APS360 Fundamentals of AI

Lisa Zhang

Lecture 2; May 9, 2019

Agenda

Last Class:

- Introduction
- Biological Neurons
- Mathematical Model of an Artificial Neuron

Today:

- Train our first neural network
- Training and test sets

Review: Biological Neuron

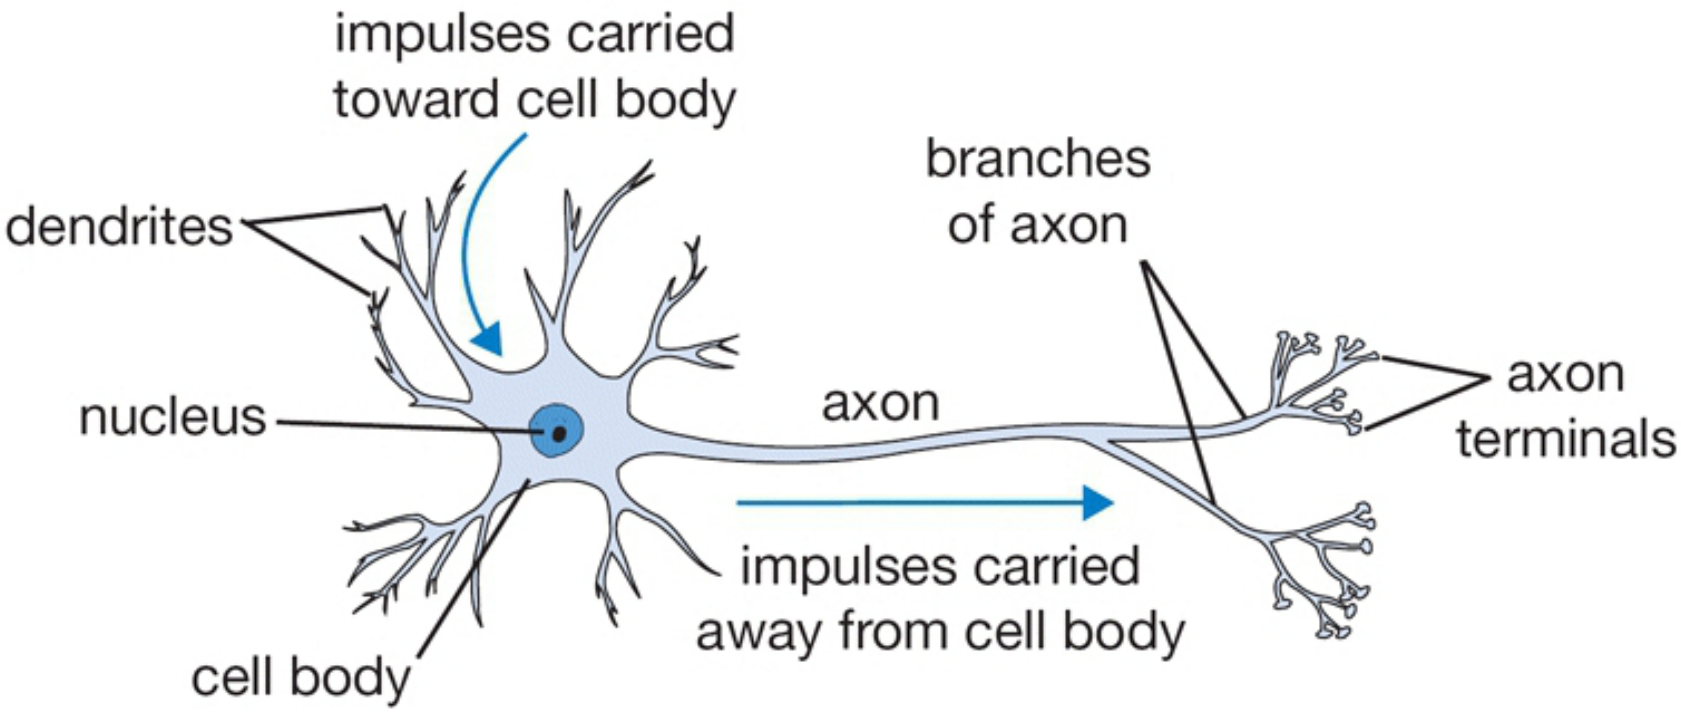

Review: Information flow

- (Axon of preivous cell)
- (Synapse)
- Dendrite
- Cell Body
- Axon
- Synapse
- (Dendrite of the next cell)

