Deep Learning

Neural-network methods for vision, sequences, language, and structured data.

A course on how deep networks are built, trained, diagnosed, and adapted, from tensors and backpropagation to convolutional networks, recurrent models, attention, Transformers, generative models, and graph neural networks. The course connects architectures to the optimization and generalization issues that make deep learning work in practice.

TOPIC MAP

Foundations

- Tensors, preprocessing, automatic differentiation.
- Computational graphs and backpropagation.
- Initialization, optimization, regularization, and dropout.
- Loss landscapes, numerical stability, and gradient flow.

Core Architectures

- Multilayer perceptrons and modular network design.
- CNNs, RNNs, GRUs, and LSTMs.
- Encoder, decoder, and sequence-to-sequence patterns.
- Graph neural networks.

Vision Models

- LeNet, AlexNet, VGG, GoogLeNet.
- Batch normalization, ResNet, DenseNet.
- Residual learning, normalization, and depth.
- Vision Transformers and patch-based representations.

Sequences and Language

- Language modeling and sequence-to-sequence learning.
- Teacher forcing, beam search, and decoding.
- Attention scoring and alignment.
- Multi-head attention, positional encoding, and Transformers.

Modern Learning Regimes

- Transfer learning and large-scale pretraining.
- Self-supervised and fewer-label learning.
- Fine-tuning, feature reuse, and representation transfer.
- Distribution shift, adversarial examples, and continual learning.

Generative Modeling

- Autoregressive and latent-variable models.
- GANs, flows, and diffusion-style models.
- Likelihood, sample quality, and diversity trade-offs.
- Evaluation issues for generated data.

Backpropagation / CNNs / RNNs / Attention / Transformers / GNNs