

# Deep Learning

Theoretical introduction and its application for face detection, recognition and camouflage.

Raffaella Lanzarotti

# Aim of the course

- Introduction to Deep Learning,
  - Theoretical
  - Practical
- We'll largely adopt the valuable material from:

<http://introtodeeplearning.com/>



# Schedule (tentative)

<b>DAY 1</b>	CLASS 1	<ul style="list-style-type: none"><li>• Introduction to Machine Learning</li></ul>
	LAB 1	<ul style="list-style-type: none"><li>• Tensor Flow</li></ul>
<b>DAY 2</b>	CLASS 2	<ul style="list-style-type: none"><li>• Deep Sequence Modelling</li></ul>
	LAB 2	<ul style="list-style-type: none"><li>• Music Generation</li></ul>
<b>DAY 3</b>	CLASS 3	<ul style="list-style-type: none"><li>• Convolutional Neural Network</li></ul>
	LAB 3	<ul style="list-style-type: none"><li>• MNIST</li></ul>
<b>DAY 4</b>	CLASS 4	<ul style="list-style-type: none"><li>• Deep Generative Models</li></ul>
	LAB 4	<ul style="list-style-type: none"><li>• Debiasing</li></ul>
<b>DAY 5</b>	CLASS 5	<ul style="list-style-type: none"><li>• Deep Reinforcement Learning</li><li>• Limits and new Frontiers</li></ul>

# Introduction to Machine Learning

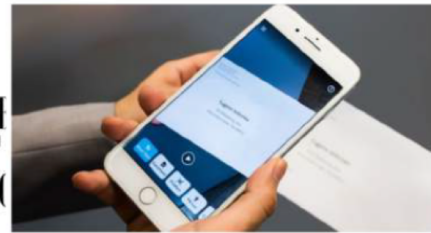
# The Rise of Deep Learning

## 'Deep Voice' Software Can Clone Anyone's Voice With Just 3.7 Seconds of Audio

Using snippets of voices, Baidu's 'Deep Voice' can generate new speech, accents, and tones.



## Let There Be Sight: How Deep Learning Is Helping the Blind 'See'



## DEEPMIND TRIUMPH



## Technology outpacing security measures

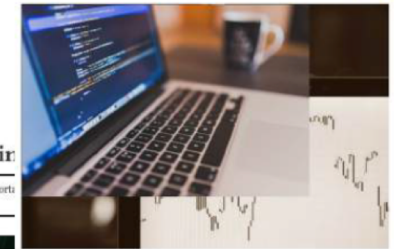
Facial Recognition | Features and Interviews

## AI beats docs in cancer spotting

A new study provides a fresh example of machine learning as an important diagnostic tool. Paul Hiegler reports.

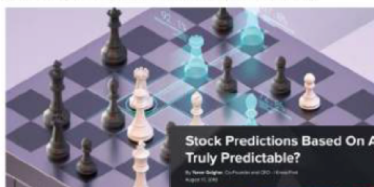


## AI Can Help In Predicting Cryptocurrency Value



## 'Creative' AlphaZero leads way for chess computers and, maybe, science

Former chess world champion Garry Kasparov likes what he sees of computer that could be used to find cures for diseases



## How an A.I. 'Cat-and-Mouse Game' Generates Believable Fake Photos

By CHAD WETZ and KEITH COLLINS JAN. 2, 2018



## Stock Predictions Based On AI: Is the Market Truly Predictable?



## Google's DeepMind aces protein folding

By Robert F. Service | Dec. 6, 2018, 12:05 PM



Complex of bacteria-infecting viral proteins modeled by CASP-13. The complex cost that were modeled individually. PROTEIN DATA BANK

## After Millions of Trials, These Simulated Humans Learned to Do Perfect Backflips and Cartwheels

By Sarah Dillman 1/15/2018 - Boston | Comment by Kenny Walker - Digital Reporter - @RandDMagazine



## Neural networks everywhere

New chip reduces neural networks' power consumption by up to 95 percent, making them practical for battery-powered devices.



## AI faces show how far AI image generation has come in just four years

AI faces on the right aren't real; they're the product of machine learning



## Automation And Algorithms: De-Risking Manufacturing With Artificial Intelligence

Sarah Goehrkke Contributor @ Manufacturing | Focus on the industrialization of additive manufacturing.

TWEET THIS

The two key applications of AI in manufacturing are pricing and manufacturability feedback.

