

CSCI 1470/2470
Spring 2022

Ritambhara Singh


January 29, 2024
Monday

Perceptron

Deep Learning



Recap



How to represent inputs and outputs

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!

Try different models

Pick a good model



Real world data tends to be complicated!



Input: X

“Temperature” “Stand Hours” “Sunny?”

Target: Y

“Profit made on selling lemonade”



$$x_1^{(1)} = 100.1 \quad x_2^{(1)} = 8 \quad x_3^{(1)} = 1$$

$$y^{(1)} = 200.0$$

$$x_1^{(2)}$$

Homework 1 - Diabetes dataset

<https://www4.stat.ncsu.edu/~boos/var.select/diabetes.html>

$$y^{(2)} = 80.5$$

$$x_1^{(3)}$$

number of samples (n) = 442

number of features (p) = 10

output (y): quantified disease progression

Linear regression task

$$y^{(3)} = 15.1$$

$$X \in \mathbb{R}^3$$

$$x_i^{(k)} = \dots$$

$$y^{(k)} = \dots$$

$$Y \in \mathbb{R}$$

(Numerical output)

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

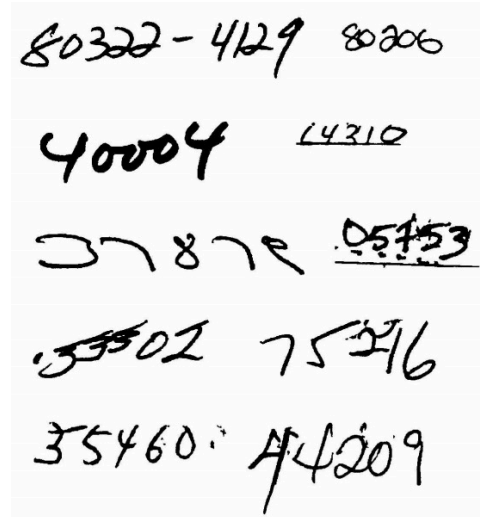
Perceptron:

- (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!



Our Problem:

Input: \mathbb{X}

3



Function: f

$$f(\mathbb{X}) \rightarrow \mathbb{Y}$$



Target: \mathbb{Y}

Which digit is it?

"3"

How does a computer know this is a three?

3



“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 |
| 155 | 182 | 163 | 74 | 75 | 62 | 93 | 17 | 110 | 210 | 180 | 154 |
| 180 | 180 | 50 | 14 | 94 | 6 | 10 | 93 | 48 | 106 | 159 | 181 |
| 206 | 109 | 5 | 124 | 131 | 111 | 120 | 204 | 166 | 15 | 56 | 180 |
| 194 | 68 | 137 | 251 | 257 | 239 | 239 | 228 | 227 | 87 | 71 | 201 |
| 172 | 106 | 207 | 233 | 233 | 214 | 220 | 239 | 228 | 98 | 74 | 206 |
| 188 | 88 | 179 | 209 | 185 | 215 | 211 | 158 | 139 | 75 | 20 | 169 |
| 189 | 97 | 165 | 84 | 10 | 168 | 134 | 11 | 31 | 62 | 22 | 148 |
| 199 | 168 | 191 | 193 | 158 | 227 | 178 | 143 | 182 | 106 | 36 | 190 |
| 205 | 174 | 155 | 252 | 236 | 231 | 149 | 178 | 228 | 43 | 95 | 234 |
| 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 | 214 | 173 | 66 | 103 | 143 | 96 | 50 | 2 | 109 | 249 | 215 |
| 187 | 196 | 235 | 75 | 1 | 81 | 47 | 0 | 6 | 217 | 255 | 211 |
| 183 | 202 | 237 | 145 | 0 | 0 | 12 | 108 | 200 | 138 | 243 | 236 |
| 195 | 206 | 123 | 207 | 177 | 121 | 123 | 200 | 175 | 13 | 96 | 218 |


| | | | | | | | | | | | |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
| 157 | 153 | 174 | 168 | 150 | 152 | 129 | 151 | 172 | 161 | 155 | 156 |
| 155 | 182 | 163 | 74 | 75 | 62 | 93 | 17 | 110 | 210 | 180 | 154 |
| 180 | 180 | 50 | 14 | 94 | 6 | 10 | 93 | 48 | 106 | 159 | 181 |
| 206 | 109 | 5 | 124 | 131 | 111 | 120 | 204 | 166 | 15 | 56 | 180 |
| 194 | 68 | 137 | 251 | 257 | 239 | 239 | 228 | 227 | 87 | 71 | 201 |
| 172 | 106 | 207 | 233 | 233 | 214 | 220 | 239 | 228 | 98 | 74 | 206 |
| 188 | 88 | 179 | 209 | 185 | 215 | 211 | 158 | 139 | 75 | 20 | 169 |
| 189 | 97 | 165 | 84 | 10 | 168 | 134 | 11 | 31 | 62 | 22 | 148 |
| 199 | 168 | 191 | 193 | 158 | 227 | 178 | 143 | 182 | 106 | 36 | 190 |
| 205 | 174 | 155 | 252 | 236 | 231 | 149 | 178 | 228 | 43 | 95 | 234 |
| 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 | 214 | 173 | 66 | 103 | 143 | 96 | 50 | 2 | 109 | 249 | 215 |
| 187 | 196 | 235 | 75 | 1 | 81 | 47 | 0 | 6 | 217 | 255 | 211 |
| 183 | 202 | 237 | 145 | 0 | 0 | 12 | 108 | 200 | 138 | 243 | 236 |
| 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?




| | | | | | | | | | | | | |
|-----|-----|-----|---|---|---|---|---|---|---|---|---|---|
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 5 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 120 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 240 | 40 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 242 | 128 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 255 | 240 | 10 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 254 | 244 | 120 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 255 | 255 | 121 | 8 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |

what the
computer sees



- 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?

