

CS540 Introduction to Artificial Intelligence

Lecture 1

Young Wu

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

June 27, 2022

Socrative

Admin

- Download the Socrative App or go to the Socrative website.
- Use Room CS540C log in with wisc ID.
- Choose " C" 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*.)

Is This Face Real

Quiz

- 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 <i>F</i>

Supervised Learning

Motivation

- Supervised learning:

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

Training and Test Sets

Motivation

- Supervised learning:

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

Simple 2D Example Diagram

Motivation

Linear Classifier

Motivation

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

$$\begin{aligned}\hat{y} &= \mathbb{1}_{\{w_1x_1 + w_2x_2 + \dots + w_mx_m + b \geq 0\}} \\ &= \mathbb{1}_{\{w^T x + b \geq 0\}}\end{aligned}$$

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

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

Brute Force LTU Learning

Motivation

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.

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

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

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

Perceptron Algorithm

Quiz

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

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 $\begin{bmatrix} w \\ b \end{bmatrix}$?

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

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 $\begin{bmatrix} w \\ b \end{bmatrix}$?

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

Perceptron Algorithm, Another One Too, Answer Quiz