# The Rise of Deep Learning

## Deep learning:

- has revolutionized many areas of machine intelligence, with particular impact on image understanding tasks
- particularly effective...
  - for unstructured data
  - to learn good representations
  - to learn good “models”

# What is Intelligence?

- The ability to process information, to inform future decisions



# What is Deep Learning?

## ARTIFICIAL INTELLIGENCE

Any technique that enables computers to mimic human behavior



## MACHINE LEARNING

Ability to learn without explicitly being programmed



## DEEP LEARNING

Extract patterns from data using neural networks

3 1 3 4 7 2  
1 7 4 2 3 5

Why deep learning?  
Why now?

# Why Deep Learning?

## Traditional ML:

- Hand engineered features
- LIMITS and PROBLEMS:
  - Time consuming
  - Brittle
  - Not scalable

## Challenge:

can we learn the **underlying features** directly from data?

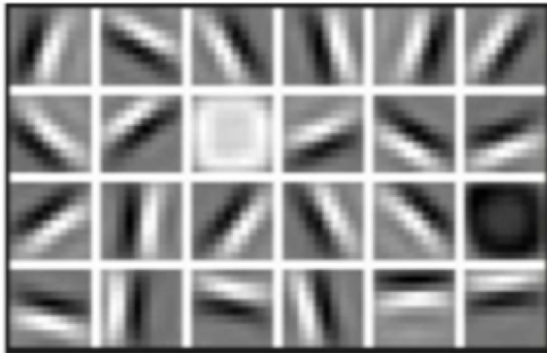
**Deep Learning**  
learns features directly from data



# Ex: Features to detect faces

- Which features characterize faces?
- They should be:
  - specific to this class
  - flexible to manage intra-class variability

Low Level Features ?



Lines & Edges

Mid Level Features?



Eyes & Nose & Ears

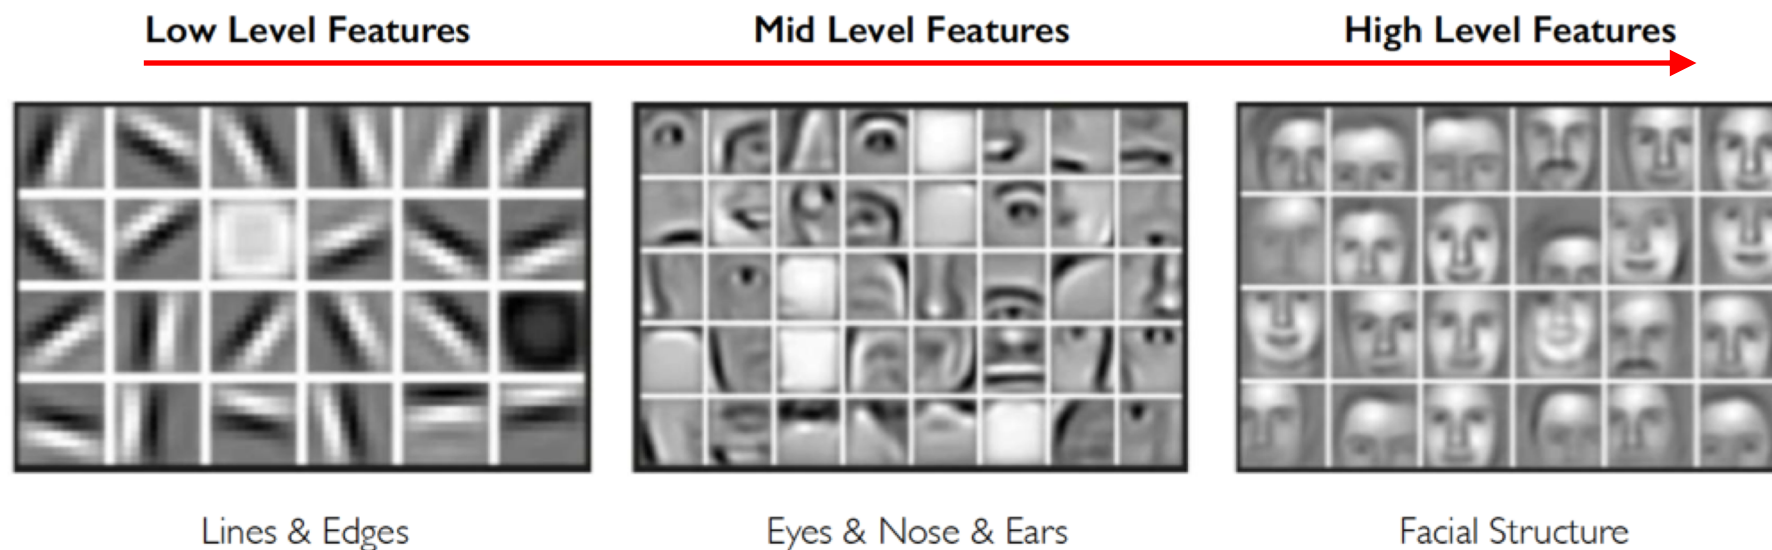
High Level Features?



