

Neural Networks

CPSC 383: Explorations in Artificial Intelligence and Machine Learning Fall 2025

Jonathan Hudson, Ph.D.
Associate Professor (Teaching)
Department of Computer Science
University of Calgary

August 27, 2025

Copyright © 2025



Neural Networks


Neurons

The basic unit of connectionist AI is the **neuron**.

It is designed to be a digital approximation of the neurons in our brains.

What are neural networks?


- Supervised machine learning technique (train to memorize model for $X \rightarrow Y$ dataset)
- Inspired by the Human Brain.
- The human brain has about 86 Billion neurons and requires 20% of your body's energy to function.
- Deep learning neural networks really popular right now
 - LLMs and generative AI! (Chat-GPT, Gemini, Co-Pilot, etc.)
- Connectionist method
 - Make network, train, hope result is useful





Neural Model


Neuron Model of Connections

- Developed to mimic the human neural system (in the brain) and its processing capabilities
- Decentralized knowledge representation and processing
👉 hopefully very efficient
- Simple components, the intelligence is in the **connections**










Single neuron

In this representation, a neuron consists of the following:


- **Input** values $x_1, x_2, x_3, \dots x_n$
- **Weights** $w_1, w_2, w_3, \dots w_n$
- A **bias** b
- An **activation function**, which in this case is

$$f(z) = \begin{cases} 1 & \text{if } z \geq 0 \\ 0 & \text{otherwise} \end{cases}$$

Exercise – What is output?



Exercise2 – What is output?



Exercise

- Suppose you are trying to train a ML model that can tell the difference between a giraffe and a horse based on height and weight.
- Can you design a single neuron that will label an animal as a giraffe if it has a height greater than 2?



Neuron Model Class

Model class

Suppose we are working with a **single** neuron.

- This is a model class for binary classification

What are the model parameters, and what types of patterns in data can we capture?

Model class

Suppose we are working with a **single** neuron.

- This is a model class for binary classification

What are the model parameters, and what types of patterns in data can we capture?

$$w_1, w_2, w_3, \dots, b$$

Patterns are hyperplanes $\rightarrow$ 2D

$$y \geq mx + b$$

Model class

2D line example

$$x_1 w_1 + x_2 w_2 - b \geq 0$$

$$x w_1 + y w_2 - b \geq 0$$


$$y w_2 \geq -x w_1 + b$$

$$y \geq -\frac{w_1}{w_2} x + \frac{1}{w_2} b$$


$$y \geq mx + b$$

Exercise

Draw a neuron that models the following linear classifier, given by the equation $y \geq \frac{1}{2}x + 1$



Exercise





$$y \geq \frac{1}{2}x + 1$$

$$2y \geq x + 2$$

$$2y - x - 2 \geq 0$$

Example

Consider the following two datasets. Can they be classified correctly by a single neuron?



Neural Networks

TensorFlow

To build neural nets, we will be using the **TensorFlow** library for Python 3:

- <https://www.tensorflow.org/guide/basics>

Keras is the high-level API for TensorFlow, and is automatically included with the installation:

- <https://www.tensorflow.org/guide/keras>

Make sure you are using TensorFlow 2 for your work in this course.

TensorBoard

TensorBoard is a visualization toolkit for TensorFlow:

- <https://www.tensorflow.org/tensorboard>
- It has way more features than we will be looking at in this course!

You will be using it a bit on your assignment to compare different models that you create for a dataset.

- Visualization will be hosted by default on port 6006

Neurons

- A **neural network** is a model class made up by connecting multiple neurons. This network is then trained to give a model that can be used to make predictions.

Question

- What do ogres and neural nets have in common?




Neural nets

A **neural network** is typically made by connecting neurons in layers:

- **Input** layers simply pass in the input features
- **Output** layers contain the neurons that output the result
- **Hidden** layers are any additional layers of neurons between the input and output

A **deep** neural network contains two or more hidden layers.

Neural nets



Neural nets

When connecting the neurons in layer k to the neurons in layer $k + 1$, we pass the output of every neuron in layer k as an input to every neuron in layer $k + 1$.

If we have four neurons in layer 3 and six neurons in layer 4, how many weights are there in the 4th layer in total?


Example

Consider the neural net on the following slide.

How many layers are there, and how many neurons in each layer?


What is the output of the neural net on the input (1,2)?

- Assume step function activations for now




Question

How does changing the number of neurons and width/depth of the net change the model class we are capturing?



Exercise

Consider the following neuron:



Neuron logic



The neuron on the previous slide computes the binary AND function!

- You can also similarly create neurons for OR, NOT, and the identity function

This gives us a nice interpretation for what the hidden layers are doing.

Example

For the XOR dataset, one neural net that can classify it would be:




Simulation

<https://playground.tensorflow.org/>

Exercise

How many neurons would we need to model the following function?

- Lines are $y \geq 2x - 1$ and $y \leq -x + 3$




Deep Neural Nets

Deep neural nets

Each layer added to the neural net allows the model to capture more complex functions.


- Width also helps, but in a different way



Deep neural nets

Increasing the number of neurons in a layer can also help to reduce the chance of getting stuck at local minima

- Effectively turns minima into saddle points



Deep neural nets


Having more layers in a neural net increases the risk of overfitting.

There is also a significant cost associated with adding layers to the net

- Even commercial nets are usually only a dozen layers deep
- Width is much less expensive than depth


Deep neural nets

As part of the training process, we need to determine how adjusting each weight affects the net's output:



Deep neural nets

As part of the training process, we need to see how adjusting each weight affects the net's output:



Deep neural nets

The amount of time needed to compute each partial derivative looks like it should be exponential in the depth of the net.

Instead, it is possible to reuse most of the computation using an algorithm called **backpropagation**.

- Based on the chain rule from calculus


But adding depth still has other problems!

History Reminder

A short history of Neural Networks

- 1957: Perceptron (Frank Rosenblatt): one layer network neural network
- **First AI Winter (General problem solvers)**
- 1988: Backpropagation (faster training)
- 1989: ALVINN: autonomous driving car
- 1989: (LeCun) Handwritten ZIP codes on mail
- **Second AI Winter (Expert systems)**
- 2012: Convolutional neural networks (vision)
- 2010: Deep learning
- 2017: Transformers (Like Chat-GPT)
- 2020s: LLM and Generative AI enabling
 - (near-human capabilities for image recognition, speech recognition, and language translation)

History



Connectionist vs Symbolic AI

Cardon, D., Cointet, J.-P. & Mazieres, A. (2018). Neurons spike back. The invention of inductive machines and the artificial intelligence controversy. Réseaux, 36(211), 173-220.

Next...data processing

Jonathan Hudson, Ph.D.
jwhudson@ucalgary.ca
<https://cspages.ucalgary.ca/~jwhudson/>



UNIVERSITY OF
CALGARY