

Python and AI

From Machine Learning Foundations to Cutting-Edge Deep Learning

FA-0379

Agentic & Generative AI · Software Development · DevOps, Cloud & Infrastructure · Data & Analytics

Introduction

Master AI with Python—explore machine learning, dive into deep learning, and build with the latest in LLMs and RAG for state-of-the-art analytics.

Artificial Intelligence (AI) is redefining industries by automating tasks, enhancing decision-making, and unlocking new possibilities. This course bridges the gap between foundational machine learning and advanced deep learning techniques, with a focus on implementing modern transformer-based architectures and retrieval-augmented generation (RAG) using LangChain. With hands-on guidance from an instructor with over 30 years of industry expertise, you'll learn how to build intelligent systems that solve real-world problems.

Learning outcomes

By the end of this course, participants will be able to:

- Understand the fundamentals of machine learning and its key techniques.
- Build and evaluate supervised, unsupervised, and reinforcement learning models using Python.
- Develop deep learning models, including Convolutional Neural Networks (CNNs) and Transformers.
- Work with language models using LangChain and implement Retrieval-Augmented Generation (RAG).
- Apply AI techniques to real-world problems with Python and popular libraries like scikit-learn, TensorFlow, PyTorch, and Hugging Face.
- Use LLMs and RAG for advanced analytics.

Prerequisites

- Good programming experience in Python.
- Basic understanding of mathematics (linear algebra, calculus, and probability).
- Familiarity with libraries such as NumPy and pandas is helpful but not mandatory.
- Access to Volve and Sleipner datasets.
- Access to Google Colab and Drive.
- Access to GPU.

Training outline

1. Day 1: Introduction to Machine Learning (ML)

Fundamentals of Machine Learning

Overview of AI, Machine Learning, and Deep Learning

1. Differences and relationships between them.
2. Applications in industry: From predictive analytics to personalization.

Types of Machine Learning

1. Supervised, unsupervised, and reinforcement learning (RL).

2. Supervised Learning

Core Concepts

1. Training data, labels, features, and target variables.
2. Linear and logistic regression.

Key Algorithms

1. Decision Trees, Random Forests, and Gradient Boosting (e.g., XGBoost, LightGBM).

Model Evaluation

1. Accuracy, precision, recall, and F1-score.
2. Cross-validation and overfitting.

3. Unsupervised Learning

Understanding Clustering and Dimensionality Reduction

1. K-means clustering.
2. Principal Component Analysis (PCA) and t-SNE.

3. Introduction to Python ML Libraries

7. scikit-learn: Building ML pipelines.
8. Hands-on exercises: Training and evaluating a simple ML model.

2. Day 2: Deep Learning (DL)

Fundamentals of Deep Learning

Introduction to Neural Networks

1. Perceptrons, activation functions, and forward/backward propagation.
2. Loss functions and optimization with gradient descent.

Frameworks for Deep Learning

1. TensorFlow and PyTorch basics.

2. **Convolutional Neural Networks (CNNs)**

Architecture of CNNs

1. Convolutional layers, pooling, and fully connected layers.
2. Applications in computer vision.

Hands-on Exercise:

1. Building a simple image classifier using TensorFlow or PyTorch.

2. **Transformers: Revolutionizing Deep Learning**

Key Concepts in Transformers

1. Self-attention mechanism and positional encoding.
2. Evolution from RNNs and LSTMs to Transformers.

Applications of Transformers

1. NLP and beyond: BERT, GPT, and image transformers.

2. **Hands-on Exercise:**

7. Building a simple Transformer for text classification using Hugging Face.

3. Day 3: Advanced Topics in AI and Language Models

Introduction to Large Language Models (LLMs)

Overview of LLMs

1. How LLMs work: Pretraining and fine-tuning.
2. Examples: GPT-3, GPT-4, and open-source LLMs.

3. LangChain for LLM Applications

What is LangChain?

1. Overview of its role in chaining LLM tasks.

Building Applications with LangChain

1. Text summarization, question answering, and chatbots.

2. Retrieval-Augmented Generation (RAG)

Key Concepts

1. Combining LLMs with external knowledge bases.
2. Retrieval techniques: Dense and sparse retrieval.

Building a RAG System with LangChain

1. Connecting to a vector database (e.g., Pinecone, FAISS).
2. Creating a real-time knowledge retrieval pipeline.

3. Deployment and Real-World Applications

Model Deployment Strategies

1. Using APIs for AI applications.
2. Hosting models with Flask or FastAPI.

AI Ethics and Future Trends

1. Responsible AI development.
2. The role of AI in shaping industries.

This intensive, three-day course equips you to navigate the world of AI with Python, whether you're tackling data challenges, building advanced deep learning models, or leveraging cutting-edge LLMs. Be ready to transform your AI aspirations into actionable expertise!