CS540 Introduction to Artificial Intelligence
Lecture 1

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

May 18, 2020

Socrative

Admin

- Download the Socrative App or go to the Socrative.
- Use Room CS540E log in with wisc ID.
- Choose "D" for the first question Q1.
- If you cannot login, private message to the person named "Questions" (also me).

Lecture Format

Admin

- Pre-recorded lectures will be posted on the course website.
- University-assigned lecture time will be used to go over examples and for participation quizzes.
- The remaining lecture time will be used as office hours.

Lecture Recording

Admin

- These BBCU sessions will not be recorded.
- All course materials will be covered in the pre-recorded lectures on YouTube.
- The screen shots of the examples and quiz questions will be posted.
- Use these sessions either as a review or preview of the materials of the week.

Grading

Admin

- Quizzes: best 10 of 12 or double exam weights.
- Math homework: best 10 of $10 + 2$.
- Programming homework: best 5 of $5 + 1$.
- Exams: one midterm and one final, 10 points each.

Quizzes

Admin

- Download Socrative, the room number is CS540E.
- Default login for Socrative is your wisc email ID.
- If someone else tries to hack your account, please email or post on Piazza.
- Quiz questions can show up any time during the lecture.
- Missing one or two questions due to technical difficulty is okay.
- If you select obviously false answers, you might lose points.

Socrative Test

Quiz

- A: Don't choose this
- B: Don't choose this
- C: Don't choose this
- D: Choose this
- E: Don't choose this

Math Homework

Admin

- Officially: due in 1 week Sunday.
- Unofficially: any time before the midterm or the final.
- Auto-graded: submit the output on Canvas.

Programming Homework

Admin

- Officially: due in 2 weeks Sunday.
- Unofficially: any time before the final.
- Solution: posted in 1 week Sunday.
- Auto-graded: submit the output on Canvas.
- Code: any language.

Favorite Programming Language

Quiz

Q2

- What is your favorite programming language (choose one)?
- A: Java
- B: Python
- C: Matlab
- D: C++
- E: Other

Midterm and Final

Admin

- Synchronous exam: morning and evening one, choose one to take.

Textbook

Admin

- Lecture slides and videos will be sufficient.
- RN is a good background reading, does not cover everything.
- SS is very theoretical, useful if you are planning to take $760, 761, 861$.

Admin

Admin

- Math and Stat Review posted under W1.
- Annotated slides will not be posted (because my handwriting is not recognizable).
- Unofficially: all homework are already posted (lots of mistakes and bugs).
- Officially: homework will be posted two to three days after the corresponding lecture.

Questions

Admin

- Questions?
- Raise your hand, or just interupt me.
- Private message to the person named "Questions".
- Do not message me.

Is This Face Real

Quiz

Q3

1

Q7

- Which face is real?
- A: Left
- B: Right
- C: Do not choose this
- D: Do not choose this
- E: Do not choose this

