

CS167: Machine Learning

Optimization for Weight Learning
Gradient Descent (GD)
Stochastic Gradient Descent (SGD)

Wednesday, April 1st, 2026





Announcements

- **Project#1**
 - due on Monday **04/06 by 11:59pm**

Today's Agenda

- Mathematical Modeling of Weight Parameters Learning
- Intuitive Understanding of Optimization (minimization/maximization)
 - Finding Moving Direction in Loss Surface (Gradient Calculation)
 - Gradient Descent
 - Stochastic Gradient Descent (SGD)

Modification of Perceptron




this step function outputs +1 if the function value is positive and -1 otherwise

It is not smooth and **not differentiable**

$$y = f(\mathbf{x}, \mathbf{w}) = o\left(\sum_{i=1, \dots, n} w_i x_i\right)$$

Make a Neuron with a Differentiable Function



this new function is referred to as an activation function. eg, sigmoid activation function


$$g\left(\sum_{i=1} w_i x_i\right) = \frac{1}{1 + \exp\left(-\sum_{i=1} w_i x_i\right)}$$

It is **smooth** and **differentiable**

$$y = f(\mathbf{x}, \mathbf{w}) = g\left(\sum_{i=1, \dots, n} w_i x_i\right)$$

Common Activation Functions

Each activation function is smooth and differentiable




Sigmoid (a.k.a. Logistic Function) Activation

Sigmoid:


$$g(u) = \frac{1}{1 + \exp^{-u}}$$

It is **smooth** and **differentiable**



Differentiation of sigmoid:

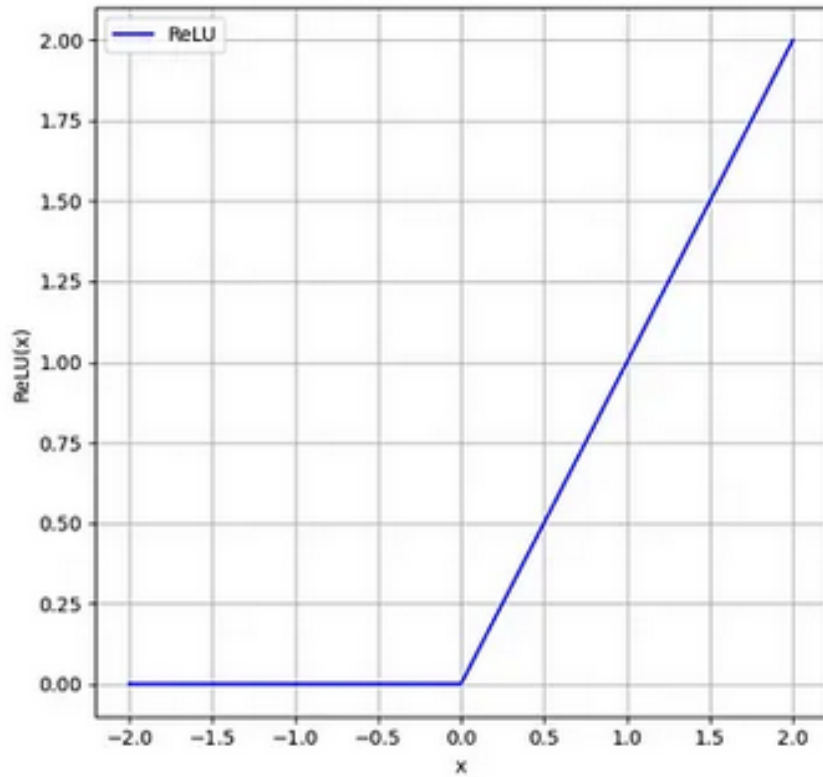
$$\frac{dg(u)}{du} = g'(u) = g(u)(1 - g(u))$$