Facial Structure

# Why Deep Learning?

**Deep Learning**  
learns features in a hierarchical manner



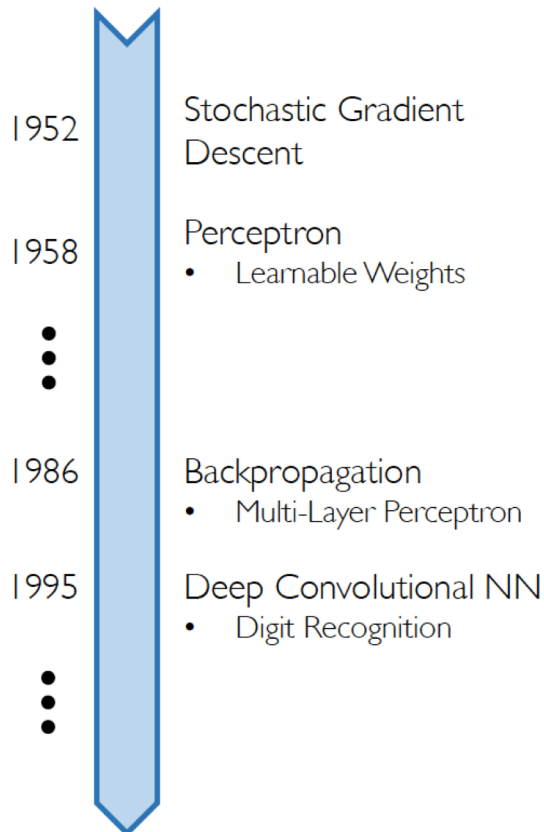
# In this course...

- We'll try to answer to this question:

HOW CAN WE GO FROM RAW DATA (e.g. pixels) TO A MORE AND MORE COMPLEX REPRESENTATION AS THE DATA FLOWS THROUGH THE MODEL?

# Why Now?

Neural Networks date back decades, so why the resurgence?



