Deep Learning Course Content

1. Introduction to Deep Learning
2. Python, Math and Tensor Foundations (Hands on)
3. Neural Networks from Scratch (Hands on)
4. PyTorch for Deep Learning (Hands on)
5. Computer Vision with CNNs (Hands on)
6. Sequence Models
7. Transformers and Attention (Hands on)
8. Fine-Tuning Large Language Models (Hands on)
9. Generative Models
10. Training at Scale and Deployment (Hands on)
11. Capstone Projects

Introduction to Deep Learning

- AI, Machine Learning and Deep Learning: Where Each Fits
- When to Use Deep Learning and When Classical ML Is Enough
- GPUs, Cloud Notebooks and the Deep Learning Workflow

Python, Math and Tensor Foundations (Hands on)

- NumPy, Tensors, Broadcasting, Matplotlib
- Linear Algebra, Calculus and Probability for Deep Learning
- Automatic Differentiation and Gradients

Neural Networks from Scratch (Hands on)

- Perceptron, Activation Functions, Loss Functions
- Forward Pass, Backpropagation, Gradient Descent
- Optimizers: SGD, Adam, AdamW, Learning Rate Schedules
- Dropout, Batch Normalization, Weight Initialization, Regularization

PyTorch for Deep Learning (Hands on)

- Tensors, nn.Module, Datasets and DataLoaders
- Training Loops, GPU Training, Mixed Precision, torch.compile
- PyTorch Lightning; Keras and TensorFlow Overview

Computer Vision with CNNs (Hands on)

- Convolution, Pooling, ResNet, EfficientNet
- Transfer Learning, Data Augmentation, Image Classification
- Object Detection with YOLO, Segmentation Overview, Vision Transformers

Sequence Models

- RNN, LSTM and GRU
- Text and Time Series with Sequence Models
- Sequence-to-Sequence Models and the Idea of Attention

Transformers and Attention (Hands on)

- Self-Attention, Multi-Head Attention, Positional Encoding
- Encoder and Decoder Models: BERT and GPT-Style Architectures
- Tokenization, Hugging Face Transformers and Datasets

Fine-Tuning Large Language Models (Hands on)

- Transfer Learning for Text and Vision
- Parameter-Efficient Fine-Tuning: LoRA and QLoRA with PEFT
- Instruction Tuning, Preference Tuning with TRL, Evaluation
- RAG Overview and When to Fine-Tune Instead

Generative Models

- Autoencoders and Variational Autoencoders
- GANs and Diffusion Models with Diffusers
- Multimodal Models Overview

Training at Scale and Deployment (Hands on)

- Experiment Tracking with Weights & Biases and MLflow
- Distributed Training: DDP, FSDP and DeepSpeed Overview
- Quantization, ONNX Export, Serving with FastAPI, vLLM and Docker

Capstone Projects

- Image Classifier, Object Detector, Transformer Text Classifier and Fine-Tuned LLM Projects
- GitHub Portfolio, Resume Building and Interview Preparation