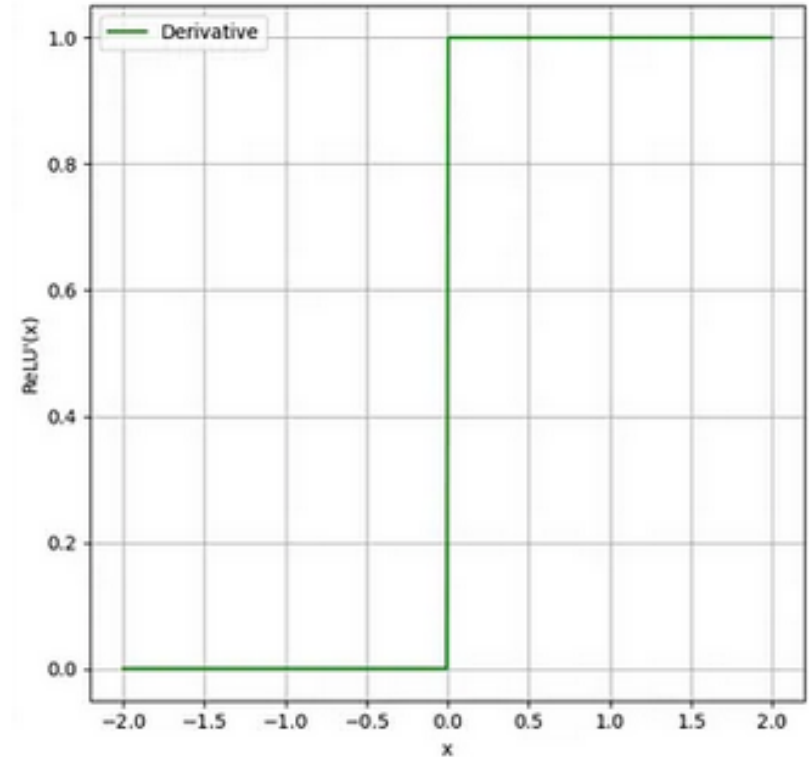
Sigmoid (a.k.a. Logistic Function) Activation

It is smooth and differentiable

ReLU Activation Function



Derivative of ReLU



Learning Weight Parameters with this Modified Neuron Model

- Instead of our simple **Perceptron Update Rule**, we can now use a better learning algorithm to learn the weight parameters ($w_0, w_1, w_2, \dots, w_n$)
- But what is this new weight parameter learning algorithm?

Learning Weight Parameters with this Modified Neuron Model

boston-housing dataset training features

| CRIM | ZN | INDUS | CHAS | NOX | RM | AGE | DIS | RAD | TAX | PTRATIO | LSTAT |
|---------|-----|-------|------|-------|-------|-------|--------|-----|-----|---------|-------|
| 0.10793 | 0.0 | 8.56 | 0 | 0.520 | 6.195 | 54.4 | 2.7778 | 5 | 384 | 20.9 | 13.00 |
| 1.34284 | 0.0 | 19.58 | 0 | 0.605 | 6.066 | 100.0 | 1.7573 | 5 | 403 | 14.7 | 6.43 |
| 4.81213 | 0.0 | 18.10 | 0 | 0.713 | 6.701 | 90.0 | 2.5975 | 24 | 666 | 20.2 | 16.42 |

$\mathbf{x}^i =$

$$\begin{bmatrix} x_0 \\ x_1 \\ x_2 \\ \vdots \\ \vdots \\ x_n \end{bmatrix}$$


training labels

MEDV
21.7
24.3
16.4


y^i

Learning Weight Parameters with this Modified Neuron Model

- Treat the problem as one of *minimization error* between a single training example's label and the network's output, given the example and weights as input



boston-housing dataset training split



| CRIM | ZN | INDUS | CHAS | NOX | RM | AGE | DIS | RAD | TAX | PTRATIO | LSTAT |
|---------|-----|-------|------|-------|-------|-------|--------|-----|-----|---------|-------|
| 0.10793 | 0.0 | 8.56 | 0 | 0.520 | 6.195 | 54.4 | 2.7778 | 5 | 384 | 20.9 | 13.00 |
| 1.34284 | 0.0 | 19.58 | 0 | 0.605 | 6.066 | 100.0 | 1.7573 | 5 | 403 | 14.7 | 6.43 |
| 4.81213 | 0.0 | 18.10 | 0 | 0.713 | 6.701 | 90.0 | 2.5975 | 24 | 666 | 20.2 | 16.42 |