## 1. Big Data

- Larger Datasets
- Easier Collection & Storage

IMAGENET



WIKIPEDIA  
The Free Encyclopedia



## 2. Hardware

- Graphics Processing Units (GPUs)
- Massively Parallelizable



## 3. Software

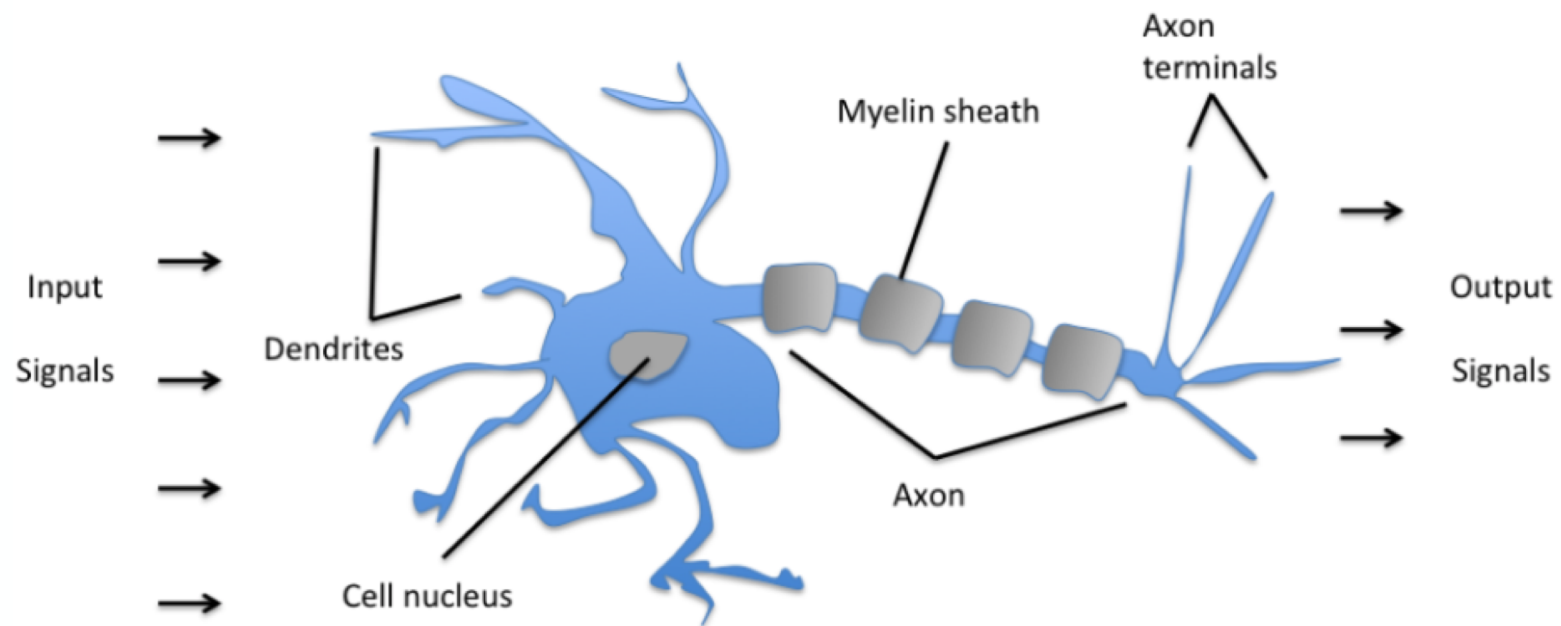
- Improved Techniques
- New Models
- Toolboxes



