CS 307

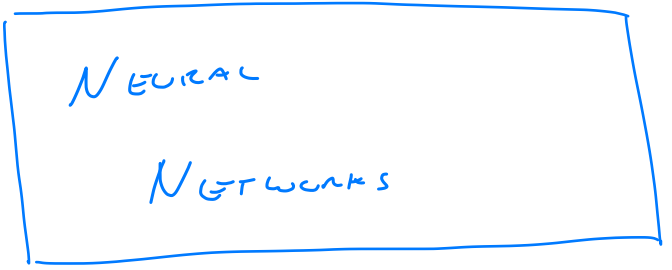

WHAT IS A "NEURAL NETWORK"?

- ARTIFICAL "ANN" / "NN"
- A FUNCTION
  - OF INPUT DATA AND PARAMETERS
  - OFTEN REPRESENTED AS A NETWORK
- PARAMETERS LEARNED FROM DATA

"WEIGHTS"

Logistic Regression

Sigmoid Function $\sigma(x) = \frac{1}{1+e^{-x}}$ $\sigma : \mathbb{R} \to [0,1]$

$p(x) = P[Y=1 \mid X=x]$

$\log\left(\frac{p(x)}{1-p(x)}\right) = \beta_0 + \beta_1 x_1 + \beta_2 x_2$

THREE PARAMETERS

TWO INPUT FEATURES

"LOGIT"

LR AS A "NETWORK"

