CSCI 1470/2470 Spring 2023

Ritambhara Singh

January 30, 2023 Monday Deep Learning

DALL-E 2 prompt "a painting of deep underwater with a yellow submarine in the bottom right corner"

Recap

+

Represent input and output as numbers

Classification – predicting categorical outputs

Regression – predicting numerical outputs



Supervised Learning Learn a function that approximates the data well

Get more data!

How to

represent

inputs and

outputs

Try different models Pick a good model

Real world data tends to be complicated!



(Image only for explaining concept, not drawn accurately)

(Numerical output)

Today's goal - Learn about the first component of deep learning model

Perceptron:

(1) Machine Learning problem – Recognizing handwritten digits

(2) Perceptron

(3) Parameters – weights and biases

Handwritten digit recognition

Motivation: ZIP codes

- In 1990s, great increase in documents on paper (mail, checks, books, etc.)
- Motivation for a ZIP code recognizer on real U.S. mail for the postal service!

80322-4129 80206

40004 (4310

J7878 05753

·5502 75316 35460: A4209

Our Problem:

Input: X Target: Y 3^{*} 3^{*} $f(X) \rightarrow Y$ How does a computer know this is a three?



Representing digits in the computer

•Numbers known as *pixel values* (a grid of discrete values that make up an image)

0 is white, 255 is black, and numbers in between are shades of gray

157	153	174	168	150	152	129	151	172	161	155	156	157	153	174	168	150	152	129	151	172	161	155	156
155	182	163	74	75	62	33	17	110	210	180	154	155	182	163	74	75	62	33	17	110	210	180	154
180	180	50	14	34	6	10	33	45	105	159	181	180	180	50	14	34	6	10	33	48	106	159	181
206	109	5	124	131	111	120	204	165	15	56	180	206	109	5	124	131	111	120	204	166	15	56	180
194	68	137	251	237	239	239	228	227	87		201	194	68	137	251	237	239	239	228	227	87	n	201
172	105	207	233	233	214	220	239	228	.98	-74	206	172	105	207	233	233	214	220	239	228	98	74	206
188	88	179	209	185	215	211	158	139	75	20	169	188	88	179	209	185	215	211	158	139	75	20	169
189	97	165	84	10	168	134	11	31	62	22	148	189	97	165	84	10	168	134	11	31	62	22	148
199	168	191	193	158	227	178	143	182	105	36	190	199	168	191	193	158	227	178	143	182	106	36	190
205	174	155	252	236	231	149	178	228	43	95	234	205	174	155	252	236	231	149	178	228	43	96	234
190	216	116	149	236	187	85	150	79	38	218	241	190	216	116	149	236	187	86	150	79	38	218	241
190	224	147	108	227	210	127	102	36	101	255	224	190	224	147	108	227	210	127	102	36	101	256	224
190	214	173	66	103	143	95	50	2	109	249	215	190	214	173	66	103	143	96	50	2	109	249	215
187	196	235	75	1	81	-17	۰	6	217	255	211	187	196	235	75	1	81	47	0	6	217	255	211
183	202	237	145	0	0	12	108	200	138	243	236	183	202	237	145	0	0	12	108	200	138	243	236
195	206	123	207	177	121	123	200	175	13	95	218	195	206	123	207	177	121	123	200	175	13	96	218

How is this different from the color image example in the last class?





• Pixel in position [15, 15] is light.

what the computer sees



Often has lighter pixels in the middle!

How does the pattern compare with digit 3?

