Machine Learing: Neural Networks

CPSC 501: Advanced Programming Techniques
Winter 2025

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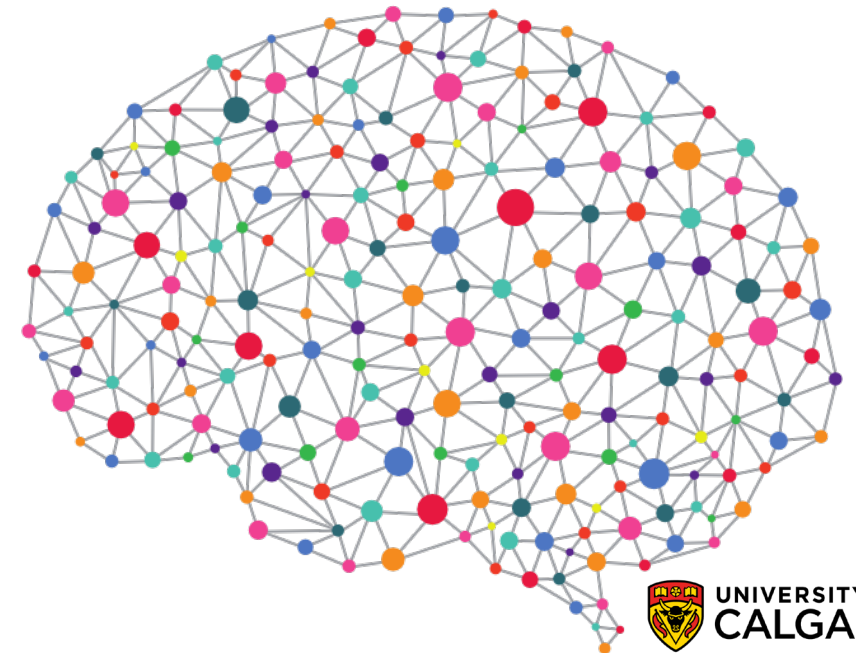University of Calgary

Friday, February 21, 2025

UNIVERSITY OF CALGARY

What are neural networks?

- Inspired by the Human Brain.

