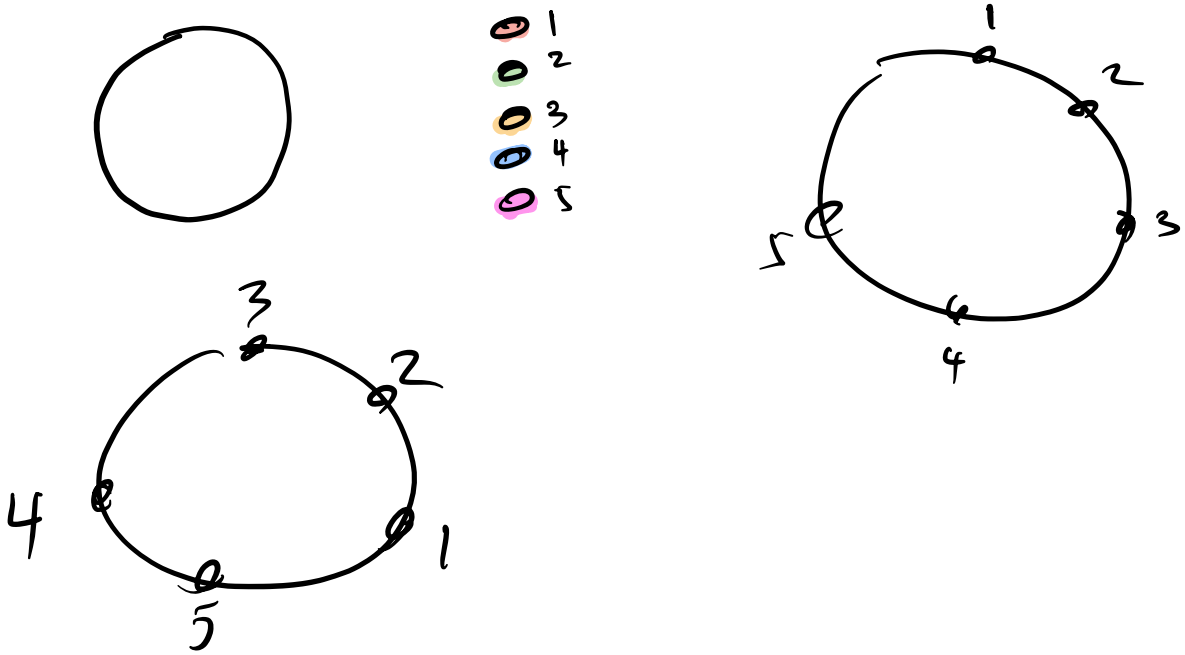

Overview:

- Large language models (LLM)
  built on problem of "next word prediction"
- Some perspective on this from brute force next letter prediction vs language like output
- next word prediction w/ large texts unreasonably hard for brute force frequency analysis.

Problem: how to reasonably do next word prediction?
Answer: neural nets (lots of them)

Neural nets give a solution to the following ill posed problem

given inputs, $x_0, x_1, \ldots$ we believe there should exist a function f

taking these to certain desired outputs $y_0, y_1, \ldots$

$f(x_0) = y_0 \quad f(x_1) = y_1 \ldots$ in some reasonable way.

and we believe if we can plug & check values for f, we'll know if it's right.

The design/architecture of a net is highly variable, depends on problem at hand.

ReLU Activation/similar for visual data (handwriting)

Sigmoid Activation for boolean/logic/classification

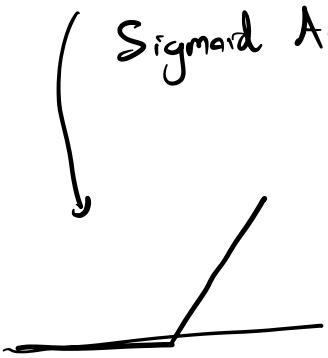

ReLU

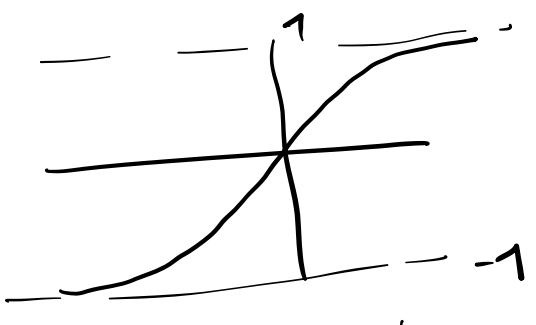

Neural nets have the flavor of linear approximation/regression, but actually very nonlinear.

Gradient descent:

method of approximating solns to eqns

$f(x) = 0$

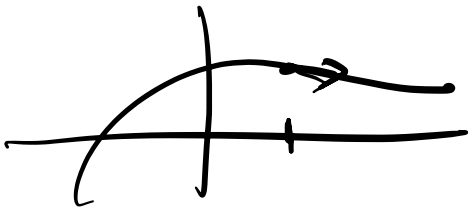

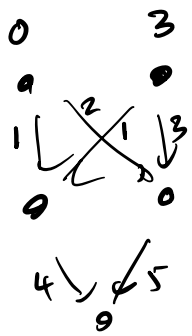

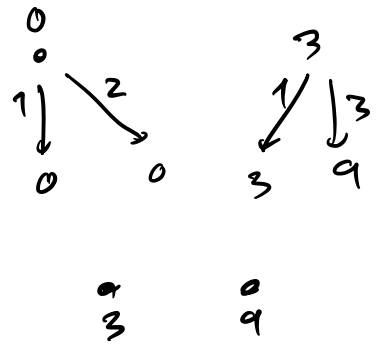

$f(x) = \sin x$ $\quad\quad$ $\sin 30 = \frac{1}{2}$

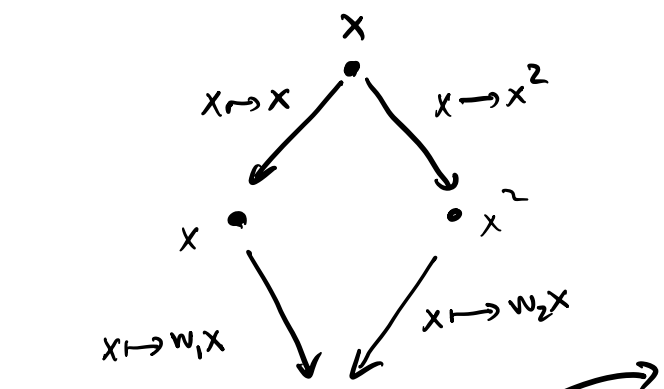

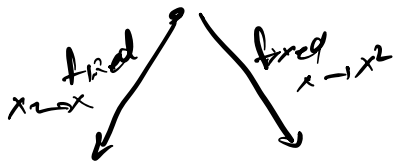

oops this was supposed to be the input not the output

$\sin(5) = .087$

$\sin(40) = 0.64$

$\sin(20) = 0.34$

0.087 $\quad$ 0.0075

0.087

$\frac{1}{2}$ $\quad$ 1

0.0

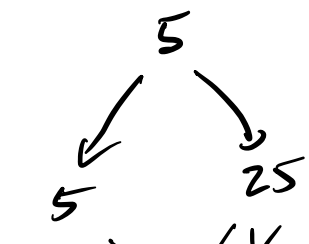

$0.005 + 0.025 = 0.03$