Neural nets unit:

- ▶ General introduction (**Today**)
- ▶ Trying out neural nets (Friday, in lab)
- ▶ How to train your perceptron (next week Monday)
- ▶ The feed-forward and back-propogation algorithms (next week Wednesday)
- ▶ Deep learning: CNNs (next week Friday and week-after Monday)
- ▶ Deep learning in practice (week-after Wednesday, in lab)

Today:

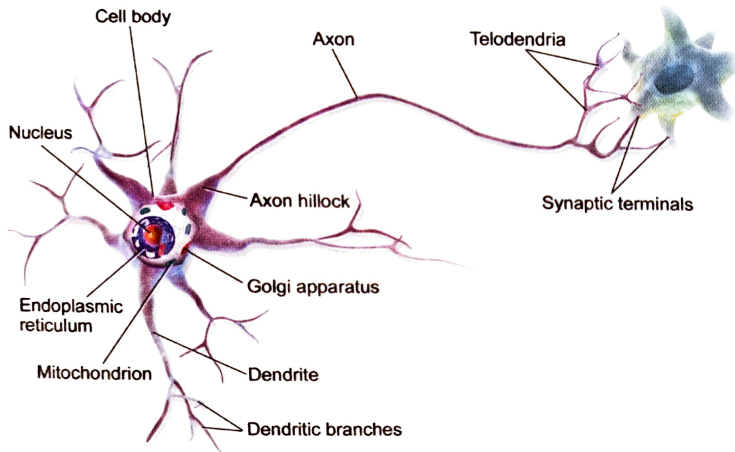- ▶ The origin of neural nets
- ▶ Perceptrons
- ▶ Multi-layer perceptrons

Figure 10-1. Biological neuron[4]

Geron, pg 282

A **perceptron** is a function $\mathbb{R}^D \to \mathbb{R}$ defined as

$$p(\boldsymbol{x}) = h(\boldsymbol{\theta} \cdot \boldsymbol{x} + b) = h\left(b + \sum_{i=0}^{D-1} \theta_i x_i\right)$$

where

- $\boldsymbol{\theta}$ is the vector of weights
- b is the bias term
- h is the activation function

Coming up:

Due Fri, Apr 4:
Implement PCA

Due Mon, Apr. 7:
Read excerpt from Geron introducing neural nets
(See Canvas)

Due Wed, Apr 9:
Read and respond to two articles about bias in algorithms
(See Canvas)

Sometime between Mar 31 and Apr 17:
Make an office-hours appointment for project check-in
(Originally the deadline was Apr 11)