Darker pixels in the middle




Can we define a set of *heuristics* (i.e. rules based on our intuition), to classify digits?

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



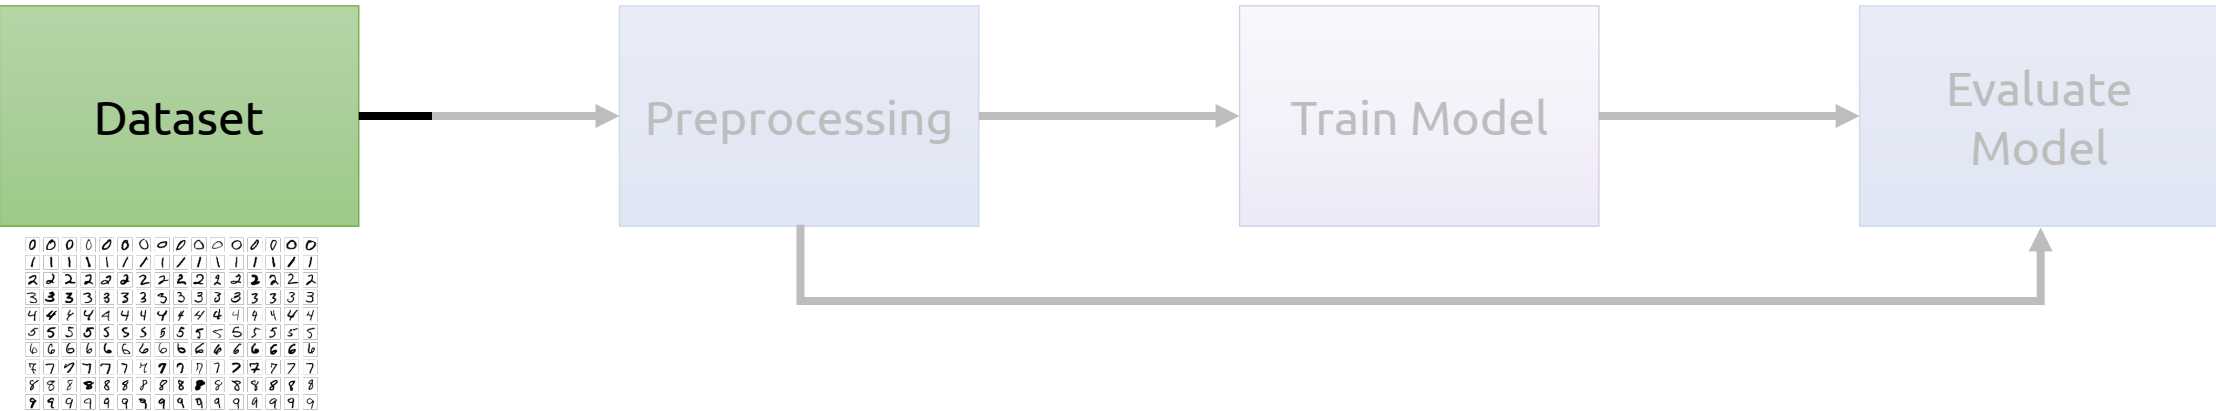
Machine Learning Pipeline for Digit Recognition



Any questions?




