

CS2109S Tutorial 8

Neural Networks

(AY 24/25 Semester 2)

April 4, 2025

(Prepared by Benson)

Contents

Perceptron

- Q1. Logic Gates
- Q2. Single vs Multi Layer Perceptron

Forward Propagation

- Q3. Forward Propagation
- Q4. Let's Activate!
- Q5. Working with Dimensions

Bonus. Perceptron Learning Algorithm

Admin Info

- ▶ For those who have reached Level 30, PS5 (Neural Networks) is *optional*. But I would still encourage you to attempt at least part of it – Pytorch is one of the most important takeaways of this course.
- ▶ Bonus submission deadlines (for EXP / bubble tea):
 - ▶ Behind Pytorch Autograd: 11 Apr (Week 12 Fri) 4pm
 - ▶ Messing with a GPT Model: 25 Apr (Reading Week Fri) 4pm

Feel free to share me your progress and I might drop some hints!

Q1. Logic Gates

- (a) Dry run the Perceptron learning algorithm with $\eta = 0.1$, all weights initialized to 0 and activation function $\text{sign}(z) = 1$ if $z \geq 0$ else 0.

AND: Hypothesis $\hat{y} = \text{sign}(b + w_1x_1 + w_2x_2)$.

Iteration	b	w_1	w_2
Initial	0	0	0
1	-0.1	0	0
2	0	1	0
3	0	1	0
4	0	1	0
5	0	1	0
6	0	1	0
7	0	1	0
8	0	1	0
9	0	1	0
10	0	1	0
11	-0.5	0.2	0.1

$b \leftarrow b + \eta(y - \hat{y})(1)$
 $b \leftarrow b + \eta(y - \hat{y})(1)$
 $b \leftarrow b + \eta(y - \hat{y})(1)$
 $b \leftarrow b + \eta(y - \hat{y})(1)$
 $b \leftarrow b + \eta(y - \hat{y})(1)$
 $b \leftarrow b + \eta(y - \hat{y})(1)$
 $b \leftarrow b + \eta(y - \hat{y})(1)$
 $b \leftarrow b + \eta(y - \hat{y})(1)$
 $b \leftarrow b + \eta(y - \hat{y})(1)$
 $b \leftarrow b + \eta(y - \hat{y})(1)$
 $0.1(0 - 1)(1) = -0.1$
 $b \leftarrow b + \eta(y - \hat{y})(1)$
 $0.1(1 - 0)(1) = 0.1$

$w_1 \leftarrow w_1 + \eta(y - \hat{y})x_1$
 $w_1 \leftarrow w_1 + \eta(y - \hat{y})x_1$
 $w_1 \leftarrow w_1 + \eta(y - \hat{y})x_1$
 $w_1 \leftarrow w_1 + \eta(y - \hat{y})x_1$
 $w_1 \leftarrow w_1 + \eta(y - \hat{y})x_1$
 $w_1 \leftarrow w_1 + \eta(y - \hat{y})x_1$
 $w_1 \leftarrow w_1 + \eta(y - \hat{y})x_1$
 $w_1 \leftarrow w_1 + \eta(y - \hat{y})x_1$
 $w_1 \leftarrow w_1 + \eta(y - \hat{y})x_1$
 $w_1 \leftarrow w_1 + \eta(y - \hat{y})x_1$
 $0.1(0 - 1)(0) = 0$
 $w_1 \leftarrow w_1 + \eta(y - \hat{y})x_1$
 $0.1(1 - 0)(1) = 0.1$

$w_2 \leftarrow w_2 + \eta(y - \hat{y})x_2 + \eta(y - \hat{y})x_2$
 $w_2 \leftarrow w_2 + \eta(y - \hat{y})x_2 \cdot (0) = 0$
 $w_2 \leftarrow w_2 + \eta(y - \hat{y})x_2$
 $w_2 \leftarrow w_2 + \eta(y - \hat{y})x_2$
 $w_2 \leftarrow w_2 + \eta(y - \hat{y})x_2$
 $w_2 \leftarrow w_2 + \eta(y - \hat{y})x_2$
 $w_2 \leftarrow w_2 + \eta(y - \hat{y})x_2$
 $w_2 \leftarrow w_2 + \eta(y - \hat{y})x_2$
 $w_2 \leftarrow w_2 + \eta(y - \hat{y})x_2$
 $w_2 \leftarrow w_2 + \eta(y - \hat{y})x_2$
 $0.1(0 - 1)(1) = -0.1$