- 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

UNIVERSITY OF
CALGARY

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

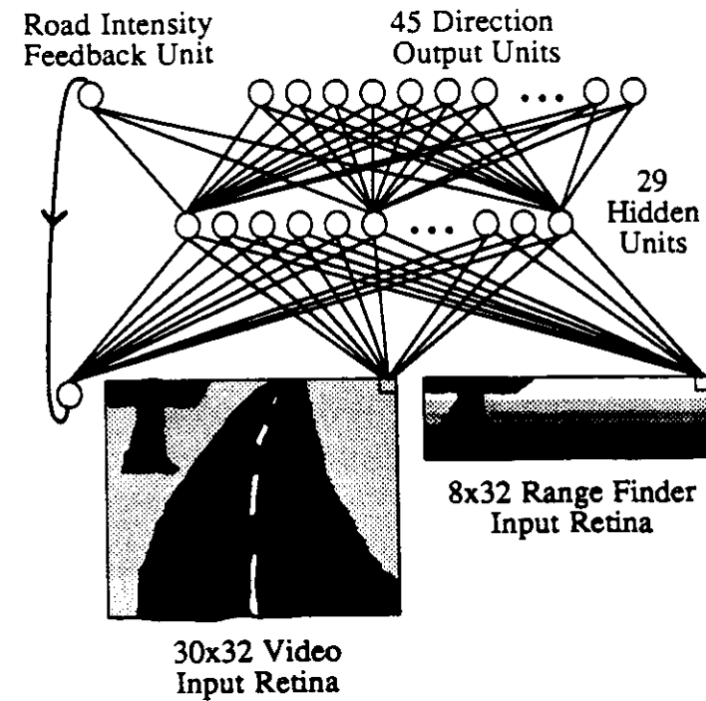- Simple components, the intelligence is in the connections
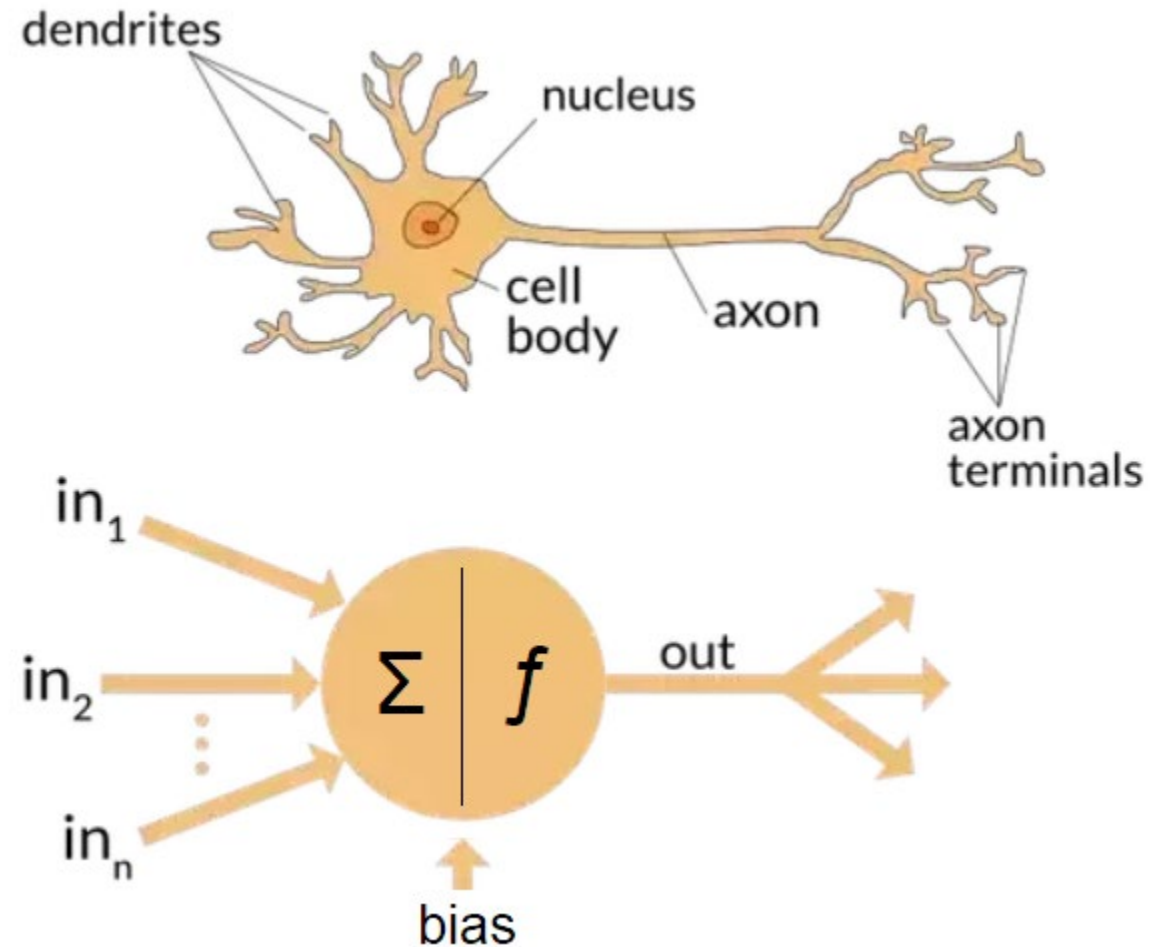
A short history of Neural Networks

- 1957: Perceptron (Frank Rosenblatt): one layer network neural network

- 1959: first neural network to solve a real world problem
  - i.e., eliminates echoes on phone lines (Widrow & Hoff)

- **First AI Winter**

- 1988: Backpropagation (Rumelhart, Hinton, Williams): learning a multi-layered network

- **Second AI Winter**

UNIVERSITY OF CALGARY

A short history of NNs

- 1989: ALVINN: autonomous driving car using NN (CMU)

Road Intensity Feedback Unit

45 Direction Output Units

29 Hidden Units

8x32 Range Finder Input Retina

30x32 Video Input Retina

UNIVERSITY OF CALGARY

A short history of NNs

- 1989: (LeCun) Successful application to recognize handwritten ZIP codes on mail using a "deep" network

- 2012-2015 convolutional neural networks (CNN)

- 2015 – AlphaGo Competition using CNN and reinforcement learning

- 2015+: near-human capabilities for image recognition, speech recognition, and language translation

- 2018 (AlphaFold) Google's protein folding prediction

- 2019 (AlphaStar) Google's StarCraft 2 AI better than 99.8% of human players (GrandMaster level)

