

Neural nets unit

- ▶ Introduction and principles (**Today**)
- ▶ Neural net language model lab (Friday)
- ▶ Begin vector semantics (next week Monday)

Today:

- ▶ Why we care about neural nets
- ▶ Principles of neural nets
 - ▶ The perceptron
 - ▶ Multi-layer perceptrons
 - ▶ Feed-forward etc
- ▶ Neural nets as language models

Machine learning is about training a function $\mathbb{R}^D \rightarrow T$ from data. Common machine learning tasks are

- ▶ *Regression*, where the target type T is $\mathbb{R}$
- ▶ *Classification*, where the target type T is a finite set (that is, a set of classes or labels)
 - ▶ *Binary classification*, where the target type T is $\{0, 1\}$
 - ▶ *Probabalistic classification*, where the target type is $[0, 1]^{|T|}$; that is, the target values are probability distributions over a finite set of classes T
- ▶ *Density estimation*, where the target type T is $[0, 1]$.

A *training algorithm* selects a model (function) from a model family by finding weights (parameters).

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

$$p(\mathbf{x}) = f(\mathbf{w} \cdot \mathbf{x} + b) = f\left(b + \sum_{i=0}^{D-1} w_i x_i\right)$$

where

- ▶ $\mathbf{w}$ is the vector of weights
- ▶ b is the bias term
- ▶ f is the activation function




Figure 6.5 The functions AND, OR, and XOR, represented with input x_1 on the x-axis and input x_2 on the y-axis. Filled circles represent perceptron outputs of 1, and white circles perceptron outputs of 0. There is no way to draw a line that correctly separates the two categories for XOR. Figure styled after [Russell and Norvig \(2002\)](#).




Figure 7.8 A simple 2-layer feedforward network, with one hidden layer, one output layer, and one input layer (the input layer is usually not counted when enumerating layers).

Jurafsky and Martin, Sec 6.3, pg 8; taken from an earlier draft




Figure 6.10 Feedforward network sentiment analysis using traditional hand-built features of the input text.




Figure 6.14 Forward inference in a feedforward neural language model. At each timestep

Coming up:

- ▶ Take neural nets quiz (Fri, Nov 14)
- ▶ Read from J&M Chapter 5 (Mon, Nov 17, before class)
- ▶ Work on stylo project (Fri, Dec 12)