AND			
x_1	x_2	y	\hat{y}
0	0	0	1010000000
0	1	0	1011010000
1	0	0	1011010000
1	1	1	1011010000

Misclassified instance
 Misclassified instance
 Misclassified instance
 Misclassified instance

Q1. Logic Gates

- (a) Dry run the Perceptron learning algorithm with $\eta = 0.1$, all weights initialized to 0 and activation function $\text{sign}(z) = 1$ if $z \geq 0$ else 0.

OR: Hypothesis $\hat{y} = \text{sign}(b + w_1x_1 + w_2x_2)$.

Iteration	b	w_1	w_2
Initial	0	0	0
1	-0.1	0	0
2	0	0	0.1
3	0	0	0.1
4	0	0	0.1
5	0	0	0.1
6	0	0	0.1
7	0	0	0.1
8	-0.2	0.2	0.1
9	-0.3	0.1	0.1
10	-0.2	0.2	0.2
11	-0.3	0.2	0.1

$$w_2 \leftarrow w_2 + \eta(y - \hat{y})x_2 + \eta(y - \hat{y})x_2$$

$$w_2 \leftarrow w_2 + \eta(y - \hat{y})x_2 \quad \text{.)(0) = 0}$$

$$0.1(0 - 1)(0) = 0$$

$$b \leftarrow b + \eta(y - \hat{y})(1)$$

$$b \leftarrow b + \eta(y - \hat{y})(1)$$

$$b \leftarrow b + \eta(y - \hat{y})(1)$$

$$b \leftarrow b + \eta(y - \hat{y})(1)$$

$$b \leftarrow b + \eta(y - \hat{y})(1)$$

$$0.1(0 - 1)(1) = -0.1$$

$$w_1 \leftarrow w_1 - \eta(y - \hat{y})x_1$$

$$w_1 \leftarrow w_1 - \eta(y - \hat{y})x_1$$

$$w_1 \leftarrow w_1 - \eta(y - \hat{y})x_1$$

$$w_1 \leftarrow w_1 - \eta(y - \hat{y})x_1$$

$$w_1 \leftarrow w_1 + \eta(y - \hat{y})x_1$$

$$0.1(0 - 1)(0) = 0$$

OR			
x_1	x_2	y	\hat{y}
0	0	0	101010
0	1	1	101111
1	0	1	101111
1	1	1	101111

Misclassified instance
Misclassified instance
Misclassified instance
Misclassified instance

Q1. Logic Gates

- (a) Dry run the Perceptron learning algorithm with $\eta = 0.1$, all weights initialized to 0 and activation function $\text{sign}(z) = 1$ if $z \geq 0$ else 0.

NOT: Hypothesis $\hat{y} = \text{sign}(b + w_1x)$.

Iteration	b	w_1	
Initial	0	0	
1	-0.1	-0.1	
2	0	-0.1	
3	-0.1	0	
4	0	0	
5	-0.1	0	
6	-0.2	0.2	0
7	-0.3	0.1	0
8	-0.2	0.2	0.1
9	-0.3	0.1	0.1
10	-0.2	0.2	0.2
11	-0.3	0.2	0.1


$b \leftarrow b + \eta(y - \hat{y})(1)$
 $b \leftarrow b + \eta(y - \hat{y})(1)$
 $0.1(1 - 0)(1) = 0.1$

$w_1 \leftarrow w_1 + \eta(y - \hat{y})x_1$
 $w_1 \leftarrow w_1 + \eta(y - \hat{y})x_1$
 $0.1(1 - 0)(0) = 0$

NOT			
x	y	\hat{y}	
0	1	1	$\text{sign}(0)0.1$
1	0	0	$\text{sign}(0)0.2$
			Misclassified instance
			Misclassified instance


Q1. Logic Gates

AND:



$$\hat{y} = \text{sign}(-0.3 + 0.2x_1 + 0.1x_2)$$

OR:




$$\hat{y} = \text{sign}(-0.1 + 0.1x_1 + 0.1x_2)$$

Q1. Logic Gates

- (b) Is it possible to model the XOR function using a single Perceptron? Refer to Figure 2 for the truth table of the XOR gate. Comment on your answer.

