


# Deep Learning: Neural Networks


Delivered by Timilehin Owolabi




# How do you fit a line to this data?




# Linear Regression?



# Exponential Regression?



# What of this?




# Logistic Regression


$$y = mx + b.$$

$$f(x) = \frac{1}{1 + e^{-x}}$$


**Linear Regression**



**Logistic Regression**




Output  $\hat{y} = \sigma(\underbrace{w^T x + b}_z)$



The shape of the line depends on the parameters  $w$  and  $b$

# How can we determine who passes?



It's now a  
classification problem



# **Cost Function: How do we determine the correctness of our line?**

$$\text{error} = |y - y_{\text{hat}}|$$

$$\text{error} = (y - y_{\text{hat}})^2$$

$$\text{error} = - (y \log(y_{\text{hat}}) + (1 - y) \log(1 - y_{\text{hat}}))$$

## **Total error = sum(error)/number of points**

## **Cost = Total Error**

# Gradient Descent

$$\text{Cost} = - (y \log(y_{\text{hat}}) + (1 - y) \log (1 - y_{\text{hat}}))$$

$$y_{\text{hat}} = \text{sigmoid}(\mathbf{w} x + \mathbf{b})$$

$$\text{Let } J = \text{Cost}$$

since  $y$  and  $x$  are constant, Cost can be seen as a function of  $\mathbf{w}$  and  $\mathbf{b}$

$$J(w, b) = - (y \log(y_{\text{hat}}) + (1 - y) \log (1 - y_{\text{hat}}))$$

So how do we reduce the cost?



