

Introduction to Neural Networks

Two separate questions:

(1) What is a neural network?
What does it do?

(2) How do we produce a good one
for our task?

This lecture addresses (1).

Future lectures address (2).

Sources:

- ★ Book: "Neural Networks From Scratch in Python"
- ★ Lots of online resources and Youtube videos that I'll share as we go.

Goals:

- * Understand the foundations of neural networks.
- * See examples of training NNs to accomplish a task
↳ "training" = "find a good NN"
- * Write our own Python code to create and train NNs (not using Machine Learning libraries)

Maybe:

- * Learn about Python ML libraries (pytorch, tensorflow)
- * Learn about specialized kinds of NNs for certain tasks.

Strong Analogy: Linear Regression

What is it?

Simplest version: you have a bunch of (x,y) points and you want to find the line the best approximates them

Usually, the (x,y) points represent inputs (x) and outputs (y) of some unknown function.

x : time since Jan 1 this year


y : temperature on my outdoor thermometer

Strong Analogy: Linear Regression

x: time since Jan 1 this year

y: temperature on my outdoor thermometer


We only have sporadic readings.



Linear Regression asks "what line is closest to these points?"


"Closest" means minimizing the sum of
$$([\text{actual } y \text{ value}] - [\text{predicted } y \text{ value}])^2$$
 over all known points.

Strong Analogy: Linear Regression




We are trying to minimize
the sum of the squares
of the lengths of the yellow lines

Strong Analogy: Linear Regression




We are trying to minimize
the sum of the squares
of the lengths of the yellow lines

Strong Analogy: Linear Regression




We are trying to minimize
the sum of the squares
of the lengths of the yellow lines

Problem: What values of m and b make the line $y = mx + b$ with the smallest error?



Demo: mlu-explain.github.io/linear-regression

Purpose: Estimate the output at new inputs.







Demo: mlu-explain.github.io/
linear-regression

This can be done with any # of input variables, and then you're finding planes and hyper planes instead of lines. Same idea though.

How do you find the m, b that make the best line?

Formulas, Linear Algebra

The problem is a lot of data isn't linear.



The four **datasets** composing Anscombe's quartet. All four sets have identical statistical parameters, but the graphs show them to be considerably different

https://commons.m.wikimedia.org/wiki/File:Anscombe%27s_quartet_3.svg

Linear Regression: Approximate data with a line

Neural Network: Approximate data with any function

Pros: Can approximate data much better

Cons: Requires a lot of computation to
produce

Big Picture:


- ★ A neural network is just a fancy function.
- ★ It takes numbers as inputs and produces numbers as outputs, just like any function from Calculus.
- ★ Finding a good neural network that approximates your data is hard.

Also: Lots of interesting things can be framed as a (numerical inputs) $\rightarrow$ (numerical outputs) function


Example function:

Input: a 28×28 grayscale image

Output: what digit it is




How is this "numeric input"?




So this is a 784-dimensional function.

* slight lie, our output will be a little different

And not really well-defined because many input pictures are not obviously a particular number



That's fine! We'll produce a neural network that takes 784 input #s (1 per pixel) and predicts the #.



MNIST data set

We'll use "training data", known (input, output) pairs, to build the NN. (just like using data to find a linear regression line)

Big Picture:

- ★ A neural network is just a fancy function.
- ★ It takes numbers as inputs and produces numbers as outputs, just like any function from Calculus.
- ★ Finding a good neural network that approximates your data is hard.

↳ We do this using lots of known (input, output) pairs.

Then we can use that good NN to predict new digits for us.


* Web applet demo

* Imagine if your job was to write an algorithm to look at the pixels manually & say what # it is! That would be so difficult.


What is a neural network?

A fancy function defined in the following way.


Example: A NN that takes 2 inputs and has 1 output
 $f(u, v) = z$



Example: A NN that takes 2 inputs and has 1 output
 $f(u, v) = z$




Example: A NN that takes 2 inputs and has 1 output
 $f(u, v) = z$




each circle is a "neuron"

Example: A NN that takes 2 inputs and has 1 output
 $f(u, v) = z$




A NN can have any # of layers, and they can be any size.

Example: A NN that takes 2 inputs and has 1 output
 $f(u, v) = z$




Every edge between neurons has a real # attached to it called its weight

Example: A NN that takes 2 inputs and has 1 output
 $f(u, v) = z$




I'm using #s to 1 decimal place for simplicity, but typically they are full floating point #s.

Example: A NN that takes 2 inputs and has 1 output
 $f(u, v) = z$




Every neuron (except the input layer) has a # attached to it called its bias.

Example: A NN that takes 2 inputs and has 1 output
 $f(u,v) = z$




The u and v get "fed forward" to the neurons in the next layer.