Hypothesis $\hat{y} = \text{sign}(b + w_1x_1 + w_2x_2)$.




XOR		
x_1	x_2	y
0	0	0
0	1	1
1	0	1
1	1	0

- ▶ The points are not **linearly separable**.
- ▶ It is impossible to model the XOR function using a single Perceptron.

Q1. Logic Gates

- (c) Model XOR function (takes 2 inputs) using a number of perceptrons that implement AND, OR, and NOT functions. Show the diagram of the final Perceptron network.

$$\text{XOR}(x_1, x_2) = \text{AND}(\text{OR}(x_1, x_2), \text{NOT}(\text{AND}(x_1, x_2)))$$



Q1. Logic Gates

- (d) If we change the ordering of data samples in Perceptron Update Rule, will the model converges to a different model weight for the AND operator? What can you conclude from the observation?

Ordering: (0, 1, 2, 3)

Iteration	b	w_1	w_2
Initial	0	0	0
1	-0.1	0	0
2	0	0.1	0.1
3	-0.1	0.1	0.1
4	-0.2	0.1	0
5	-0.1	0.2	0.1
6	-0.2	0.2	0
7	-0.3	0.1	0
8	-0.2	0.2	0.1
9	-0.3	0.1	0.1
10	-0.2	0.2	0.2
11	-0.3	0.2	0.1

Ordering: (0, 2, 3, 1)

Iteration	b	w_1	w_2
Initial	0	0	0
1	-0.1	0	0
2	0	0.1	0.1
3	-0.1	0.1	0
4	-0.2	0	0
5	-0.1	0.1	0.1
6	-0.2	0.1	0
7	-0.1	0.2	0.1
8	-0.2	0.2	0
9	-0.3	0.1	0
10	-0.2	0.2	0.1
11	-0.3	0.1	0.1
12	-0.2	0.2	0.2
13	-0.3	0.2	0.1


Ordering: (0, 2, 1, 3)

Iteration	b	w_1	w_2
Initial	0	0	0
1	-0.1	0	0
2	0	0.1	0.1
3	-0.1	0.1	0.1
4	-0.2	0	0.1
5	-0.1	0.1	0.2
6	-0.2	0	0.2
7	-0.3	0	0.1
8	-0.2	0.1	0.2
9	-0.3	0.1	0.1
10	-0.2	0.2	0.2
11	-0.3	0.1	0.2

Q1. Logic Gates

- (d) If we change the ordering of data samples in Perceptron Update Rule, will the model converges to a different model weight for the AND operator? What can you conclude from the observation?
- ▶ Reordering data points could help the model converge much faster.
 - ▶ Manipulating the ordering could direct the model to a different weight. There is no guarantee to converge to the same model even for such a simple gate function.

Q1. Logic Gates

- (e)  With regards to Figure 1, does your proposed model have high bias and high variance?
(Recap: *What is bias? What is variance?*)
- ▶ **Low bias:** It classifies all data points correctly.
 - ▶ **Low variance:** It is a simple linear model, the simplest model to perform the logic gates given.

Q1. Logic Gates

📌 Does the Perceptron learning algorithm *always* converge on linearly separable datasets? Make your guess!


- A. Yes, always!
- B. Yes, if the learning rate is small enough.
- C. Yes, if the parameters are ordered properly.
- D. No, it can be stuck in a loop regardless.

Q2. Single vs Multi Layer Perceptron

After training both networks, you obtain a mean squared error of 1000 on the training set and a mean squared error of 2000 on the validation set for the single-layer perceptron, and a mean squared error of 800 on the training set and a mean squared error of 1200 on the validation set for the multi-layer perceptron.

- (a) What might be the reasons for the difference in performance between the single-layer perceptron and the multi-layer perceptron?
 - ▶ More layers \Rightarrow able to learn more complex (non-linear) patterns.
 - ▶ The single-layer perceptron is limited to a linear classifier.
- (b) How might you modify the single-layer perceptron to improve its performance, and what are the advantages and disadvantages of doing so?
 - ▶ Add transformed features (polynomial $(x_1)^2$, interaction x_1x_2).
- (c) What techniques could you use to improve the performance of the multilayer perceptron?
 - ▶ Underfitting: Increase the number of hidden layers (parameters).
 - ▶ Overfitting: Regularization.