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

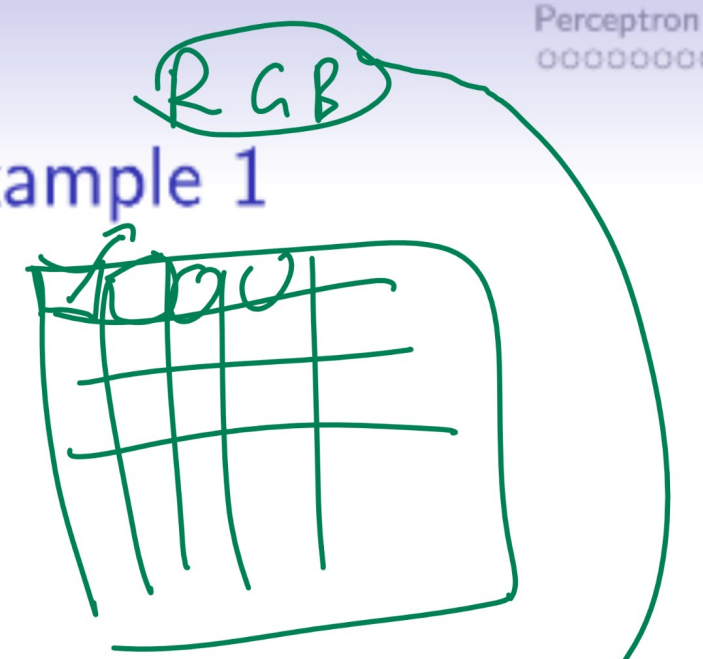

Data	images of cats and dogs
Features (Input)	height, length, eye color, ...
	pixel intensity
Output	cat or dog

Supervised Learning Example 2

Motivation

Data	medical records
Features (Input)	scan, blood, and other test results
Output	disease or no

Data	patient information
Features (Input)	age, pre-existing conditions, ...
Output	likelihood of death

Supervised Learning Example 3

Motivation

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

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

Supervised Learning Example 4

Motivation

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

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

Supervised Learning Example 5

Motivation

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

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

Supervised Learning Example 6

Motivation

Data	handwritten letters
Features (Input)	pixel, stroke
Output	δ or σ, φ or ψ

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

Supervised Learning Example 7

Motivation

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

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

Supervised Learning

Motivation

- Supervised learning:

Data	Features (Input)	Output	-
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	