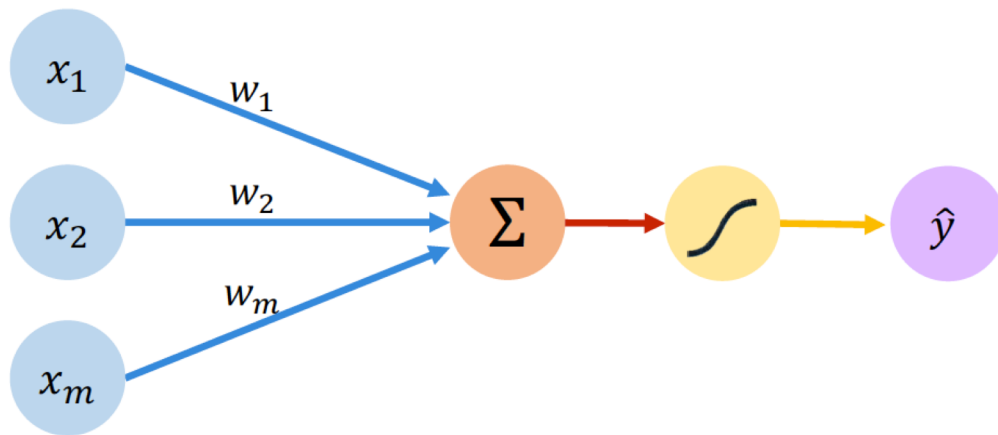
# The Perceptron

The structural building block of deep learning

# Biological Inspiration



# The Perceptron: Forward Propagation



Output

Linear combination of inputs

$$\hat{y} = g \left( \sum_{i=1}^m x_i w_i \right)$$

Non-linear activation function

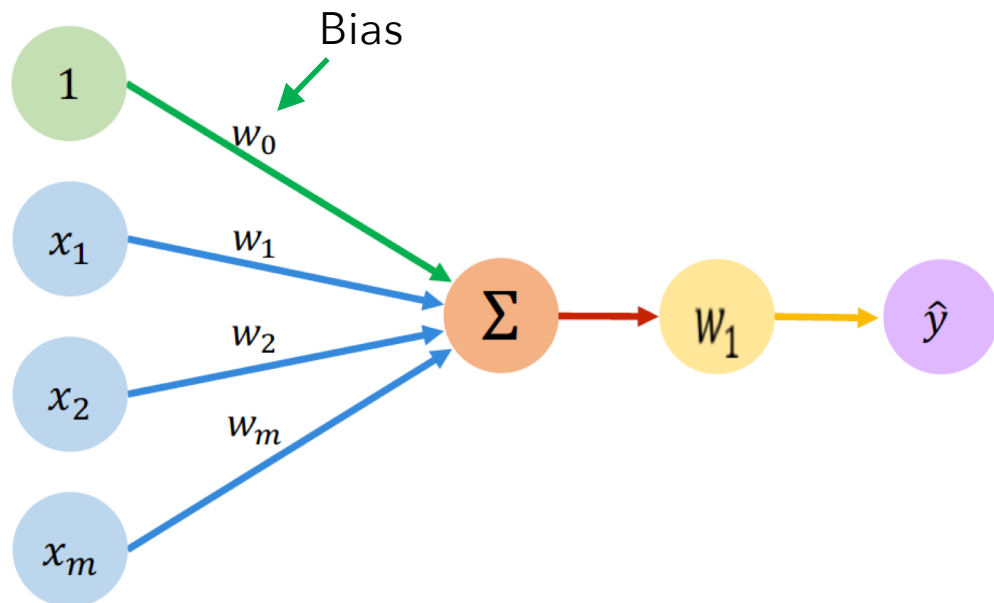
Inputs

Sum

Non-Linearity

Output

# The Perceptron: Forward Propagation



Output

Linear combination of inputs

$$\hat{y} = g \left( w_0 + \sum_{i=1}^m x_i w_i \right)$$

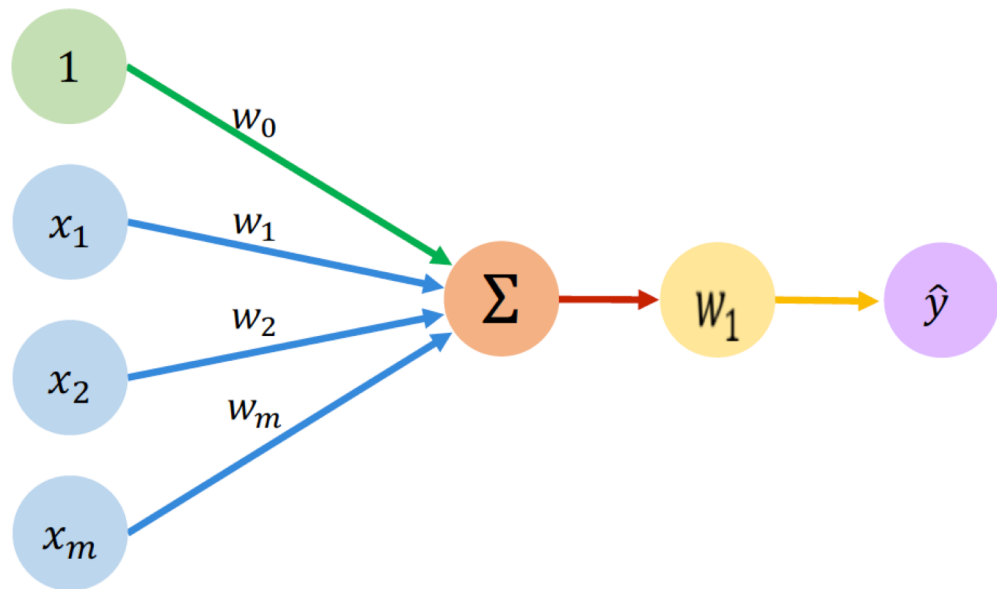
Non-linear activation function

Bias

Inputs    Weights    Sum    Non-Linearity    Output



# The Perceptron: Forward Propagation



$$\hat{y} = g \left( w_0 + \sum_{i=1}^m x_i w_i \right)$$

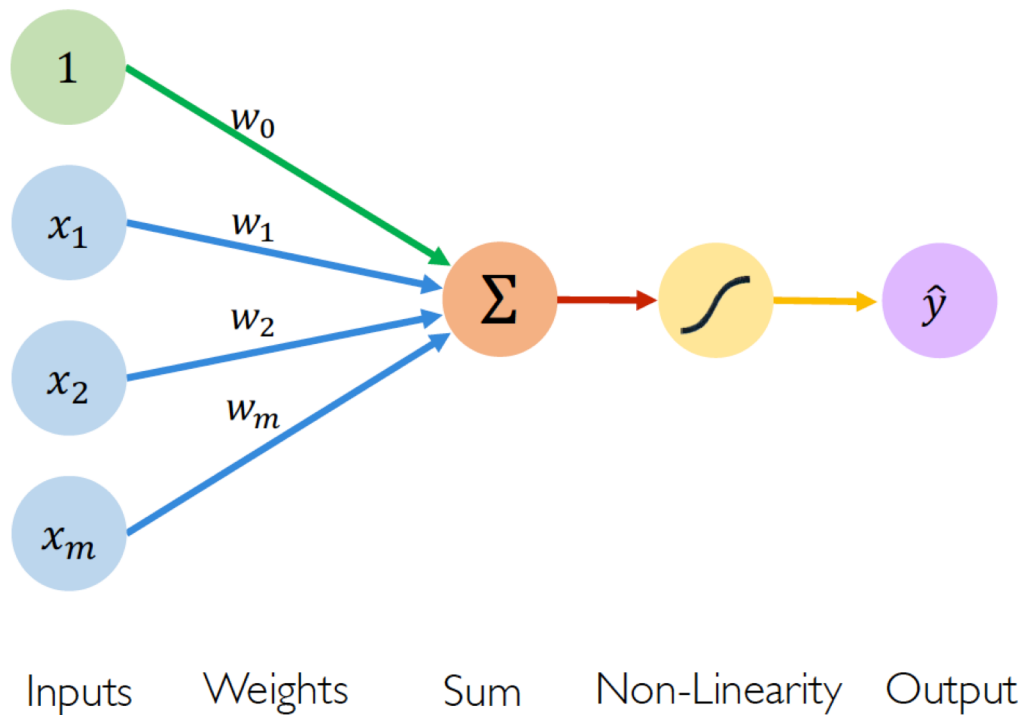
$$\hat{y} = g ( w_0 + \mathbf{X}^T \mathbf{W} )$$