- 2022 MidJourney, Chat-GPT – generative AI for text and images go mainstream

UNIVERSITY OF CALGARY

Basic data structures

Directed, weighted graph:

- Nodes represent a function (activation function f_{act}) with n arguments, if n links lead into the node, producing one result
  - Input nodes: take values from outside
  - Output nodes: represent activation values for different concepts to detect
  - Inner nodes: usually organized in layers (hidden layers)
- Labeled weighted links
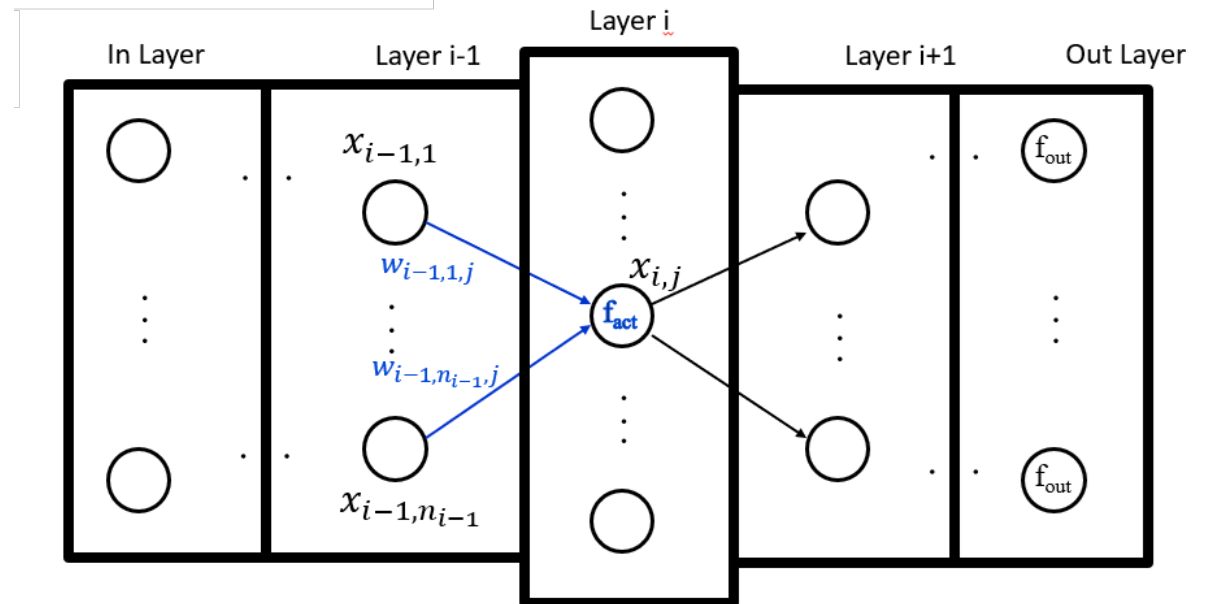
Semantics

- Whole **net** represents **a decision function f** linking input nodes to output nodes

1. We connect a bunch of nodes together between these two end points in a directed graph
   - We give the connections weights and nodes trigger sub-f_{act} functions
   - This makes the middle of the f function 'complicated'
2. Then we give it inputs with expected outputs,
   - **If it is wrong**, then we **change the weights** inside graph
3. We do this until the function seems to be right a lot in the future
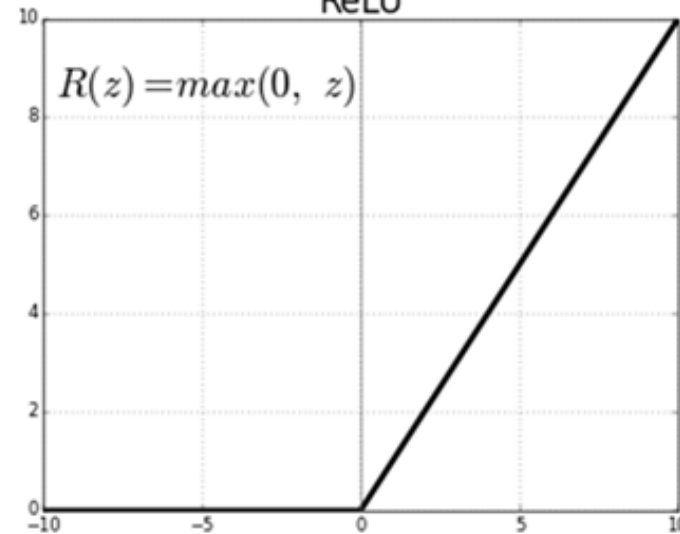
Activation Functions

