

Complete Artificial Intelligence for Beginners

Course Duration: 15 days (8 hours/day)

Pre-requisite: Basic knowledge of any Object-Oriented Programming (Python is recommended but any other language is good)

Courseware: Unofficial pdf/ppt

Lab: Koenig DC (VM) for lab performance

Course Objective

The course aims to provide a comprehensive foundation in Python programming, machine learning, and deep learning. It covers essential concepts, practical skills, and hands-on experience, enabling participants to proficiently analyze, model, and solve real-world data science problems.

Day 1: Introduction to AI and Python Refresher

Session 1: Introduction to AI

- What is AI?
- History and evolution of AI
- AI applications in various fields
- Introduction to machine learning and deep learning

Session 2: Python Basics

- Python installation and setup
- Basic syntax and data types
- Control structures: if statements, loops
- Functions and modules

Day 2: Python for Data Analysis & Pre-processing

Session 1: Python Libraries for Data Science

- Introduction to NumPy, Pandas, Matplotlib, Seaborn
- Basic operations with NumPy
- Data manipulation with Pandas

Session 2: Data Cleaning and EDA

- Data cleaning techniques
- Handling missing and categorical data
- Exploratory Data Analysis (EDA) with visualizations

Day 3: Data Preprocessing and Transformation

Session 1: Data Preprocessing

- Importance of data preprocessing
- Handling missing values
- Label encoding and one-hot encoding

Session 2: Data Transformation

- Data normalization techniques: Standard Scaler, MinMax Scaler
- Splitting data: Train, test, and validation sets

Day 4: Introduction to Machine Learning

Session 1: Machine Learning Fundamentals

- Types of machine learning
- Lifecycle of a data science project
- Overview of supervised and unsupervised learning

Session 2: Supervised Learning - Regression

- Simple Linear Regression
- Multiple Linear Regression
- Evaluation metrics: R-squared, RMSE

Day 5: Supervised Learning - Classification

Session 1: Logistic Regression

- Concept and intuition
- Confusion matrix and performance metrics

Session 2: Advanced Classification Techniques

- Support Vector Machine (SVM)
- Decision Trees
- Random Forests

Day 6: Feature Selection and Dimensionality Reduction

Session 1: Feature Selection

- Importance of feature selection
- Univariate feature selection techniques
- Recursive Feature Elimination (RFE)

Session 2: Dimensionality Reduction

- Principal Component Analysis (PCA)

- Hands-on PCA implementation

Day 7: Unsupervised Learning and Clustering

Session 1: Clustering Techniques

- Introduction to clustering
- K-means clustering
- Elbow method

Session 2: Natural Language Processing (NLP) Basics

- Introduction to NLP
- Tokenization, stop words, stemming, lemmatization

Day 8: Introduction to Deep Learning

Session 1: Deep Learning Fundamentals

- Need and applications of deep learning
- Working of Artificial Neural Networks (ANN)
- Introduction to TensorFlow and Keras

Session 2: Building Neural Networks

- Keras model building: Construct, compile, evaluate
- Activation functions
- Loss functions and optimization techniques

Day 9: Neural Network Basics and Shallow Neural Networks

Session 1: Neural Network Basics

- Machine learning problem setup
- Neural network mindset
- Vectorization for efficient computation

Session 2: Shallow Neural Networks

- Building a neural network with one hidden layer
- Forward propagation and backpropagation

Day 10: Deep Neural Networks

Session 1: Deep Neural Networks

- Building and training deep neural networks
- Computation in deep learning

Session 2: Practical Aspects of Deep Learning

- Initialization methods
- Regularization techniques to prevent overfitting

Day 11: Optimization and Hyperparameter Tuning

Session 1: Optimization Algorithms

- Advanced optimization techniques
- Random minibatching and learning rate decay

Session 2: Hyperparameter Tuning

- Batch normalization
- Grid search and randomized search

Day 12: Convolutional Neural Networks (CNNs)

Session 1: Foundations of CNNs

- Understanding pooling and convolutional layers
- Building deep CNNs for image classification

Session 2: Advanced CNN Techniques

- Advanced tricks and methods in deep CNNs
- Transfer learning with pretrained models

Day 13: Object Detection and Face Recognition

Session 1: Object Detection

- Using CNNs for object detection tasks

Session 2: Face Recognition & Neural Style Transfer

- Applying CNNs for face recognition
- Implementing neural style transfer for art generation

Day 14: Recurrent Neural Networks (RNNs) and NLP

Session 1: Recurrent Neural Networks

- Introduction to RNNs
- Variants of RNNs for sequential data modeling

Session 2: NLP with Deep Learning

- NLP applications with deep learning models
- Word embeddings for text analysis

Day 15: Advanced Sequence Models and Transformers

Session 1: Sequence Models & Attention Mechanism

- Enhancing sequence models with attention mechanisms
- Speech recognition and audio data processing

Session 2: Transformer Networks

- Understanding transformer networks
- Applications of transformers in NLP