Cost is minimum when  $dJ(w, b) = 0$

# Gradient Descent

Recap:  $\hat{y} = \sigma(w^T x + b)$ ,  $\sigma(z) = \frac{1}{1+e^{-z}}$  ←

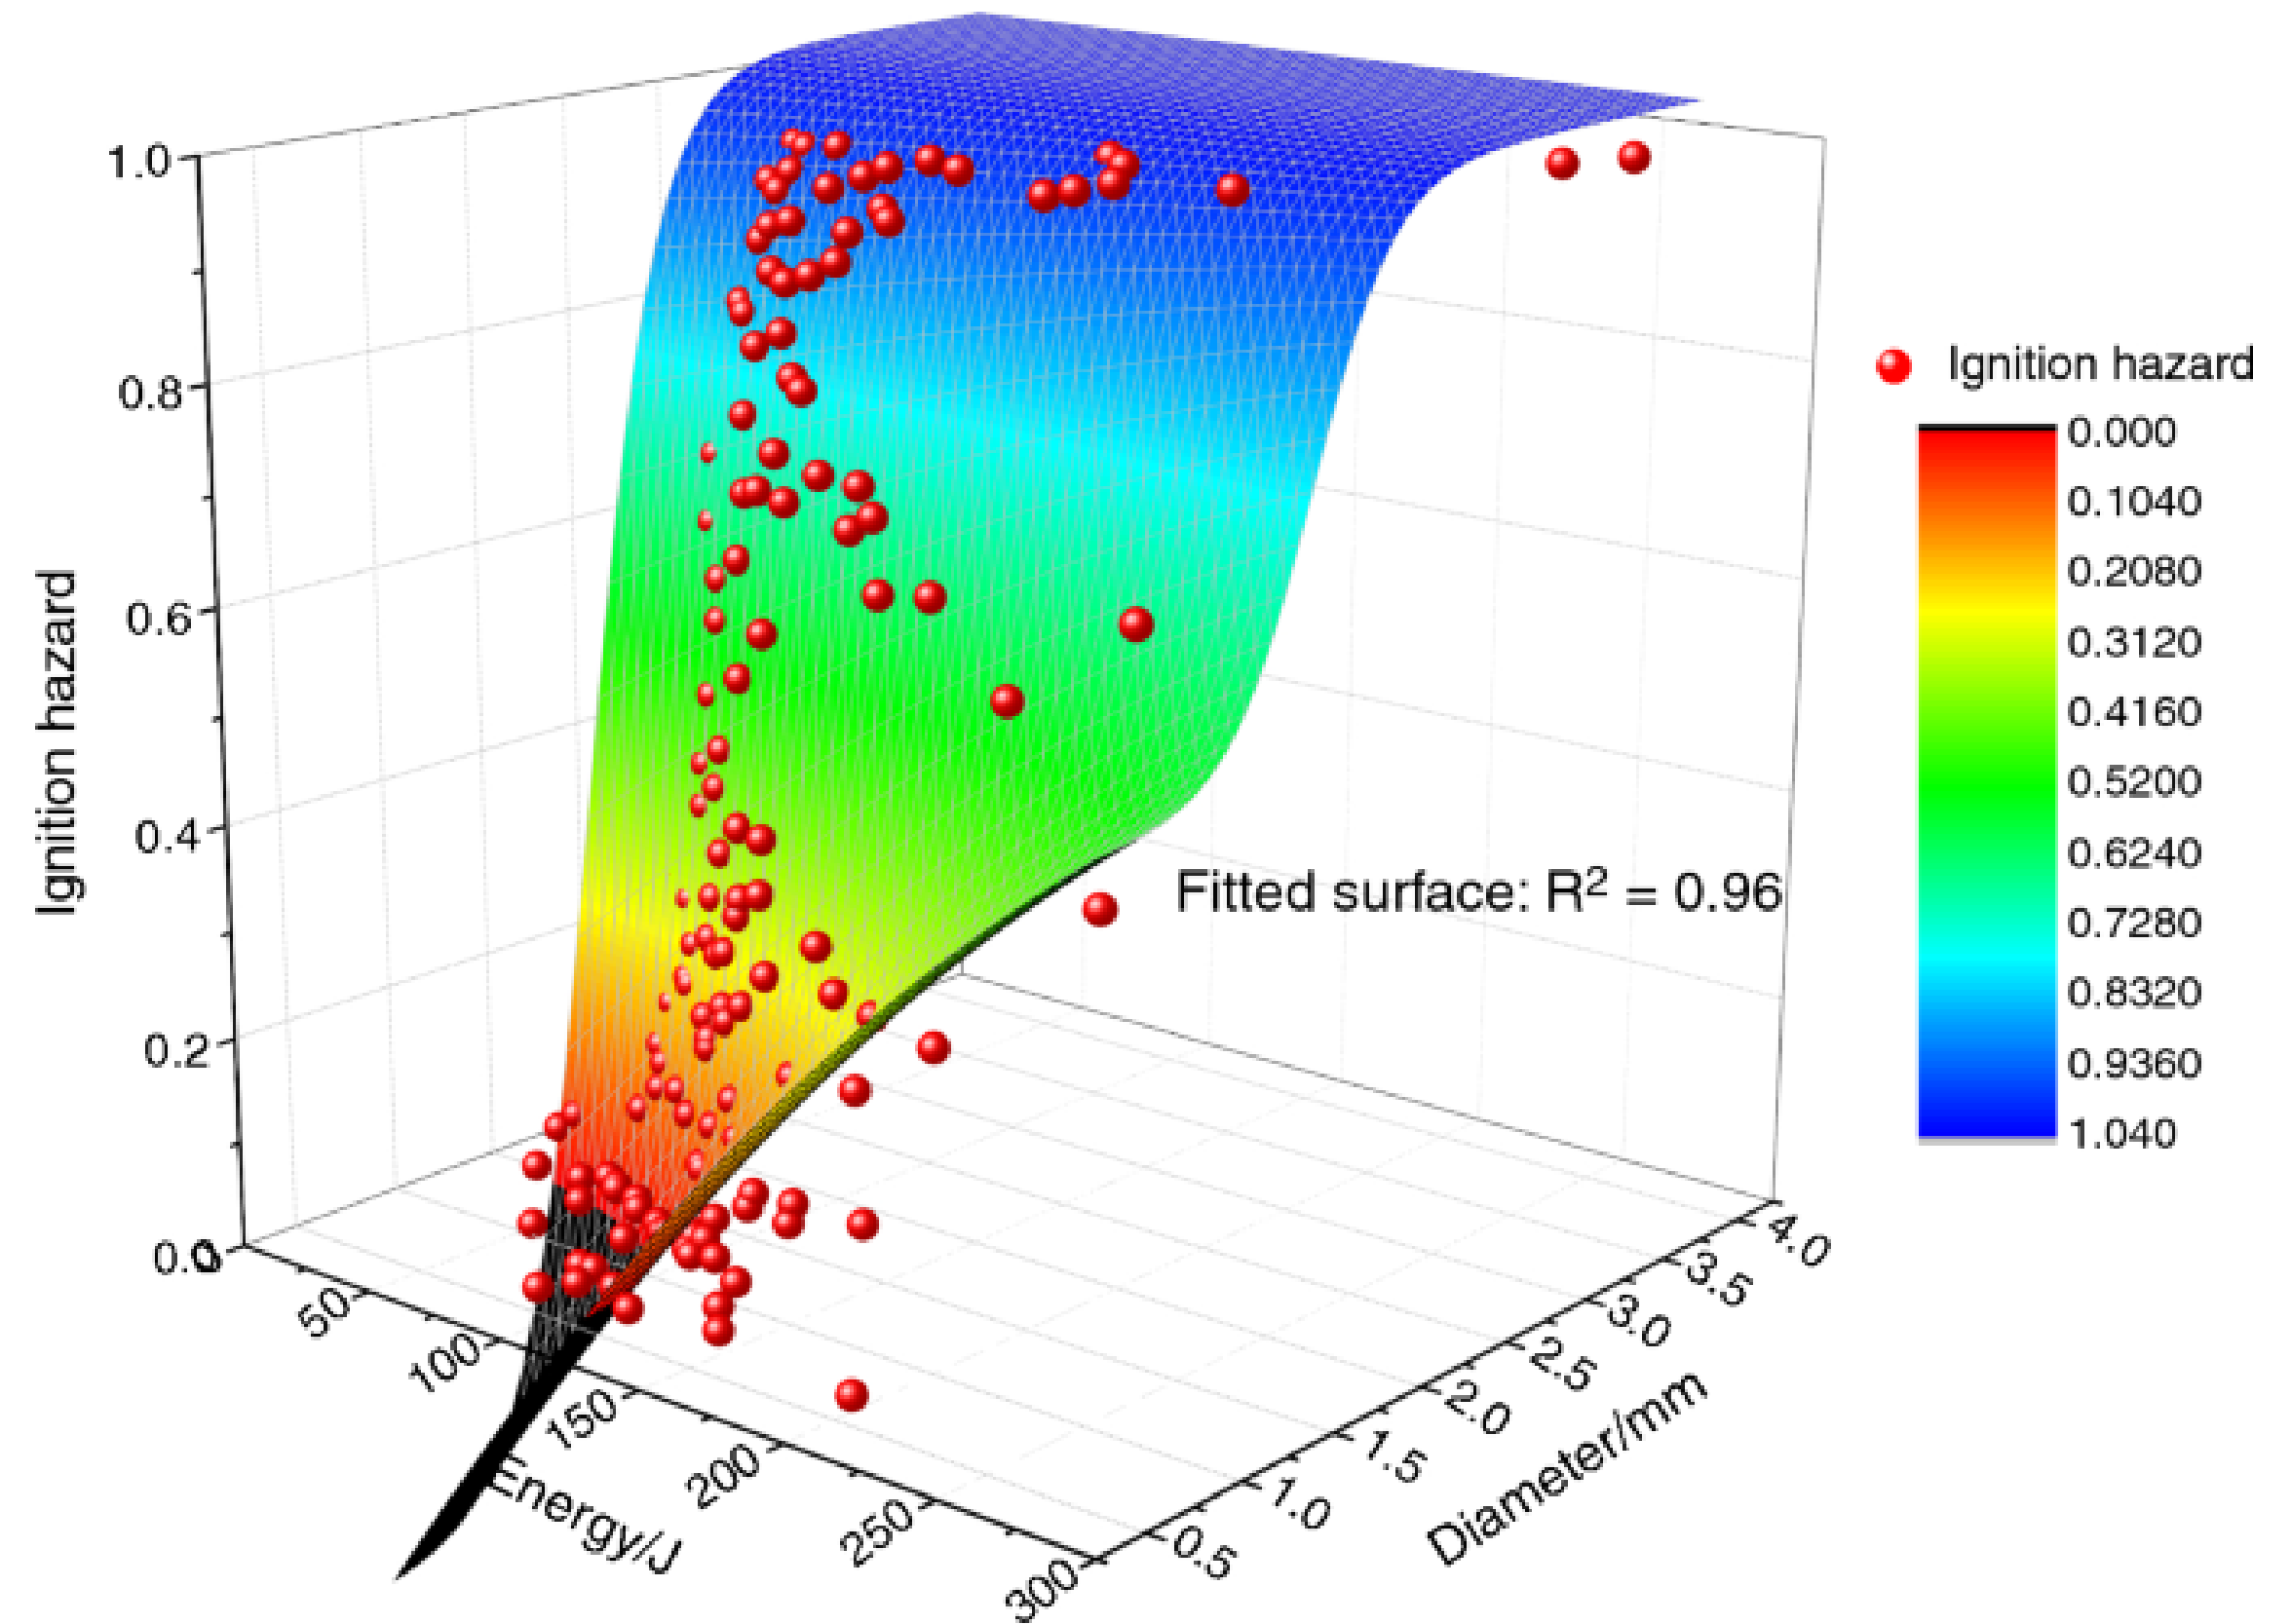
$$\underline{J(w, b)} = \frac{1}{m} \sum_{i=1}^m \mathcal{L}(\underline{\hat{y}^{(i)}}, \underline{y^{(i)}}) = -\frac{1}{m} \sum_{i=1}^m y^{(i)} \log \hat{y}^{(i)} + (1 - y^{(i)}) \log(1 - \hat{y}^{(i)})$$

Want to find  $w, b$  that minimize  $J(w, b)$



# But what if we have 2 independent variables?

$$y = w_1x_1 + w_2x_2 + b$$



Our Cost function will become  
 $J(w_1, w_2, b)$

For each variable, we add a **w**  
and a **b**

Solving for the minimum of  $J(w_1, w_2, \dots, w_x, b)$  by  
partial differentiation will become very  
computationally expensive

