3-Month AI Course Syllabus

Course Overview

- **Duration:** 12 weeks (60 sessions)
- **Level:** Beginner → Intermediate/Advanced
- **Prerequisites:**
  - Comfortable with basic computer usage
  - Some prior exposure to Python programming is helpful
- **Tools Used:**
  - Python
  - Jupyter Notebook / Google Colab
  - scikit-learn
  - Pandas, NumPy
  - TensorFlow/Keras & PyTorch (basics to intermediate)
  - Matplotlib, Seaborn
  - NLP libraries (NLTK, SpaCy)
  - OpenCV
  - MLflow or similar tools for model tracking (intro level)

SYLLABUS DETAILS

Week 1 – Foundations of AI & Python Refresher

- Day 1: Introduction to AI
  - AI definitions & concepts
  - Types of AI (Narrow, General, Super AI)
  - AI use cases in various industries
- Day 2: AI vs ML vs DL
  - Differences explained
- Day 3: Python Programming Refresher (Part 1)
  - Syntax, variables, data types
- Day 4: Python Programming Refresher (Part 2)
  - Loops, functions, modules
- Day 5: Jupyter Notebooks / Google Colab Basics
- Day 6: Data Structures in Python
  - Lists, Dictionaries, Tuples
- Day 7: Numpy Basics
  - Arrays, operations, indexing

Week 2 – Data Handling & Visualization

- Day 8: Pandas Basics
  - DataFrames, Series
- Day 9: Importing Data
  - CSV, Excel, JSON
- Day 10: Data Cleaning
  - Handling missing data
- Day 11: Data Transformation
  - GroupBy, merging, pivot tables
- Day 12: Data Visualization with Matplotlib
- Day 13: Data Visualization with Seaborn
- Day 14: Exploratory Data Analysis (EDA) workflow

Week 3 – Machine Learning Fundamentals

- Day 15: Machine Learning Introduction
  - Supervised, Unsupervised, Reinforcement Learning
- Day 16: ML Pipeline Overview
  - Data preparation → Training → Testing → Deployment
- Day 17: Regression vs Classification
- Day 18: Model Evaluation Metrics
  - MSE, RMSE, R^2, confusion matrix, precision, recall
- Day 19: Overfitting vs Underfitting
- Day 20: Cross-validation
- Day 21: Hyperparameter tuning basics

Week 4 – Regression Algorithms

- Day 22: Simple Linear Regression
- Day 23: Multiple Linear Regression
- Day 24: Polynomial Regression
- Day 25: Ridge and Lasso Regression
- Day 26: Practical Regression Project
- Day 27: Model Evaluation for Regression
- Day 28: Assignment: Regression Model on Real Dataset

Week 5 – Classification Algorithms

- Day 29: Logistic Regression
- Day 30: K-Nearest Neighbors (KNN)
- Day 31: Decision Trees
- Day 32: Random Forest
- Day 33: Naive Bayes
- Day 34: Support Vector Machines (SVM)
- Day 35: Model Evaluation for Classification (ROC, AUC)

Week 6 – Unsupervised Learning

- Day 36: Introduction to Unsupervised Learning
- Day 37: Clustering Algorithms Overview
- Day 38: K-Means Clustering
- Day 39: Hierarchical Clustering
- Day 40: DBSCAN
- Day 41: Dimensionality Reduction Concepts
- Day 42: PCA (Principal Component Analysis)

Week 7 – Working with Real-World Data

- Day 43: Data Cleaning in Real-Life Datasets
- Day 44: Outlier Detection
- Day 45: Feature Engineering Techniques
- Day 46: Feature Scaling & Encoding

- Day 47: Handling Imbalanced Data
- Day 48: Pipelines in scikit-learn
- Day 49: Mini Project – Data Pipeline & Model

Week 8 – Introduction to Deep Learning

- Day 50: Introduction to Neural Networks
- Day 51: Perceptron and Multi-layer Perceptrons
- Day 52: Activation Functions (ReLU, Sigmoid, Softmax)
- Day 53: Forward and Backpropagation
- Day 54: TensorFlow Basics
- Day 55: Building a Simple Neural Network in Keras
- Day 56: Overfitting and Regularization Techniques (Dropout, Early Stopping)

Week 9 – Computer Vision Fundamentals

- Day 57: What is Computer Vision?
- Day 58: Basics of Image Processing with OpenCV
- Day 59: Convolutional Neural Networks (CNNs)
- Day 60: Pooling Layers
- Day 61: Building a CNN with Keras
- Day 62: Data Augmentation for Images
- Day 63: Transfer Learning & Pre-trained Models (VGGNet, ResNet)

Week 10 – Natural Language Processing (NLP)

- Day 64: Introduction to NLP
- Day 65: Text Preprocessing
  - Tokenization, stop words, stemming, lemmatization
- Day 66: Feature Extraction in NLP
  - Bag of Words, TF-IDF
- Day 67: Text Classification Basics
- Day 68: Word Embeddings
  - Word2Vec, GloVe
- Day 69: Building a Simple Sentiment Analysis Model
- Day 70: Introduction to Transformer Models

Week 11 – Advanced Topics & Deployment

- Day 71: Time Series Analysis Introduction
- Day 72: Time Series Forecasting Techniques
- Day 73: Introduction to Reinforcement Learning
- Day 74: AI Ethics & Responsible AI
- Day 75: Explainable AI (XAI)
- Day 76: Model Deployment Basics
- Day 77: Introduction to MLflow or similar tracking tools

Week 12 – Projects & Final Assessment

- Day 78–79: Final Project Work
- Day 80: Project Review & Presentations
- Day 81: Course Revision & Q&A
- Day 82: Career Paths in AI
- Day 83: Building a Portfolio
- Day 84: Final Assessment Exam

Skills You'll Gain

- Solid understanding of ML &DL fundamentals
- Practical experience with real-world datasets
- Ability to build supervised & unsupervised models
- Knowledge of CNNs, NLP, and advanced techniques
- Exposure to model deployment basics
- Completed portfolio projects for your resume

Suggested Final Project Ideas

- House Price Prediction (Regression)
- Customer Churn Prediction
- Image Classification on custom dataset
- Sentiment Analysis on Social Media Data
- Object Detection with pre-trained models
- Recommendation System
- Time Series Forecasting of sales data
- Text Summarization using transformers

Resources

- scikit-learn official docs
- TensorFlow/Keras documentation
- PyTorch tutorials
- Kaggle for datasets & competitions
- Books:
  - "Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow" – Aurélien Géron
  - "Deep Learning with Python" – François Chollet
  - "Python Machine Learning" – Sebastian Raschka
