Deep Learning Fundamentals: Models & Core Concepts

Welcome to Week 2! Today we'll explore the foundations of deep learning, compare it with traditional machine learning, and understand key concepts that power modern AI systems.

by S MM

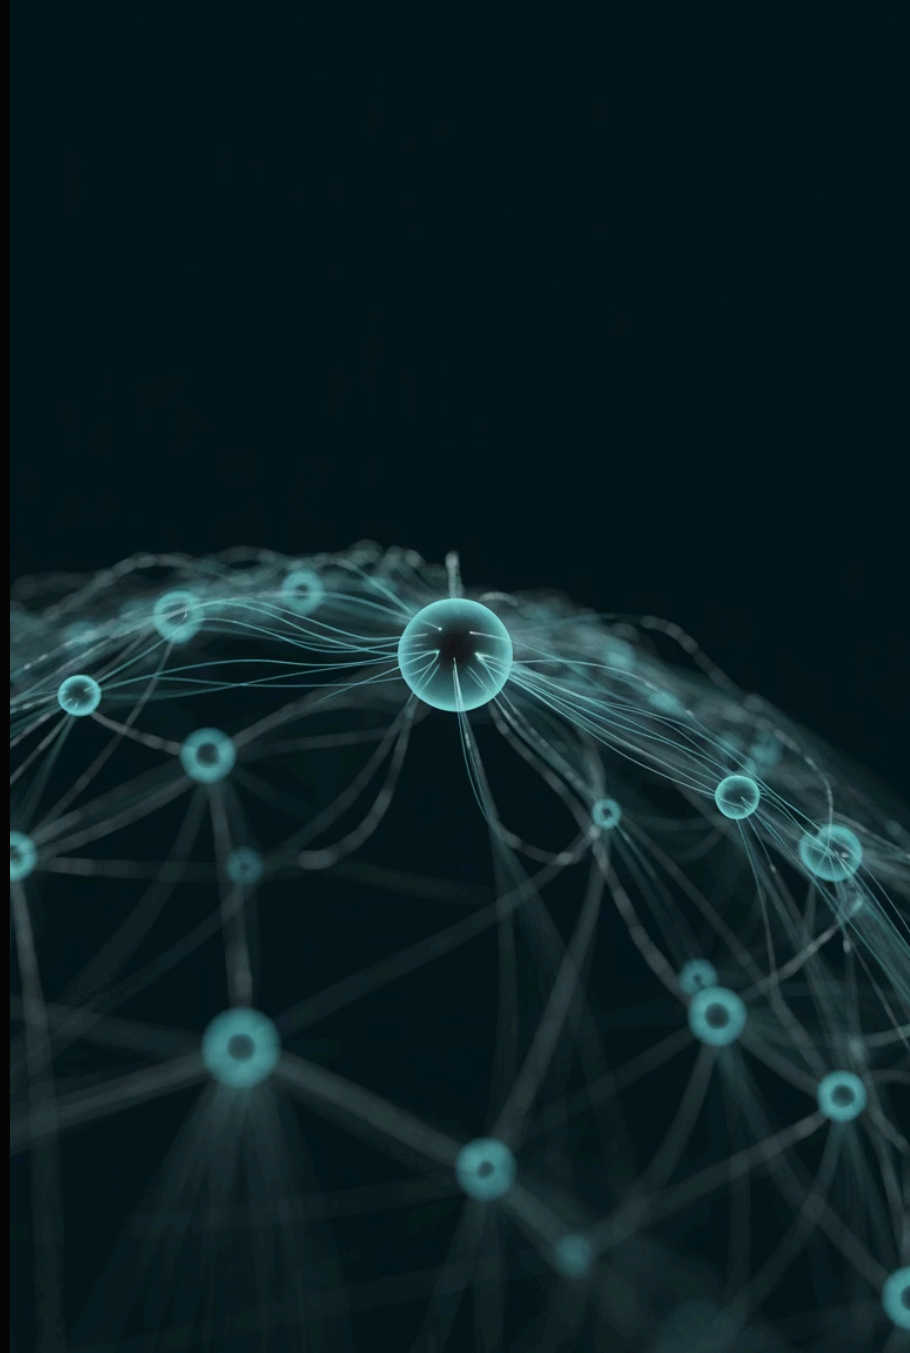

Deep Learning vs. Machine Learning

Machine Learning

- Requires feature engineering
- Often uses structured data
- Linear models, trees, SVMs

