

Introduction to Deep Learning

Introduction to Deep Learning

Key concepts, building blocks, and practical considerations

1. What Is Deep Learning?

Aspect	Description
Definition	A sub?field of machine learning that uses artificial neural networks with many layers (hence “deep”) to automatically learn hierarchical representations from data.
Why “deep”?	Depth refers to the number of hidden layers. More layers enable the model to capture increasingly abstract features (e.g., edges ? shapes ? objects).
Relationship to AI	Deep learning is a technique within artificial intelligence (AI) and a subset of machine learning (ML).
Typical Applications	Image & video analysis, speech & natural?language processing, recommendation systems, game playing, drug discovery, autonomous driving, etc.

2. Core Components of a Deep Neural Network (DNN)

Component	Role	Common Choices
Neuron (unit)	Computes a weighted sum of inputs, adds a bias, and applies a non?linear activation.	Linear, sigmoid, tanh, ReLU, Leaky ReLU, ELU, Swish
Layer	A collection of neurons that operate on the same input vector.	Dense (fully connected), Convolutional, Recurrent, Transformer, Pooling, Normalization
Loss Function	Quantifies the discrepancy between predictions and ground truth; guides learning.	MSE, Cross?entropy, Hinge, KL?divergence, CTCLoss
Optimizer	Updates weights using gradients of the loss.	SGD, Momentum, Adam, RMSprop, AdamW

Introduction to Deep Learning

Component	Role	Common Choices
Regularization	Prevents overfitting and improves generalization.	L1/L2 weight decay, Dropout, BatchNorm, Data augmentation, Early stopping
Training Loop	Repeatedly: forward pass → compute loss → backpropagation → weight update.	Mini-batch training, epochs, learning rate schedules

3. Popular Network Architectures

Architecture	Main Idea	Typical Use Cases
Multilayer Perceptron (MLP)	Stacks of dense layers; universal function approximator.	Tabular data, simple classification/regression.
Convolutional Neural Network (CNN)	Convolutions exploit spatial locality; weight sharing reduces parameters.	Image classification, object detection, video analysis, medical imaging.
Recurrent Neural Network (RNN)	Processes sequences step-by-step, maintaining hidden state.	Language modeling, time series forecasting, speech recognition.
Long Short-Term Memory (LSTM) / GRU	RNN variants that mitigate vanishing gradients via gated cells.	Machine translation, caption generation, music modeling.
Transformer	Self-attention mechanism replaces recurrence; enables parallel processing of sequences.	NLP (BERT, GPT), vision transformers, multimodal models.
Autoencoder	Learns compact latent representation by reconstructing input.	Dimensionality reduction, denoising, anomaly detection.
Generative Adversarial Network (GAN)	Two networks (generator & discriminator) contest in a zero-sum game.	Image synthesis, style transfer, data augmentation.
Graph Neural Network (GNN)	Operates on graph-structured data using message passing.	Social network analysis, molecular property prediction.

Introduction to Deep Learning

4. Training Deep Networks – Practical Tips

1. Data Matters Most

- Gather diverse, high-quality labeled data.
- Apply preprocessing (normalization