Learning Weight Parameters with this Modified Neuron Model

- Treat the problem as one of *minimization error* between a single training example's label and the network's output, given the example and weights as input

$$\mathbf{x}^i = \begin{bmatrix} x_0 \\ x_1 \\ x_2 \\ \vdots \\ \vdots \\ x_n \end{bmatrix}$$

$$\mathbf{x}^1 = \begin{bmatrix} 1.0 \\ 0.11 \\ 0.0 \\ 8.56 \\ 0 \\ 0.52 \\ 6.20 \\ 54.4 \\ 2.78 \\ 5.0 \\ 384.0 \\ 20.9 \\ 13.0 \end{bmatrix}$$

$$\mathbf{x}^2 = \begin{bmatrix} 1.0 \\ 1.34 \\ 0.0 \\ 19.58 \\ 0 \\ 0.61 \\ 6.07 \\ 100.0 \\ 1.76 \\ 5.0 \\ 403.0 \\ 14.7 \\ 6.43 \end{bmatrix}$$

$$\mathbf{x}^3 = \begin{bmatrix} 1.0 \\ 4.81 \\ 0.0 \\ 18.1 \\ 0 \\ 0.71 \\ 6.70 \\ 90.0 \\ 2.60 \\ 24.0 \\ 666.0 \\ 20.2 \\ 16.42 \end{bmatrix}$$

take a training
example

boston-housing dataset training split

| CRIM | ZN | INDUS | CHAS | NOX | RM | AGE | DIS | RAD | TAX | PTRATIO | LSTAT |
|---------|-----|-------|------|-------|-------|-------|--------|-----|-----|---------|-------|
| 0.10793 | 0.0 | 8.56 | 0 | 0.520 | 6.195 | 54.4 | 2.7778 | 5 | 384 | 20.9 | 13.00 |
| 1.34284 | 0.0 | 19.58 | 0 | 0.605 | 6.066 | 100.0 | 1.7573 | 5 | 403 | 14.7 | 6.43 |
| 4.81213 | 0.0 | 18.10 | 0 | 0.713 | 6.701 | 90.0 | 2.5975 | 24 | 666 | 20.2 | 16.42 |

Learning Weight Parameters with this Modified Neuron Model

- Treat the problem as one of *minimization error* between a single training example's label and the network's output, given the example and weights as input

$$\mathbf{x}^i = \begin{bmatrix} x_0 \\ x_1 \\ x_2 \\ \vdots \\ \vdots \\ x_n \end{bmatrix}$$

$$\mathbf{x}^1 = \begin{bmatrix} 1.0 \\ 0.11 \\ 0.0 \\ 8.56 \\ 0 \\ 0.52 \\ 6.20 \\ 54.4 \\ 2.78 \\ 5.0 \\ 384.0 \\ 20.9 \\ 13.0 \end{bmatrix}$$

$$\mathbf{x}^2 = \begin{bmatrix} 1.0 \\ 1.34 \\ 0.0 \\ 19.58 \\ 0 \\ 0.61 \\ 6.07 \\ 100.0 \\ 1.76 \\ 5.0 \\ 403.0 \\ 14.7 \\ 6.43 \end{bmatrix}$$

$$\mathbf{x}^3 = \begin{bmatrix} 1.0 \\ 4.81 \\ 0.0 \\ 18.1 \\ 0 \\ 0.71 \\ 6.70 \\ 90.0 \\ 2.60 \\ 24.0 \\ 666.0 \\ 20.2 \\ 16.42 \end{bmatrix}$$

y^i

take a training ground truth labels

training labels

MEDV
21.7
24.3
16.4


$$y^1 = 21.7$$

$$y^2 = 24.3$$

$$y^3 = 16.4$$


Learning Weight Parameters with this Modified Neuron Model

- Treat the problem as one of *minimization error* between a single training example's label and the network's output, given the example and weights as input



Learning Weight Parameters with this Modified Neuron Model

- Treat the problem as one of *minimization error* between a single training example's label and the network's output, given the example and weights as input




Plug-in a training example

boston-housing dataset training split

| CRIM | ZN | INDUS | CHAS | NOX | RM | AGE | DIS | RAD | TAX | PTRATIO | LSTAT |
|---------|-----|-------|------|-------|-------|-------|--------|-----|-----|---------|-------|
| 0.10793 | 0.0 | 8.56 | 0 | 0.520 | 6.195 | 54.4 | 2.7778 | 5 | 384 | 20.9 | 13.00 |
| 1.34284 | 0.0 | 19.58 | 0 | 0.605 | 6.066 | 100.0 | 1.7573 | 5 | 403 | 14.7 | 6.43 |
| 4.81213 | 0.0 | 18.10 | 0 | 0.713 | 6.701 | 90.0 | 2.5975 | 24 | 666 | 20.2 | 16.42 |