Deep Learning

- Automatic feature extraction
- Excels with unstructured data
- Multiple processing layers

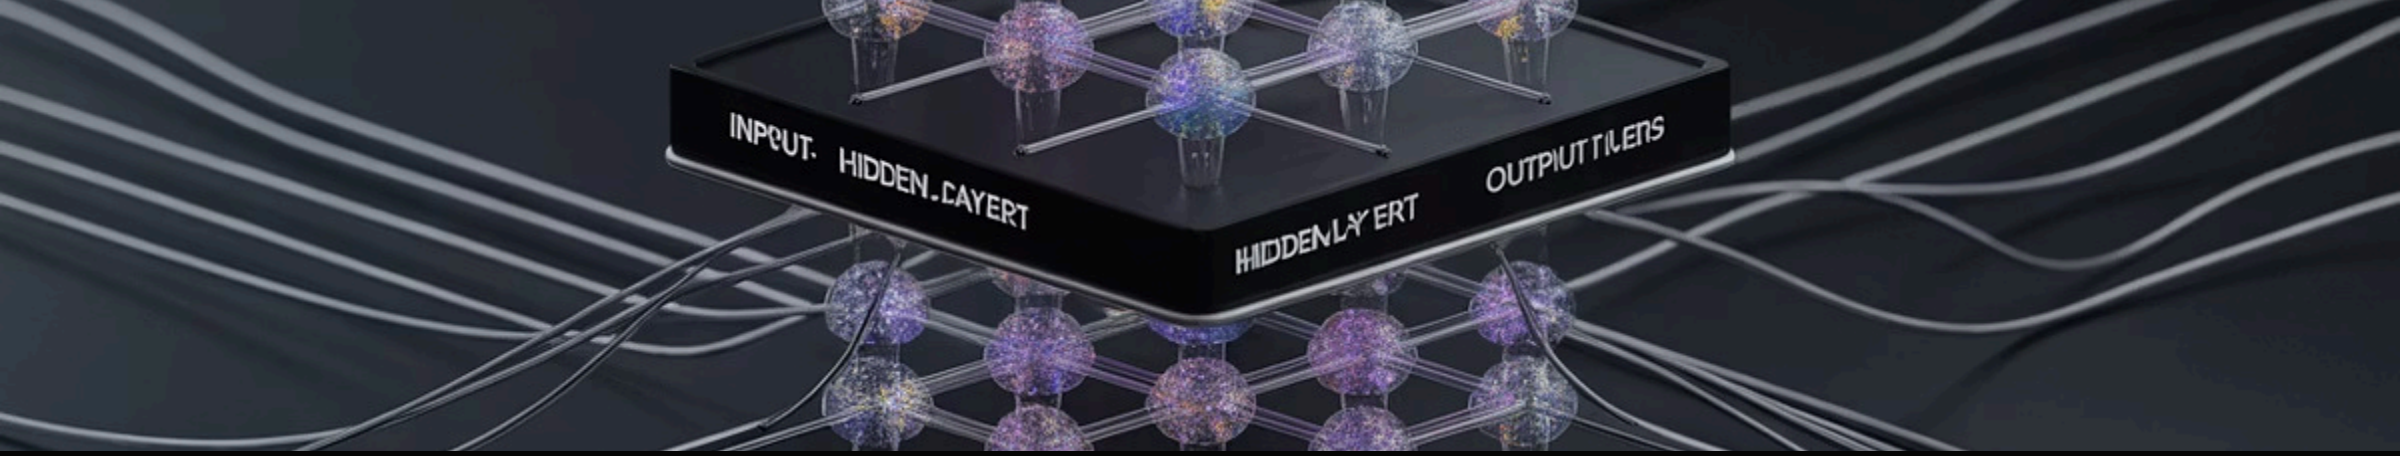

Neural Network Architecture

Input Layer

Receives raw data. Each neuron represents one feature in your dataset.