Sigmoid

$$\phi(z) = \frac{1}{1 + e^{-z}}$$
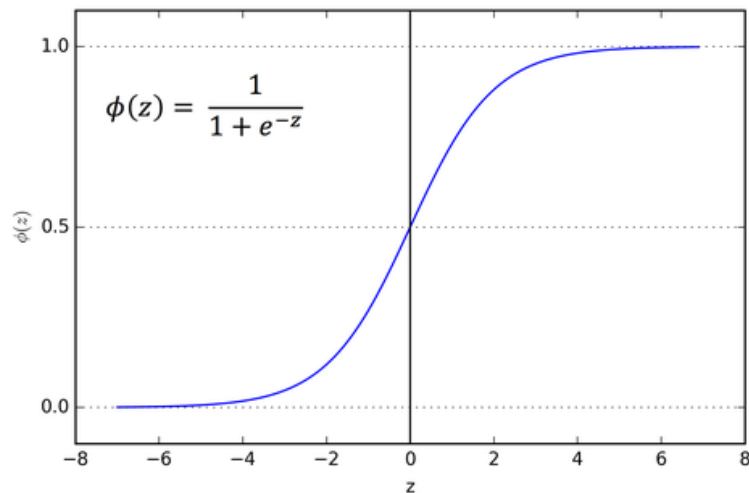
Step sigmoid
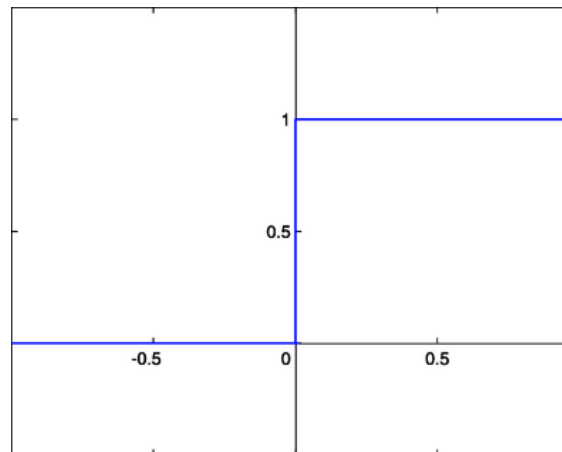
Linear Function