Learning Weight Parameters with this Modified Neuron Model

- Treat the problem as one of *minimization error* between a single training example's label and the network's output, given the example and weights as input




Learning Weight Parameters with this Modified Neuron Model

- Treat the problem as one of *minimization error* between a single training example's label and the network's output, given the example and weights as input



Learning Weight Parameters with this Modified Neuron Model


- Treat the problem as one of *minimization error* between a single training example's label and the network's output, given the example and weights as input



MEDV
21.7
24.3
16.4

Learning Weight Parameters with this Modified Neuron Model

- Treat the problem as one of *minimization error* between a single training example's label and the network's output, given the example and weights as input




For example, here, we used **Mean Squared Error (MSE)** to measure the discrepancy. We could use any other measure to find the discrepancy between prediction and ground-truth

$$(y^i - f(\mathbf{x}^i, \mathbf{w}))^2$$

Learning Weight Parameters with this Modified Neuron Model

- Treat the problem as one of *minimization error* between a single training example's label and the network's output, given the example and weights as input



$$Error(\mathbf{x}^i, y^i, \mathbf{w}) = (y^i - f(\mathbf{x}^i, \mathbf{w}))^2$$

Learning Weight Parameters with this Modified Neuron Model

- Treat the problem as one of *minimization error* between a single training example's label and the network's output, given the example and weights as input

$$Error(\mathbf{x}^i, y^i, \mathbf{w}) = (y^i - f(\mathbf{x}^i, \mathbf{w}))^2$$

- If we consider a collection of training examples and sum the above error over all examples

$$E(\mathbf{w}) = \sum_i Error(\mathbf{x}^i, y^i, \mathbf{w}) = \sum_i (y^i - f(\mathbf{x}^i, \mathbf{w}))^2$$

Learning Weight Parameters with this Modified Neuron Model

- If we consider a collection of training examples and Sum the above error term over all examples

$$E(\mathbf{w}) = \sum_i Error(\mathbf{x}_i, y_i, \mathbf{w}) = \sum_i (y_i - f(\mathbf{x}_i, \mathbf{w}))^2$$

- Minimize errors using an optimization algorithm:
 - Gradient Descent (GD)
 - Stochastic Gradient Descent (SGD)

$E(\mathbf{w})$ term is also known as *loss function*

Today's Agenda


- Mathematical Modeling of Weight Parameters Learning
- Intuitive Understanding of Optimization (minimization/maximization)
 - Finding Moving Direction in Loss Surface (Gradient Calculation)
 - Gradient Descent
 - Stochastic Gradient Descent (SGD)

Optimization

- **minimization**: trying to find the subset of values for attributes that gives you the minimum value in the objective function
- The term objective function is generalized term which leaves room for the function to be something that we want to either **minimize** or **maximize**. The other terms used for the minimizing setting are as follows:
 - loss function
 - error function
 - cost function

Optimization Intuition

- **minimization**: trying to find the subset of values for attributes that gives you the minimum value in the objective function




Optimization Intuition

- **minimization**: trying to find the subset of values for attributes that gives you the minimum value in the objective function

A smooth 3D surface (each point correspond to a loss value)


You are viewing a 3D surface from a third-person perspective.




Play with the code I uploaded on Blackboard to generates different surfaces

Optimization Intuition

- **minimization**: trying to find the subset of values for attributes that gives you the minimum value in the objective function
- **How to reach to the minimum?**
 - we can start at an arbitrary point on the surface and gradually explore the surface until we reach the minimum value






This visualization is called a contour map. It appears as if you are viewing a 3D surface from a bird's-eye perspective.

iteratively move in direction $-\nabla E$

$$E(\mathbf{w}) = \sum_i \text{Error}(\mathbf{x}_i, y_i, \mathbf{w}) = \sum_i (y_i - f(\mathbf{x}_i, \mathbf{w}))^2$$

Adapted from K. Hauser's slide




This visualization is called a contour map. It appears as if you are viewing a 3D surface from a bird's-eye perspective.