Hidden Layers

Process information through weighted connections. More layers enable learning complex patterns.

Output Layer

Produces final predictions or classifications based on processed information.

Training Process: Backpropagation

Forward Pass

Data flows through network. Model makes prediction.

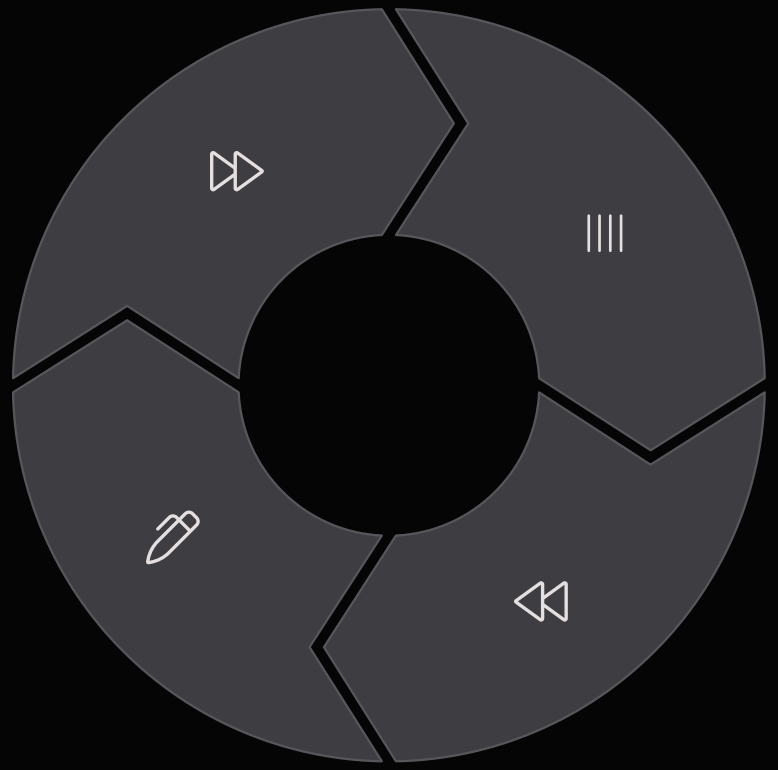

Error Calculation

Compare prediction with actual values. Measure the difference.

Backward Pass

Error propagates backward. Gradients calculated for each weight.

Weight Update

Adjust weights using gradient descent to minimize future errors.

Activation Functions

ReLU

Returns x if positive, 0 if negative. Solves vanishing gradient problem.

Sigmoid

Squashes values between 0 and 1. Used in binary classification.

Softmax

Converts values to probabilities. Perfect for multi-class classification.