Simple 2D Example Diagram

Motivation

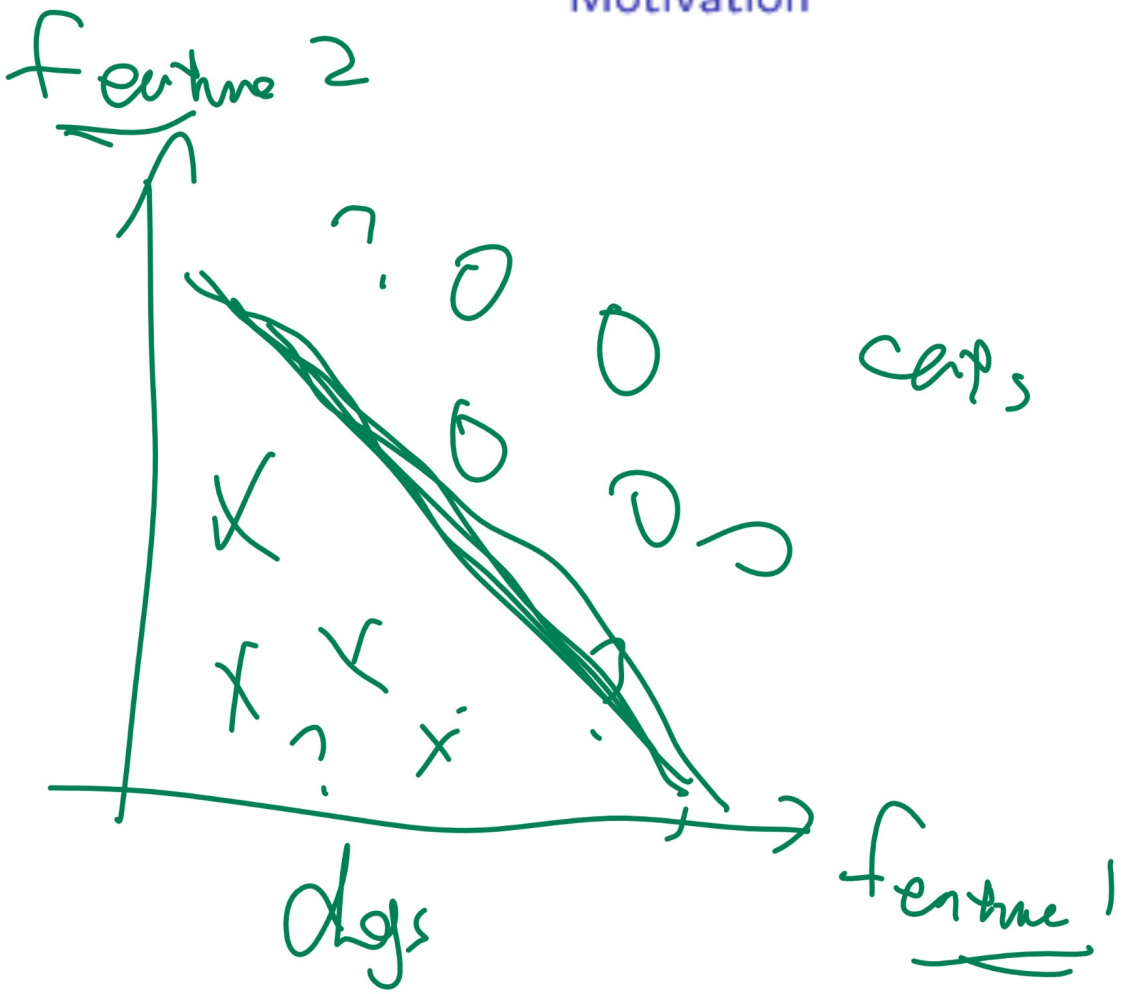

Linear Classifier

Motivation

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

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

