CS540 Introduction to Artificial Intelligence

Lecture 2

Young Wu

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

May 23, 2022

Two-thirds of the Average Game

Quiz

A7

- Pick an integer between 0 and 100 (including 0 and 100) that is the closest to two-thirds of the average of the numbers other people picked.

Make Up Lectures

Admin

- Next Monday is Memorial Day.
- A : No make up lecture
- B : Make up lecture next Wednesday
- C : Make up lecture next next Wednesday

Supervised Learning

Motivation

$y \in \{w^T x + b\}$

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

Loss Function Diagram

Motivation

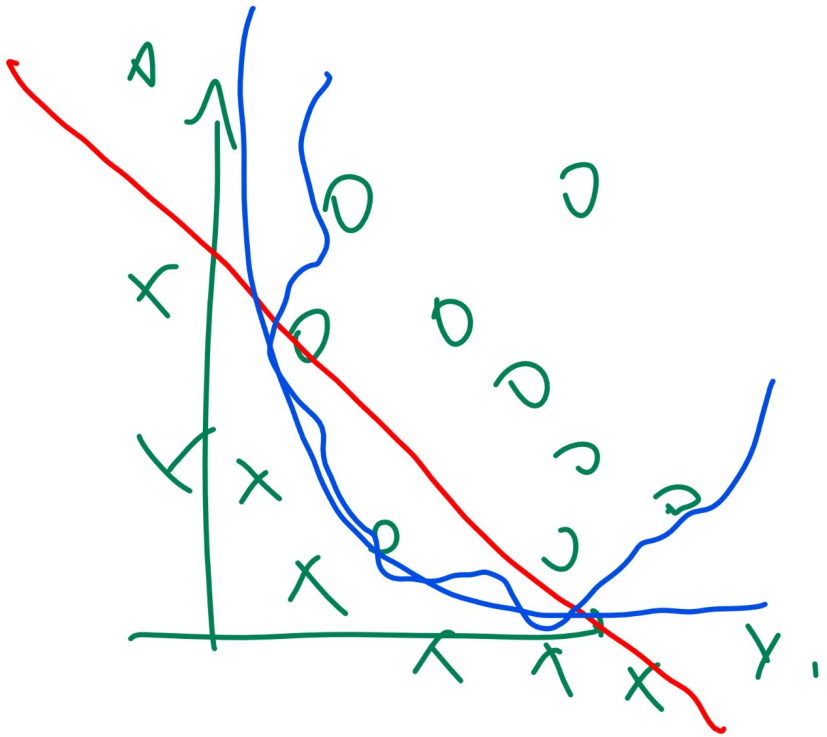

Zero-One Loss Function

Motivation

- An objective function is needed to select the "best" $\hat{f}$. An example is the zero-one loss.

$$\hat{f} = \underset{f}{\operatorname{argmin}} \sum_{i=1}^{n} \mathbb{1}_{\{f(x_i) \neq y_i\}}$$

- $\underset{f}{\operatorname{argmin}}$ objective (f) outputs the function that minimizes the objective.
- The objective function is called the cost function (or the loss function), and the objective is to minimize the cost.

Squared Loss Function

Motivation

- Zero-one loss counts the number of mistakes made by the classifier. The best classifier is the one that makes the fewest mistakes.
- Another example is the squared distance between the predicted and the actual y value:

$$\hat{f} = \operatorname*{argmin}_f \frac{1}{2} \sum_{i=1}^{n} \left(f\left(x_i\right) - y_i\right)^2$$

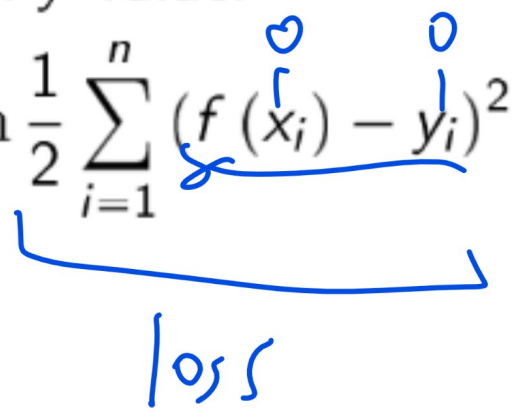

Loss Functions Equivalence

Quiz

- Which one of the following functions is not equivalent to the squared error for binary classification?

$$C = \sum_{i=1}^{n} \left(f\left(x_i\right) - y_i\right)^2, f\left(x_i\right) \in \{0,1\}, y_i \in \{0,1\}$$

