AI in Depth: From Fundamentals of Deep Learning to Real-World Applications

Course Overview

This 5-day mini course, *AI in Depth: From Fundamentals of Deep Learning to Real-World Applications*, offers an intensive, hands-on introduction to deep neural networks (DNNs). Participants will start with foundational concepts and small “baby applications” to quickly see deep learning in action, before progressing to advanced architectures such as convolutional neural networks, recurrent models, and transformers. Each day combines theory with practical coding exercises, enabling participants to build and train models on image, text, and multimodal datasets. Real-world case studies in either healthcare, finance, misinformation detection, and materials science provide context for applied learning, while final team-based projects allow learners to integrate their skills. The course also emphasizes responsible AI, encouraging participants to consider issues of fairness, explainability, and ethics alongside technical expertise.

Target Audience

This course is designed for advanced undergraduates, graduate students, and early-career professionals in data science, computer science, mathematics, or engineering. It is also suitable for industry professionals seeking an intensive introduction to deep learning, as well as researchers or students from other domains (e.g., social sciences, biology, physics) who wish to apply AI techniques to their data.

Prerequisites:
- Basic programming knowledge (preferably Python).
- Familiarity with basic machine learning concepts.

5-Day Course Outline

The course runs for 5 days, with two 90-minute morning sessions.

Day 1 – Foundations of Deep Learning (with intro applications)

Morning Session 1: What is AI & How Do Neural Networks Work?

- Introduction to AI & deep learning breakthroughs
- Neural network fundamentals: perceptron, layers, activation functions
- Baby Application Demo: Digit recognition with a trained MNIST model
- Interactive activity: students pick up a digit and the network predicts it

Morning Session 2: Building Your First Neural Network

- Forward and backward propagation with a small dataset
- Hands-on: build a 2-layer neural net in Python
- Baby Application Extension: MNIST dataset. Handwritten digit recognition

Day 2 – Architectures & Training

Morning Session 1

- Key architectures: MLPs, CNNs, RNNs, Transformers (overview)
- Training dynamics: loss functions, gradient descent, regularization, batch norm

Morning Session 2

- Hands-on: Training a CNN on an image dataset. Classification Cats vs Dogs
- Evaluation metrics & visualization of learning curves

Day 3 – Advanced Architectures

Morning Session 1

- Deep dives: RNNs, LSTMs, GRUs, and Transformers
- Attention mechanisms & large language models

Morning Session 2

- Hands-on: Text classification with RNN/Transformer

Day 4 – Applications Across Domains

Morning Session 1

- Applications of DNNs in computer vision, NLP, and multimodal AI
- Real-world case studies: healthcare, finance, misinformation detection, materials science

Morning Session 2

- Hands-on: Transfer learning with a pretrained model (ResNet, BERT)
- Student exercise: adapt a pretrained model to a small dataset using one of the real-world case study

Day 5 – AI in Practice: Challenges & Future Directions

Morning Session 1

- Responsible AI: ethics, fairness, bias, explainability
- Scaling deep learning: GPUs, TPUs, distributed training
- The future: generative AI, multimodal learning, foundation models

Morning Session 2

- Hands-on mini-projects: teams work on a dataset (vision, text, or tabular)
- Final wrap-up: presentations, reflections, and next steps