- 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

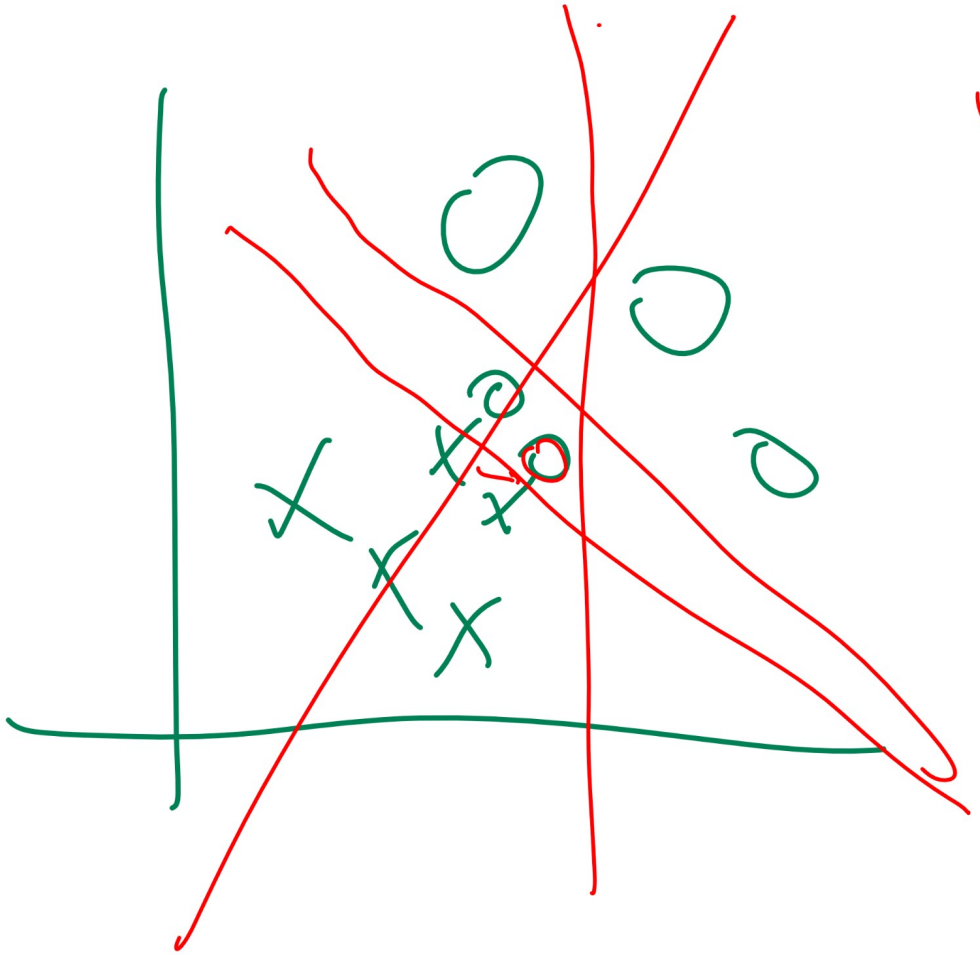

features

$w_1 x_1 + w_2 x_2 + b \geq 0$

weights bias

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, 0 Example

Description

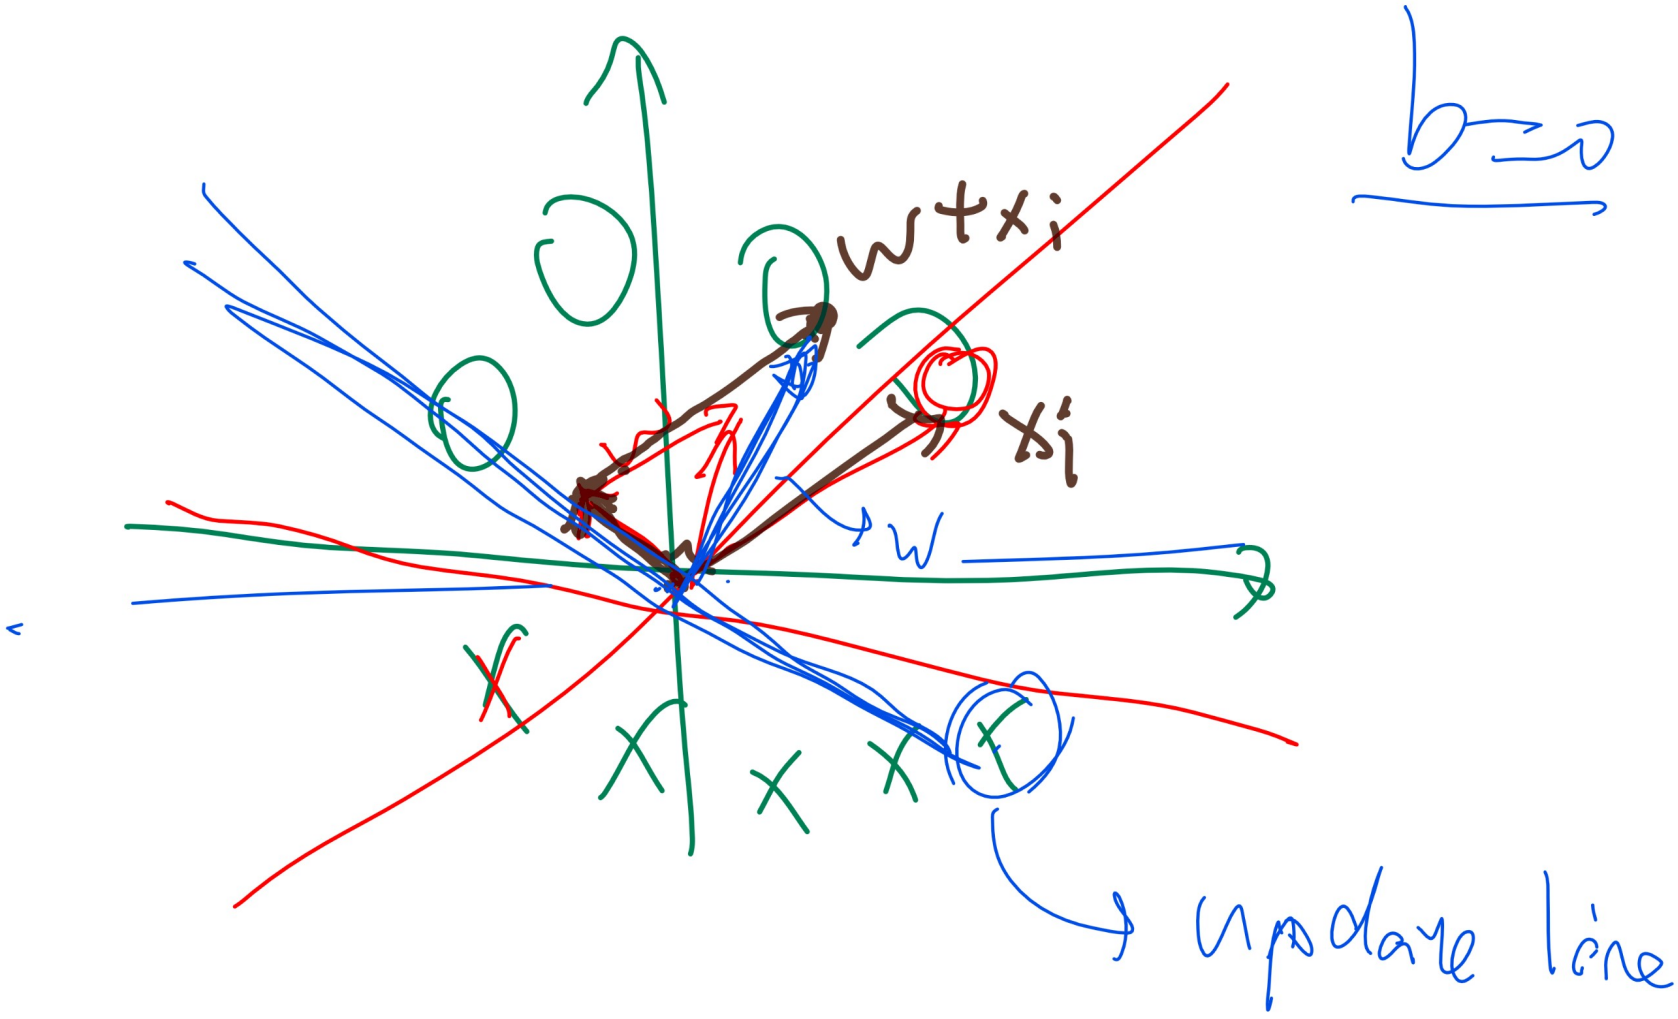