**What if we have 3 independent variables?**

$$J(w_1, w_2, w_3, b)$$


**A 1920x1080 image has 2,073,600 x 3 variables**

**We can classify the image using the same principles**

**A pixel is just a number**

$$J(w_1, w_2, \dots, w_{2073600}, b)$$

# Gradient Descent



Repeat {

$$w := w - \alpha \frac{dJ(w)}{dw}$$

learning rate

"dw"

$$w := w - \alpha dw$$

$\frac{dJ(w)}{dw} = ?$

$J(w, b)$

$$w := w - \alpha \frac{\partial J(w, b)}{\partial w}$$

$$b := b - \alpha \frac{\partial J(w, b)}{\partial b}$$

$\frac{\partial J(w, b)}{\partial w}$

$\frac{\partial J(w, b)}{\partial b}$

"partial derivative"

$\partial$

$\partial$


$dw$

$db$

Andrew Ng




# We will be dealing with 28x28 images

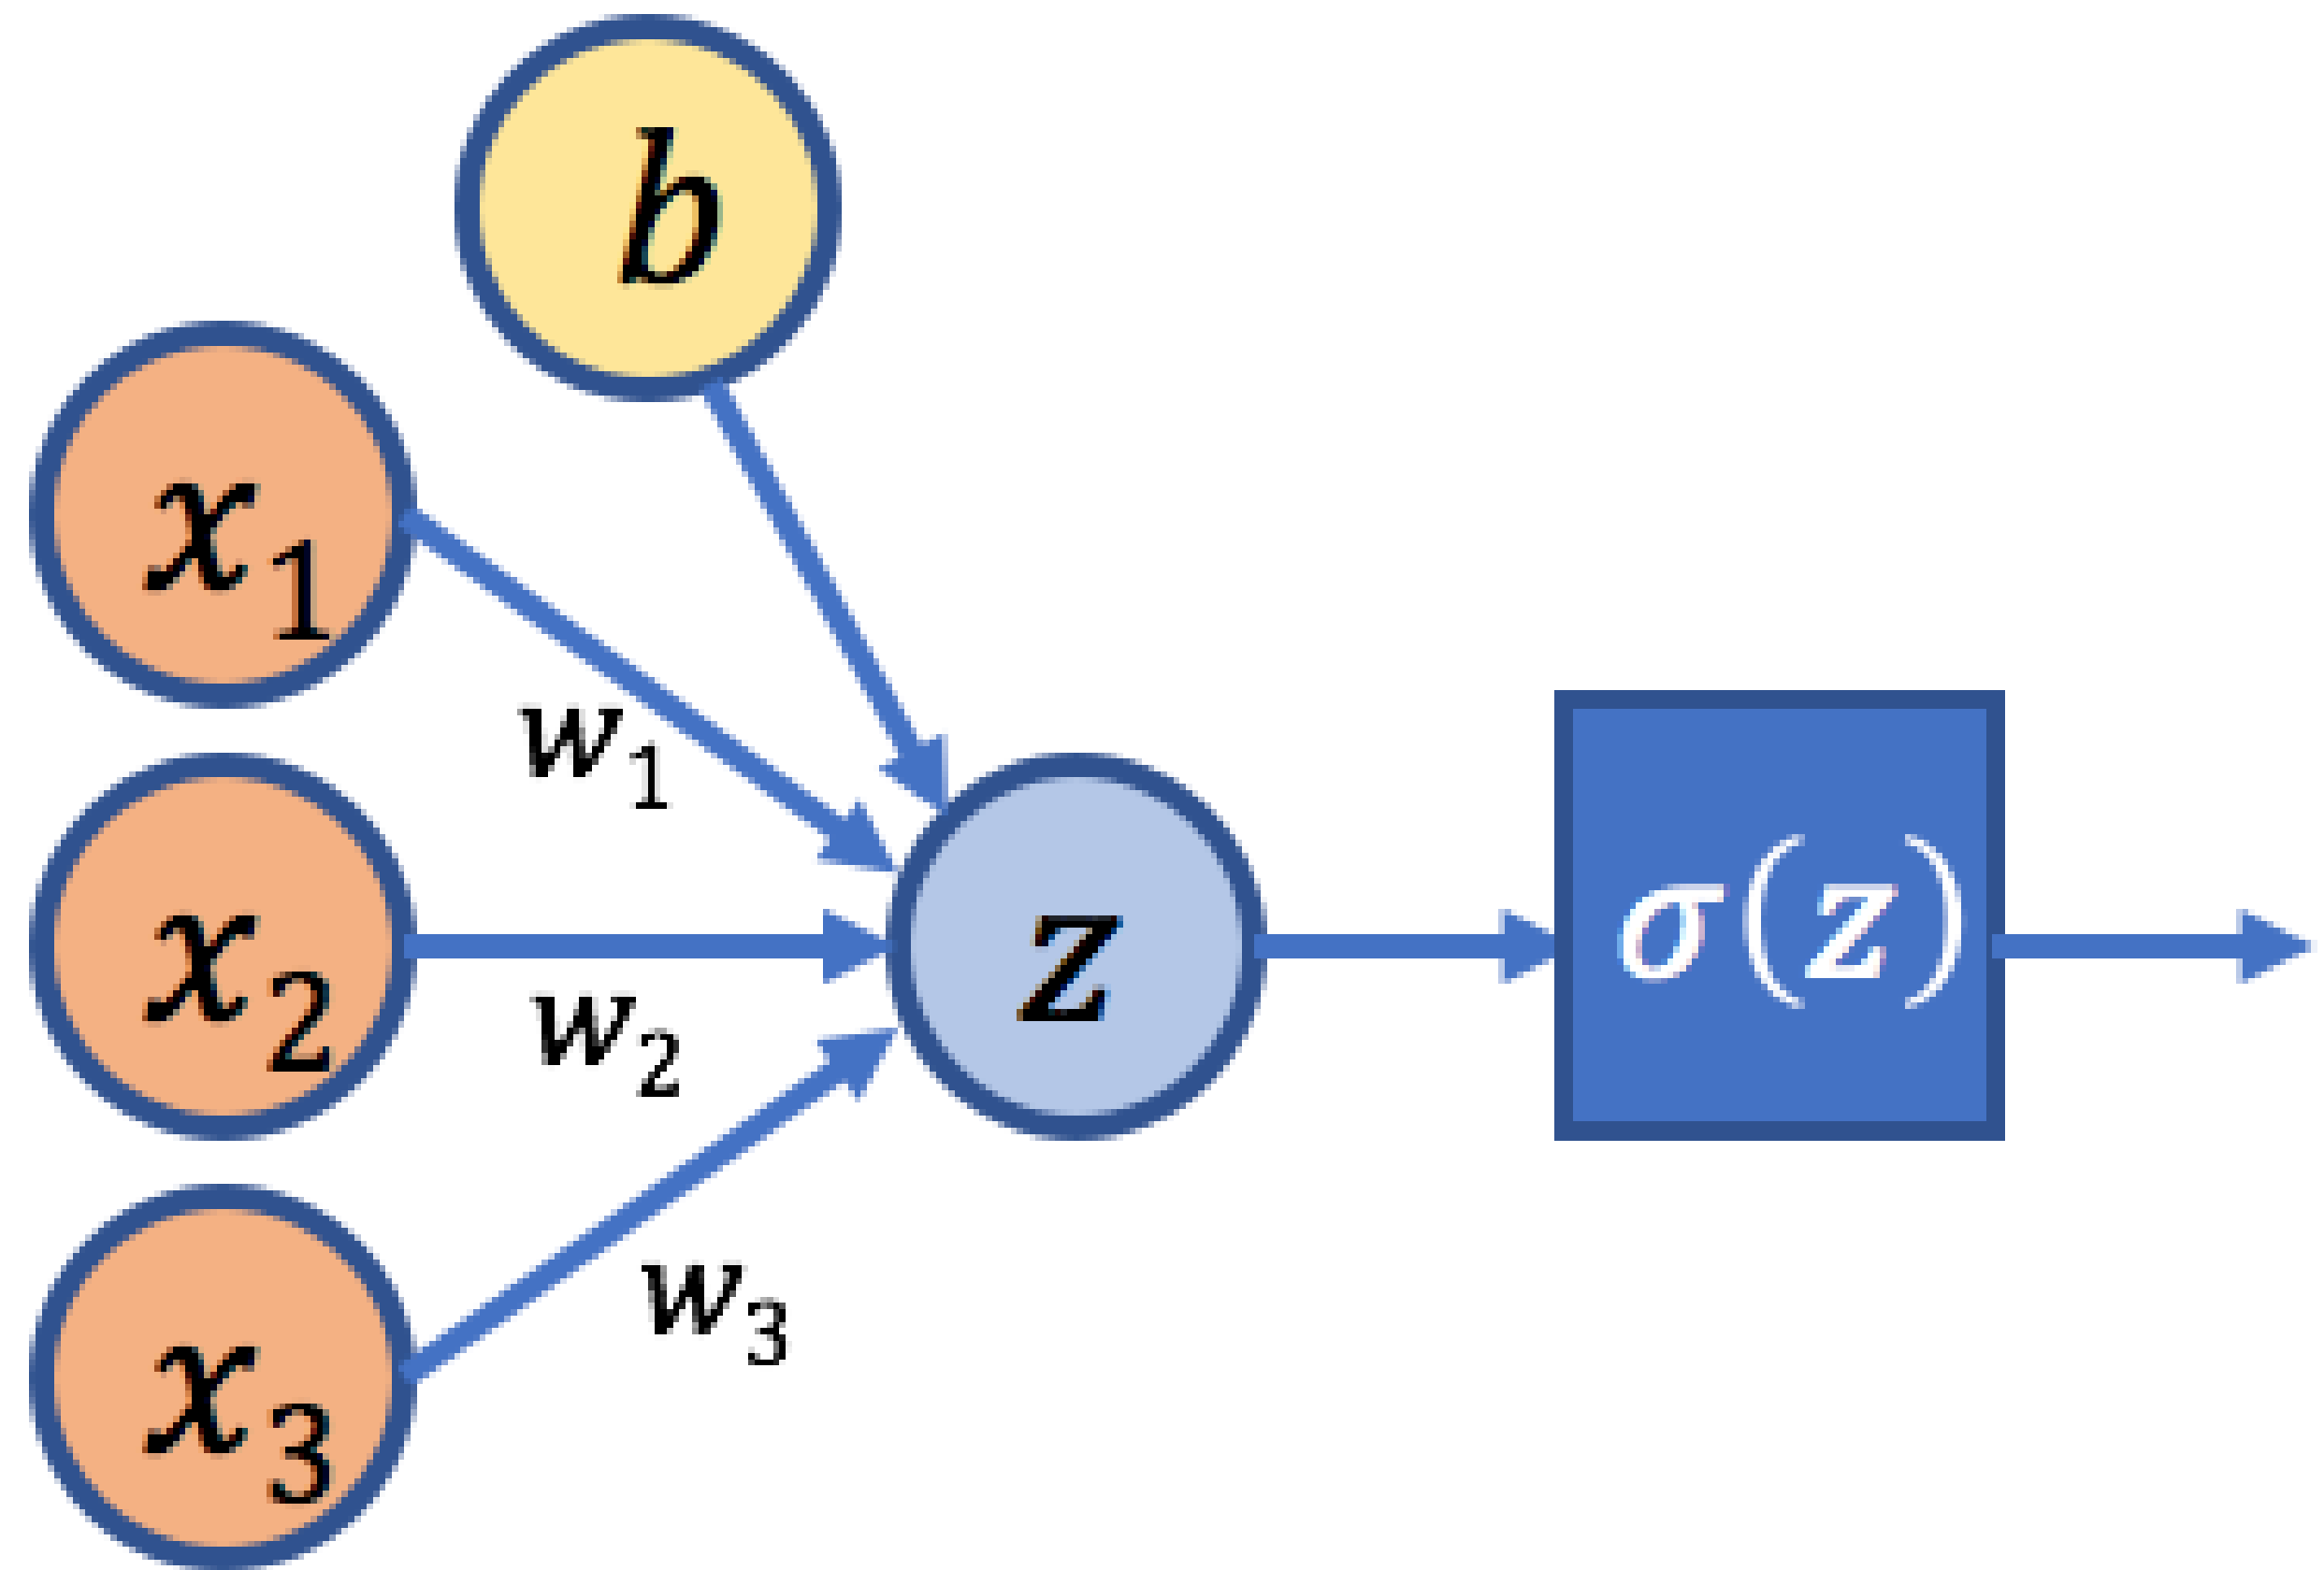





# We flatten the image



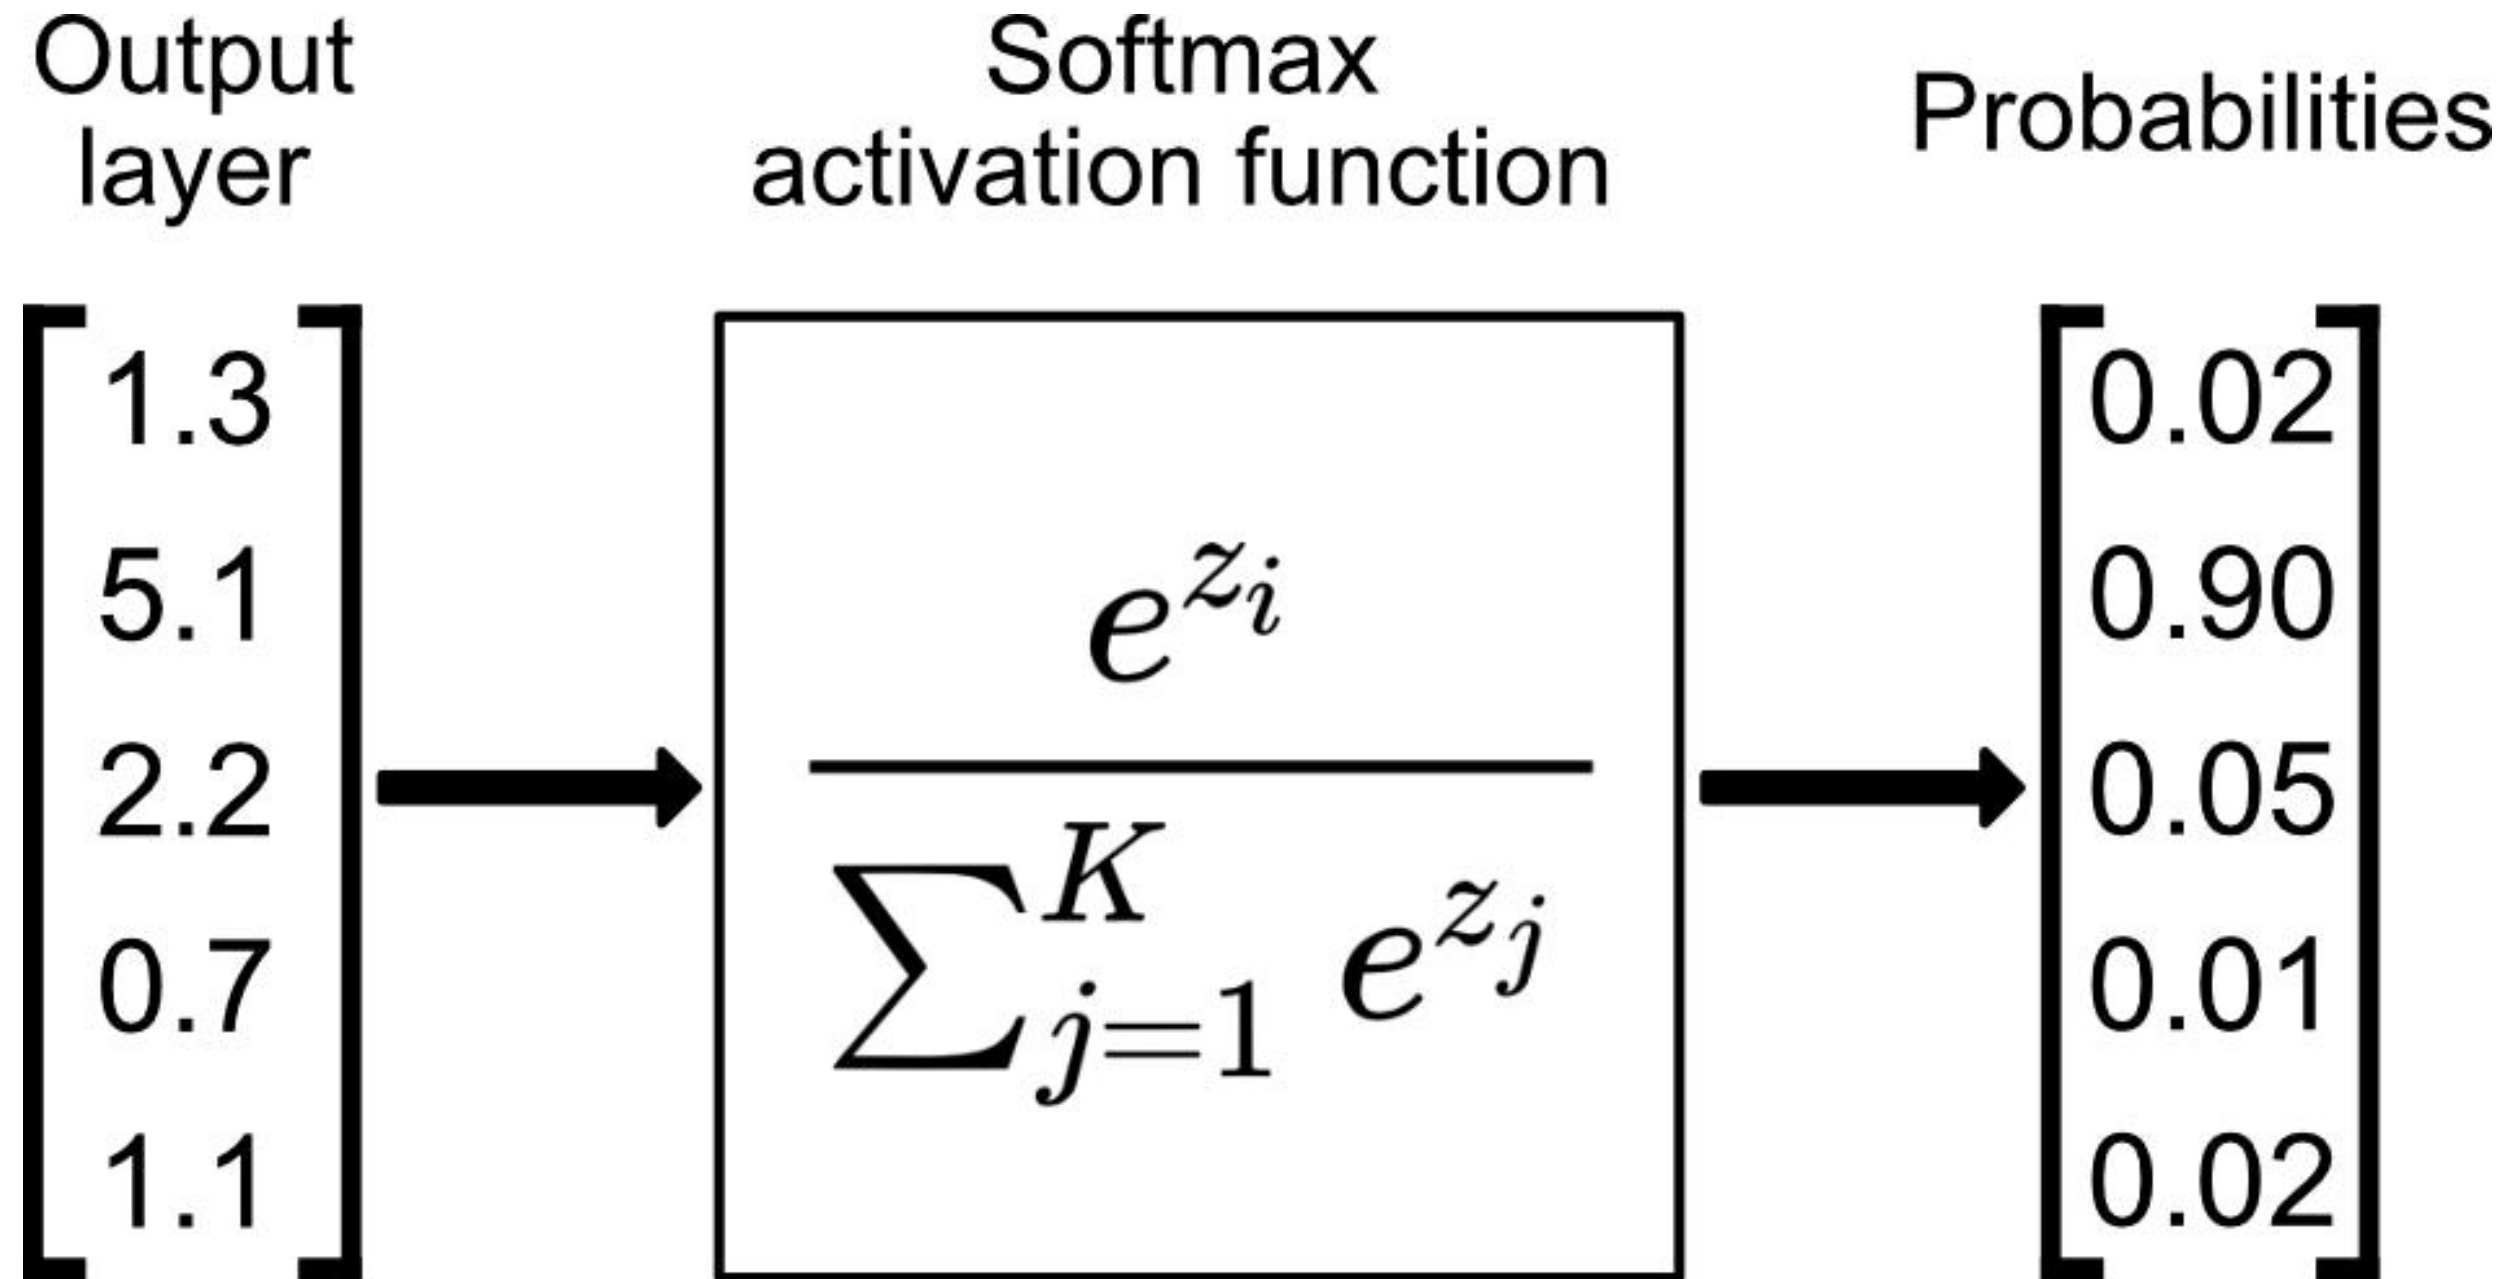
Each pixel of an image can be treated as an independent variable