Recap: Forward Propagation




$$\mathbf{a}^{[1]} = \begin{bmatrix} a_1^{[1]} \\ a_2^{[1]} \end{bmatrix} = \begin{bmatrix} W_{01}^{[1]} & W_{11}^{[1]} & W_{21}^{[1]} \\ W_{02}^{[1]} & W_{12}^{[1]} & W_{22}^{[1]} \end{bmatrix} \times \begin{bmatrix} 1 \\ x_1^{[0]} \\ x_2^{[0]} \end{bmatrix} = \mathbf{W}^\top \times \mathbf{x}^{[0]}$$

Q3. Forward Propagation

Suppose there is a data input $\mathbf{x} = (2, 3)^\top$ and the actual output label is $\mathbf{y} = (0.1, 0.9)^\top$. The weights for the network are

$$\mathbf{W}^{[1]} = \begin{bmatrix} 0.1 & 0.1 \\ -0.1 & 0.2 \\ 0.3 & -0.4 \end{bmatrix}, \mathbf{W}^{[2]} = \begin{bmatrix} 0.1 & 0.1 \\ 0.5 & -0.6 \\ 0.7 & -0.8 \end{bmatrix}$$

Calculate $\mathbf{a}^{[1]}$, $\hat{\mathbf{y}}^{[2]}$ and $L(\hat{\mathbf{y}}^{[2]}, \mathbf{y})$.



Activation Func: $\text{ReLU}(x) = \max(0, x)$.

$$\mathbf{a}^{[1]} = \text{ReLU}((\mathbf{W}^{[1]})^\top \mathbf{X}) = \text{ReLU} \left(\begin{bmatrix} 0.1 & -0.1 & 0.3 \\ 0.1 & 0.2 & -0.4 \end{bmatrix} \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \right) = \text{ReLU} \left(\begin{bmatrix} 0.8 \\ -0.7 \end{bmatrix} \right) = \begin{bmatrix} 0.8 \\ 0 \end{bmatrix}$$

$$\hat{\mathbf{y}}^{[2]} = \text{ReLU}((\mathbf{W}^{[2]})^\top \mathbf{a}^{[1]}) = \text{ReLU} \left(\begin{bmatrix} 0.1 & 0.5 & 0.7 \\ 0.1 & -0.6 & -0.8 \end{bmatrix} \begin{bmatrix} 1 \\ 0.8 \\ 0 \end{bmatrix} \right) = \text{ReLU} \left(\begin{bmatrix} 0.5 \\ -0.38 \end{bmatrix} \right) = \begin{bmatrix} 0.5 \\ 0 \end{bmatrix}$$

$$L(\hat{\mathbf{y}}^{[2]}, \mathbf{y}) = \frac{1}{2} \left((\hat{y}_1^{[2]} - y_1)^2 + (\hat{y}_2^{[2]} - y_2)^2 \right) = \frac{1}{2} \left((0.5 - 0.1)^2 + (0 - 0.9)^2 \right) = 0.485$$

Q4. Let's Activate!

We can define a neural network as follows:

$$\hat{y} = g((\mathbf{W}^{[L]})^\top \dots g(\mathbf{W}^{[2]})^\top \cdot g(\mathbf{W}^{[1]})^\top \mathbf{x})$$

where $\mathbf{W}^{[l] \in \{1, \dots, L\}}$ is a weight matrix. You're given the following weight matrices:

$$\mathbf{W}^{[3]} = \begin{bmatrix} 1.2 & -2.2 \\ 1.2 & 1.3 \end{bmatrix}, \mathbf{W}^{[2]} = \begin{bmatrix} 2.1 & -0.5 \\ 0.7 & 1.9 \end{bmatrix}, \mathbf{W}^{[1]} = \begin{bmatrix} 1.4 & 0.6 \\ 0.8 & 0.6 \end{bmatrix}$$

You are given $g(z) = \text{SiLU}(z) = \frac{z}{1 + e^{-z}}$ between all layers except the last layer.

