

Lesson 34 Summary: MLP & Activations

Data Science with Python – Key Concepts

Data Science Program

Multi-Layer Perceptron & Activation Functions

Architecture

Input layer
Hidden layers
Output layer

Activations

ReLU: $\max(0, x)$
Sigmoid: $1/(1+e^{-x})$
Tanh: $(e^x - e^{-x}) / (e^x + e^{-x})$

Why Non-Linear?

Learn complex patterns
XOR solvable
Universal approx.

ReLU Benefits

Fast computation
No vanishing grad
Sparse activation

Output Layer

Binary: Sigmoid
Multi-class: Softmax
Regression: Linear

Finance Apps

Price prediction
Risk scoring
Pattern detection

MLPs use non-linear activations to learn complex patterns

Network architecture:

- **Input layer:** Receives features
- **Hidden layers:** Learn representations
- **Output layer:** Produces predictions

Each layer transforms data for the next

Why Non-Linear Activations?

The key insight:

- **Linear only:** Collapses to single layer
- **Non-linear:** Enables complex functions
- **Universal approximation:** Can learn any function

Without non-linearity, depth provides no benefit

Most popular hidden activation:

- **Formula:** $f(x) = \max(0, x)$
- **Benefit:** Fast, no vanishing gradient
- **Issue:** Dead neurons (always 0)

Leaky ReLU fixes dead neuron problem

Classic activations:

- **Sigmoid:** $\sigma(x) = 1/(1 + e^{-x})$ outputs 0-1
- **Tanh:** Outputs -1 to 1, zero-centered
- **Issue:** Vanishing gradients at extremes

Sigmoid still used for binary output layers

Output Layer Activations

Depends on task:

- **Binary classification:** Sigmoid
- **Multi-class:** Softmax (probabilities sum to 1)
- **Regression:** Linear (no activation)

Match activation to your prediction task

Theoretical foundation:

- **Theorem:** MLP can approximate any function
- **Requirement:** Sufficient hidden neurons
- **Practice:** Finding right architecture is art

More layers often better than more neurons

Deep learning in finance:

- **Price prediction:** Regression with linear output
- **Credit scoring:** Classification with sigmoid
- **Portfolio:** Multi-output with softmax

Choose architecture based on prediction task

Activation Functions:

Function	Range	Use
ReLU	$[0, \infty)$	Hidden layers
Sigmoid	$(0, 1)$	Binary output
Tanh	$(-1, 1)$	Hidden (legacy)
Softmax	$(0, 1)$ sum=1	Multi-class output
Linear	$(-\infty, \infty)$	Regression output

ReLU for hidden, task-specific for output