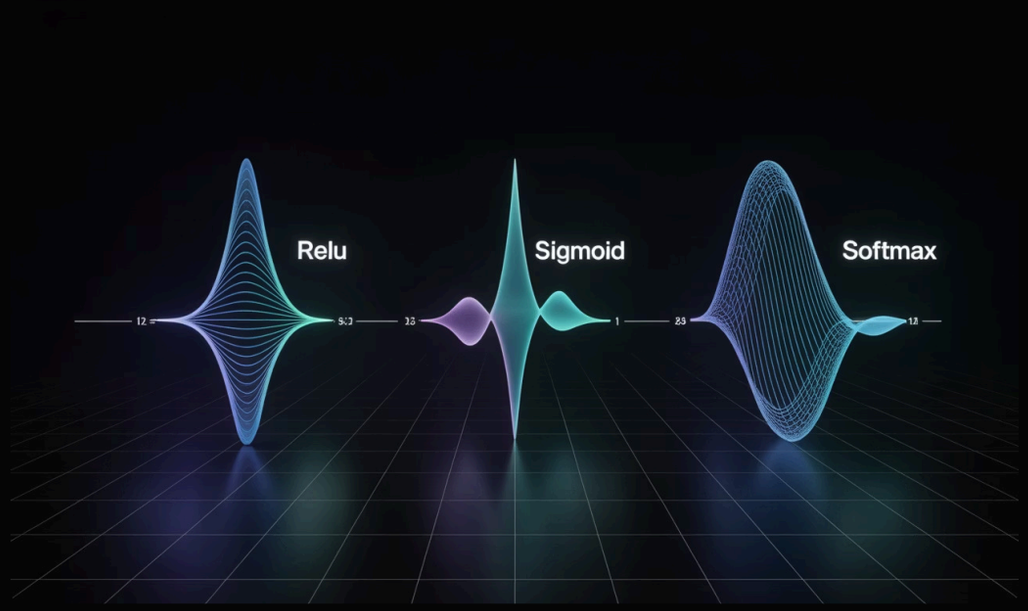

Model Fitting Challenges

Underfitting

Model is too simple. Cannot capture data patterns. High bias.

Solution: More complex model, fewer regularization constraints.

Overfitting

Model memorizes training data. Poor generalization. High variance.

Solution: Dropout layers, more training data, regularization.