	255	255	255	255	255	253	254	245	255
	255	255	251	255	255	255	254	235	252
	255	252	255	250	255	245	255	253	234
	253	255	255	255	251	254	255	255	235
	255	255	252	255	249	255	239	243	255
	255	250	255	245	255	255	254	244	254
	255	255	255	255	249	255	255	255	244
	249	255	253	255	233	255	249	245	239
	255	255	255	250	255	254	251	243	251
	245	240	244	240	239	244	255	244	248
	242	128	140	150	130	128	110	245	246
	240	240	4	5	4	3	2	118	120
	240	5	4	2	0	0	0	4	2
	0	0	0	0	0	0	0	0	0
						2	.	-	
Can w	e det	ine a	a set	of he	uristic	s (i.e.	rules k	based	on ou
				1	2				
Intuiti	on), t	<mark>O Cla</mark>	assity	/ digit	S?				

Let's define some rules (heuristic) for classifying "7"



Digit is a 7 if $P_1 >$ 128 and $P_2 >$ 128 and $P_3 >$ 128

But what if...



Slanted digit?

An Improved Heuristic!



Digit is a 7 if $P_1 >$ 128 and $P_2 >$ 128 and ($P_3 >$ 128 or $P_4 >$ 128)

Not so fast...



Digit shifted up?

Heuristics...

- Not as simple as we think!
- Distortions, overlappings, underlinings, etc.
- Cannot rely on a set of exact rules

Let's do some machine learning!

Distorted numbers

م 3->9 *Q 9* 6->0 9->8 **9** 4->9 6->1 9->4 9->1 **L 3** 6->1 3->5 **> 9 0** 3->2 **9**->5 **6**->0 **4** 4->3 **9** 9->7 **2**->7 **%) % % % 3** 3->8 **>** 8->5 6->5 3->8 9->8 **9 6 0 6 7** 9->8 6->3 0->2 6->5 9->5 **9** 0->7 **9** 9->7 2->8 S 4->9 2->8

Machine Learning Pipeline for Digit Recognition





Machine Learning Pipeline for Digit Recognition



MNIST

- Modified National Institute of Standards and Technology database
- Handwritten digits
- •0 9 (10 *classes*)
- 70,000 images

G F Л qq Ø

Machine Learning Pipeline for Digit Recognition



Machine Learning Pipeline for Digit Recognition



Train, validation, and test sets

- Train set used to adjust the parameters of the model
- Validation set used to test how well we're doing as we develop
 - Prevents overfitting
- Test set used to evaluate the model once the model is done



MNIST

- Training set 60,000 images
- Test set 10,000 images
- No explicit validation set

What do you suggest we do here?

G F Л qØ P

Machine Learning Pipeline for Digit Recognition



Machine Learning Pipeline for Digit Recognition





Our simplified problem:



Perceptron

(Our first deep learning model unit)

Biological motivation

- Loosely inspired by neurons, basic working unit of the brain
- Serve to transmit information between cells



<u>here</u>

https://en.wikipedia.org/wiki/Depolarization

The Perceptron





Biological Neuron

Artificial Neuron (Perceptron)

Input

• Input: a vector of numbers $\mathbf{x} = [x_1, x_2, \dots x_n]$



What was x_i for lemonade stand example?

What is x, MINIST image?

x is represented by a 28 * 28 matrix of pixel values, flattened into a one-dimensional vector (size 784) (more on this later)

Predicting with a Perceptron

- [•]1. Multiply each input x_i by its corresponding weight w_i , sum them up.
- 2. Add the bias b





Predicting with a Perceptron

- [•]1. Multiply each input x_i by its corresponding weight w_i , sum them up.
- 2. Add the bias b

3. If the result value is greater than 0, return 1, otherwise return 0

$$f_{\Phi}(x) = \begin{cases} 1, & \text{if } b + \sum_{i=0}^{n} w_{i} x_{i} > 0\\ 0, & \text{otherwise} \end{cases}$$



How is perceptron different from linear regression?



Performs binary classification!

- w and b are parameters of the perceptron
 - Parameters: values we adjust during learning
 - Let $\Phi = \{w \cup b\}$ (the set of all parameters)



Go to www.menti.com and use the code 5592 2451

• Weights — the importance of each input to determining the output

b

Σ

output

- Weight near 0 imply this input has little influence on the output
- Negative weight means?

