

Deep Learning

How neural networks learn to predict, generate, and act, with architectures for images, sequences, and graphs.

COURSE CONTENT

Neural-network foundations	Tensors, automatic differentiation, linear regression, softmax classification, and multilayer perceptrons.
Training and generalization	Backpropagation, initialization, gradient-based optimization, weight decay, dropout, and overfitting.
Convolutional networks	Convolutions, pooling, LeNet, residual networks, and batch normalization.
Sequences and attention	RNNs, LSTMs, GRUs, encoder-decoder models, attention, and Transformers.
Generative models	Variational autoencoders, diffusion, text-to-image generation, guidance, and sample-quality evaluation.
Reinforcement learning	Markov decision processes, Q-learning and DQN, policy gradients, actor-critic methods, and PPO.
Learning with fewer labels	Data augmentation, transfer learning, contrastive and semi-supervised learning, and meta-learning.
Graph neural networks	Message passing, graph convolutions, GraphSAGE, and graph attention.

LEARNING FOCUS

Relate architectures and learning objectives to training behavior, data structure, and decision-making.

Mathematical explanations, architecture diagrams, and Python/PyTorch examples.