Perceptron Algorithm Diagram, 1 Example

Description

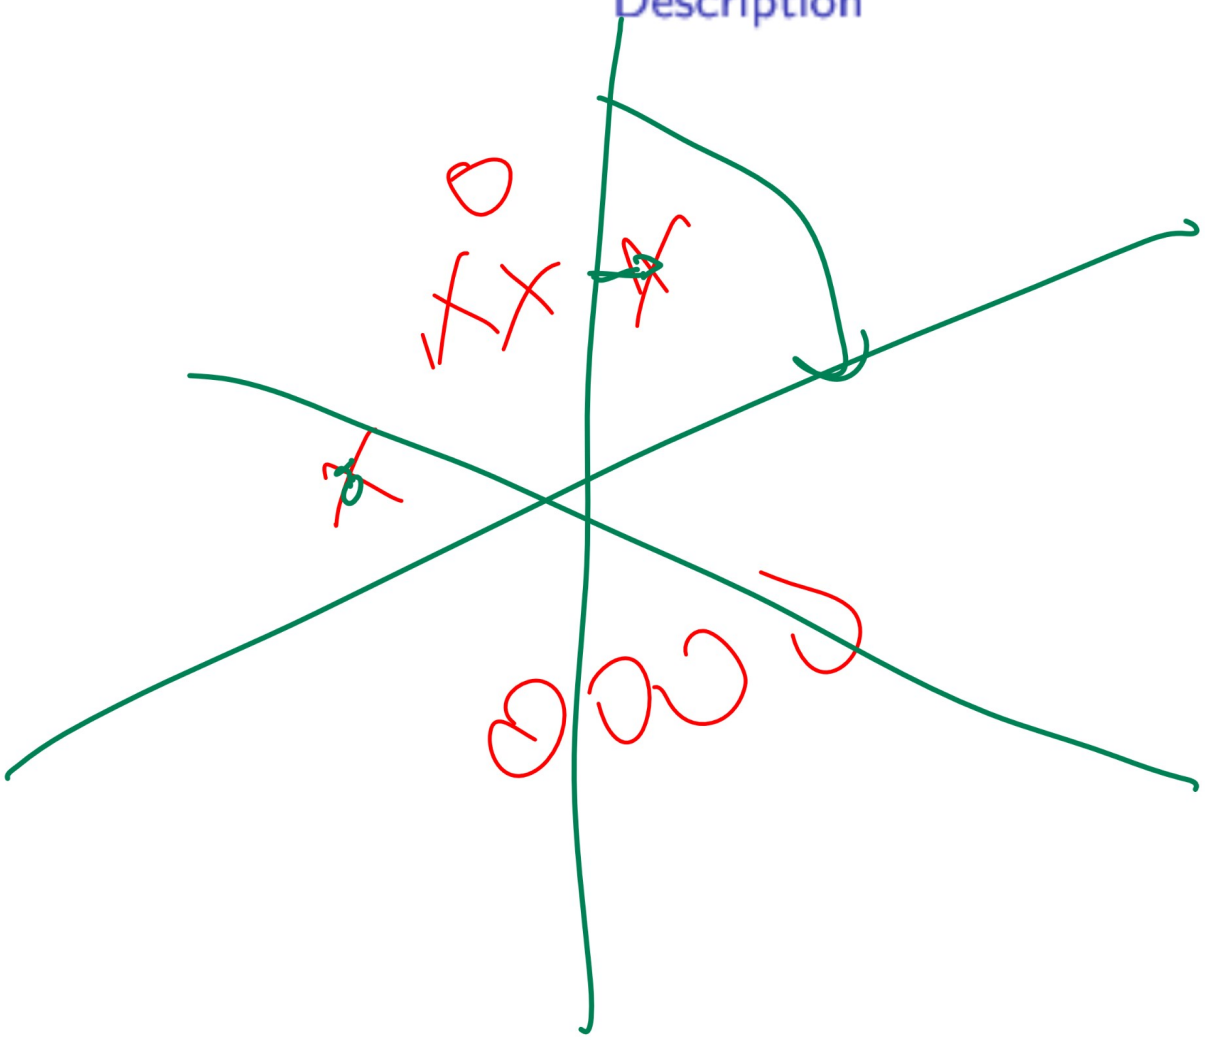

Perceptron Algorithm

Definition

- Update weights using the following rule.

$$w = w - \alpha(a_i - y_i)x_i$$
$$b = b - \alpha(a_i - y_i)$$
$$a_i = \mathbb{1}_{\{w^T x_i + b \geqslant 0\}}$$

predicted 0 1

learning rate

Perceptron Algorithm

Quiz

- Spring 2017 Final Exam Q3
- 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 $\begin{bmatrix} w \\ b \end{bmatrix}$?

