CS540 Introduction to Artificial Intelligence

Lecture 1

Young Wu

Based on lecture slides by Jerry Zhu, Yingyu Liang, and Charles Dyer

May 23, 2022

Socrative

Admin

wu489@wisc.edu

ID = "WU489"

- Download the Socrative App or go to the Socrative website.
- Use Room CS540E log in with wisc ID.
- Choose "E" for the first question $Q1$.

Logistics, Grading

Admin

- Everything is on the course website.
- Talk about these and answer questions at the end of the lecture.

What is AI

Motivation

Is This Face Real

Quiz

- Which face is real?
- A : Left
- B :
- C :
- D :
- E : Right
- (Do not choose B, C, D.)

Q2 ~ Q5

Is This Face Real

Quiz

A5

- How is the real face different from the fake one?
- Short answer.
- (Write something random if you do not know: Do not leave it blank.)

Socrative

Admin

- Submit a regrade request if you missed any questions or selected the incorrect answer by mistake.

Generative Adversarial Network

Motivation

- Generative Adversarial Network (GAN):

1. Generative part: input random noise and output fake images.
2. Discriminative part: input real and fake images and output labels real or fake.
3. The two parts compete with each other.

Supervised Learning Example 1

Motivation

Data	images of cats and dogs
Features (Input)	height, length, eye color, ...
Labels (Output)	cat or dog

Data	images of 1000 object classes
Features (Input)	pixel information ...
Labels (Output)	turtle or rifle

Supervised Learning Example 2

Motivation

Data	handwritten characters
Features (Input)	pixel intensity, stroke, ...
Labels (Output)	δ or σ, φ or ψ

Data	voice recording
Features (Input)	signal, sound (phoneme), ...
Labels (Output)	recognize speech or wreck a nice beach

Supervised Learning Example 3

Motivation

Data	medical records
Features (Input)	scan, blood, and test results, ...
Labels (Output)	cancer or no cancer

Data	patient information
Features (Input)	age, pre-existing conditions, ...
Labels (Output)	cancer or no cancer

Supervised Learning Example 4

Motivation

Data	emails
Features (Input)	word count, capitalization, ...
Labels (Output)	spam or ham

Data	comments
Features (Input)	word count, capitalization, ...
Labels (Output)	offensive or not

Supervised Learning Example 5

Motivation

Data	face images
Features (Input)	edges, corners, ...
Labels (Output)	face or non-face

Data	self-driving car data
Features (Input)	color, distance (depth), movement, ...
Labels (Output)	road or car or pedestrian

Supervised Learning Example 6

Motivation

Data	book or movie reviews
Features (Input)	word count, capitalization, ...
Labels (Output)	positive or negative

Data	financial transactions
Features (Input)	amount, frequency, ...
Labels (Output)	fraud or not

Supervised Learning Example 7

Motivation

Data	painting
Features (Input)	appearance, price, ...
Labels (Output)	art or garbage

Data	essay
Features (Input)	length, key words, ...
Labels (Output)	A+ or F

Supervised Learning

Motivation

$y \approx \hat{f}(x)$

x_{ij} → ith instance, jth feature

- Supervised learning:

Data	Features	Labels	-
Sample	$\{(x_{i1}, \ldots, x_{im})\}_{i=1}^{n}$	$\{y_i\}_{i=1}^{n}$	find "best" $\hat{f}$
-	observable	known	-
New	$(x'_1, \ldots, x'_m)$	y'	guess $\hat{y} = \hat{f}(x')$
-	observable	unknown	-

Training and Test Sets

Motivation

- Supervised learning:

Data	Features	Labels	-
Training	$\{(x_{i1}, \ldots, x_{im})\}_{i=1}^{n'}$	$\{y_i\}_{i=1}^{n'}$	find "good" $\hat{f}$
-	observable	known	-
Validation	$\{(x_{i1}, \ldots, x_{im})\}_{i=n'}^{n}$	$\{y_i\}_{i=n'}^{n}$	find "best" $\hat{f}$
-	observable	known	-
Test	$(x'_1, \ldots, x'_m)$	y'	guess $\hat{y} = \hat{f}(x')$
-	observable	unknown	-

evaluate how good $\hat{f}$ is.