iteratively move in direction $-\nabla E$

$$E(\mathbf{w}) = \sum_i \text{Error}(\mathbf{x}_i, y_i, \mathbf{w}) = \sum_i (y_i - f(\mathbf{x}_i, \mathbf{w}))^2$$

Adapted from K. Hauser's slide




This visualization is called a contour map. It appears as if you are viewing a 3D surface from a bird's-eye perspective.

iteratively move in direction $-\nabla E$

$$E(\mathbf{w}) = \sum_i \text{Error}(\mathbf{x}_i, y_i, \mathbf{w}) = \sum_i (y_i - f(\mathbf{x}_i, \mathbf{w}))^2$$

Adapted from K. Hauser's slide



This visualization is called a contour map. It appears as if you are viewing a 3D surface from a bird's-eye perspective.

iteratively move in direction $-\nabla E$

$$E(\mathbf{w}) = \sum_i \text{Error}(\mathbf{x}_i, y_i, \mathbf{w}) = \sum_i (y_i - f(\mathbf{x}_i, \mathbf{w}))^2$$

Adapted from K. Hauser's slide




This visualization is called a contour map. It appears as if you are viewing a 3D surface from a bird's-eye perspective.

iteratively move in direction $-\nabla E$

$$E(\mathbf{w}) = \sum_i \text{Error}(\mathbf{x}_i, y_i, \mathbf{w}) = \sum_i (y_i - f(\mathbf{x}_i, \mathbf{w}))^2$$

Adapted from K. Hauser's slide




This visualization is called a contour map. It appears as if you are viewing a 3D surface from a bird's-eye perspective.

iteratively move in direction $-\nabla E$

$$E(\mathbf{w}) = \sum_i \text{Error}(\mathbf{x}_i, y_i, \mathbf{w}) = \sum_i (y_i - f(\mathbf{x}_i, \mathbf{w}))^2$$

Adapted from K. Hauser's slide




This visualization is called a contour map. It appears as if you are viewing a 3D surface from a bird's-eye perspective.

iteratively move in direction $-\nabla E$

$$E(\mathbf{w}) = \sum_i \text{Error}(\mathbf{x}_i, y_i, \mathbf{w}) = \sum_i (y_i - f(\mathbf{x}_i, \mathbf{w}))^2$$

Adapted from K. Hauser's slide




This visualization is called a contour map. It appears as if you are viewing a 3D surface from a bird's-eye perspective.

iteratively move in direction $-\nabla E$

$$E(\mathbf{w}) = \sum_i \text{Error}(\mathbf{x}_i, y_i, \mathbf{w}) = \sum_i (y_i - f(\mathbf{x}_i, \mathbf{w}))^2$$

Adapted from K. Hauser's slide



This visualization is called a contour map. It appears as if you are viewing a 3D surface from a bird's-eye perspective.

iteratively move in direction $-\nabla E$

$$E(\mathbf{w}) = \sum_i \text{Error}(\mathbf{x}_i, y_i, \mathbf{w}) = \sum_i (y_i - f(\mathbf{x}_i, \mathbf{w}))^2$$

Adapted from K. Hauser's slide

Today's Agenda


- Mathematical Modeling of Weight Parameters Learning
- Intuitive Understanding of Optimization (minimization/maximization)
 - Finding Moving Direction in Loss Surface (Gradient Calculation)
 - Gradient Descent (GD)
 - **Stochastic Gradient Descent (SGD)**

Stochastic Gradient Descent (SGD)

- Keep doing the **Gradient Descent**, but instead of using all the training samples, *use small subset of training samples* picked randomly when computing the **gradient vector**
 - divide the entire training data into **mini batches**
 - calculate the **gradient vector** based on that batch $\nabla E(\mathbf{w})$
 - adjust (or update) the values of the weights based on the **gradient vector** to that batch

$$\mathbf{w}^{\text{new}} = \mathbf{w}^{\text{old}} - \eta \nabla E(\mathbf{w})$$

Stochastic Gradient Descent (SGD)



Useful Online Resources for Gradient Descent

- Another [mathematical derivation for Gradient Descent](#)
- One more [mathematical derivation for Gradient Descent](#)
- [Google course's Gradient Descent](#)
- [Visualization of Gradient Descent](#)
- [Visual explanation of Gradient Descent and other optimizers](#)