriculum provides ample time for theory, practice, and real-world project work, equipping learners to either enter the AI workforce or pursue further specialization.