LET $X_0 = 1$

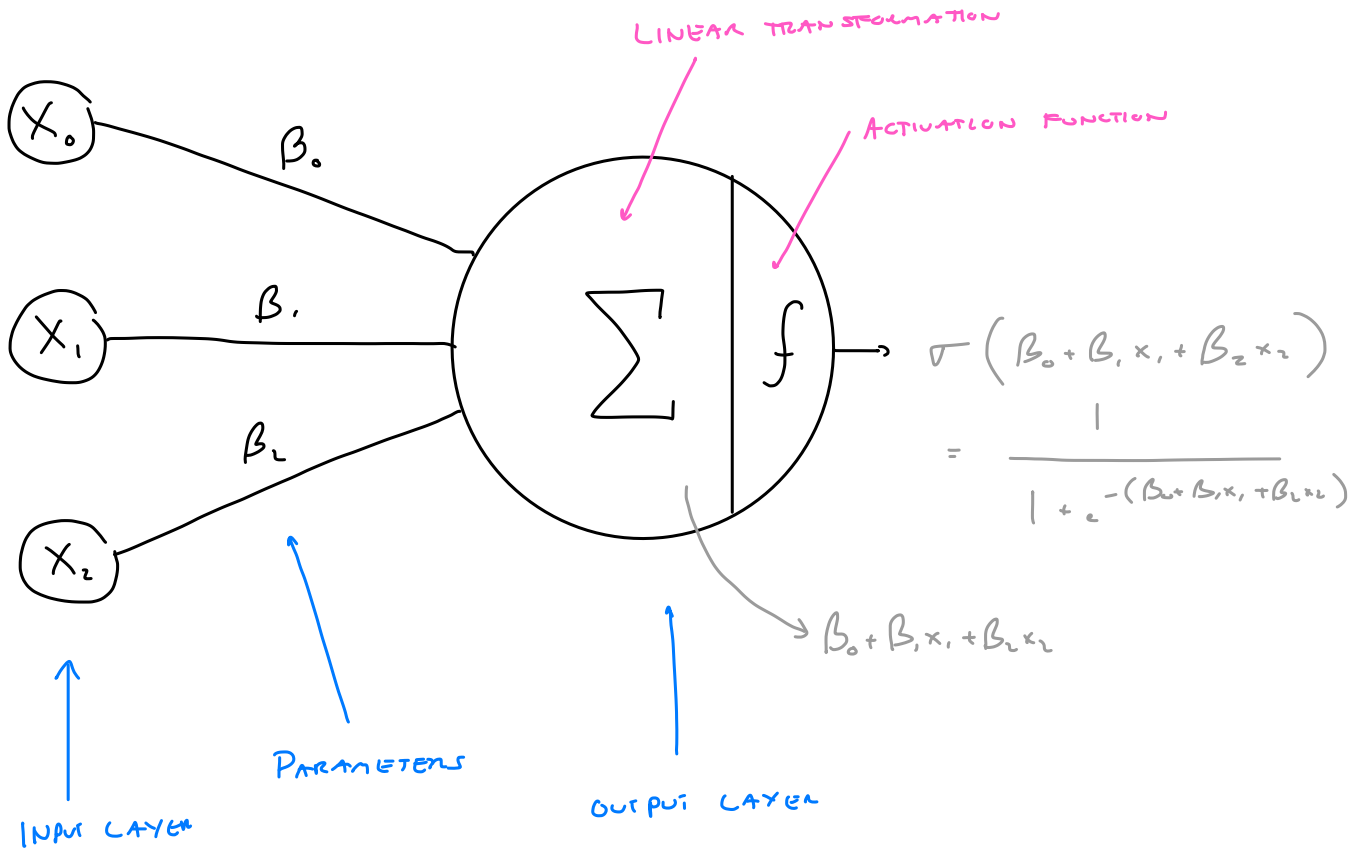

NN NOMENCLATURE

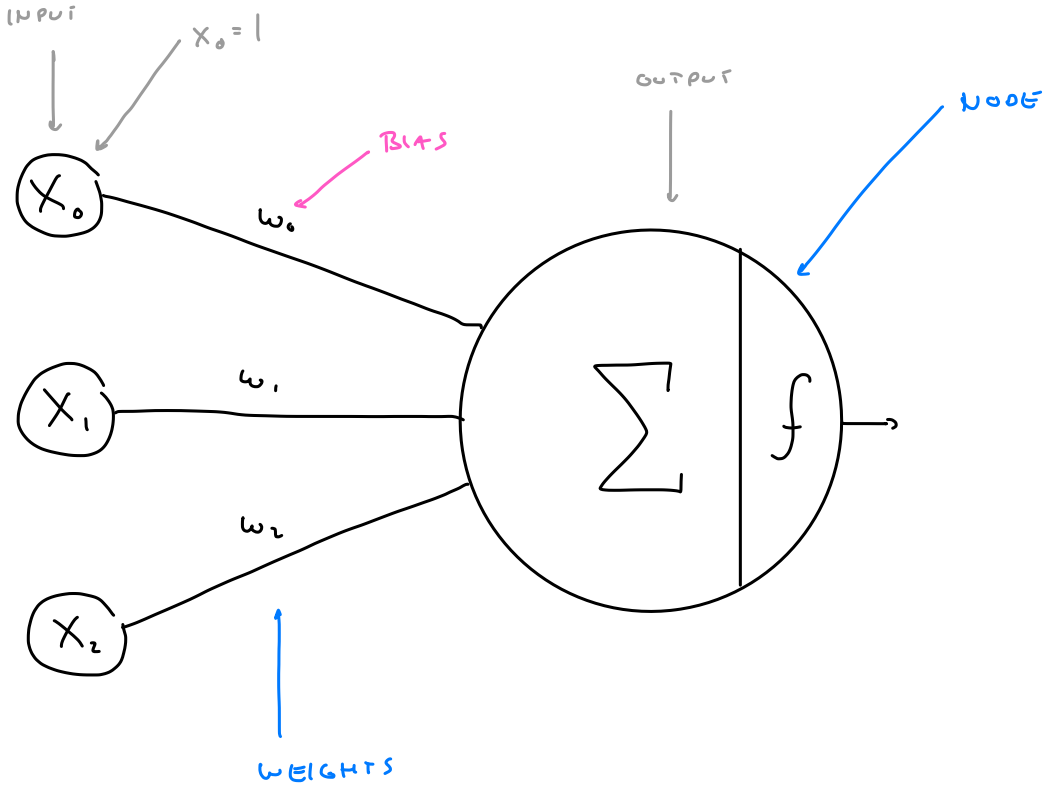

ReLU

Rectified Linear Unit

Activation

$$f(x) = \max(0, x) = \begin{cases} x & \text{if } x > 0 \\ 0 & \text{otherwise} \end{cases}$$