where:  $\mathbf{X} = \begin{bmatrix} x_1 \\ \vdots \\ x_m \end{bmatrix}$  and  $\mathbf{W} = \begin{bmatrix} w_1 \\ \vdots \\ w_m \end{bmatrix}$

Using Linear Algebra...

Inputs    Weights    Sum    Non-Linearity    Output

# The Perceptron: Forward Propagation

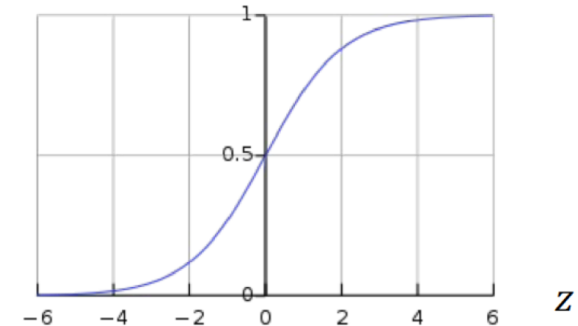


## Activation Functions

$$\hat{y} = g(w_0 + \mathbf{X}^T \mathbf{W})$$

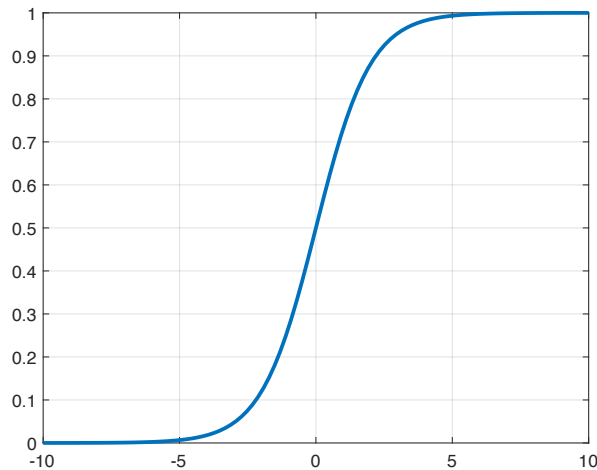
- Example: sigmoid function

$$g(z) = \sigma(z) = \frac{1}{1 + e^{-z}}$$



# Sigmoid

Near-0 gradient

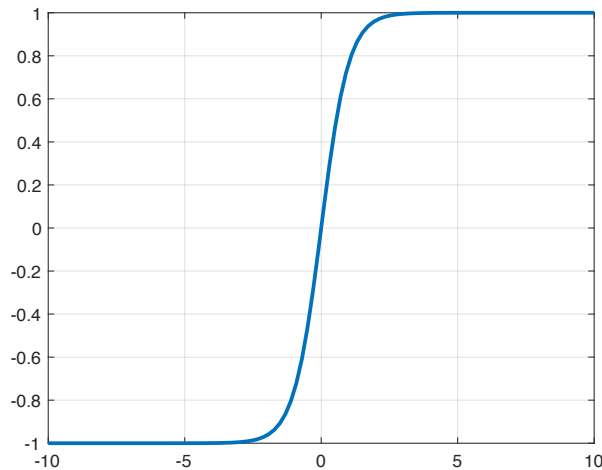


Near-0 gradient

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

- Historically the most used for binary classification
- Useful for modelling probability, because it collapse the input between 0 and 1
- It suffers from the vanishing gradient problem
- Non-zero centered output that may cause zig-zagging