Is it possible to replace the whole neural network with just one matrix in both cases **with** and **without** non-linear activations $g(z)$?

Q4. Let's Activate!

Without non-linear activations:

$$\begin{aligned}\hat{y} &= (\mathbf{W}^{[3]})^\top (\mathbf{W}^{[2]})^\top (\mathbf{W}^{[1]})^\top \mathbf{x} \\ &= \begin{bmatrix} 1.2 & 1.2 \\ -2.2 & 1.3 \end{bmatrix} \begin{bmatrix} 2.1 & 0.7 \\ -0.5 & 1.9 \end{bmatrix} \begin{bmatrix} 1.4 & 0.8 \\ 0.6 & 0.6 \end{bmatrix} \mathbf{x} \\ &= \left(\begin{bmatrix} 1.2 & 1.2 \\ -2.2 & 1.3 \end{bmatrix} \begin{bmatrix} 2.1 & 0.7 \\ -0.5 & 1.9 \end{bmatrix} \begin{bmatrix} 1.4 & 0.8 \\ 0.6 & 0.6 \end{bmatrix} \right) \mathbf{x} \\ &= \begin{bmatrix} 4.56 & 3.408 \\ -6.82 & -3.658 \end{bmatrix} \mathbf{x} \\ &= \mathbf{M}^\top \mathbf{x}\end{aligned}$$

\therefore It is possible to replace the whole neural network with $M = \begin{bmatrix} 4.56 & -6.82 \\ 3.408 & -3.658 \end{bmatrix}$.

Q4. Let's Activate!

With non-linear activations: Suppose $\hat{y} = \mathbf{M}^\top \mathbf{x}$. We expect \hat{y} to be doubled if we double \mathbf{x} .

► When $\mathbf{x} = \begin{bmatrix} 1 & 0 \end{bmatrix}^\top$:

$$\begin{aligned}\hat{y} &= g\left((\mathbf{W}^{[3]})^\top g\left((\mathbf{W}^{[2]})^\top g\left((\mathbf{W}^{[1]})^\top \mathbf{x}\right)\right)\right) \\ &= \begin{bmatrix} 3.0571 \\ -5.2727 \end{bmatrix}\end{aligned}$$

Intuition: Prove that it is impossible since the network is no longer “linear”.

► When $\mathbf{x} = \begin{bmatrix} 2 & 0 \end{bmatrix}^\top$:

$$\begin{aligned}\hat{y} &= g\left((\mathbf{W}^{[3]})^\top g\left((\mathbf{W}^{[2]})^\top g\left((\mathbf{W}^{[1]})^\top \mathbf{x}\right)\right)\right) \\ &= \begin{bmatrix} 7.7257 \\ -13.2458 \end{bmatrix} \neq 2 \cdot \begin{bmatrix} 3.0571 \\ -5.2727 \end{bmatrix}\end{aligned}$$

The outputs are not linear \Rightarrow a contradiction! \therefore No such \mathbf{M} exists.

Q4. Let's Activate!

Conclusion:

- ▶ Without non-linear activations, the entire network **collapses to a simple linear model**.


$$\begin{aligned}\hat{y} &= (\mathbf{W}^{[L]})^\top \dots (\mathbf{W}^{[2]})^\top (\mathbf{W}^{[1]})^\top \mathbf{x} \\ &= \left((\mathbf{W}^{[L]})^\top \dots (\mathbf{W}^{[2]})^\top (\mathbf{W}^{[1]})^\top \right) \mathbf{x}\end{aligned}$$

- ▶ Non-linear activation functions let the network model non-linear relationships in the data. Increasing the depth of the network will help the model learn more complex relationships.

Q5. Working with Dimensions

You're building a self-driving car program that takes in **grayscale** images of size 32×32 where 32 is the image height and width. There are 4 classes your simplified program has to classify: {car, person, traffic light, stop sign}. You start off experimenting with a Multi-layer Perceptron composed of three linear layers of the form $\mathbf{y} = \mathbf{W}^\top \mathbf{x}$, where $\mathbf{x} \in R^d$ is the input vector, \mathbf{W} is the weight matrix, and \mathbf{y} is the network output.