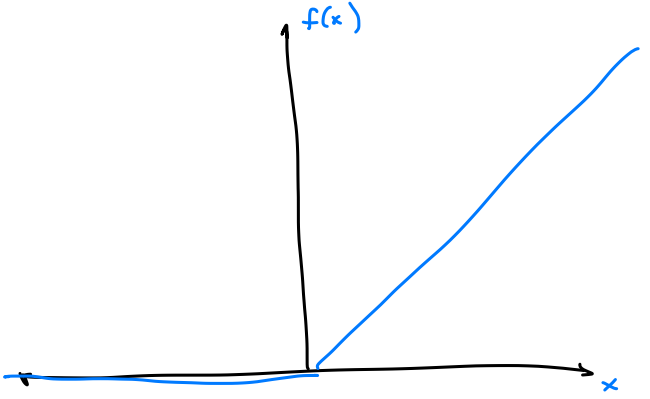

GOING "DEEP"

EVERY CONNECTION REPRESENTS A WEIGHT PARAMETER

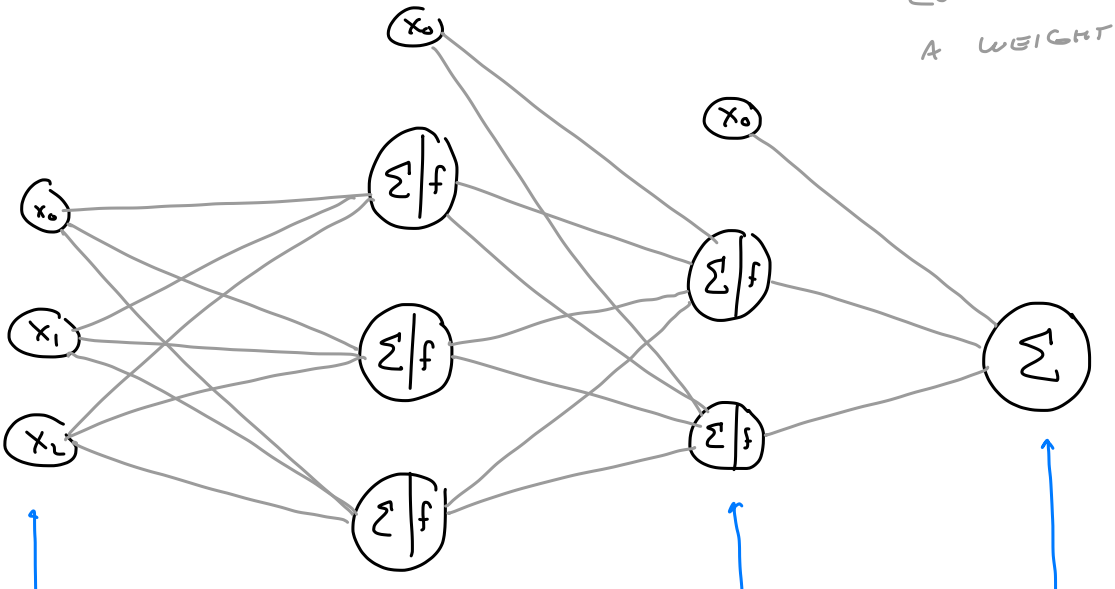

INPUT

HIDDEN
3 NODES
+ BIAS

HIDDEN
2 NODES
+ BIAS

OUTPUT

- SINGLE OUTPUT FOR BINARY
- ONE-DER FOR MULTICLASS

WHAT DO WE CONTROL?

- HOW MANY LAYERS?
- HOW MANY NODES PER LAYER?

} NETWORK ARCHITECTURE

- WHICH ACTIVATION FUNCTION? USE RELU FOR HIDDEN
- HOW TO FIT / OPTIMIZE?
  - METHOD
  - LEARNING RATE

DEEP LEARNING = BIG NETWORKS

How to train NNs?

"How to learn weights from data?"