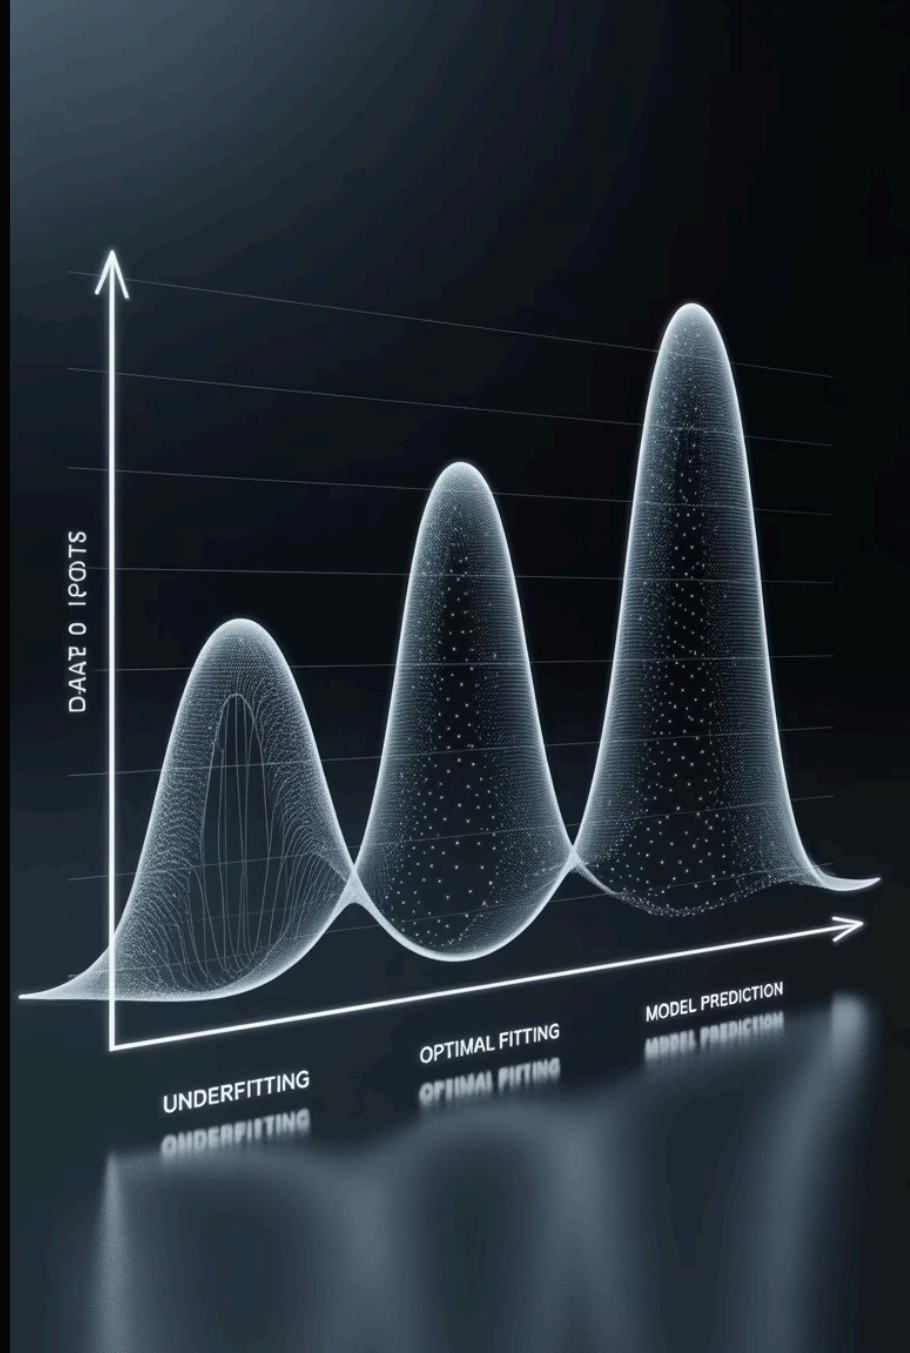

Dropout Layers

How Dropout Works

- Randomly deactivates neurons during training
- Typically drops 20-50% of connections
- Forces network to learn redundant representations
- Prevents co-adaptation of neurons

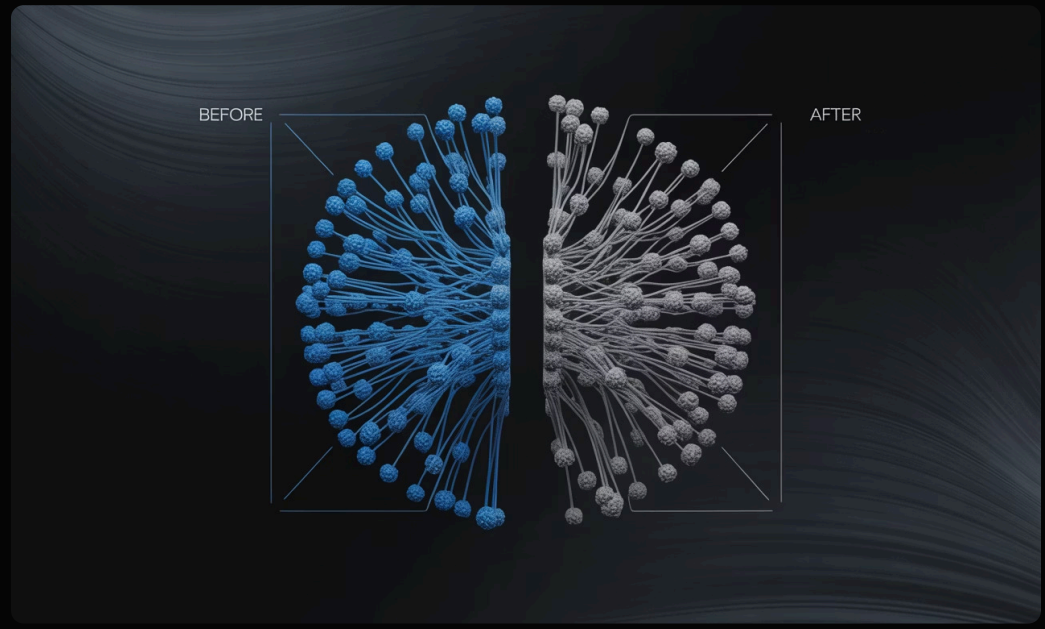

During testing, all neurons are active but outputs are scaled proportionally.

Popular Deep Learning Frameworks

TensorFlow

Google's framework. Production-ready. Excellent deployment options. TensorFlow.js for web.

PyTorch

Facebook's framework. Dynamic computation graphs. Popular in research. Pythonic design.

Keras

High-level API. Works with multiple backends. Rapid prototyping. Beginner-friendly.