Machine Learning Pipeline for Digit Recognition

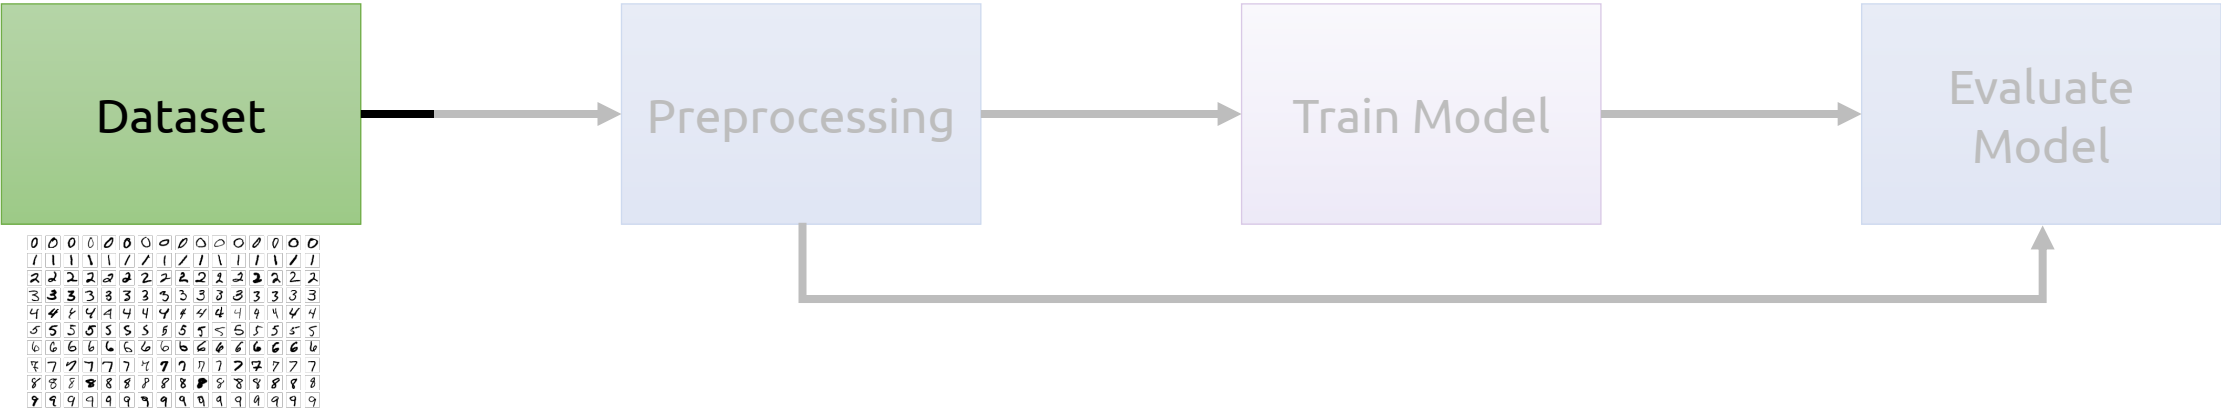


MNIST

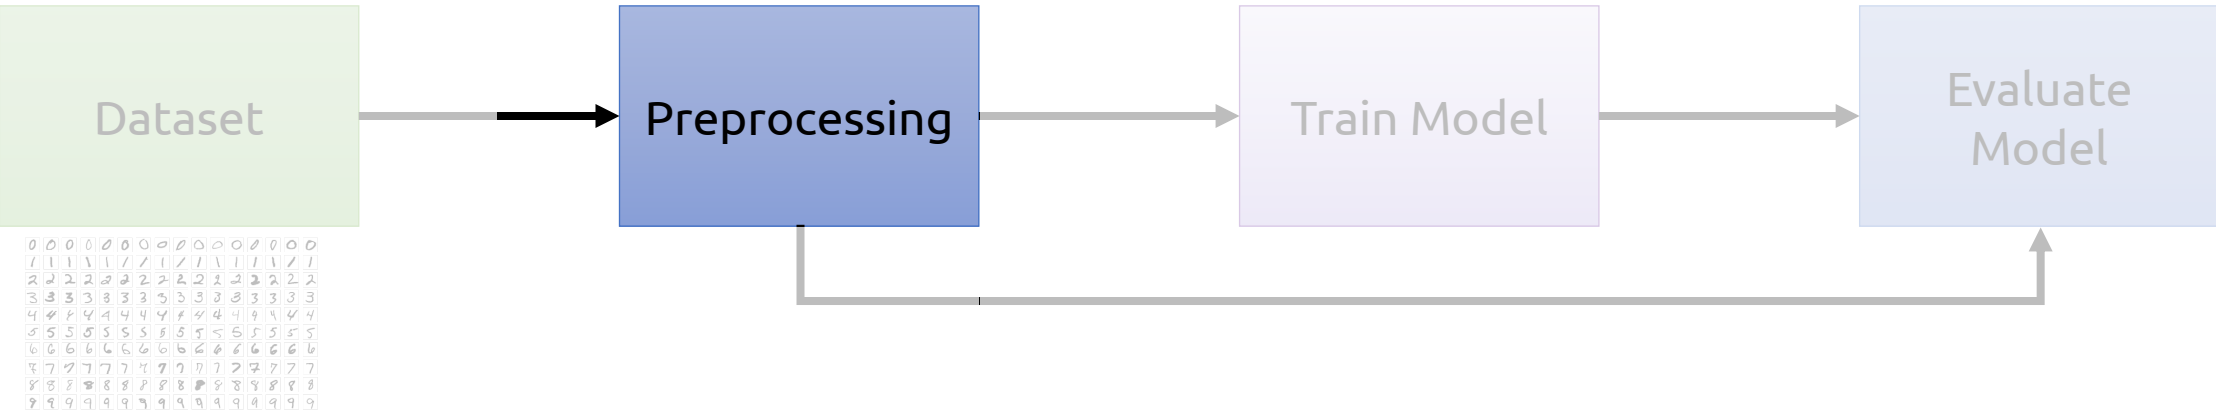
- **Modified National Institute of Standards and Technology** database
- Handwritten digits
- 0 — 9 (10 **classes**)
- 70,000 images



