


Course 2

Linear Regression:

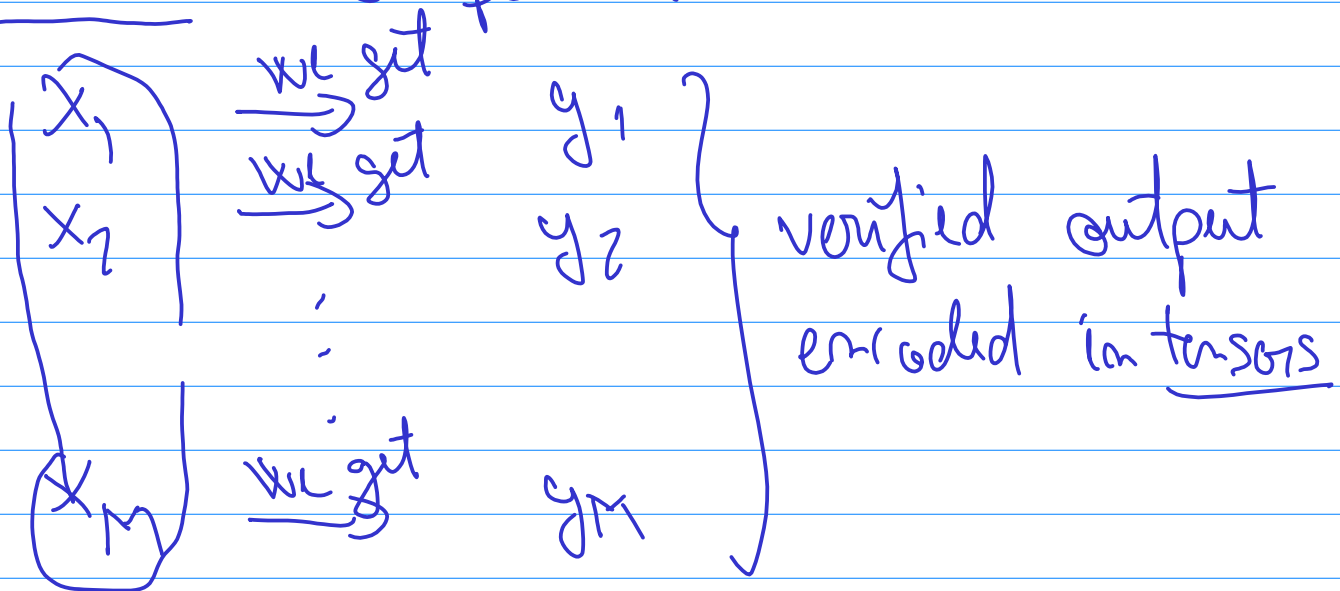
fitting data with a linear (affine) function.



Data: (x_i, y_i) , $i=1 \dots N$


Find the line $y = mx + m$ such that
 $[mx_i + m \approx y_i]$

DATA: For parameters



Objects \rightarrow convert to
real values/vectors/matrices/tensors

Data as a straight line:



$$y = \overset{\text{weight}}{\textcircled{m}}x + \overset{\text{bias}}{m}$$

↓ ↓
parameters

Samples x_i

compute the values $y_i = m x_i + m$

$\Rightarrow (x_i, y_i)$

DATA
input

output

x_1	y_1
x_2	y_2
\vdots	

80%

training set
Use to train the model
using THIS DATA

x_{80}	y_{80}
\vdots	
x_{100} ---	y_{100}

20%

test set
HOT USED FOR
TRAINING

WHY? We need to test the model on data it "hasn't seen before."

Plotting the data helps us understand better what's happening.


MODEL FOR LINEAR REG:


→ Parameters: weight (slope)
bias (translation)

→ "Compute" the model

"forward" pass through the model

FORWARD






$$x \longrightarrow x \cdot w + b \longrightarrow \text{output}$$

→ For training we need to differentiate the "formula" that is used in the forward function.

→ automated differentiation

→ `requires_grad = True` keeps track of all computations such that the derivative can be computed automatically



Loss function: measure the "difference" between model results and expected results.


$$L_1 \text{ loss} = \sum_{i=1}^N |\text{Model}(x_i) - y_i|$$

result of model for x_i *Data for x_i*

Loss "small" \Rightarrow model predicts "well"

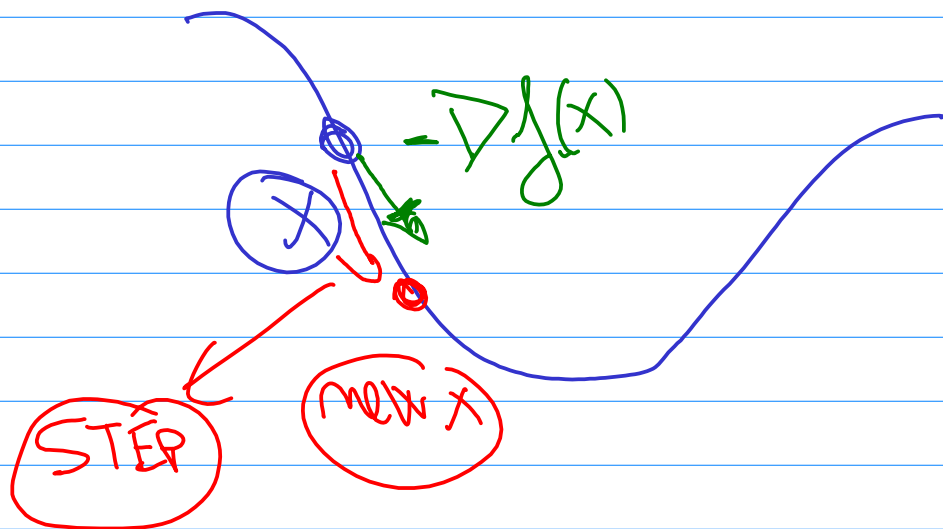
$$\text{MSE Loss} = \sum_{i=1}^N (\text{Model}(x_i) - y_i)^2$$

OBJECTIVE: make Loss function small! Use on optimization alg.



$$\text{Loss} := \sum (\text{output}_i - y_i)^2$$

∇Loss



Observations: \rightarrow Do as many epochs as you can (reasonable in time)
 \rightarrow Decreasing learning rate can

improvi results after an initial run.