ReLU

$$R(z) = max(0, \ z)$$
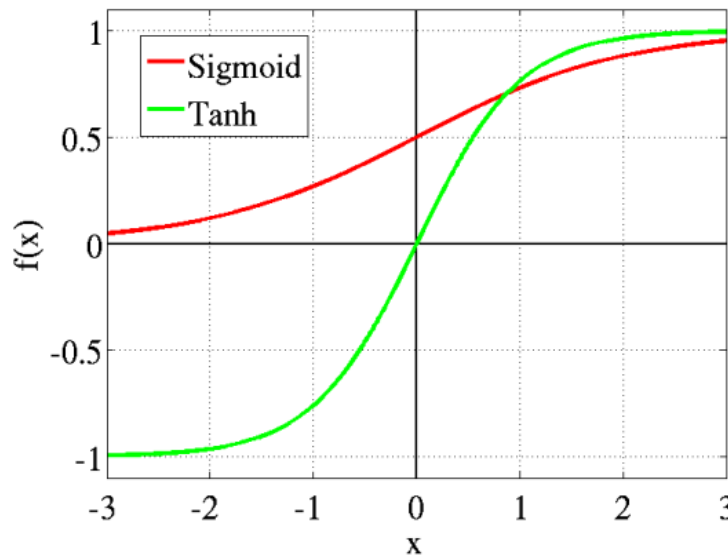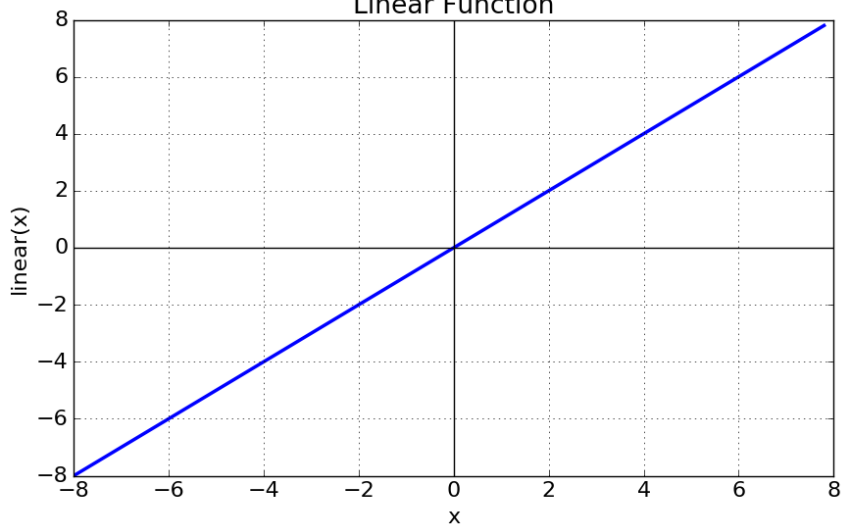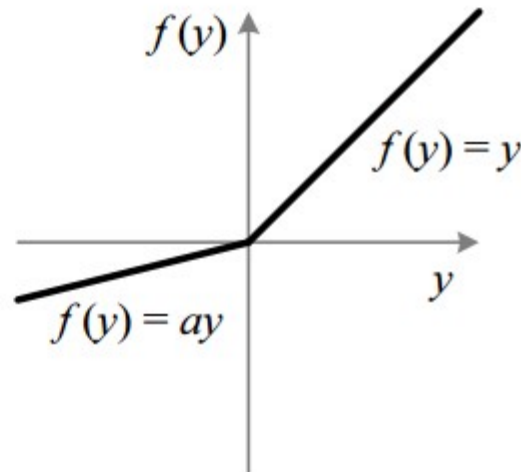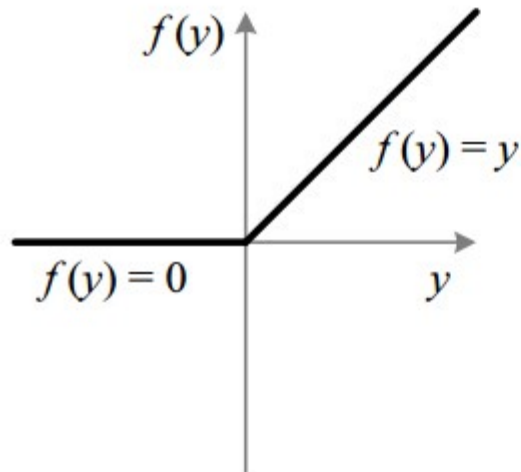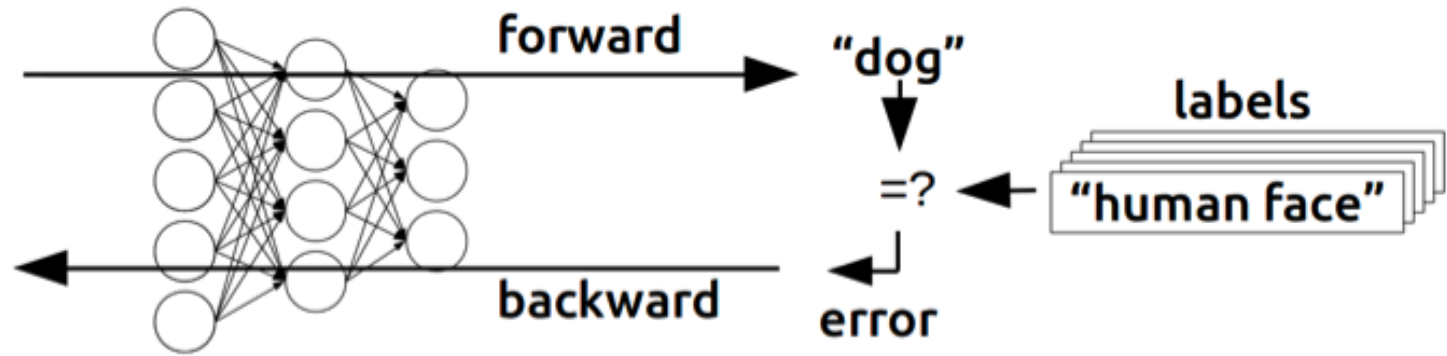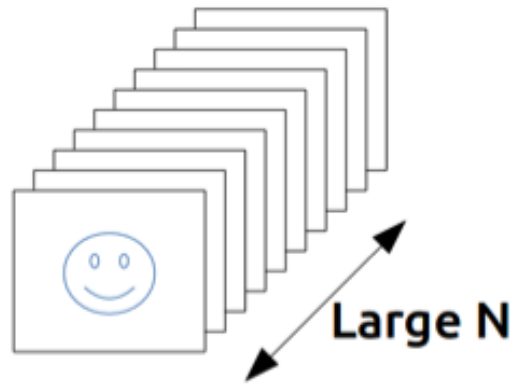
Leaky ReLU Activation Function