Review: Artificial Neuron

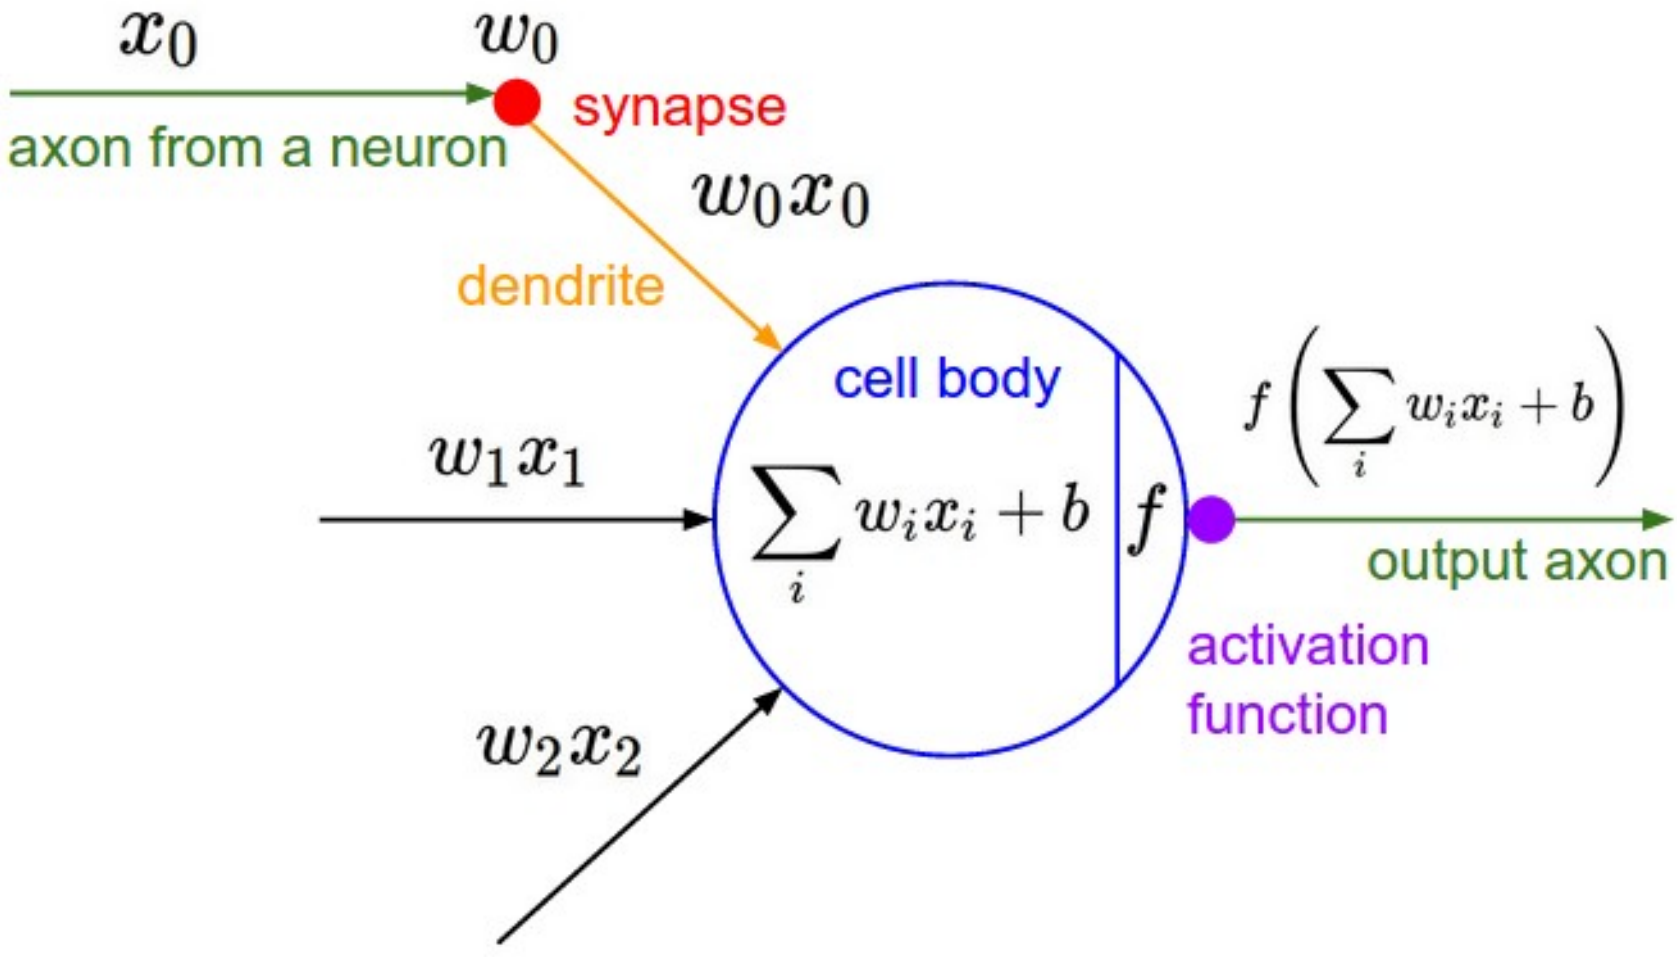

Review: Artificial Neural Network

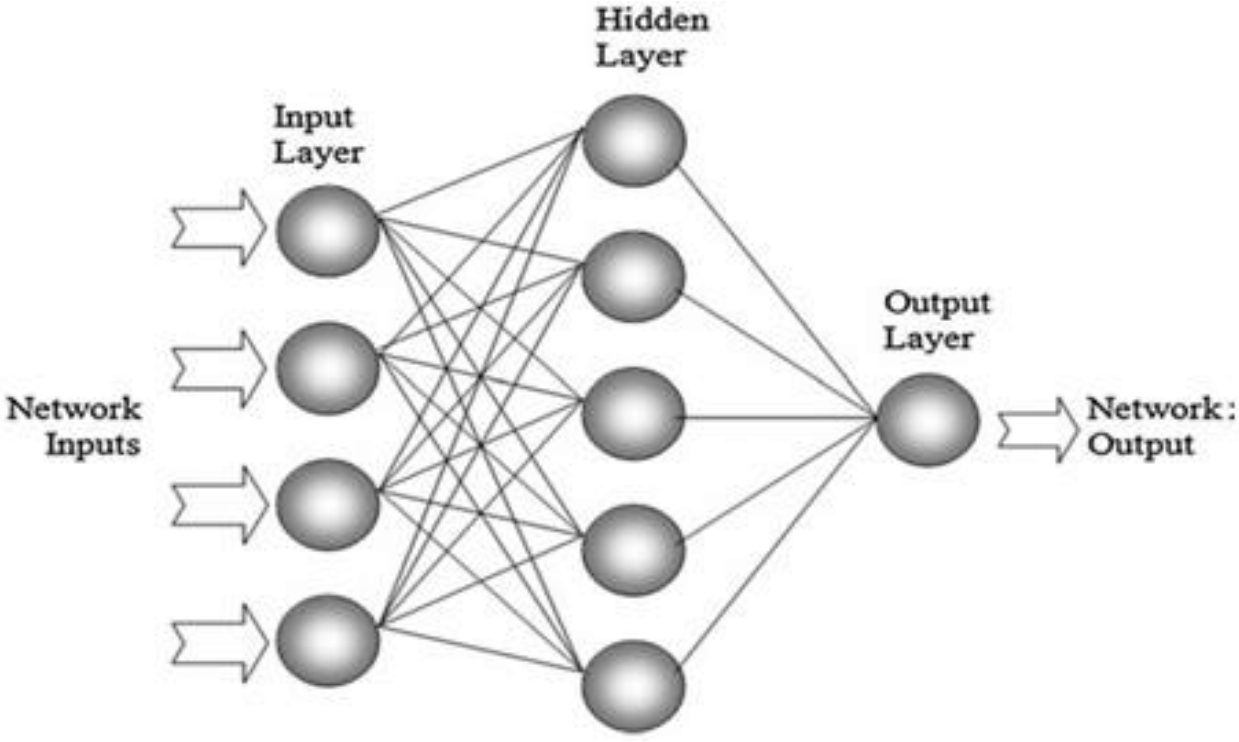

- **fully-connected, feed-forward network**
- $x_1, \ldots$ = the neurons **activation** of input layer neurons
- h_1 = the neuron activation of a hidden layer neuron
- y = the neuron activation of the output layer neuron

Machine Learning Models

When we describe models (like neural networks), we usually:

- first describe how to **make predictions**
- then describe how to **train the model**

This seems a little backward, but it is difficult to understand how to *train* a network without first describing how to *use* that network.