 $f_{\Phi}(x) = \begin{cases} 1, & \text{if } b + \sum_{i=0}^{n} w_i x_i > 0\\ 0, & \text{otherwise} \end{cases}$

Option 1: Increasing input will increase output

Option 2: Increasing input will decrease output

Option 3: Decreasing input will decrease output

• Weights — the importance of each input to determining the output

- Weight near 0 imply this input has little influence on the output
- Negative weight means increasing the input will decrease the output





• **Bias** — What do we need this for?



Bias: Geometric Explanation

• the bias is essentially the **b** term in **y** = **mx+b**



Bias: Conceptual Explanation

• **Bias** — the *a priori* likelihood of the positive class

- Ensures that even if all inputs are 0, there will be some result value
- Just because all inputs are 0, it does not mean there are no 1's in the world
- Maybe there just happen to be more, say, 0's than 1's



Bias as special type of weight

• Another way to think of bias is to represent it as an extra weight for an input/feature that is always 1



Bias as special type of weight

 Another way to think of bias is to represent it as an extra weight for an input/feature that is always 1

$$[x_0, x_1, x_2, \dots x_n] \cdot [w_0, w_1, w_2, \dots w_n] + b$$

$$= [x_0, x_1, x_2, \dots x_n, 1] \cdot [w_0, w_1, w_2, \dots w_n, b]$$

Recall

$$\mathbf{a} = [a_1, a_2, \dots, a_n]$$
 and $\mathbf{b} = [b_1, b_2, \dots, b_n]$ with vector space n

the dot product is

$$\mathbf{a} \cdot \mathbf{b} = \sum_{i=1}^{n} a_i b_i = a_1 b_1 + a_2 b_2 + \dots + a_n b_n$$

Simplifying some notation...

- Recall: the dot product of two vectors of length *n* is $\mathbf{a} \cdot \mathbf{b} = \sum_{i=1}^{n} a_i b_i$
- We can rewrite the perceptron function accordingly:

1

Any questions?

$$f_{\Phi}(x) = \begin{cases} 1, & \text{if } b + \sum_{i=0}^{n} w_i x_i > 0 \\ 0, & \text{otherwise} \end{cases}$$

$$f_{\Phi}(x) = \begin{cases} 1, & \text{if } b + w \cdot x > 0 \\ 0, & \text{otherwise} \end{cases}$$

• In modern deen learning narlance $h \perp w \cdot v$ is known as a *linear unit*

A Binary Perceptron for MNIST

- Inputs $[x_1, x_2, ..., x_n]$ are all positive
 - n = 784 (28 * 28 pixel values)
- *output* is either 0 or 1
 - 0 → input is not the digit type we're looking for
 - 1 → input *is* the digit type we're looking for



Training a perceptron



0. set the parameters $\Phi = \{w \cup b\}$ to 0

1. Iterate over training set several times,

feeding in each training example into the model,

producing an output, and adjusting the parameters according to whether that output was right or wrong

2. Stop once we either
(a) get every training
example right or
(b) after N iterations, a
number set by the
programmer.



The Perceptron Learning Algorithm

- •1. set *w*'s to 0.
- 2. for N iterations, or until the weights do not change:
 - a) for each training example \mathbf{x}^k with label y^k

i. if
$$y^k - f(\mathbf{x}^k) = 0$$
 continue

ii. else for all weights
$$w_i$$
, $\Delta w_i = (y^k - f(\mathbf{x}^k)) x_i^k$

- b = bias
- w = weights
- N =maximum number of training iterations
- $\mathbf{x}^k = \mathbf{k}^{\text{th}}$ training example

- $y^k = label$ for the kthexample
- w_i = weight for the ith input where $i \le n$
- *n* = number of pixels per image
- $x_i^k = i^{th}$ input of the example where $i \le n$