Example: A NN that takes 2 inputs and has 1 output
 $f(u,v) = z$




To calculate the value for a neuron, we multiply the value of each connected neuron from the last layer by the weight of the connection, add these all up, and add the bias

Example: A NN that takes 2 inputs and has 1 output
 $f(u,v) = z$




$$w_1 = (0.4)(0.1) + (-0.7)(-0.8) + 0.1$$
$$= 0.7$$

Example: A NN that takes 2 inputs and has 1 output
 $f(u,v) = z$




$$\begin{aligned} W_2 &= (0.4)(-0.3) + (-0.7)(0.1) + -0.2 \\ &= -0.39 \end{aligned}$$

Example: A NN that takes 2 inputs and has 1 output
 $f(u,v) = z$




$$\begin{aligned} w_3 &= (0.4)(0.7) + (-0.7)(0.2) + 0.1 \\ &= 0.24 \end{aligned}$$

Example: A NN that takes 2 inputs and has 1 output
 $f(u,v) = z$




$$\begin{aligned} w_4 &= (0.4)(-0.5) + (-0.7)(-0.6) + -0.8 \\ &= -0.58 \end{aligned}$$

Example: A NN that takes 2 inputs and has 1 output
 $f(u,v) = z$




Then, the neurons from the first hidden layer feed forward to the next hidden layer.

Example: A NN that takes 2 inputs and has 1 output
 $f(u,v) = z$




$$\begin{aligned}
 W &= (0.7)(0.3) + (-0.39)(0.8) + (0.24)(-0.2) + (-0.58)(0.3) + -0.4 \\
 &= -0.724
 \end{aligned}$$


Example: A NN that takes 2 inputs and has 1 output
 $f(u,v) = z$




Example: A NN that takes 2 inputs and has 1 output
 $f(u,v) = z$



Example: A NN that takes 2 inputs and has 1 output
 $f(u,v) = z$




Example: A NN that takes 2 inputs and has 1 output
 $f(u,v) = z$



So, for this particular neural network, with these **weights** and **biases**, the inputs $(0.4, -0.7)$ produce output -0.601

$$f(0.4, -0.7) = -0.601$$


Example: A NN that takes 2 inputs and has 1 output
 $f(u,v) = z$



So, for this particular neural network, with these **weights** and **biases**, the inputs (0.4, -0.7) produce output -0.601

- * NNs can have any # of layers
- * The layers can have any # of neurons in them
- * If every neuron in a layer is connected to every neuron in the next layer, then the network is "fully connected"


* But we're missing a really key component so far.



$$\begin{aligned}
 & (w_{11}x_{11}y_{11} + w_{11}x_{12}y_{21} + w_{11}x_{13}y_{31} + w_{12}x_{21}y_{11} + w_{12}x_{22}y_{21} + w_{12}x_{23}y_{31} + w_{13}x_{31}y_{11} + w_{13}x_{32}y_{21} + w_{13}x_{33}y_{31} + w_{14}x_{41}y_{11} + w_{14}x_{42}y_{21} + w_{14}x_{43}y_{31}) \cdot u \\
 & + (w_{21}x_{11}y_{11} + w_{21}x_{12}y_{21} + w_{21}x_{13}y_{31} + w_{22}x_{21}y_{11} + w_{22}x_{22}y_{21} + w_{22}x_{23}y_{31} + w_{23}x_{31}y_{11} + w_{23}x_{32}y_{21} + w_{23}x_{33}y_{31} + w_{24}x_{41}y_{11} + w_{24}x_{42}y_{21} + w_{24}x_{43}y_{31}) \cdot v \\
 & + b_1x_{11}y_{11} + b_1x_{12}y_{21} + b_1x_{13}y_{31} + b_2x_{21}y_{11} + b_2x_{22}y_{21} + b_2x_{23}y_{31} + b_3x_{31}y_{11} + b_3x_{32}y_{21} + b_3x_{33}y_{31} + b_4x_{41}y_{11} + b_4x_{42}y_{21} + b_4x_{43}y_{31} + c_1y_{11} + c_2y_{21} + c_3y_{31} + d_1
 \end{aligned}$$


As we've built them NNs are still just (big) linear functions, so this is no better than linear regression.

The problem is even more obvious when there's only one input and one output.



$$\Rightarrow y = 0.026x - 0.25$$


Just a line defined in a complicated way!



$$\Rightarrow y = 0.026x - 0.25$$

$$\Rightarrow (0.026)(0.3) - 0.25 = -0.2422 \checkmark$$

Understanding check: Let's feed forward $x = 0.3$



$$A = (0.2)(0.3) + 1.0 = 1.06$$

$$B = (-0.6)(0.3) - 1.0 = -1.18$$

$$C = (0.7)(0.3) - 0.5 = -0.29$$

$$D = 0.244$$

$$E = -1.001$$

$$F = 0.079$$

$$y = -0.2422$$