Training our first neural network:

Here is how we will train our articial neural network:

1. Make a prediction for some input data, whose output we already know.
2. Compare the predicted output to the *ground truth* (actual output).
3. Adjust the *weights/biases* to make the prediction close to the ground truth.
4. Repeat steps 1-3 for some number of iterations.

Task for the day

- Input: An 28x28 pixel image
- Output: Whether the digit is a **small** digit (0, 1, or 2)
  - output=1 means that the digit is small
  - output=0 means that the digit is not small

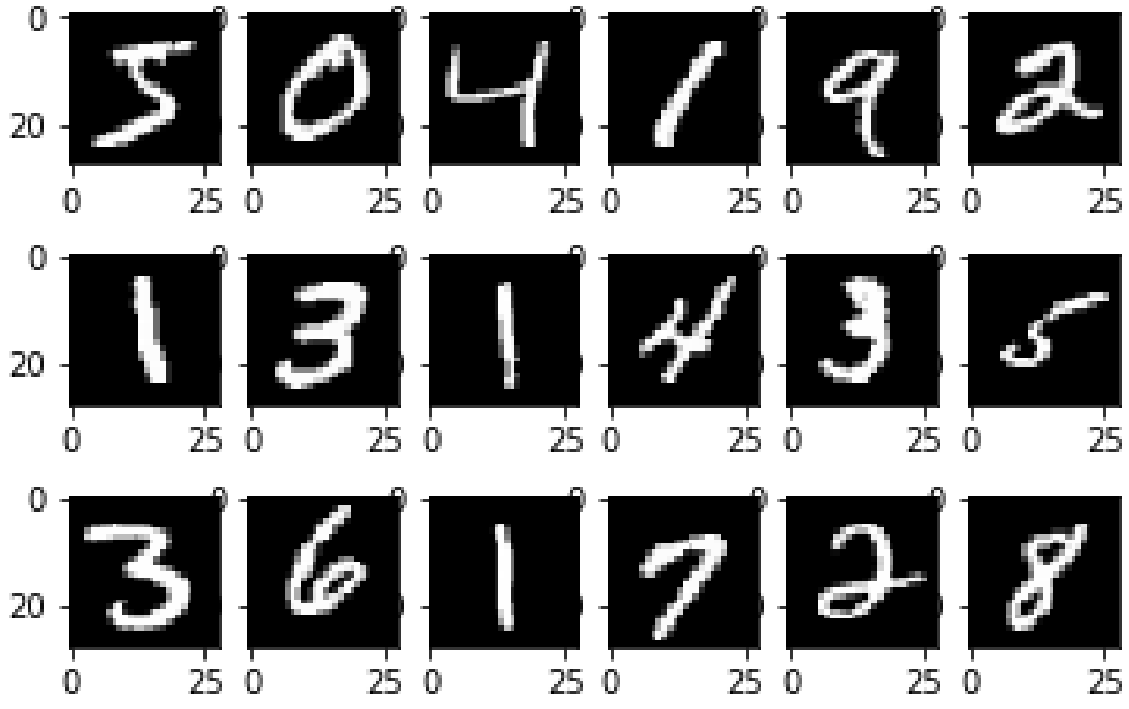

Task for the day

- Input: An 28x28 pixel image
- Output: Whether the digit is a **small** digit (0, 1, or 2)
  - output=1 means that the digit is small
  - output=0 means that the digit is not small