# Tanh

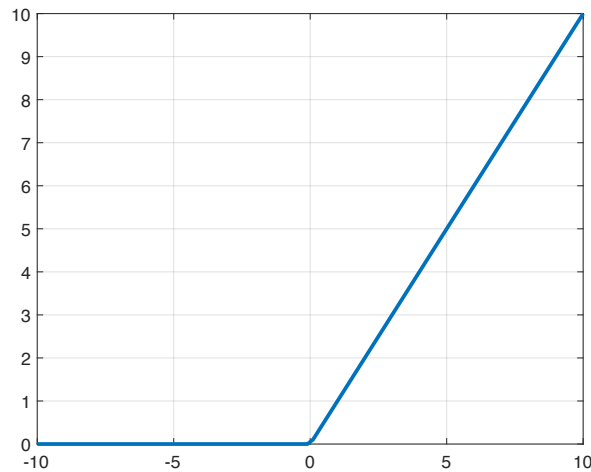


$$\tanh(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}}$$

$$\tanh(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}} = 2\sigma(2x) - 1$$

- It suffers from the vanishing gradient problem
- Output is zero centered, thus it has better gradient properties than sigmoid
- It is a scaled version of Sigmoid:

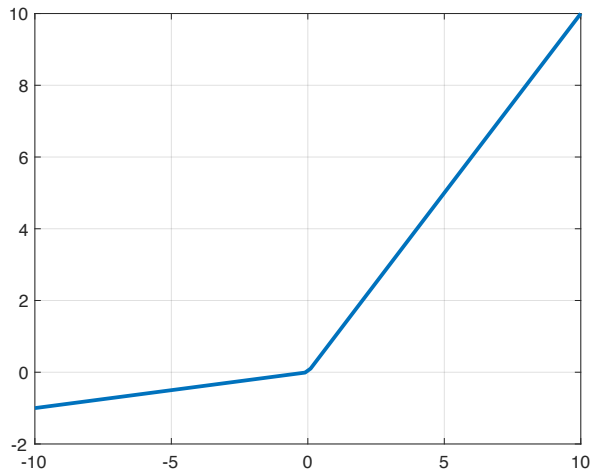
# ReLU (Rectified Linear Unit)



$$f(x) = \max(0, x)$$

- Very popular and simple: it thresholds values below 0
- It allows for fast convergence of the optimization function
- The weight may irreversibly die

# Leaky ReLU



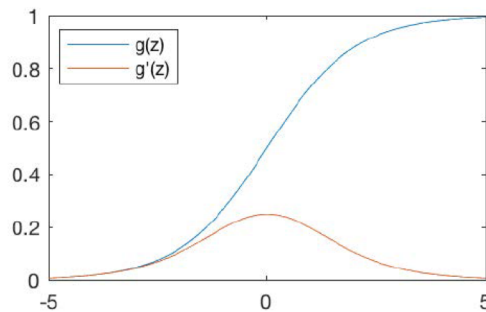
- It is aimed to fix the dying ReLU problem
- In a variant (called parametric ReLU) the slope for negative values can be learnt

$$f(x) = \begin{cases} \alpha x & \text{if } x < 0 \\ x & \text{if } x \geq 0 \end{cases}$$

