Module 7: Deep Learning

Data Science with Python – BSc Course

Why Deep Learning?

Citadel and Two Sigma collectively manage over $100 billion and invest heavily in neural networks for alpha generation. JPMorgan's LOXM uses deep learning to execute trades optimally, minimizing market impact.
From simple perceptrons in the 1950s to networks that process satellite imagery of parking lots, credit card transactions at millisecond speed, and social media sentiment at scale — deep learning is transforming finance.
Traditional models assume linear relationships or simple decision boundaries. Neural networks learn complex, non-linear patterns that even experts struggle to define manually.

Deep learning captures complexity that traditional models miss

Why This Matters

- **Alternative data processing:** Satellite images, ship tracking, social media sentiment
- **High-frequency trading:** Pattern recognition in order flow, millisecond execution
- **Credit scoring:** Non-linear relationships that logistic regression misses
- **Language understanding:** Earnings calls, analyst reports, SEC filings (foundation for NLP)

Neural networks handle complexity at scales traditional ML cannot match

What You'll Learn

By the end of this module, you will be able to:

- Understand the perceptron as the building block of neural networks
- Build multi-layer perceptrons with appropriate activation functions
- Implement backpropagation and gradient descent for training
- Prevent overfitting with dropout, early stopping, and regularization
- Apply neural networks to financial prediction problems

From single neurons to networks that learn complex patterns

Lesson Roadmap

Lesson	Topic	Narrative Arc	Focus
L33	Perceptron	"The Artificial Neuron"	Biology to math
L34	MLPs & Activations	"Stacking Legos"	Layers, ReLU, sigmoid
L35	Backpropagation	"Learning from Mistakes"	Gradient descent, training
L36	Overfitting Prevention	"Too Good to Be True"	Dropout, early stopping, L2

Module Story: This module tells one complete story: how to build a brain from scratch. L33 creates a single neuron. L34 stacks neurons into powerful networks. L35 teaches the network to learn from its mistakes. L36 prevents the network from learning TOO well. By the end, you have a complete deep learning toolkit.

Four lessons, one toolkit — from single neuron to regularized deep networks

Key Concepts

- **Perceptron & Linear Boundaries** – Single neuron as a linear classifier
- **Multi-Layer Perceptrons** – Stack layers to learn non-linear patterns
- **Activation Functions** – ReLU, sigmoid, tanh — introducing non-linearity
- **Backpropagation & Gradient Descent** – How networks learn from errors
- **Overfitting Prevention** – Dropout, early stopping, L1/L2 regularization

Master these and you understand how modern AI systems learn

Finance Application

Scenario: Non-Linear Default Predictor
Using the techniques from this module, you will:

- Build a neural network that captures interaction effects between borrower features
- Learn patterns that logistic regression cannot represent (e.g., income$\times$debt interactions)
- Apply dropout and early stopping to prevent overfitting on historical data
- Compare performance against traditional logistic regression

This is where deep learning shines — capturing the complex, non-linear relationships that drive real-world outcomes.

When linear models plateau, neural networks break through

Who Uses This?

- **Quantitative Funds** – Citadel, Two Sigma, D.E. Shaw use neural networks for alpha generation
- **Trading** – JPMorgan's LOXM executes trades optimally using deep learning
- **Alternative Data** – Orbital Insight analyzes satellite imagery with neural networks
- **Fintech** – Zest AI uses deep learning for fair lending and credit scoring

Deep learning powers the quantitative revolution in finance

What's Next: Module 8 – NLP & Text Analysis

Neural networks can process numbers — stock prices, balance sheets, transaction amounts. But finance is full of text: 300,000+ financial news articles published daily, earnings calls, analyst reports, SEC filings.

Module 8 introduces Natural Language Processing: teaching machines to read, understand, and extract insights from text. Sentiment analysis, document classification, embeddings — all built on neural network foundations.

The deep learning techniques you learn here are the building blocks for language understanding.

Numbers tell part of the story — text tells the rest

Let's Begin!

First up: L33 – Perceptron

From a single neuron to networks that learn the world.