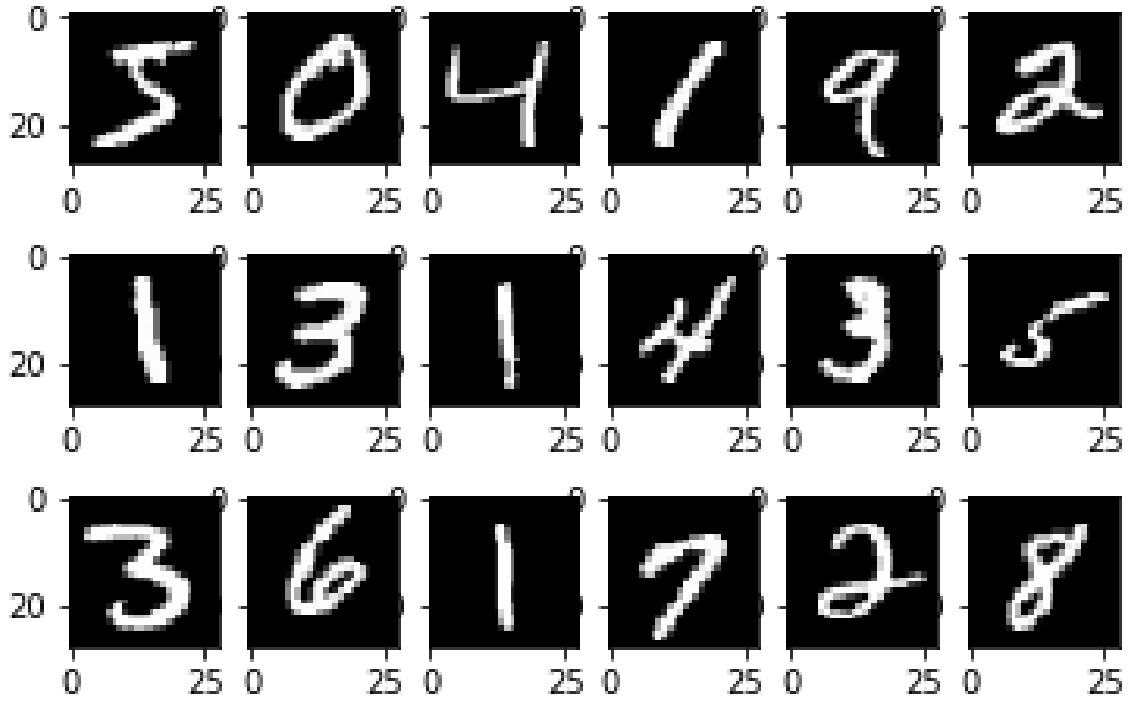

Is this a supervised or unsupervised learning problem?

Task for the day

- Input: An 28x28 pixel image
- Output: Whether the digit is a **small** digit (0, 1, or 2)
  - output=1 means that the digit is small
  - output=0 means that the digit is not small

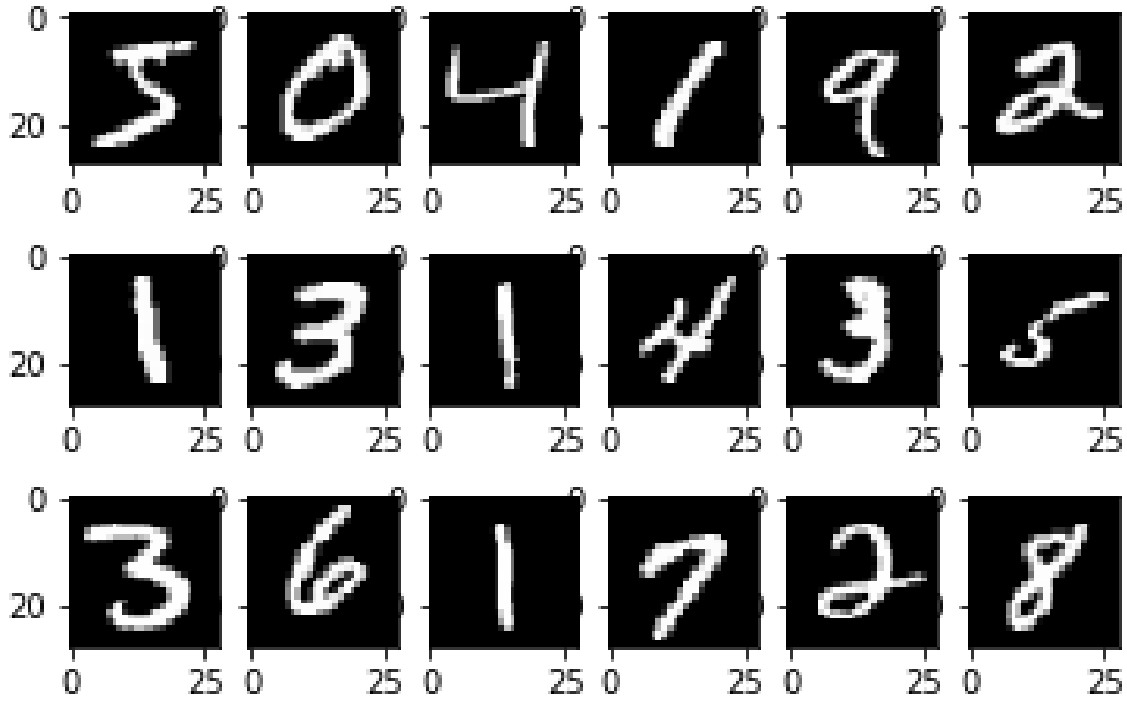

Is this a supervised or unsupervised learning problem?

Is this a regression or classification problem?

Let's write some code!