12-Month AI Course Syllabus

Course Overview

- **Duration:** 12 months (48 weeks)
- **Level:** Beginner → Advanced
- **Prerequisites:**
  - Basic computer knowledge
  - Familiarity with math (high school level algebra, some statistics)
  - Optional but helpful: basic Python
- **Key Tools & Libraries:**
  - Python, Jupyter Notebook/Colab
  - NumPy, Pandas, Matplotlib, Seaborn, Plotly
  - Scikit-learn
  - TensorFlow, Keras, PyTorch
  - NLTK, SpaCy, Hugging Face Transformers
  - OpenCV
  - Flask, FastAPI, Streamlit
  - MLflow, Docker, Kubernetes basics
  - Spark, Hadoop basics (for Big Data modules)
  - Cloud (AWS, GCP, Azure)
  - Git, GitHub
  - SQL, NoSQL basics

Detailed Syllabus

Month 1 – Introduction to AI & Python Foundations

Weeks 1-4

- What is AI, ML, DL?
- AI history and trends
- AI applications in various industries
- Python basics:
  - Syntax, variables, data types
  - Lists, tuples, dictionaries
  - Loops, functions, file handling
- Working with Jupyter Notebook / Google Colab
- Version control basics (Git, GitHub)

Month 2 – Data Science & Visualization

Weeks 5-8

- NumPy deep dive
- Pandas for data manipulation
- Data loading from CSV, Excel, JSON, APIs
- Data cleaning:
  - Handling missing data
  - Duplicate handling
- Exploratory Data Analysis (EDA)
- Data visualization:
  - Matplotlib

  - Seaborn
  - Plotly

Month 3 – Statistics & Probability for ML

Weeks 9-12

- Descriptive statistics
- Measures of central tendency
- Variance, standard deviation
- Probability concepts:
  - Conditional probability
  - Bayes' theorem
- Hypothesis testing
- Correlation vs causation
- Sampling methods
- Outlier detection basics

Month 4 – Core Machine Learning: Regression & Classification

Weeks 13-16

- ML pipeline overview
- Simple Linear Regression
- Multiple Linear Regression
- Polynomial Regression
- Regularization:
  - Lasso
  - Ridge
- Logistic Regression
- K-Nearest Neighbors (KNN)
- Model evaluation:
  - MSE, RMSE, MAE
  - Confusion matrix
  - Precision, Recall, F1-score, ROC, AUC

Month 5 – Machine Learning Algorithms Continued

Weeks 17-20

- Decision Trees
- Random Forests
- Naive Bayes
- Support Vector Machines (SVM)
- Ensemble methods:
  - Bagging
  - Boosting (AdaBoost, Gradient Boosting, XGBoost, LightGBM)
- Cross-validation techniques
- Hyperparameter tuning:
  - Grid Search
  - Randomized Search

Month 6 – Unsupervised Learning & Feature Engineering

Weeks 21-24

- Unsupervised Learning concepts
- Clustering:
  - K-Means
  - Hierarchical Clustering
  - DBSCAN
- Dimensionality Reduction:
  - PCA
  - t-SNE
- Anomaly Detection
- Feature Engineering:
  - Encoding categorical data
  - Feature scaling
  - Binning, transformations
- Data pipelines in scikit-learn

Month 7 – Time Series & Advanced Topics

Weeks 25-28

- Time Series fundamentals
- Trend, seasonality, noise
- Stationarity
- AR, MA, ARIMA, SARIMA models
- Forecasting with machine learning
- Introduction to Prophet library

Month 8 – Deep Learning Fundamentals

Weeks 29-32

- Introduction to Neural Networks
- Perceptron, activation functions
- Loss functions & optimizers
- Backpropagation concept
- Building neural networks:
  - TensorFlow
  - Keras
- Introduction to PyTorch

Month 9 – Computer Vision

Weeks 33-36

- Understanding images as arrays
- CNN architecture:
  - Convolutions
  - Pooling
- Transfer Learning:
  - Using pre-trained models
- Data augmentation
- Object detection basics:
  - YOLO
  - SSD
- OpenCV basics

- Mini-project: Image classification

Month 10 – Natural Language Processing (NLP)

Weeks 37-40

- NLP basics:
  - Tokenization
  - Stopwords
  - Stemming, Lemmatization
- Text vectorization:
  - Bag of Words
  - TF-IDF
- Word embeddings:
  - Word2Vec
  - GloVe
- Named Entity Recognition (NER)
- Sentiment analysis
- Transformers basics:
  - Hugging Face
  - BERT overview
- Chatbot fundamentals

Month 11 – Model Deployment & MLOps

Weeks 41-44

- Introduction to deployment:
  - Flask
  - FastAPI
  - Streamlit
- Model serialization (pickle, joblib)
- Docker fundamentals
- MLflow for experiment tracking
- Cloud deployment:
  - AWS Sagemaker
  - GCP AI Platform
- CI/CD basics for ML pipelines
- Monitoring ML models in production

Month 12 – Advanced AI Topics & Capstone Project

Weeks 45-48

- Reinforcement Learning basics
- Generative AI concepts:
  - GANs introduction
  - Diffusion models (overview)
- Explainable AI (XAI)
- AI Ethics and Fairness
- Big Data overview:
  - Hadoop, Spark basics
- Capstone Project:
  - End-to-end ML/DL solution
  - Presentation and reporting

Capstone Project (Last Month Focus)

Students choose a project topic to build a fully functioning AI solution:

- Data gathering
- Data cleaning and EDA
- Model building and evaluation
- Deployment (if feasible)
- Documentation & presentation

Example Capstone Projects:

- Predictive maintenance for factories
- Customer churn prediction
- Object detection in retail stores
- NLP chatbot for customer support
- Real-time fraud detection
- Recommendation systems
- Medical image classification
- Time-series forecasting for financial markets

Skills Gained

- Solid understanding of ML and DL theory and practice
- Python programming and data handling
- Advanced model tuning and optimization
- Computer Vision and NLP projects
- MLOps and deployment techniques
- Exposure to real-world projects and cloud services
- Foundation for research or advanced studies in AI

Recommended Resources

- **Books:**
  - Hands-On Machine Learning – Aurélien Géron
  - Deep Learning with Python – François Chollet
  - Python Machine Learning – Sebastian Raschka
- **Online platforms:**
  - Kaggle
  - Coursera
  - Udemy
  - YouTube channels: StatQuest, Data School
- **Documentation:**
  - scikit-learn
  - TensorFlow
  - PyTorch
  - Hugging Face

This 12-month curriculum is highly thorough, designed to turn beginners into proficient AI practitioners ready for industry roles or further specialization.