Machine Learning Pipeline for Digit Recognition

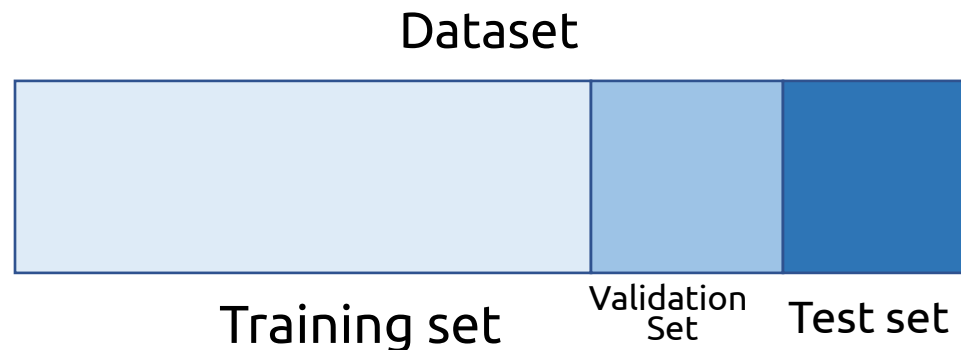


Machine Learning Pipeline for Digit Recognition



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



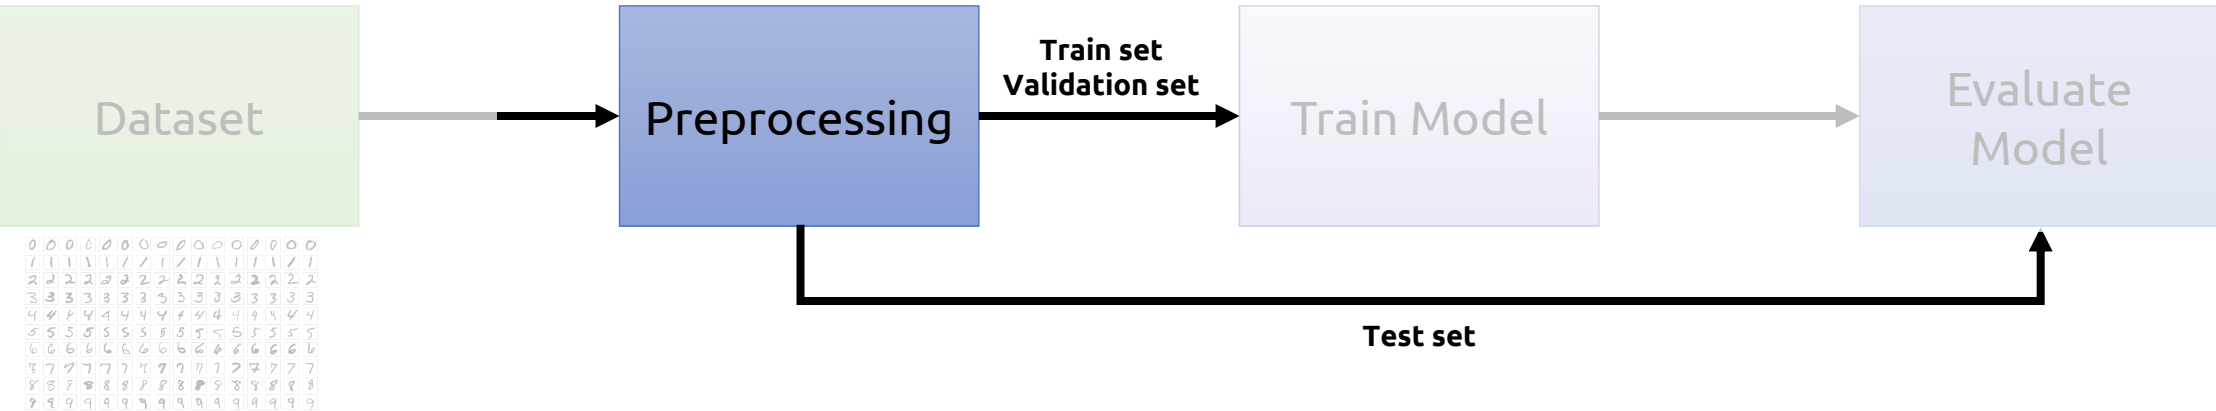
MNIST

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

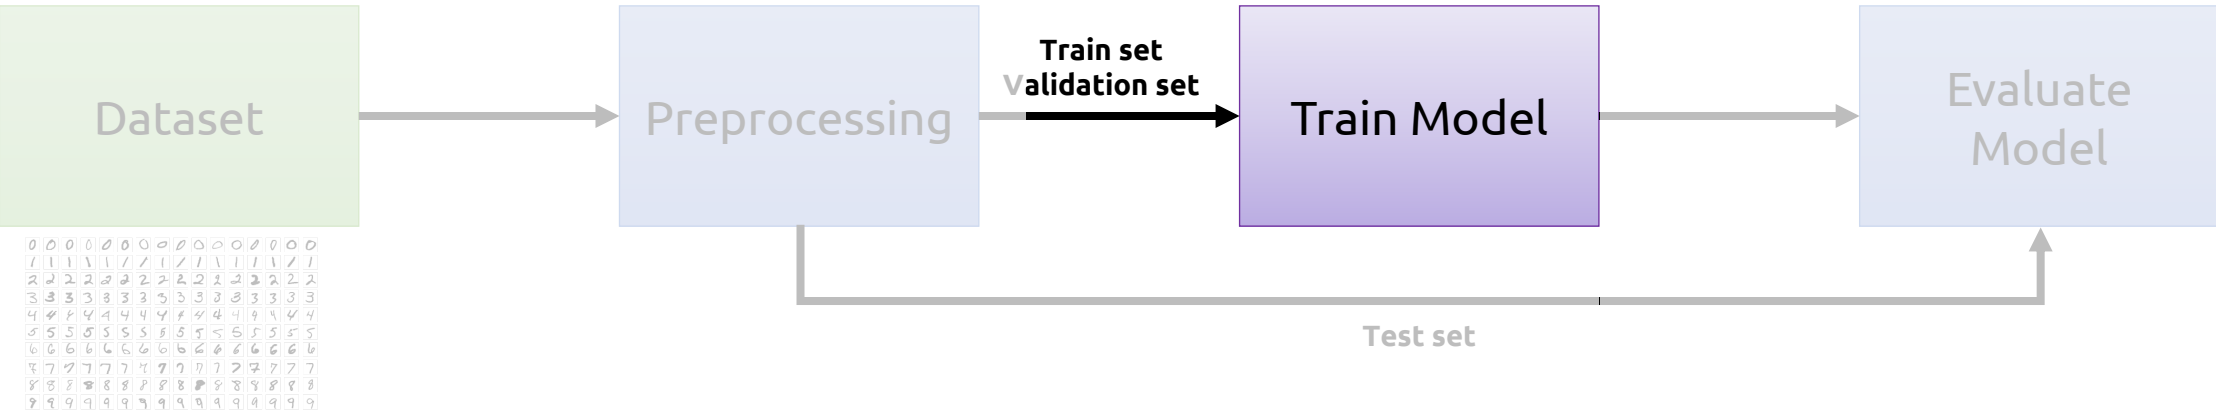
What do you suggest we do here?



Machine Learning Pipeline for Digit Recognition



Machine Learning Pipeline for Digit Recognition




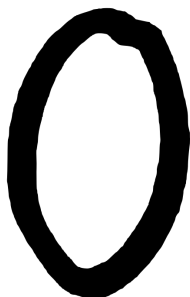
Our Problem:

Classifying MNIST digits requires predicting 1 of 10 possible values

Input: \mathbb{X}

Pixel Grid

$x^{(1)} =$ 
28x28 pixels

$x^{(2)} =$ 

What is our input space?

What is our output space?

What is our prediction task?