$$\hat{y} = a_i = \mathbb{1}\{w^T x + b \geq 0\} \leftarrow \text{prediction}$$

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

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

Perceptron Algorithm, Answer

Quiz

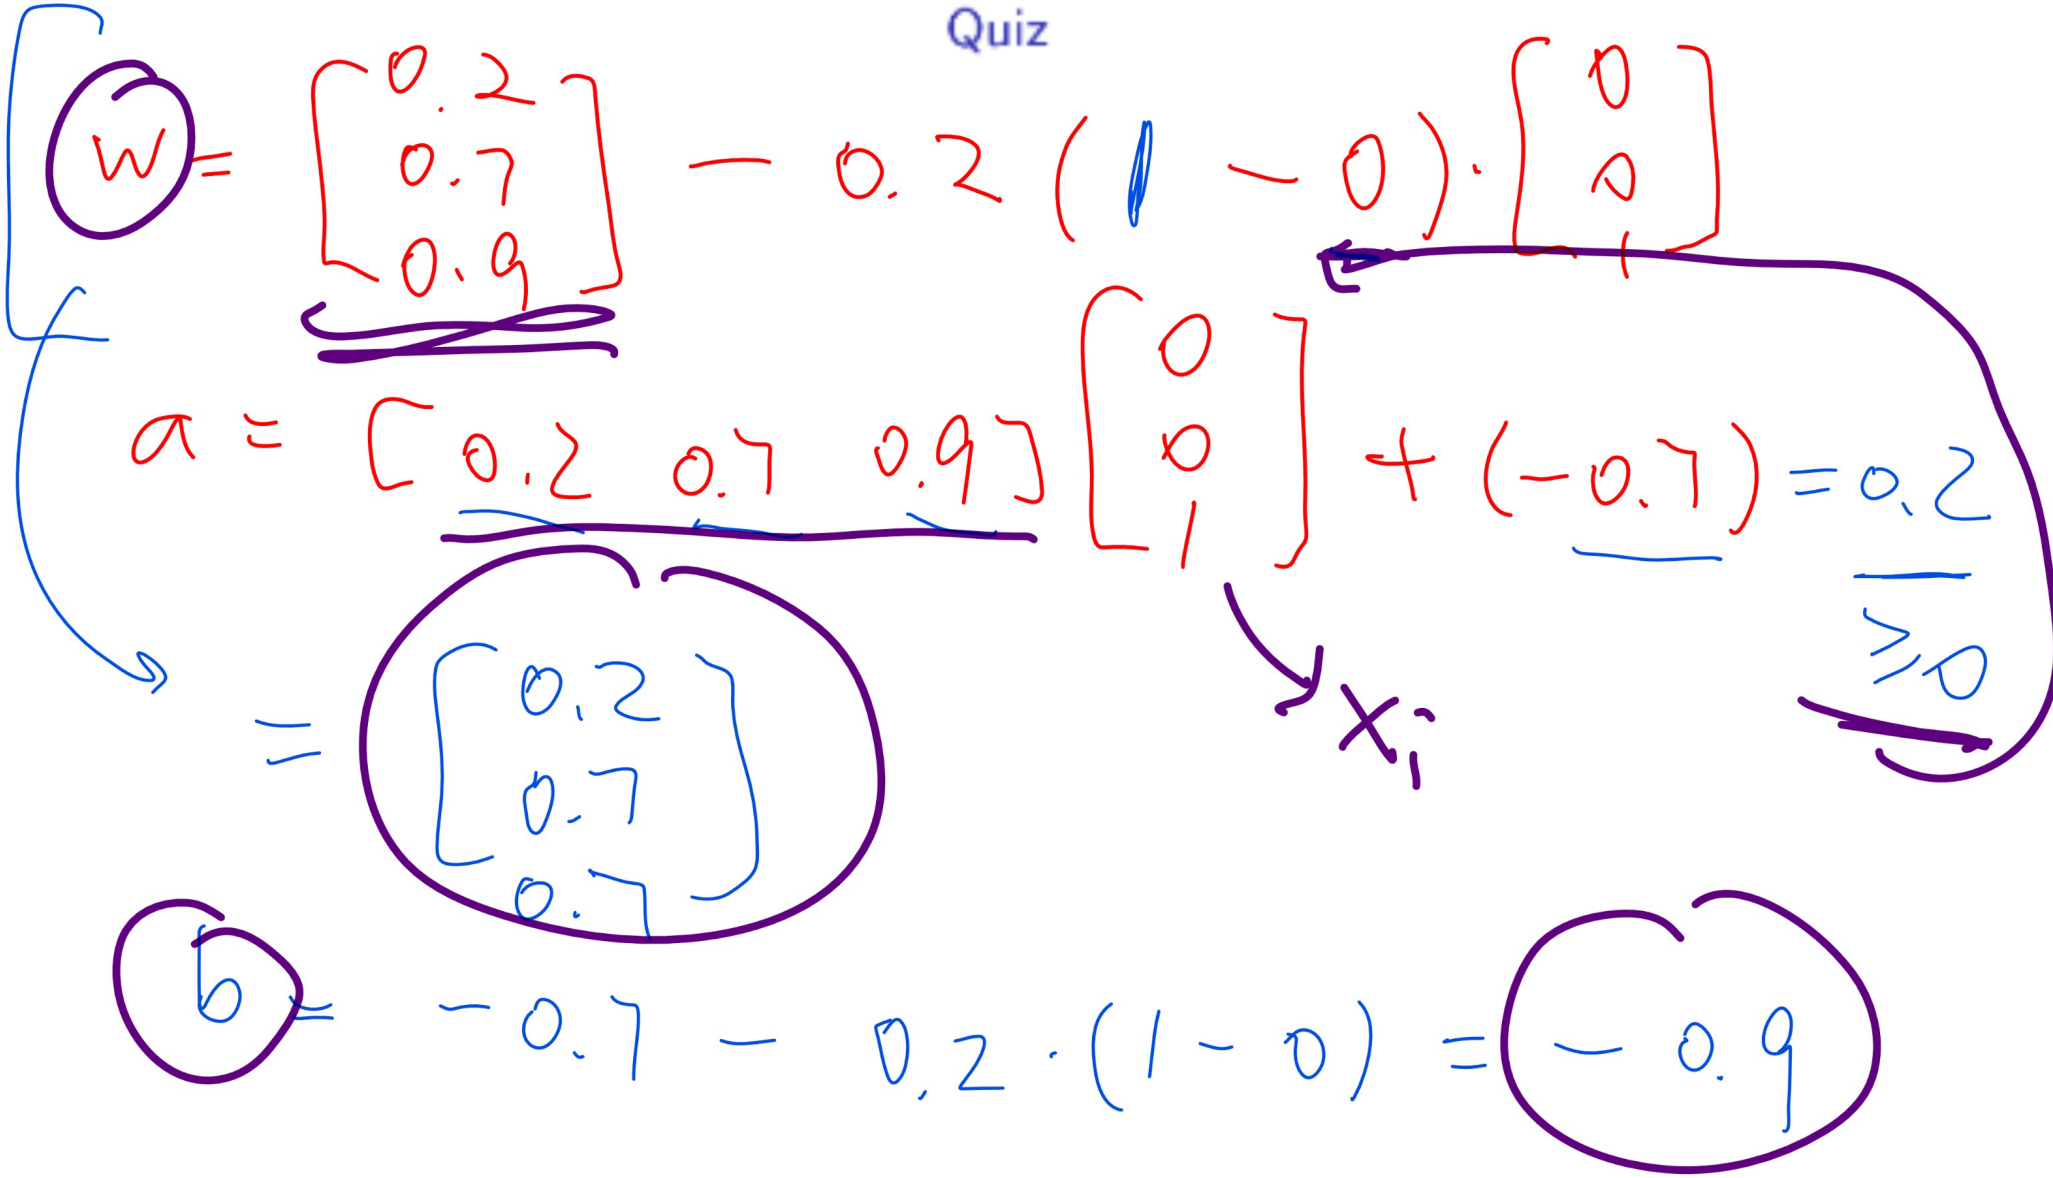

Perceptron Algorithm, Another One

Quiz

Q8

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

$w_1 x_1 + w_2 x_2 + b \cdot 1 \geq 0$

$0.2 \cdot 0 + (-0.3) \cdot 1 + 0.4$

$= 0.1 \geq 0$

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

$a_i = 1$

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

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

$b = b - \alpha(a_i - y_i) \cdot 1$

Perceptron Algorithm, Another One, Answer

Quiz