- $A: \sum \mathbb{1}_{\{f(x_i) \neq y_i\}}$
- $B: \sum \mathbb{1}_{\{f(x_i) = y_i\}}$
- $C: \sum \left|f\left(x_i\right) - y_i\right|$
- $D: \sum \max\left\{0, 1 - f\left(x_i\right) y_i\right\}$
- $E: \sum \frac{1}{2} \max\left\{0, 1 - \left(2 \cdot f\left(x_i\right) - 1\right)\left(2 \cdot y_i - 1\right)\right\}$

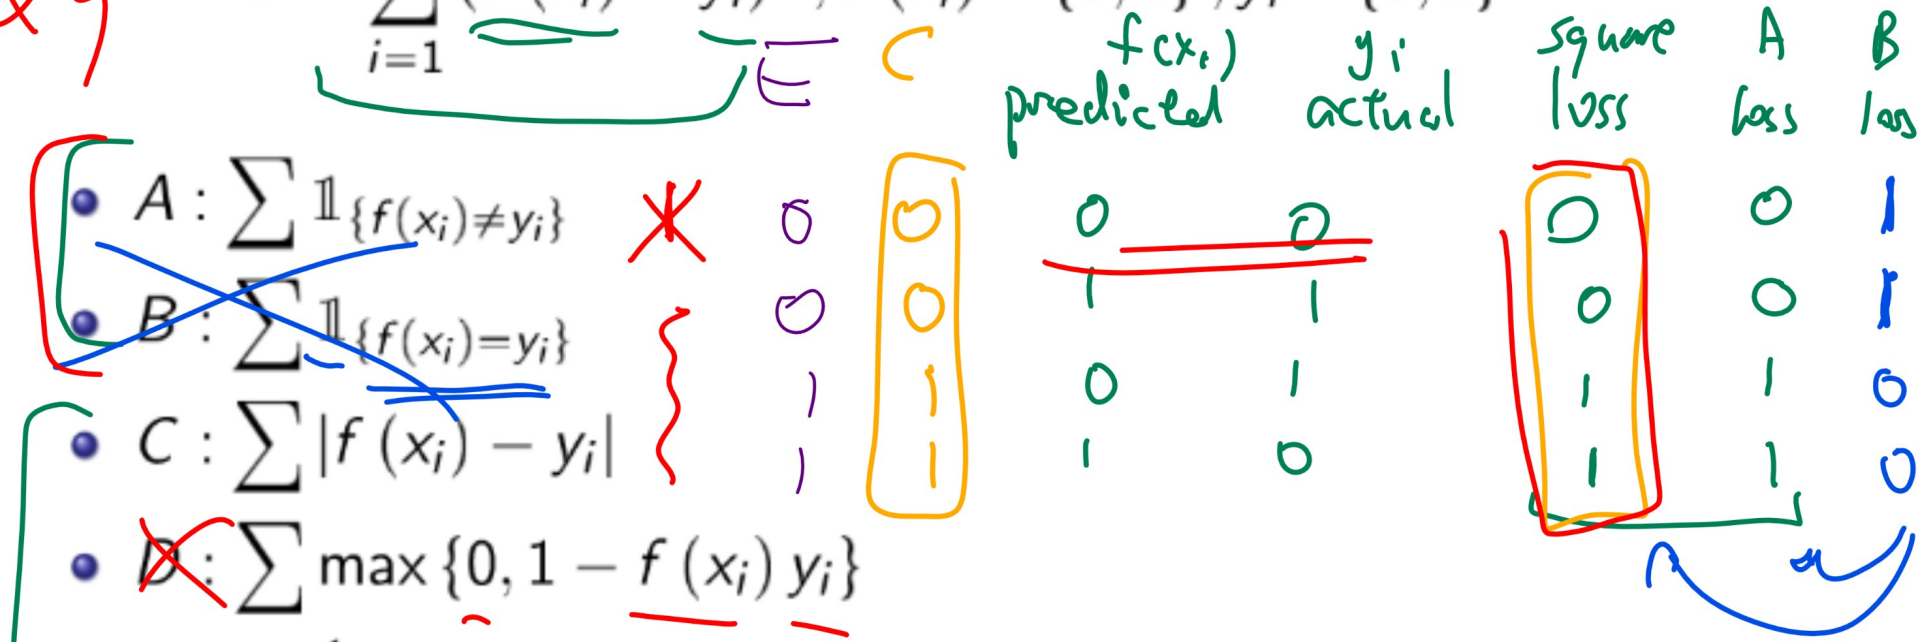

Loss Functions Equivalence, Answer

Quiz

Function Space Diagram

Motivation

Hypothesis Space

Motivation

- There are too many functions to choose from.
- There should be a smaller set of functions to choose $\hat{f}$ from.