Simple 2D Example Diagram

Motivation

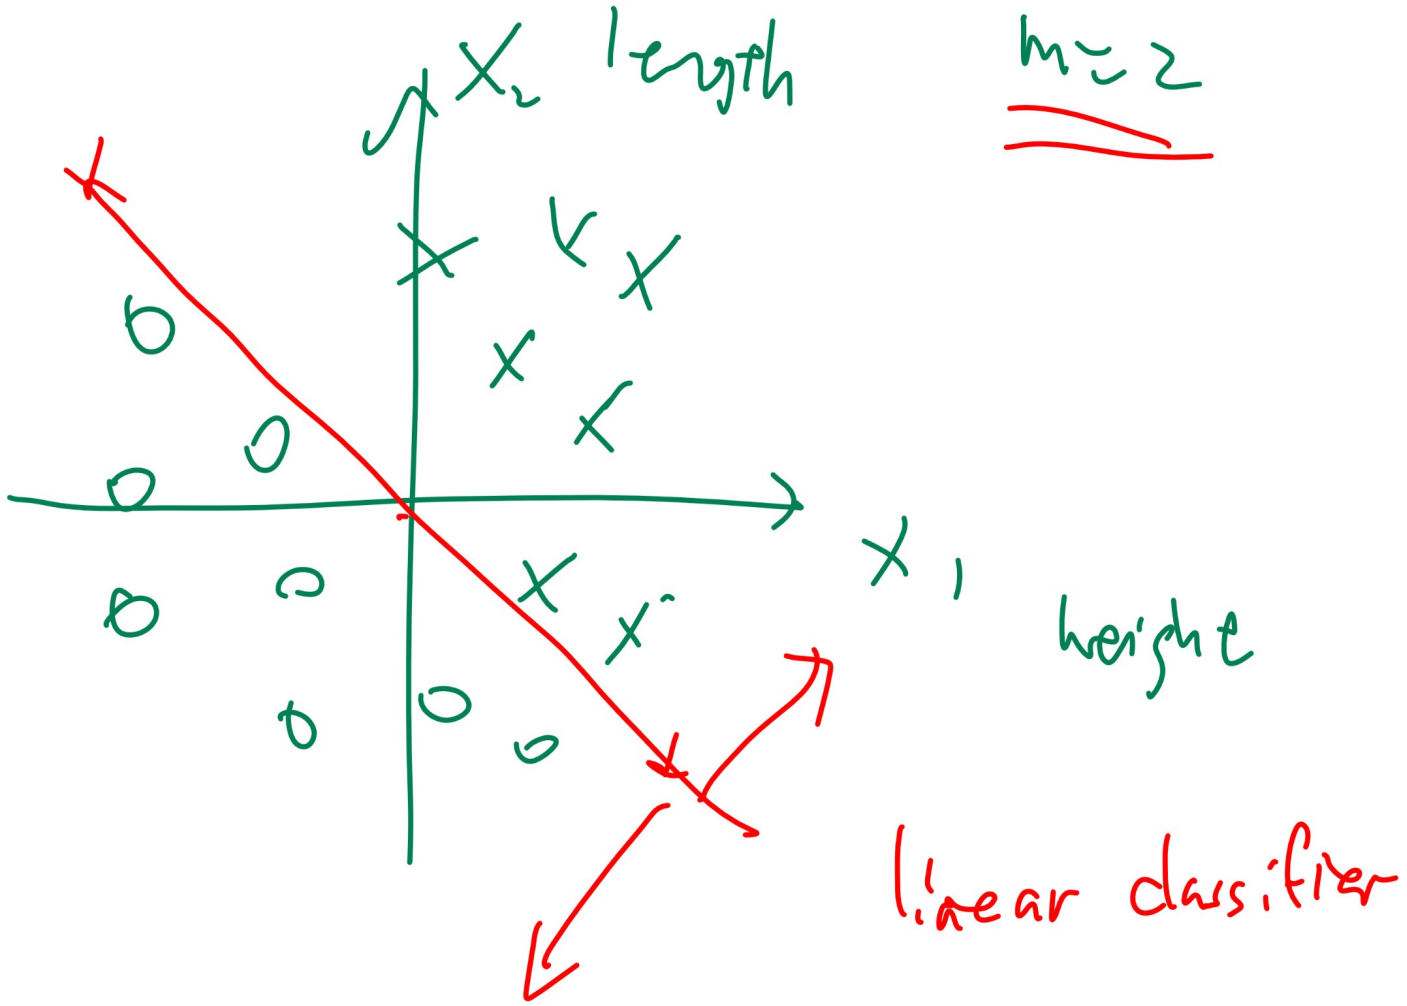

Linear Classifier

Motivation

$\hat{y}$ is 0 or 1

- One possible guess is in the form of a linear classifier,

$$\hat{y} = \mathbb{1}_{\{w_1 x_1 + w_2 x_2 + \ldots + w_m x_m + b \geq 0\}} = \mathbb{1}_{\{w^T x + b \geq 0\}}$$

$= \begin{cases} 1 & w_1 x_1 + \ldots + b \geq 0 \\ 0 & \text{—— c0} \end{cases}$

$\mathbb{1}_E$

- The $\mathbb{1}$ (open number 1) is the indicator function.

$$\mathbb{1}_E = \begin{cases} 1 & \text{if } E \text{ is true} \\ 0 & \text{if } E \text{ is false} \end{cases}$$

weight

$$w = \begin{pmatrix} w_1 \\ w_2 \\ \vdots \\ w_m \end{pmatrix}$$

features

$$X = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_m \end{pmatrix}$$

bias

$$(w_1, w_2 \ldots w_m) \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} + b$$

Brute Force LTU Learning

Motivation

linear threshold unit.