→ Function: f →

$f(\mathbb{X}) \rightarrow \mathbb{Y}$

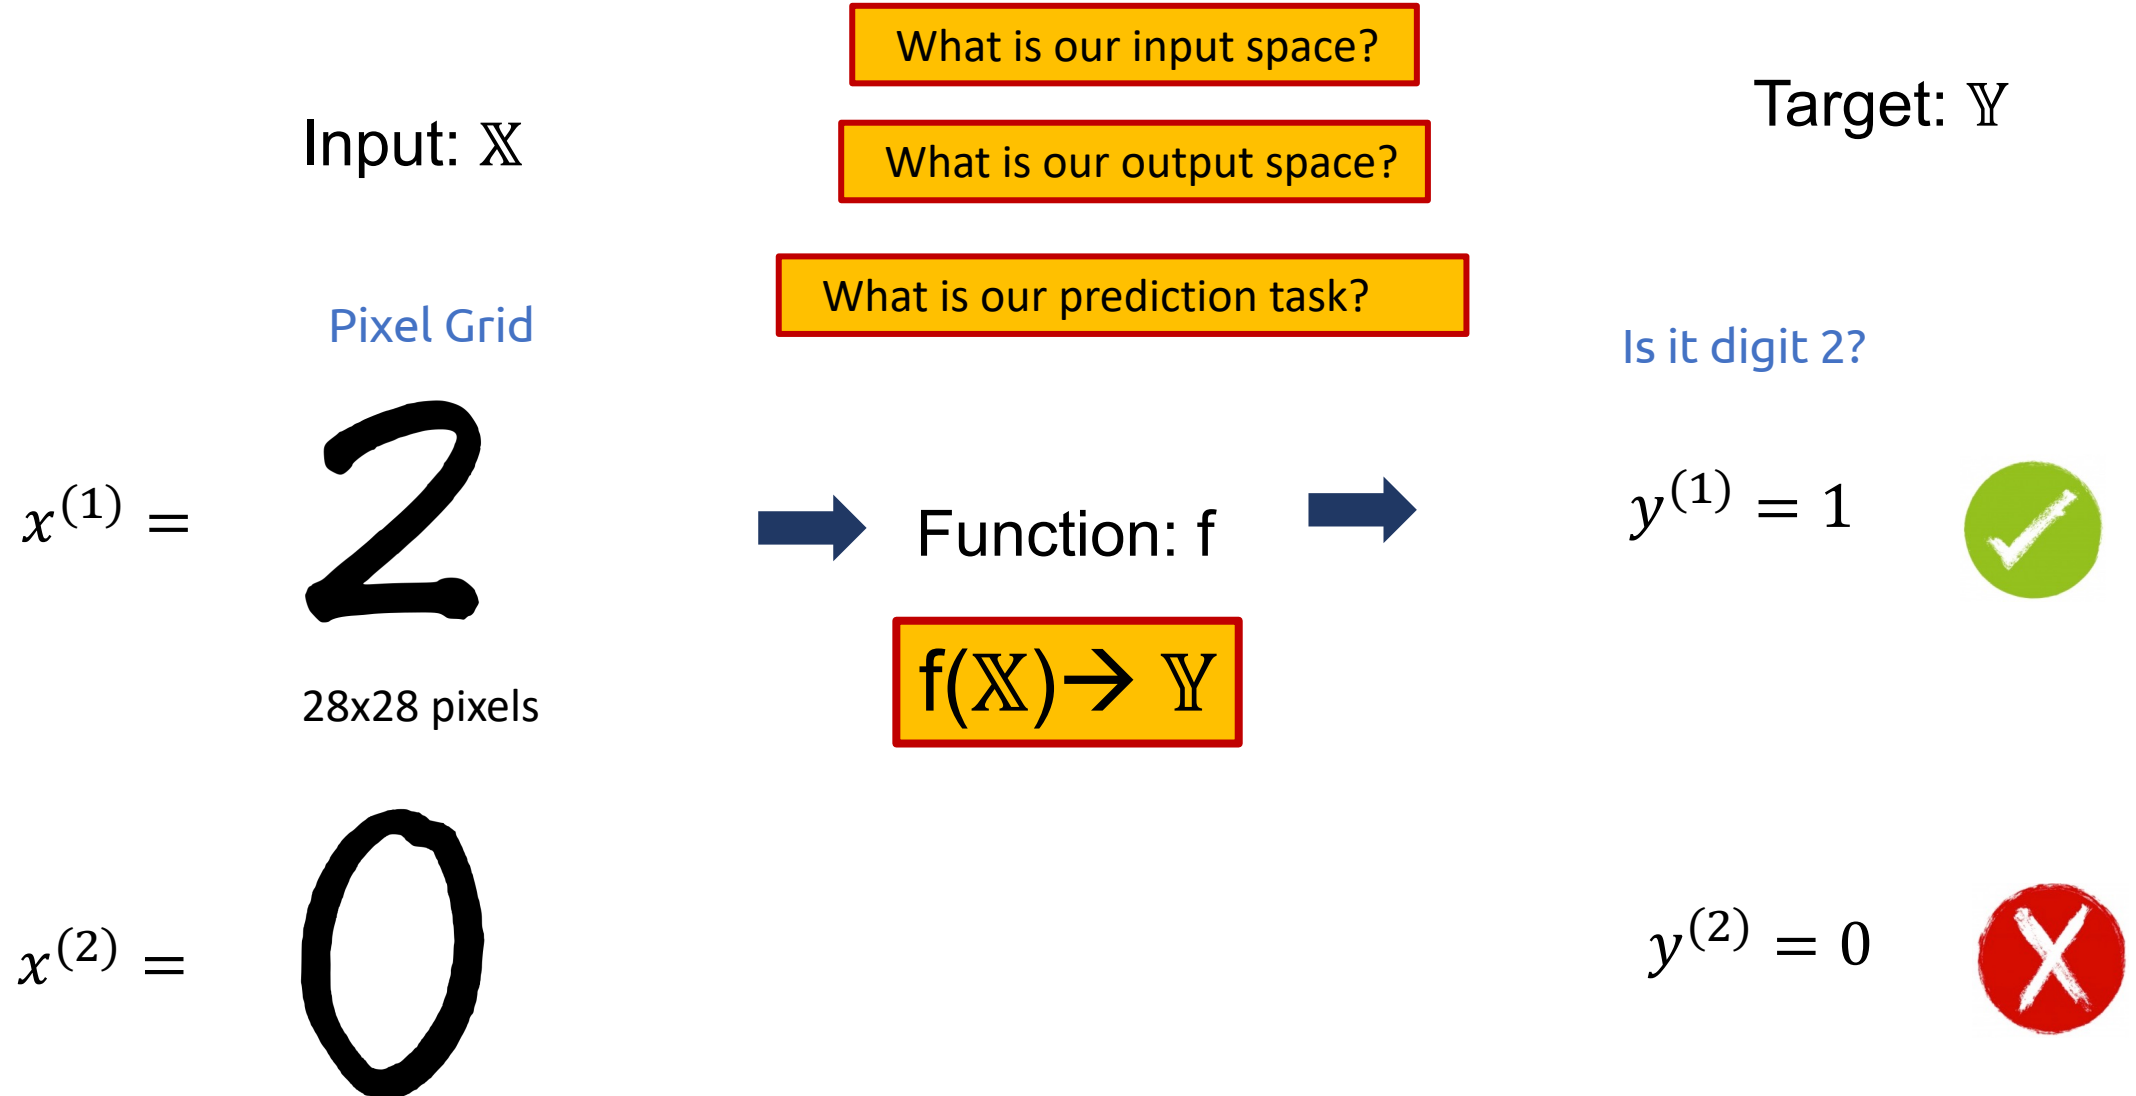
Target: \mathbb{Y}

Which digit is it?

$y^{(1)} = \text{"2"}$

$y^{(2)} = \text{"0"}$

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




(Depolarization is a change within a cell, during which the cell undergoes a shift in electric charge distribution)

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

The Perceptron




Biological Neuron



Artificial Neuron (Perceptron)

Input

- Input: a vector of numbers $x = [x_1, x_2, \dots, x_n]$



What was x_i for
lemonade stand
example?


What is x_i 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


Does this look familiar?



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?

Threshold value = 0

Performs binary classification!