$$\hat{f} = \underset{f \in \mathcal{H}}{\operatorname{argmin}} \frac{1}{2} \sum_{i=1}^{n} \left(f\left(x_i\right) - y_i \right)^2$$

- The set $\mathcal{H}$ is called the hypothesis space.

Activation Function

Motivation

- Suppose $\mathcal{H}$ is the set of functions that are compositions between another function g and linear functions.

$$\left(\hat{w}, \hat{b}\right) = \underset{w,b}{\operatorname{argmin}} \frac{1}{2} \sum_{i=1}^{n} (a_i - y_i)^2$$

where $a_i = g\left(w^T x + b\right)$

- g is called the activation function.

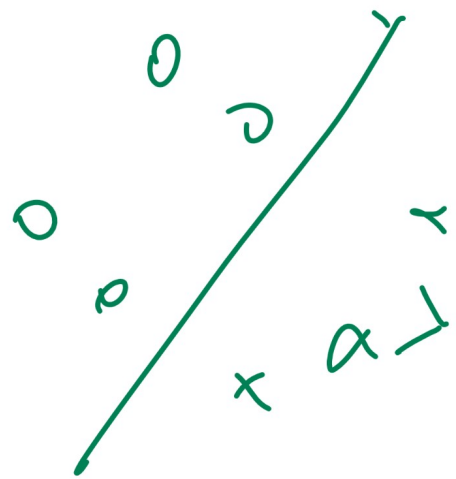

Linear Threshold Unit

Motivation

- One simple choice is to use the step function as the activation function:

$$g(\boxed{\cdot}) = \mathbb{1}_{\{\boxed{\cdot} \geqslant 0\}} = \begin{cases} 1 & \text{if } \boxed{\cdot} \geqslant 0 \\ 0 & \text{if } \boxed{\cdot} < 0 \end{cases}$$

- This activation function is called linear threshold unit (LTU).

Sigmoid Activation Function

Motivation

- When the activation function g is the sigmoid function, the problem is called logistic regression.

$$g\left(\boxed{\cdot}\right) = \frac{1}{1 + \exp\left(-\boxed{\cdot}\right)}$$

- This g is also called the logistic function.