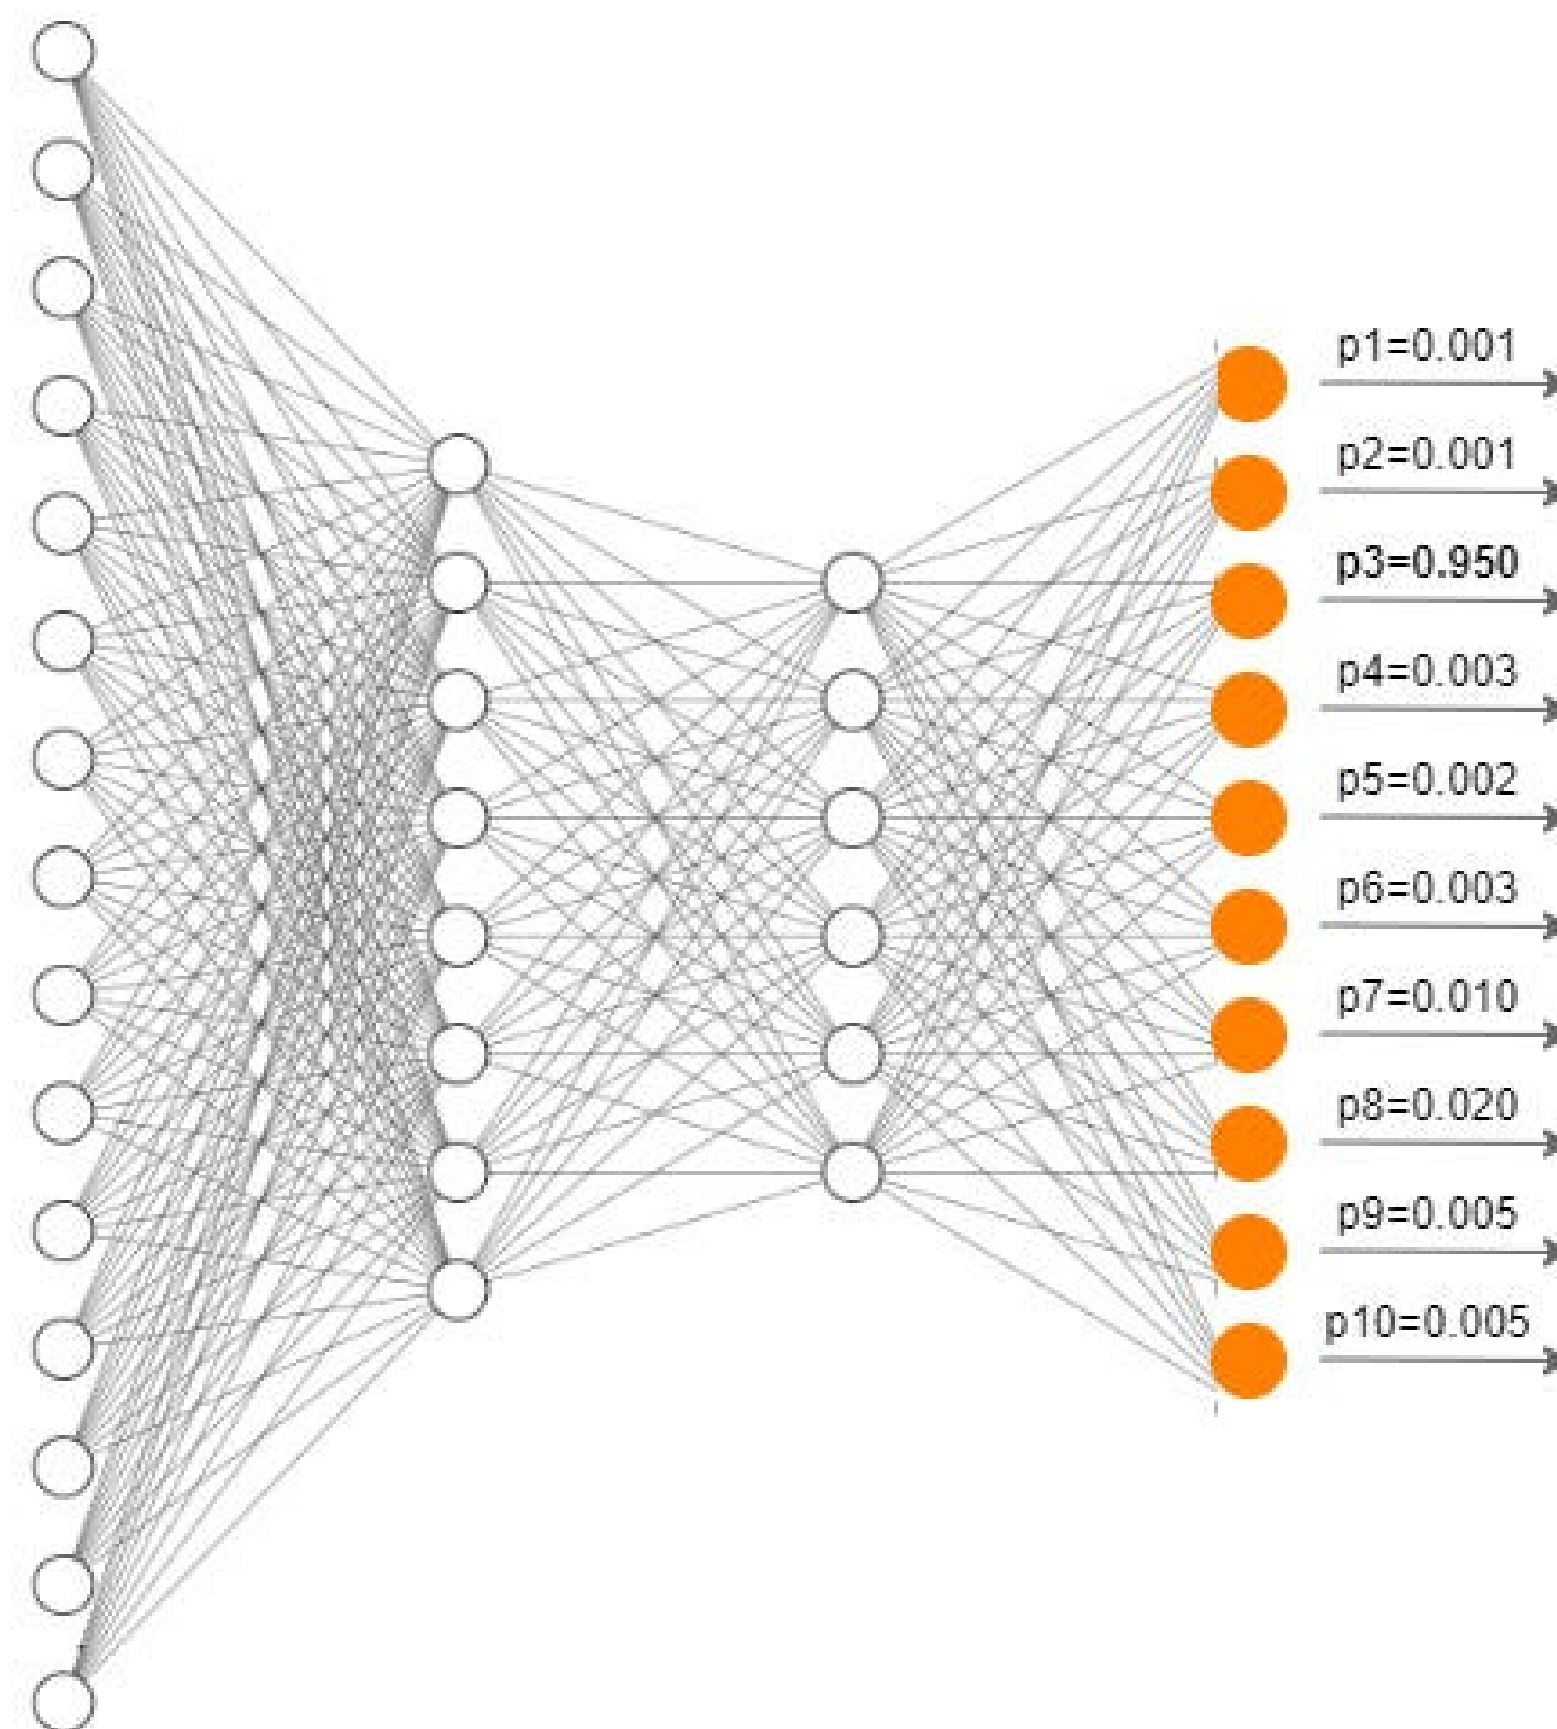



# What are hidden layers?



# Softmax Function



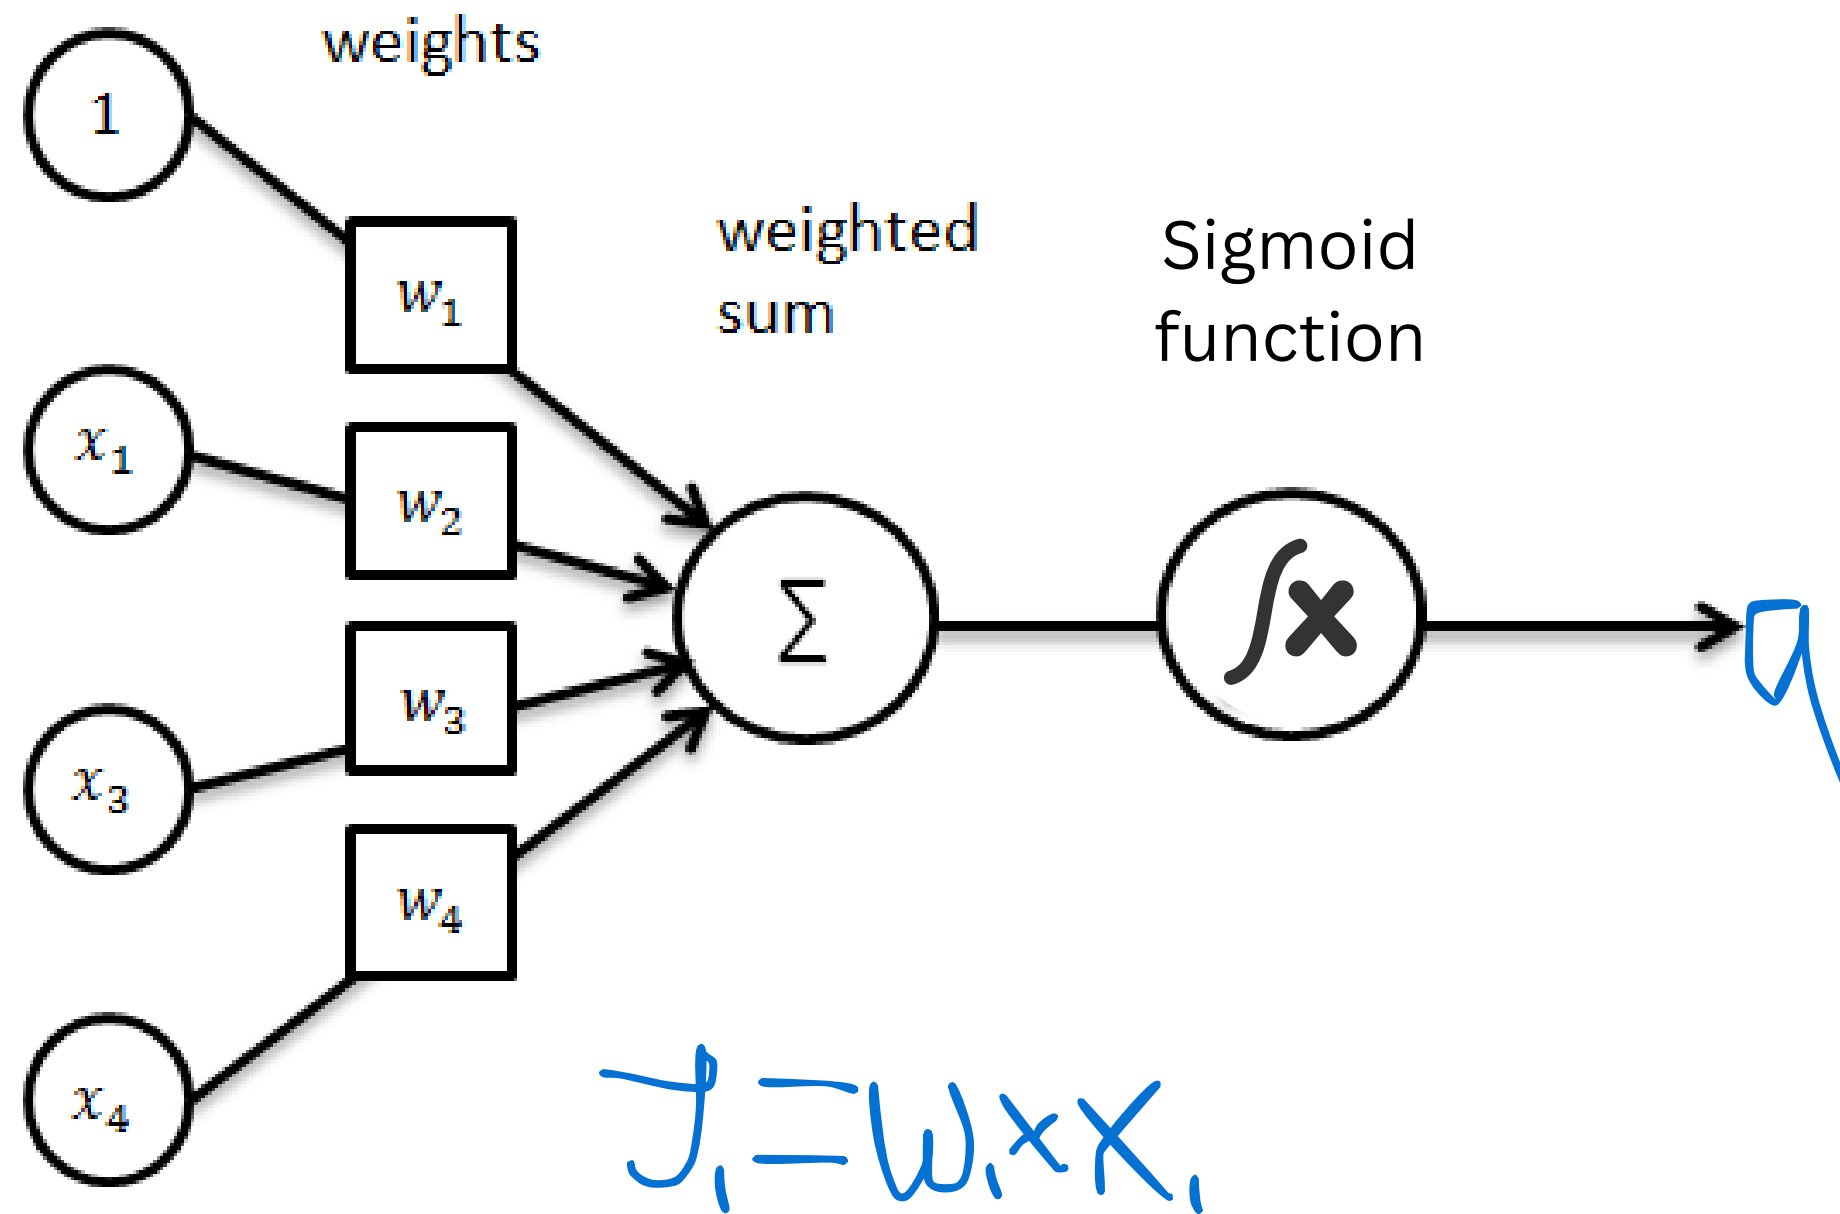


$p1=0.001$   
 $p2=0.001$   
 $p3=0.950$   
 $p4=0.003$   
 $p5=0.002$   
 $p6=0.003$   
 $p7=0.010$   
 $p8=0.020$   
 $p9=0.005$   
 $p10=0.005$

10 Classes (mutually exclusive)	
0	0
1	1
2	2
3	3
4	4
5	5
6	6
7	7
8	8
9	9

# Forward Propagation


inputs



$$\begin{bmatrix} w_1 & w_2 & w_3 & w_4 \end{bmatrix} \begin{bmatrix} 1 \\ x_1 \\ x_3 \\ x_4 \end{bmatrix} = z$$


$$z = y_1 + y_2 + y_3 + y_4 + b$$

$$a = \sigma(z)$$



$$\begin{matrix} \blacksquare & \blacksquare & \blacksquare \\ W3 \end{matrix} \cdot \begin{matrix} \blacksquare \\ \blacksquare \\ \blacksquare \\ H2 \end{matrix} = \begin{matrix} \blacksquare \\ Z3 \end{matrix}$$

# Gradient Descent





# Backpropagation

$$\underbrace{\mathbf{w}^{(t+1)}}_{\text{position of next iteration}} = \underbrace{\mathbf{w}}_{\text{position of previous step}} - \underbrace{\alpha \nabla f_i(\mathbf{w}^{(t)})}_{\text{step}}$$

learning rate

observation  $i$

Now do this  
for all **w**  
and **b**

Momentum

ReLU

Batch normalization

Regularisation

Attention

CNN

RNN

Transformers

Latent Diffusion

GANs

# Advantages and Disadvantages

## Neural Networks: Advantages and Disadvantages

Advantages	Disadvantages
1) The ability to learn by themselves	1) The black box nature and uncertain prediction rates