Parameters

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



Parameters

Join at [menti.com](https://www.menti.com) | use code 3340 9640


- **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?

Option 1: Increasing input will increase output

Option 2: Increasing input will decrease output

Option 3: Decreasing input will decrease output


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



Parameters

- **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

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




Any questions?



Parameters

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


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



Bias: Geometric Explanation

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

only the
line with
bias can fit
the data



$$f(x) = 2x$$


$$f(x) = \frac{3}{2}x$$

$$f(x) = x + 1$$

Bias: Conceptual Explanation

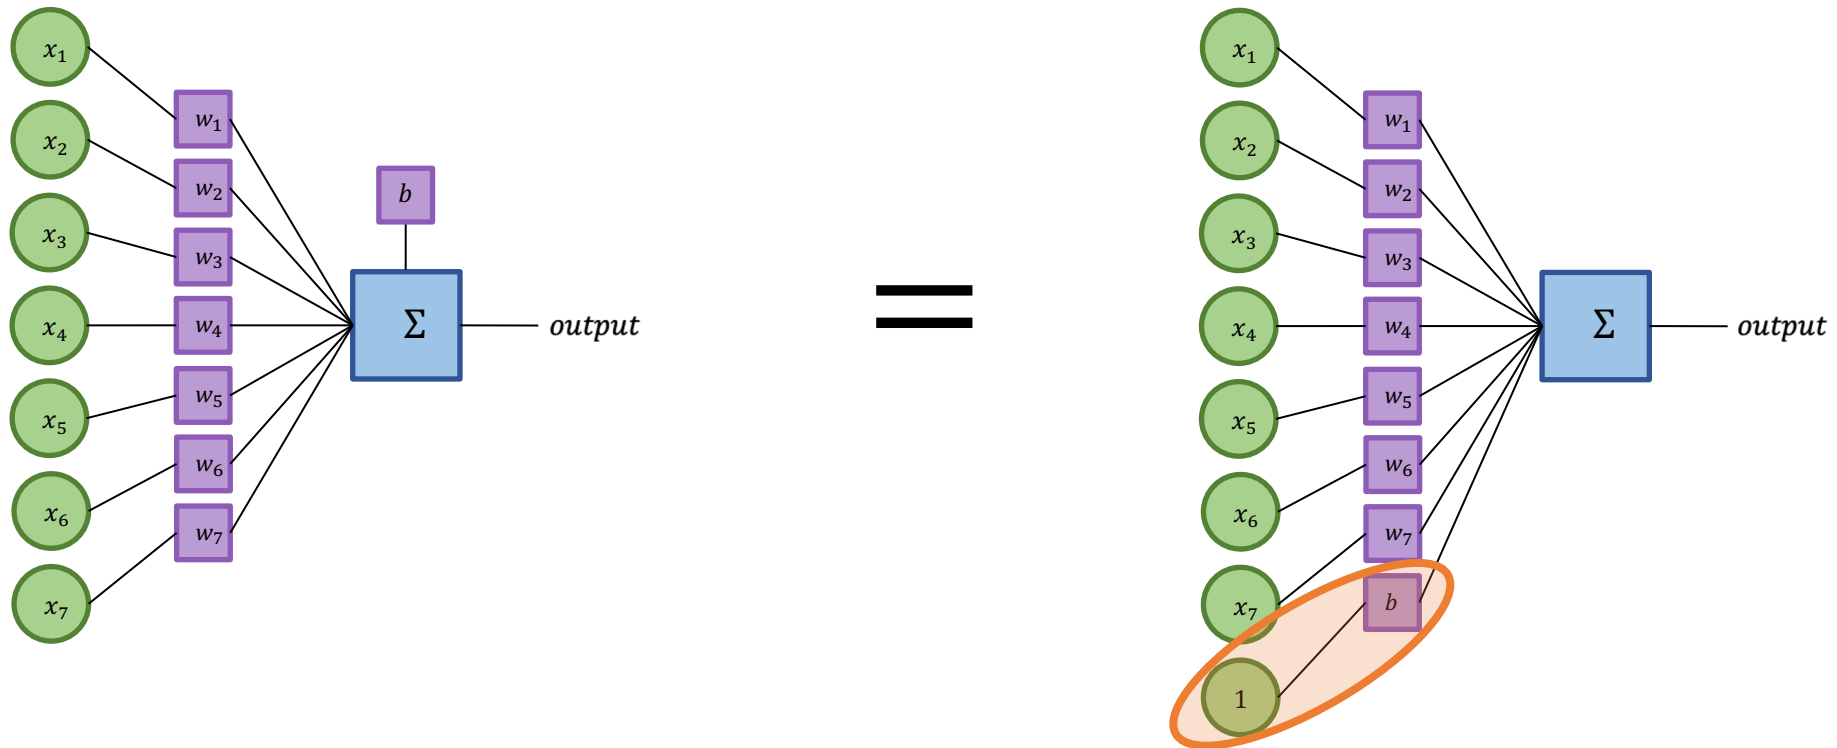
- **Bias** — the *a priori* likelihood of the positive class
 - Ensures that even if all inputs are 0, there will be some result value
 - Just because all inputs are 0, it does not mean there are no 1's in the world
 - Maybe there just happen to be more, say, 0's than 1's

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



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

$$\begin{aligned} & [x_0, x_1, x_2, \dots x_n] \cdot [w_0, w_1, w_2, \dots w_n] + b \\ = & [x_0, x_1, x_2, \dots x_n, 1] \cdot [w_0, w_1, w_2, \dots w_n, b] \end{aligned}$$

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:

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

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


Any questions?