Perceptron Algorithm

Description

- Initialize random weights.
- Evaluate the activation function at one instance x_i to get $\hat{y}_i$.
- If the prediction $\hat{y}_i$ is 0 and actual y_i is 1, increase the weights by x_i.
- If the prediction $\hat{y}_i$ is 1 and actual y_i is 0, decrease the weights by x_i.
- Repeat for all data points and until convergent.

Perceptron Algorithm Diagram

Description

Perceptron Algorithm

Definition

- Update weights using the following rule.

$$w = w - \alpha (a_i - y_i) x_i$$

$$b = b - \alpha (a_i - y_i)$$

$$a_i = \mathbb{1}_{\{w^T x_i + b \geq 0\}}$$

predicted value

random instance

actual label in training set

learning rate.

GU $\hat{y}_i = a_i$ → activation value

PI

Perceptron Algorithm

Quiz

- Let the learning rate be $\alpha = 0.2$. Currently $w = \begin{bmatrix} 0.2 & 0.7 & 0.9 \end{bmatrix}^T$, $b = -0.7$, and $x_i = \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^T$ and $y_i = 0$. What is the updated weights $\begin{bmatrix} w \\ b \end{bmatrix}$?

plane in 3D

$$\begin{bmatrix} w \\ b \end{bmatrix} = \begin{bmatrix} w \\ b \end{bmatrix} - \alpha (a_i - y_i) \begin{bmatrix} x_i \\ 1 \end{bmatrix} = \begin{bmatrix} 0.2 \\ 0.7 \\ 0.9 \\ -0.7 \end{bmatrix} - 0.2 \begin{bmatrix} 0 \\ 0 \\ 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 0.2 \\ 0.7 \\ 0.7 \\ -0.9 \end{bmatrix}$$

$$a_i = \mathbb{1}(w^T x_i + b) = \mathbb{1}\left(\begin{bmatrix} 0.2 \\ 0.7 \\ 0.9 \end{bmatrix}^T \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} - 0.7 = \mathbb{1}(0.2 \geq 0) = 1 \right)$$

Perceptron Algorithm, Answer

Quiz

Perceptron Algorithm, Another One

Quiz

- Let the learning rate be $\alpha = 0.1$. Currently $w = \begin{bmatrix} 0.2 \\ -0.3 \end{bmatrix}, b = 0.4$, and $x_i = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$ and $y_i = 1$. What is the updated weights $\begin{bmatrix} w \\ b \end{bmatrix}$?
- $A: \begin{bmatrix} 0.2 \\ -0.3 \\ 0.4 \end{bmatrix}, B: \begin{bmatrix} 0.2 \\ -0.2 \\ 0.5 \end{bmatrix}, C: \begin{bmatrix} 0.2 \\ -0.4 \\ 0.3 \end{bmatrix}$
- $D: \begin{bmatrix} 0.2 \\ -0.2 \\ 0.3 \end{bmatrix}, E: \begin{bmatrix} 0.2 \\ -0.4 \\ 0.5 \end{bmatrix}$

Q6

like in CJ

$-0.3 + 0.4$

$0.1 \geq 0$

$a_i = 1$

$\begin{bmatrix} w \\ b \end{bmatrix} - \alpha (a_i - y_i) \begin{bmatrix} x_i \\ 1 \end{bmatrix}$

$a_i = \mathbb{1}\{w^T x_i + b\}$

Perceptron Algorithm, Another One, Answer

Quiz

Perceptron Algorithm, Another One Too

Quiz

- Let the learning rate be $\alpha = 0.1$. Currently $w = \begin{bmatrix} 0.2 \\ -0.3 \end{bmatrix}$, $b = 0.4$, and $x_i = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$ and $y_i = 0$. What is the updated weights $\begin{bmatrix} w \\ b \end{bmatrix}$?

- $A: \begin{bmatrix} 0.2 \\ -0.3 \\ 0.4 \end{bmatrix}$, $B: \begin{bmatrix} 0.2 \\ -0.2 \\ 0.5 \end{bmatrix}$, $C: \begin{bmatrix} 0.2 \\ -0.4 \\ 0.3 \end{bmatrix}$

- $D: \begin{bmatrix} 0.2 \\ -0.2 \\ 0.3 \end{bmatrix}$, $E: \begin{bmatrix} 0.2 \\ -0.4 \\ 0.5 \end{bmatrix}$

Q7

$\mathbb{1}\{w^T x + b \geq 0\}$

$\begin{bmatrix} w \\ b \end{bmatrix} - \alpha(a_i - y_i)\begin{pmatrix} x_i \\ 1 \end{pmatrix}$

0.1

1 0

back @ 6:40

Perceptron Algorithm, Another One Too, Answer

Quiz