Sigmoid Function Diagram

Motivation

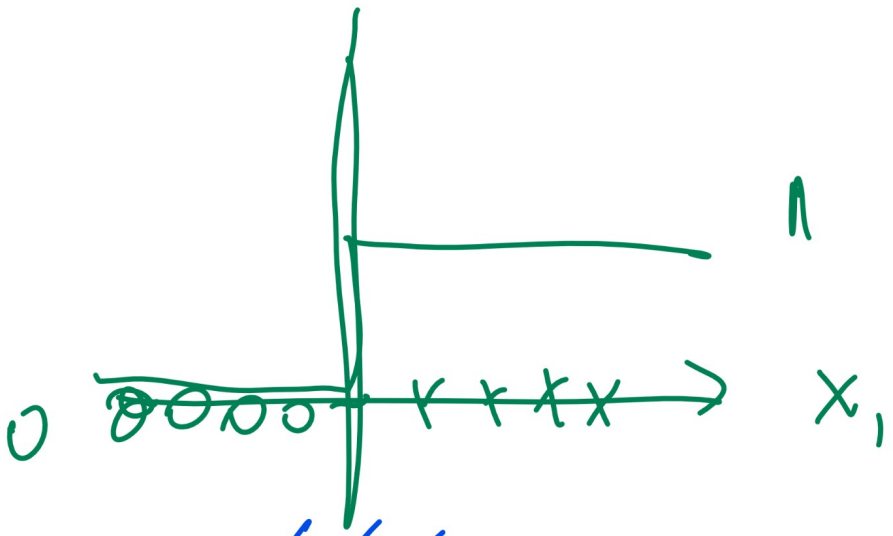

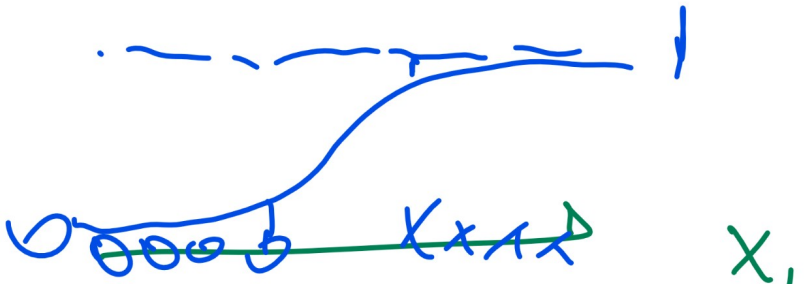

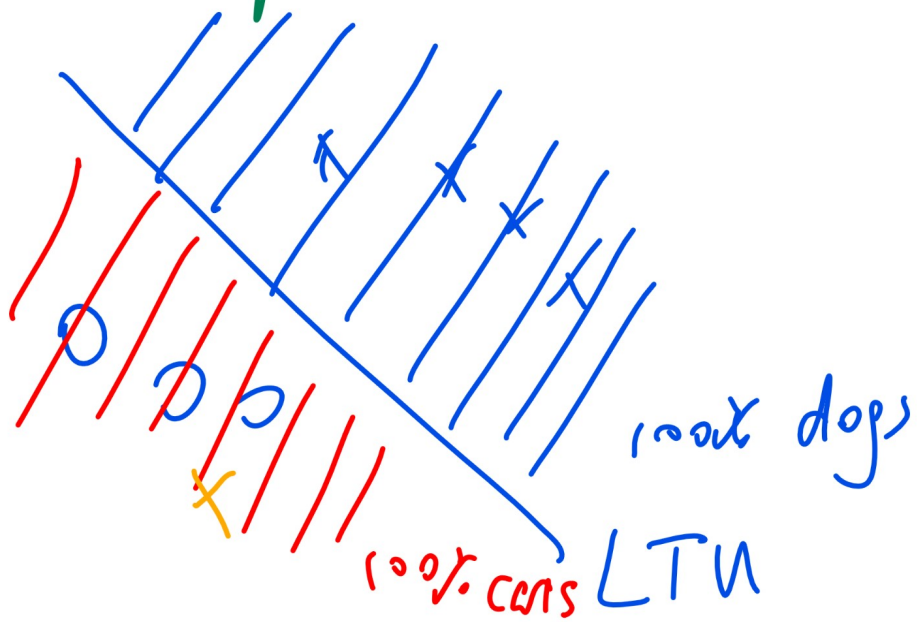

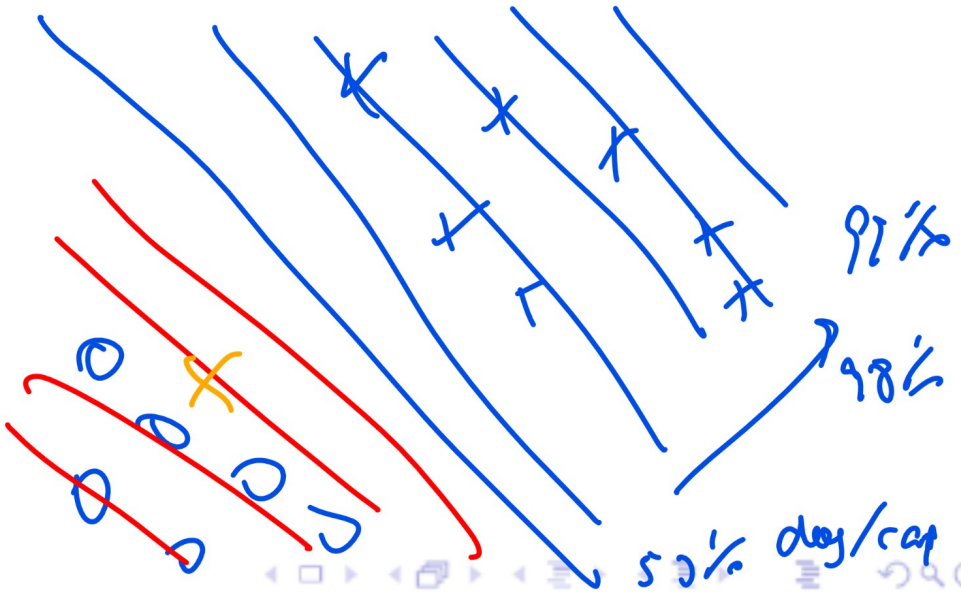

Cross-Entropy Loss Function

Motivation

- The cost function used for logistic regression is usually the log cost function.