- It is an attempt to solve the dying ReLU problem

- The leak helps to increase the range of the ReLU function. Usually, the value of **a** is 0.01 or so.
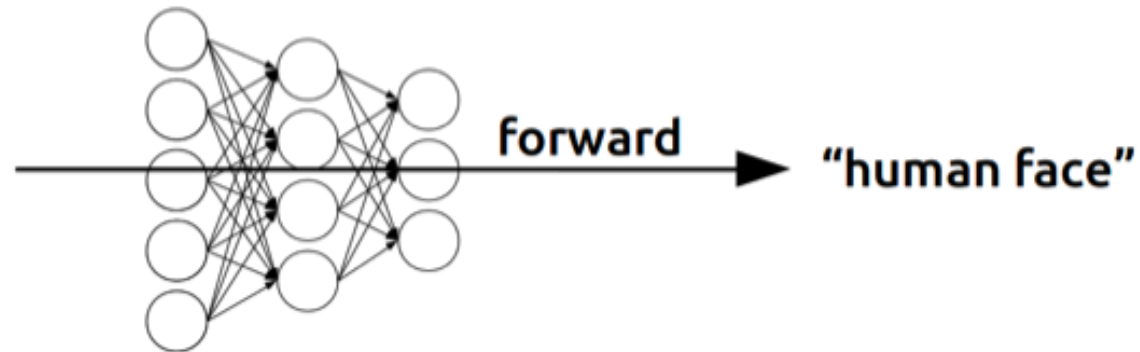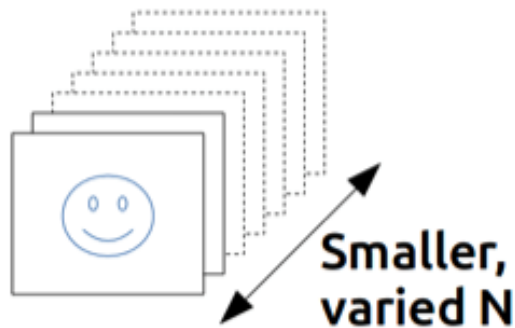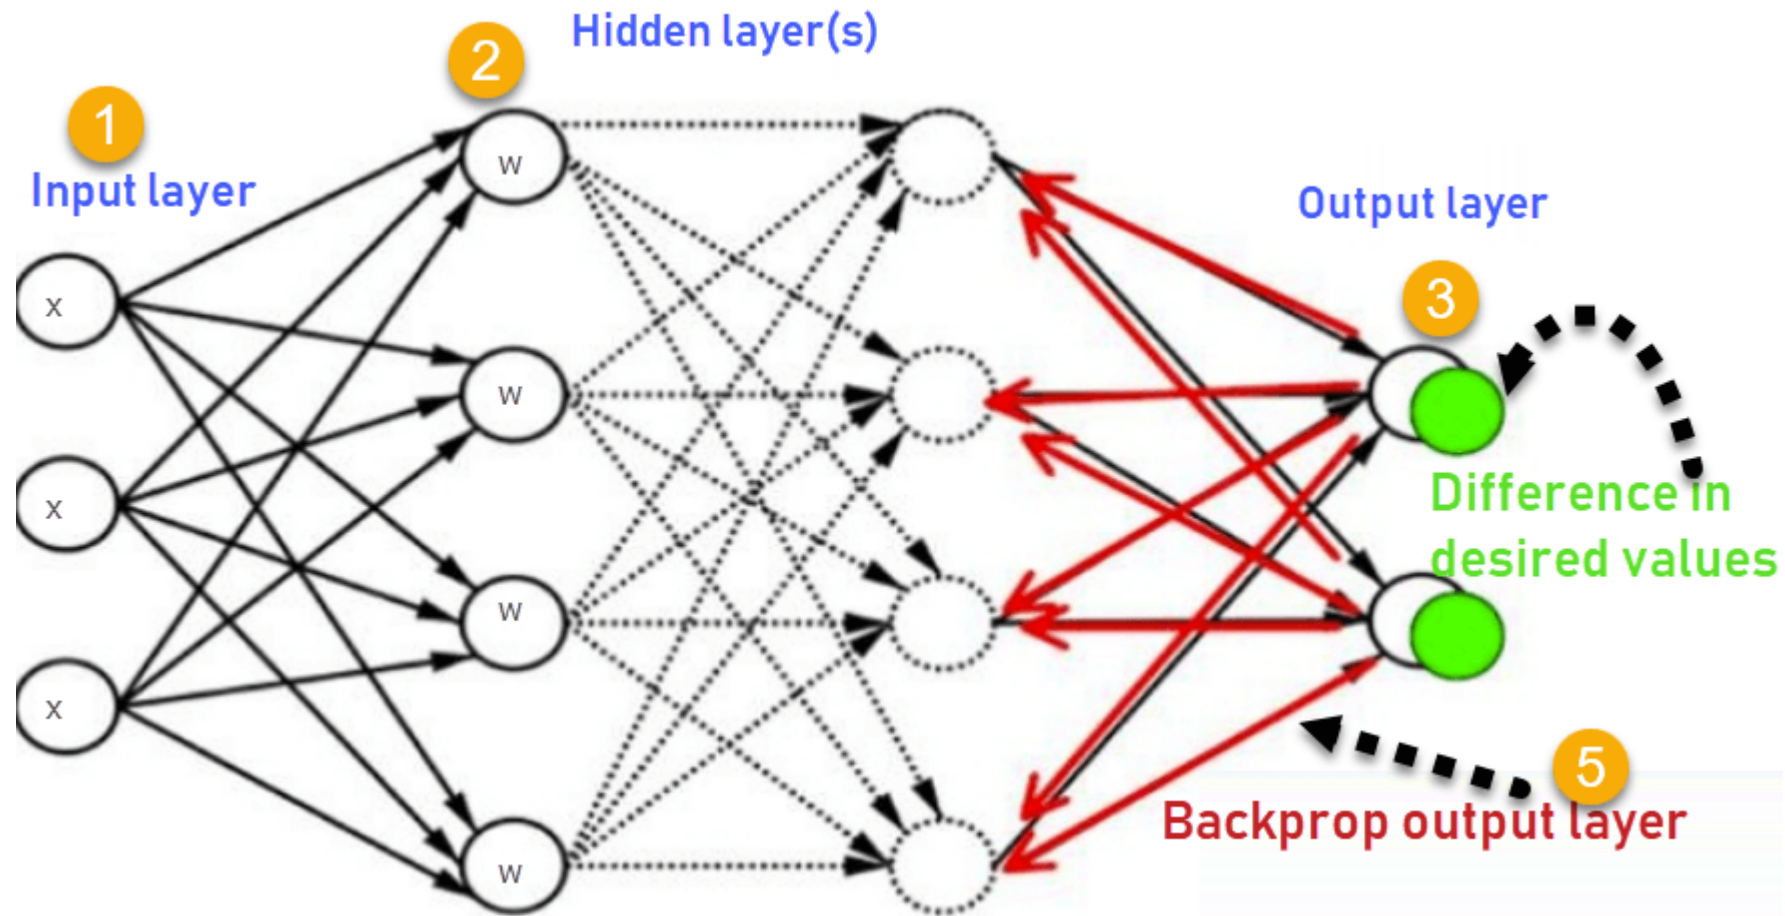
Learning: Backpropagation

From: http://www.andreykurenkov.com/writing/ai/a-brief-history-of-neural-nets-and-deep-learning/

Back-Propagation

Stochastic Gradient Descent

- Gradient descent -> follow slope to best

- Stochastic -> random influence

- 1950s, Frank Rosenblatt used SGD to optimize his perceptron

- 2014, Adam (for "Adaptive Moment Estimation") was published
  - Most machine learning libraries are dominated by Adam-type optimizers

UNIVERSITY OF CALGARY

Onward to ... Tensorflow

Jonathan Hudson
jwhudson@ucalgary.ca
https://pages.cpsc.ucalgary.ca/~jwhudson/

UNIVERSITY OF
CALGARY