

Intelligent Applications Advanced

Course Description

An accelerated, hands-on program covering advanced deep learning, scalable machine learning, reinforcement learning, natural language processing, computer vision, responsible AI, and intelligent application development. Participants build and evaluate practical AI solutions using Python-based frameworks.

Learning Objectives

- Design and implement advanced neural network architectures.
- Apply scalable machine learning and model optimization techniques.
- Develop reinforcement learning agents and intelligent decision systems.
- Build NLP and transformer-based applications.
- Implement computer vision workflows for images and video.
- Assess AI fairness, explainability, privacy, and governance.
- Integrate AI components into a working application prototype.

Prerequisites

- Working knowledge of Python programming and data handling.
- Familiarity with core machine learning concepts and evaluation metrics.
- Basic experience with NumPy, Pandas, Scikit-learn, and neural networks.

Day 1: Advanced Neural Networks and Scalable Machine Learning

Module 1: Advanced Deep Learning Architectures

- Advanced intelligent application architectures and neural network fundamentals review
- Deep feedforward networks, CNNs, RNNs, LSTM and GRU
- Attention mechanisms and introduction to transformers
- Activation functions, loss functions, optimizers, regularization and batch normalization

Lab: Train and evaluate a CNN image classifier with TensorFlow or PyTorch.

Module 2: Generative Models and Unsupervised Learning

- Autoencoders and variational autoencoders
- Generative Adversarial Networks: generators and discriminators
- Introduction to diffusion models
- Feature extraction, dimensionality reduction and practical use cases

Lab: Implement an autoencoder for reconstruction or anomaly detection.

Module 3: Machine Learning at Scale

- Large-scale dataset challenges and big data analytics
- Distributed computing, Apache Spark and Spark MLlib
- Parallel processing, feature pipelines and distributed training
- GPU acceleration and scalable AI infrastructure

Lab: Build a machine learning pipeline using PySpark MLlib.

Module 4: Advanced Model Optimization and Evaluation

- Adam, RMSProp and learning-rate scheduling

- Hyperparameter search, cross-validation and benchmarking
- Bias-variance trade-off, generalization and inference optimization
- Explainability with SHAP and LIME

Lab: Tune a model and interpret predictions using explainability methods.

Day 1 Outcome: Develop deep learning models, apply generative methods, and optimize scalable ML pipelines.

Day 2: Reinforcement Learning and Natural Language Processing

Module 5: Reinforcement Learning and Intelligent Decision-Making

- Agents, environments, states, actions and reward signals
- Markov Decision Processes and Bellman equations
- Q-Learning and Deep Q-Networks
- Exploration versus exploitation and practical RL use cases

Lab: Implement a Q-Learning agent in a Gymnasium environment.

Module 6: Advanced Reinforcement Learning and Multi-Agent Systems

- Policy gradients and actor-critic architectures
- Deep reinforcement learning concepts
- Cooperative and competitive multi-agent systems
- Reward engineering and deployment challenges

Lab: Explore a policy-gradient agent in a configured simulation.

Module 7: Advanced Natural Language Processing

- NLP pipelines, preprocessing and tokenization
- Word embeddings and contextual representations
- Sentiment analysis, text classification and named entity recognition
- Semantic similarity, information extraction and evaluation

Lab: Build a pretrained-model text classification or sentiment analysis application.

Module 8: Transformers and Intelligent Language Applications

- Self-attention, encoder and decoder architectures
- BERT, GPT, transfer learning and fine-tuning concepts
- Prompt engineering, embeddings and vector similarity search
- Retrieval-Augmented Generation and application integration

Lab: Create a document question-answering prototype with semantic retrieval.

Day 2 Outcome: Implement RL agents and develop transformer-based language applications.

Day 3: Computer Vision, Responsible AI, and Integrated Project

Module 9: Advanced Computer Vision and Perception

- Computer vision workflows and image preprocessing
- Object detection with YOLO and image segmentation
- Video analysis, object tracking and action recognition
- Industrial and automation use cases

Lab: Detect objects in images or recorded video using a pretrained YOLO model.

Module 10: 3D Vision, Augmented Reality, and Intelligent Perception

- Depth estimation, stereo vision and 3D reconstruction

- Pose estimation and spatial understanding
- Augmented reality concepts and AI perception integration

Practical exercise: Explore pose or depth estimation using a pretrained model.

Module 11: AI Ethics, Governance, and Responsible AI

- Bias, fairness, transparency and explainability
- Privacy, data protection and secure AI practices
- Governance, accountability, human oversight and model monitoring

Practical exercise: Assess fairness and explainability for a sample AI model.

Module 12: Integrated Intelligent Application – Capstone Project

- Define a business use case and select suitable AI models
- Prepare data and design an intelligent application architecture
- Integrate text processing, semantic search and a pretrained language model
- Expose inference functionality through FastAPI
- Test outputs, evaluate performance and demonstrate the prototype

Capstone lab: Build an intelligent document analysis and decision-support application with a query interface, semantic retrieval and an API.

Module 13: Course Review and Knowledge Assessment

- Review deep learning, generative AI, scalable ML and reinforcement learning
- Review NLP, transformers and computer vision
- Discuss implementation risks and application design trade-offs
- Capstone presentation, feedback and final assessment

Activity: Demonstration, knowledge assessment and Q&A.

Day 3 Outcome: Develop vision capabilities, apply responsible AI principles, and present an integrated AI application.

Tools and Technologies

Python, Jupyter Notebook / Google Colab, Scikit-learn, TensorFlow, PyTorch, PySpark MLlib, Gymnasium, Stable-Baselines3, Hugging Face Transformers, Sentence Transformers, OpenCV, YOLO, FastAPI, SHAP.

Training Deliverables

- Practical labs for deep learning, distributed ML, reinforcement learning, NLP and computer vision.
- An integrated intelligent application capstone prototype.
- Course review, discussion and assessment.

Implementation note: Preconfigured datasets, pretrained models and ready-to-use lab environments are recommended to support the breadth of this accelerated course.