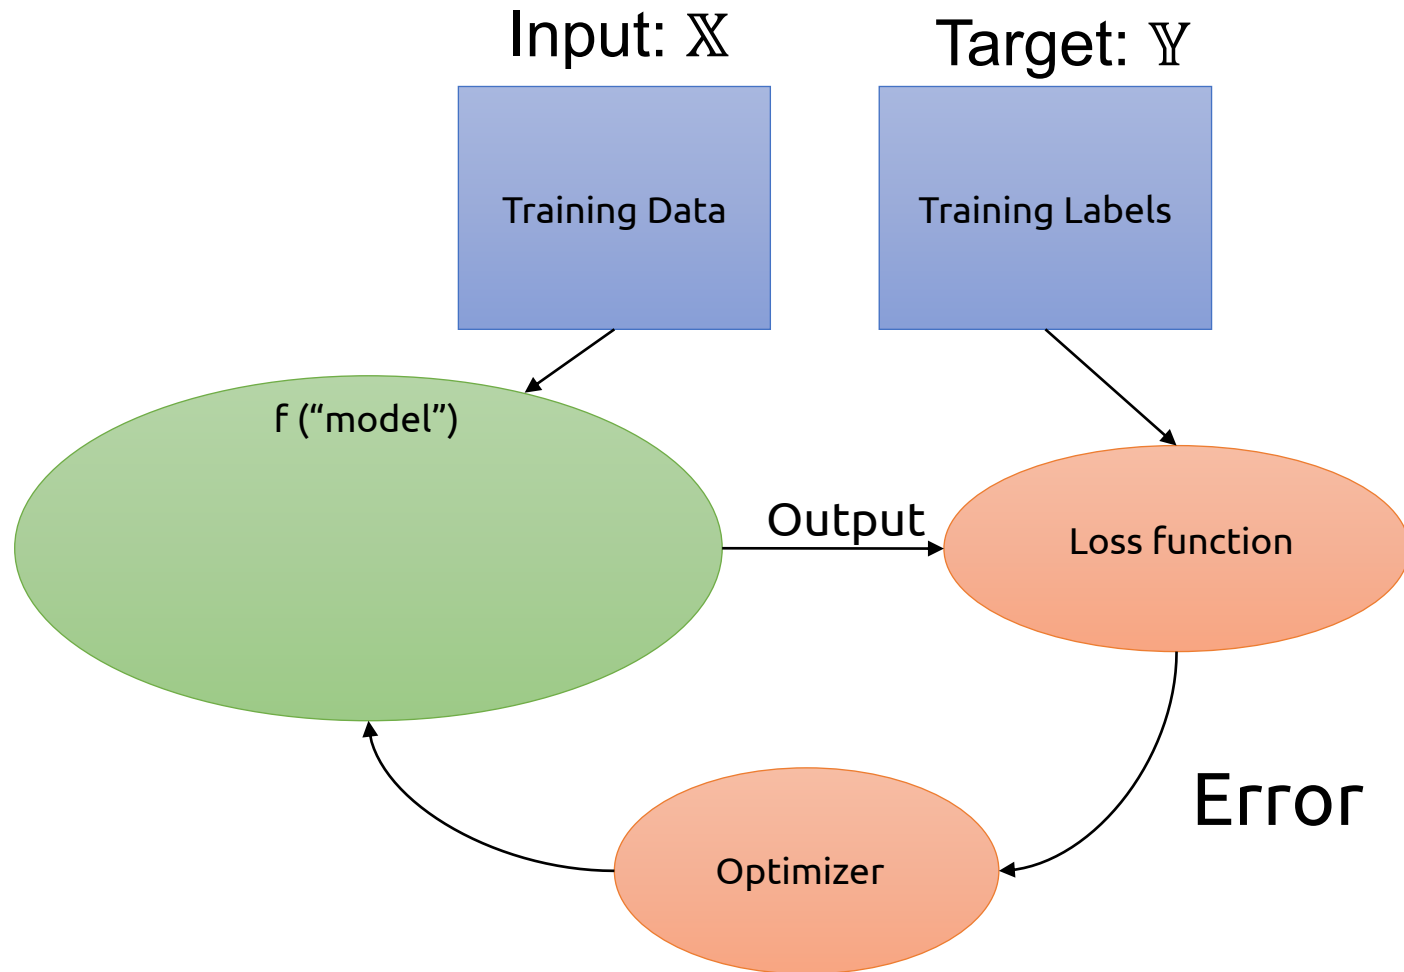
- In modern deep learning parlance, $b + \mathbf{w} \cdot \mathbf{x}$ is known as a ***linear unit***

A Binary Perceptron for MNIST

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



Training a perceptron (Next Class)




The Perceptron Learning Algorithm (Next class)

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 = k^{th} training example

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

Recap



MNIST Data

Representing
handwritten digits

Handwritten digit
prediction task

Machine learning
pipeline



Biological motivation

Perceptron equation

Perceptron parameters and
training

