Neural Networks

Outline

- Questions about projects?
- Recap of a singular neuron model
- Notation and terminology
- Compute graphs
- Optimization
- Backpropagation
- This will be our most math heavy week
- Next week we'll rely on PyTorch to compute all derivatives

Recap: A Single Neuron

- Take five minutes to draw
 - Whatever will help you remember (no correct or incorrect drawings)
 - You'll keep a running drawing log the rest of the semester

Fully-Connected (Feed-Forward) Network



Fully-Connected (Feed-Forward) Network



Fully-Connected (Feed-Forward) Network



Simplified (easier to draw) diagram

Vectorized Equations

Data

MNIST Dataset Example

- MNIST includes 60,000 training images
- Each image is grayscale and 28x28 pixels in size
- Each output is a one-hot encoding of the digits 0 through 9
- What is the shape of *X*?

• What is the shape of *Y*?

MNIST Neural Network

• What is the shape of $Z^{[1]}$?

• What is the shape of $A^{[1]}$?

• What is the shape of $Z^{[2]}$?

• What is the shape of $A^{[2]}$?



MNIST Neural Network

- Imagine we have a two-layer network
- The hidden layer has 17 neurons
- What is the shape of $W^{[1]}$?
- What is the shape of $b^{[1]}$?
- What is the shape of $W^{[2]}$?
- What is the shape of $b^{[2]}$?



Compute Graph

Optimization with Binary Cross Entropy Loss

Sigmoid Activation Functions

Backpropagation $W^{[2]}$

Backpropagation $b^{[2]}$

Backpropagation $W^{[1]}$

Backpropagation $b^{[1]}$

Parameter Updates