- Optimization problem!
  - Define loss function
  - Backpropagation (chain rule)
  - SGD / ADAM

1. LEARN TOY BINARY DATA

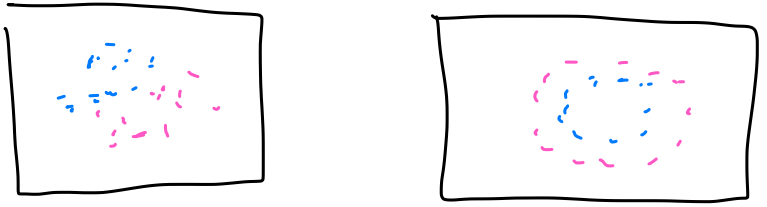

2. LEARN MULTICLASS DATA

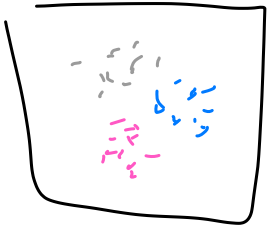

3. LEARN TO CLASSIFY IMAGES!

IMAGES AS FEATURE DATA

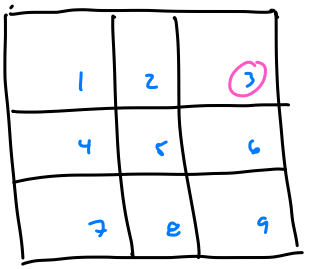

"FLATTEN"

X_1	X_2	X_3	. . .	X_9
1	2	3	. . .	9

3×3 PIXEL IMAGE

FOR SIMPLICITY, ASSUME GREYSCALE IMAGES

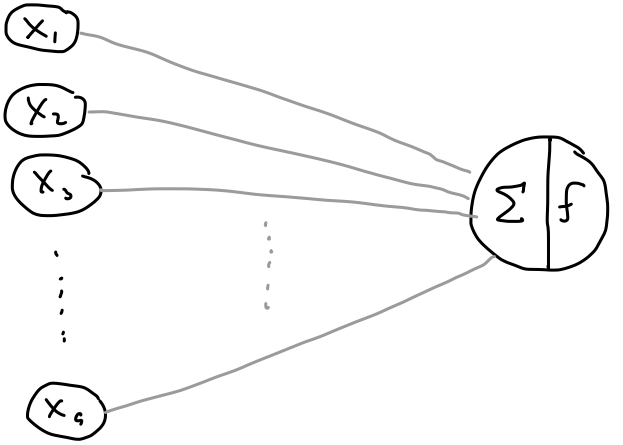

Conv Net

Input

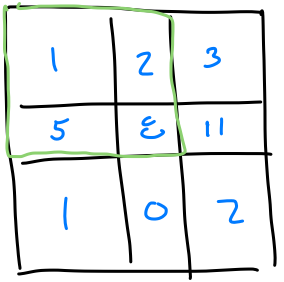

3x3

Kernel

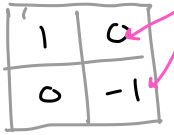

2x2

Output

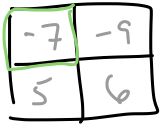

2x2

$$1(1) + 0(2) + 0(5) - 1(6) = -7$$

PADDING

INPUT

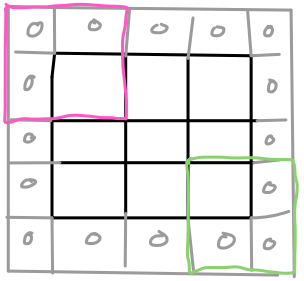

3×3

5×5

AFTER PADDING

KERNEL

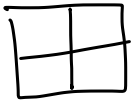

2×2

OUTPUT

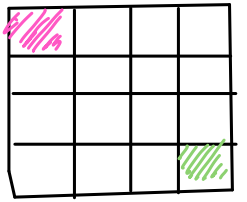

4×4

Pooling

MAX / AVERAGE

Input

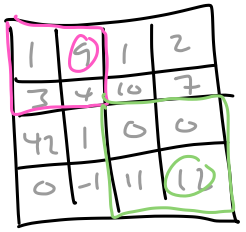

4x4

Kernel

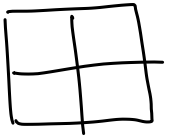

2x2

STRIDE: 2

often set to kernel size

Output