Faysal Aziz

Machine Learning and AI

3-day intensive course using python and various libraries

FA-0295 · 3-day

DevOps, Cloud & Infrastructure · Data & Analytics

Introduction

This course is crafted to transition you from understanding basic Python programming to applying it in machine learning and artificial intelligence. We will explore the core concepts of machine learning, dive into deep learning fundamentals, and get hands-on experience with leading Python libraries in this domain. Our objective is to empower you with the knowledge and skills needed to design, implement, and deploy machine learning and AI solutions.

Learning outcomes

By the end of this course, participants will:

1. Develop a comprehensive understanding of the principles, algorithms, and mathematical foundations underlying machine learning and deep learning.
2. Gain hands-on experience with Python's most powerful libraries for machine learning and AI, including Scikit-learn for classical machine learning models, XGBoost for boosting techniques, TensorFlow for deep learning, and PyTorch for both research and production needs.
3. Learn to preprocess data effectively, including normalization, standardization, handling missing values, and feature engineering to make datasets ready for machine learning models.
4. Master the art of selecting and applying the appropriate machine learning algorithms for various types of data and predictive modeling problems, including regression, classification, clustering, and more.
5. Acquire skills to design, build, and train advanced neural network architectures, including Convolutional Neural Networks (CNNs) for image processing and Recurrent Neural Networks (RNNs) for time series analysis and natural language processing.
6. Understand and implement best practices for model evaluation, including cross-validation techniques, performance metrics, and strategies for dealing with overfitting and underfitting.

7. Gain proficiency in tuning models using hyperparameter optimization techniques to achieve the best performance.
8. Learn to implement end-to-end machine learning and deep learning projects, from data collection and preprocessing to model training, evaluation, and deployment.
9. Explore advanced topics in AI, such as transfer learning, generative adversarial networks (GANs), and reinforcement learning, to stay ahead in the field.
10. Develop critical thinking and problem-solving skills applicable to real-world AI challenges, preparing for further research or industry projects.

Prerequisites

- Strong foundation in Python: Proficiency in Python programming, including familiarity with syntax, control structures, data types, and functions. Experience with Python's data science stack (NumPy, Pandas) is highly beneficial.
- Mathematical background: A good grasp of basic mathematics, including algebra and calculus. Understanding of statistics and probability is crucial for interpreting data and model outcomes.
- Basic knowledge of machine learning concepts: Familiarity with core machine learning concepts and terminology, such as supervised vs. unsupervised learning, overfitting, and model evaluation metrics.
- Software and tools: Comfort with using Python development environments (like Jupyter Notebooks or PyCharm) and managing packages (conda or pip). Basic experience with any version control system (e.g., Git) is a plus.
- Analytical skills: Ability to think critically about problems and data. Experience with data manipulation and visualization tools will be helpful.
- Curiosity and willingness to experiment: Machine learning and AI are rapidly evolving fields. An eagerness to learn new concepts, experiment with different algorithms, and apply creative solutions to problems is essential.

Training outline

1. Machine Learning Fundamentals with Scikit-learn and XGBoost

Introduction to Machine Learning

1. Overview of machine learning, types of learning, and algorithm categories.
2. Setting up the development environment.

Data Preprocessing and Visualization

1. Handling missing values, categorical data, feature scaling.
2. Visualizing data for machine learning insights.

Supervised Learning with Scikit-learn

1. Regression models (Linear Regression, Decision Trees).
2. Classification models (Logistic Regression, SVM, Random Forests).
3. Model evaluation and hyperparameter tuning.

Ensemble Learning and XGBoost

1. Introduction to ensemble methods and boosting.
2. Implementing XGBoost for performance improvement.

2. Deep Learning Fundamentals with TensorFlow and Keras

Introduction to Neural Networks and Deep Learning

1. Understanding neural networks, activation functions, and network architectures.
2. TensorFlow and Keras overview.

Building and Training Neural Networks

1. Designing a neural network with TensorFlow and Keras.
2. Optimizers, loss functions, and regularization techniques.

Convolutional Neural Networks (CNNs)

1. Introduction to CNNs and their applications in image recognition.
2. Building a CNN with TensorFlow and Keras.

Recurrent Neural Networks (RNNs) and LSTM

1. Understanding RNNs and their use in sequence modeling.
2. Implementing an LSTM model for time series prediction or text generation.

3. Advanced Models and Techniques with PyTorch

Introduction to PyTorch

1. PyTorch basics, tensors, and autograd.
2. Building blocks of neural networks in PyTorch.

Implementing Deep Learning Models with PyTorch

1. Designing and training deep learning models for classification and regression tasks.
2. Advanced neural network architectures: Transfer Learning, GANs.

Natural Language Processing (NLP) with PyTorch

1. Tokenization, word embeddings, and sequence models for NLP.
2. Building an NLP model with PyTorch for sentiment analysis or text classification.

Project Work and Capstone Challenge

1. Applying the learned concepts in a capstone project.
2. Guidance on further learning paths and resources for deep diving into machine learning and AI topics.

This course combines theoretical knowledge with practical, hands-on exercises and projects, ensuring that participants not only understand the concepts but also know how to apply them in real-world scenarios.

© 2026 Faysal Aziz · Kuala Lumpur