2) The ability to work with insufficient data	2) Long training processes and limited data efficiency
3) The ability of parallel processing	3) Economically and computationally expensive



**Pytorch vs TensorFlow**

<b>Feature</b>	<b>PyTorch</b>	<b>TensorFlow</b>
Ease of use	Relatively easy to learn, but can be complex for large projects	It can be difficult to understand, but it has a large community of users and resources
Speed	Faster than TensorFlow but not as fast as some other deep learning frameworks	Slower than PyTorch but faster than some other deep-learning frameworks
Flexibility	Very flexible, allowing for a wide variety of deep learning tasks	Less flexible than PyTorch, but still capable of a wide range of tasks

# Credits and Materials

3blue1brown

Andrew Ng

Lightning AI

Awesome Course on Deep Learning by Lightning AI: [https://www.youtube.com/watch?v=6Py-tIEiXKw&list=PLaMu-SDt\\_RB4Ly0xb0qsQVpLwRQcjtOb-&pp=iAQB](https://www.youtube.com/watch?v=6Py-tIEiXKw&list=PLaMu-SDt_RB4Ly0xb0qsQVpLwRQcjtOb-&pp=iAQB)

Andrew Ng's Course: <https://www.coursera.org/learn/neural-networks-deep-learning/home/welcome>

Want to learn more theory 🤔💧?: [https://www.youtube.com/playlist?list=PLZHQObOWTQDNU6R1\\_67000Dx\\_ZCJB-3pi](https://www.youtube.com/playlist?list=PLZHQObOWTQDNU6R1_67000Dx_ZCJB-3pi)

Thank  
you!