$$C(f) = -\sum_{i=1}^{n} \left(y_i \log\left(f\left(x_i\right)\right) + \left(1 - y_i\right) \log\left(1 - f\left(x_i\right)\right)\right)$$

- It is also called the cross-entropy loss function.

logistic

Logistic Regression Objective

Motivation

- The logistic regression problem can be summarized as the following.

$$\left(\hat{w}, \hat{b}\right) = \underset{w,b}{\operatorname{argmin}} - \sum_{i=1}^{n} \left(y_i \log\left(a_i\right) + \left(1 - y_i\right) \log\left(1 - a_i\right)\right)$$

where $a_i = \dfrac{1}{1 + \exp\left(-z_i\right)}$ and $z_i = w^T x_i + b$

Optimization Diagram

Motivation

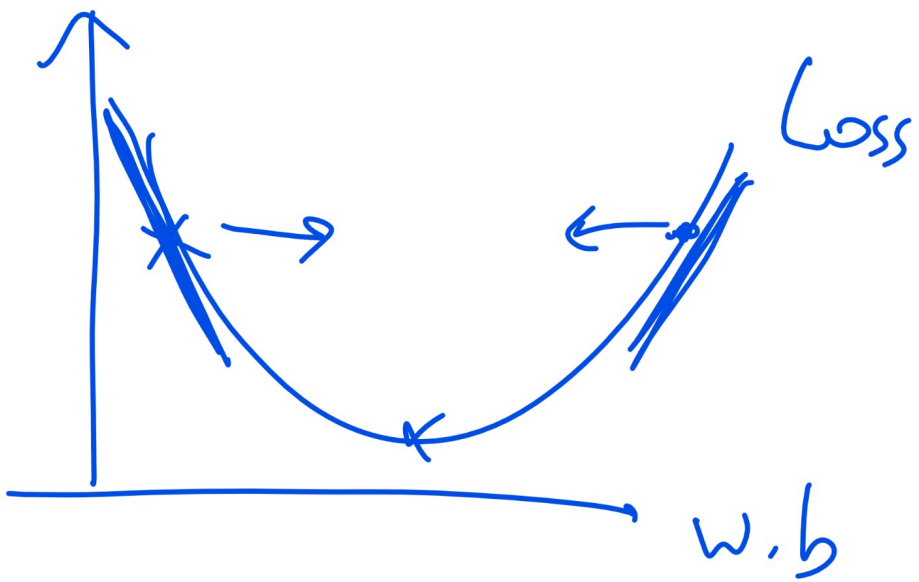

gradient
↓
derivative

Logistic Regression

Description

- Initialize random weights.
- Evaluate the activation function. → a_i
- Compute the gradient of the cost function with respect to each weight and bias.
- Update the weights and biases using gradient descent.
- Repeat until convergent.

Gradient Descent Step

Definition

- For logistic regression, use chain rule twice.

$$w = w - \alpha \sum_{i=1}^{n} (a_i - y_i) x_i$$

$$b = b - \alpha \sum_{i=1}^{n} (a_i - y_i)$$

$$a_i = g\left(w^T x_i + b\right), g\left(\boxed{\cdot}\right) = \frac{1}{1 + \exp\left(-\boxed{\cdot}\right)}$$

gradient descent

- α is the learning rate. It is the step size for each step of gradient descent.

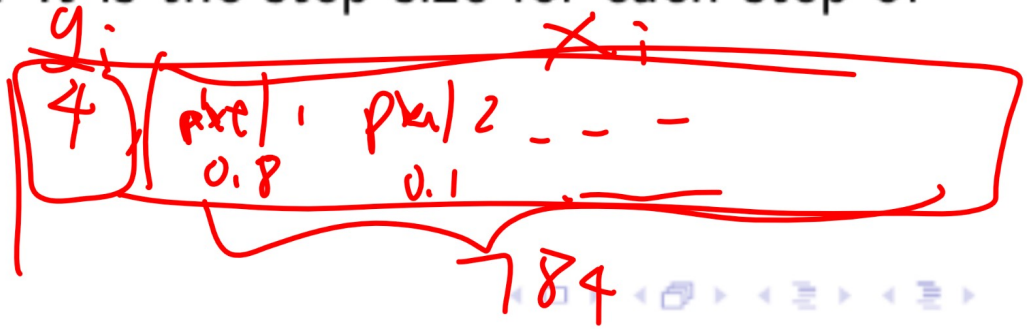

Perceptron Algorithm

Definition

- Update weights using the following rule.