Layer	Input dim	Weight Matrix dim	Output dim
Linear layer 1	<u>1024</u> \times 1	<u>1024</u> \times <u>512</u>	512 \times 1
Linear layer 2	512 \times 1	<u>512</u> \times 128	<u>128</u> \times 1
Linear layer 3	128 \times 1	<u>128</u> \times 4	<u>4</u> \times 1



Bonus. Perceptron Learning Algorithm

Prove that the Perceptron Learning Algorithm always converges on a **linearly separable dataset**! You may assume the initial weights are 0 and $y^{(i)} \in \{-1, 1\}$.

(Hint: Let \mathbf{w}^ be a **unit** weight vector that linearly separates the data, we have $y^{(i)}(\mathbf{w}^* \cdot \mathbf{x}^{(i)}) \geq \gamma$ since the margin is positive. Also, denote $R = \max_i \|\mathbf{x}^{(i)}\|$. Prove that after k iterations, (1) $\mathbf{w} \cdot \mathbf{w}^* \geq k\eta\gamma$ and (2) $\|\mathbf{w}\|^2 \leq k\eta^2 R^2$. Finally, show that $k \leq \frac{R^2}{\gamma^2}$, i.e. the algorithm never takes more than this number of iterations.)*

Solution:

- Consider the k -th iteration with initial weights \mathbf{w}_0 and misclassified instance $(\mathbf{x}^{(i)}, y^{(i)})$. The update rule gives $\mathbf{w} = \mathbf{w}_0 + \eta y^{(i)} \mathbf{x}^{(i)}$. Then

$$\begin{aligned}\mathbf{w} \cdot \mathbf{w}^* &= (\mathbf{w}_0 + \eta y^{(i)} \mathbf{x}^{(i)}) \cdot \mathbf{w}^* \\ &= \mathbf{w}_0 \cdot \mathbf{w}^* + \eta y^{(i)} (\mathbf{x}^{(i)} \cdot \mathbf{w}^*) \\ &\geq \mathbf{w}_0 \cdot \mathbf{w}^* + \eta \gamma\end{aligned}$$

This implies $\mathbf{w} \cdot \mathbf{w}^*$ increases by at least $\eta \gamma$ in each iteration. After k iterations, $\mathbf{w} \cdot \mathbf{w}^* \geq k \eta \gamma$.

- Since \mathbf{w}^* is a unit vector, we have $\mathbf{w} \cdot \mathbf{w}^* = \|\mathbf{w}\| \|\mathbf{w}^*\| \cos \theta = \|\mathbf{w}\| \cos \theta \leq \|\mathbf{w}\|$. Hence $\|\mathbf{w}\|^2 \geq (\mathbf{w} \cdot \mathbf{w}^*)^2 \geq k^2 \eta^2 \gamma^2$.

Bonus. Perceptron Learning Algorithm

Solution:

- ▶ Since point $(\mathbf{x}^{(i)}, y^{(i)})$ is misclassified, we also have $y^{(i)}(\mathbf{w}_0 \cdot \mathbf{x}^{(i)}) \leq 0$. Then

$$\begin{aligned}\|\mathbf{w}\|^2 &= \|\mathbf{w}_0 + \eta y^{(i)} \mathbf{x}^{(i)}\|^2 \\&= \|\mathbf{w}_0\|^2 + \eta^2 (y^{(i)})^2 \|\mathbf{x}^{(i)}\|^2 + 2\eta y^{(i)} (\mathbf{w}_0 \cdot \mathbf{x}^{(i)}) \\&= \|\mathbf{w}_0\|^2 + \eta^2 \|\mathbf{x}^{(i)}\|^2 + 2\eta y^{(i)} (\mathbf{w}_0 \cdot \mathbf{x}^{(i)}) \\&\leq \|\mathbf{w}_0\|^2 + \eta^2 \|\mathbf{x}^{(i)}\|^2 \\&\leq \|\mathbf{w}_0\|^2 + \eta^2 R^2\end{aligned}$$

This implies $\|\mathbf{w}\|^2$ increases by at most $\eta^2 R^2$ in each iteration. After k iterations, $\|\mathbf{w}\|^2 \leq k\eta^2 R^2$.

- ▶ Combining both inequalities, we have $k^2 \eta^2 \gamma^2 \leq k\eta^2 R^2 \Rightarrow k \leq \frac{R^2}{\gamma^2}$.