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

Geometry

Learning Rate Diagram

Definition

Other Non-linear Activation Function

Discussion

- Activation function: $g\left(\boxed{\cdot}\right) = \tanh\left(\boxed{\cdot}\right) = \dfrac{e^{\boxed{\cdot}} - e^{-\boxed{\cdot}}}{e^{\boxed{\cdot}} + e^{-\boxed{\cdot}}}$
- Activation function: $g\left(\boxed{\cdot}\right) = \arctan\left(\boxed{\cdot}\right)$
- Activation function (rectified linear unit): $g\left(\boxed{\cdot}\right) = \boxed{\cdot}\mathbb{1}_{\left\{\boxed{\cdot}\geqslant 0\right\}}$
- All these functions lead to objective functions that are convex and differentiable (almost everywhere). Gradient descent can be used.

Gradient Descent

Quiz

- What is the gradient descent step for w if the objective (cost) function is the squared error?

cross entropy

$$C = \frac{1}{2}\sum_{i=1}^{n}(a_i - y_i)^2, a_i = g\left(w^T x_i + b\right), g(z) = \frac{1}{1+e^{-z}}$$

Chain Rule.

$$w_j = w_j - \alpha \frac{\partial C}{\partial w_j} = w - \alpha \sum_{i=1}^{n} \frac{\partial C}{\partial a_i} \cdot \frac{\partial a_i}{\partial w_j}$$

partial derivative: $\frac{1}{2} \cdot 2(a_i - y_i) \cdot (1 - a_i)\, a_i\, x_j$

Gradient Descent, Answer

Quiz

$$a_i = \frac{1}{1 + e^{-x_1 w_1 - x_2 w_2 \dots x_n w_n - b}}$$

$$\frac{\partial a_i}{\partial w_j} = + \frac{1}{(1 + e^{-(w^T x + b)})^2} \left(e^{-(w^T x + b)} \right) \cdot x_j$$

$$= \frac{e^{-(w^T x + b)} + 1 - 1}{(1 + e^{-(w^T x + b)})^2} x_j$$

$$= \left(1 - \frac{1}{1 + e^{-(w^T x + b)}} \right) \left(\frac{1}{1 + e^{-(w^T x + b)}} \right) x_j$$

$$= (1 - a_i)\, a_i\, x_i$$

Gradient Descent, Answer Too

Quiz

Gradient Descent

Quiz

- What is the gradient descent step for w if the objective (cost) function is the squared error?

$$C = \frac{1}{2}\sum_{i=1}^{n}\left(a_i - y_i\right)^2, a_i = g\left(w^T x_i + b\right), g'\left(z\right) = g\left(z\right)\cdot\left(1 - g\left(z\right)\right)$$

- $A : w = w - \alpha\sum\left(a_i - y_i\right)$
- $B : w = w - \alpha\sum\left(a_i - y_i\right)x_i$
- $C : w = w - \alpha\sum\left(a_i - y_i\right)a_i x_i$
- $D : w = w - \alpha\sum\left(a_i - y_i\right)\left(1 - a_i\right)x_i$
- $E : w = w - \alpha\sum\left(a_i - y_i\right)a_i\left(1 - a_i\right)x_i$

Gradient Descent, Another One, Answer

Quiz

Gradient Descent, Another One Too

Quiz

- What is the gradient descent step for w if the activation function is the identity function?

$$C = \frac{1}{2}\sum_{i=1}^{n}(a_i - y_i)^2, a_i = w^T x_i + b$$

- $A: w = w - \alpha \sum (a_i - y_i)$
- $B: w = w - \alpha \sum (a_i - y_i)\, x_i$
- $C: w = w - \alpha \sum (a_i - y_i)\, a_i x_i$
- $D: w = w - \alpha \sum (a_i - y_i)\,(1 - a_i)\, x_i$
- $E: w = w - \alpha \sum (a_i - y_i)\, a_i\,(1 - a_i)\, x_i$

Gradient Descent, Another One Too, Answer

Quiz

Convexity Diagram

Discussion

Questions?

Admin

10/12

- Missed lectures and quizzes. ✓
- Math used in the course. ✓
- Homework due dates. ✓
- Discussions and sharing solutions. ✓