$$\alpha = 0.1$$


# Common Activation Functions

Sigmoid Function

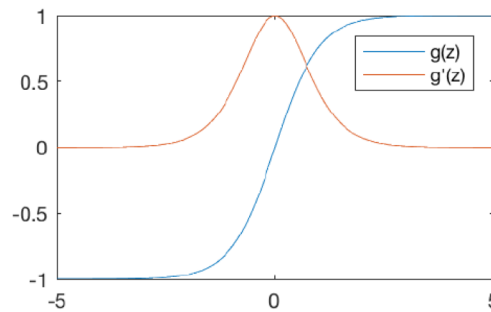


$$g(z) = \frac{1}{1 + e^{-z}}$$

$$g'(z) = g(z)(1 - g(z))$$


 `tf.math.sigmoid(z)`

Hyperbolic Tangent

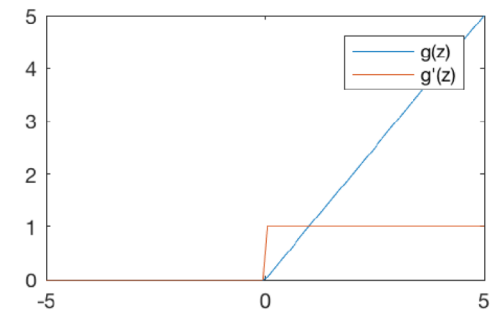


$$g(z) = \frac{e^z - e^{-z}}{e^z + e^{-z}}$$

$$g'(z) = 1 - g(z)^2$$


 `tf.math.tanh(z)`

Rectified Linear Unit (ReLU)



$$g(z) = \max(0, z)$$

$$g'(z) = \begin{cases} 1, & z > 0 \\ 0, & \text{otherwise} \end{cases}$$

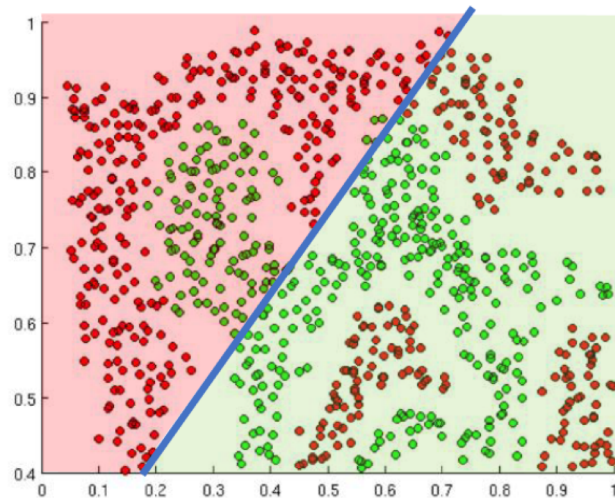
 `tf.nn.relu(z)`

 TensorFlow code blocks

NOTE: All activation functions are non-linear

# Importance of Activation Functions

The purpose of activation functions is to **introduce non-linearities** into the network

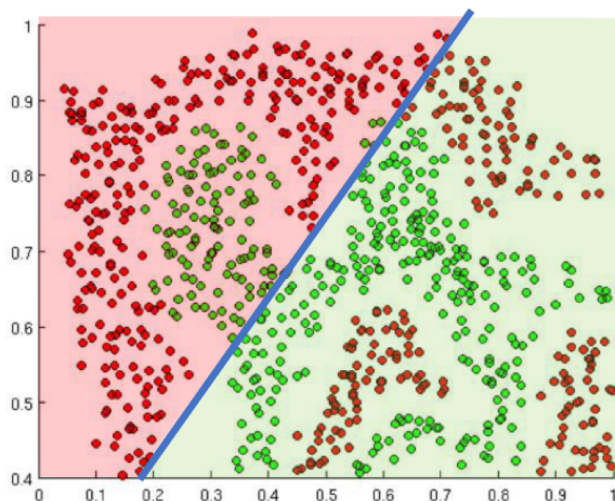


Linear Activation functions produce linear decisions no matter the network size

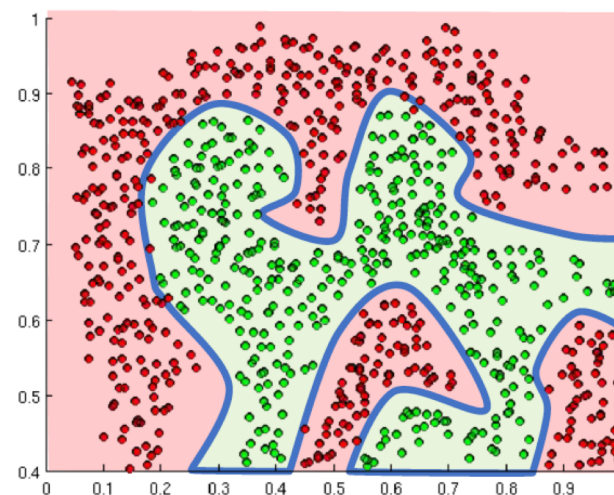


# Importance of Activation Functions

The purpose of activation functions is to *introduce non-linearities* into the network

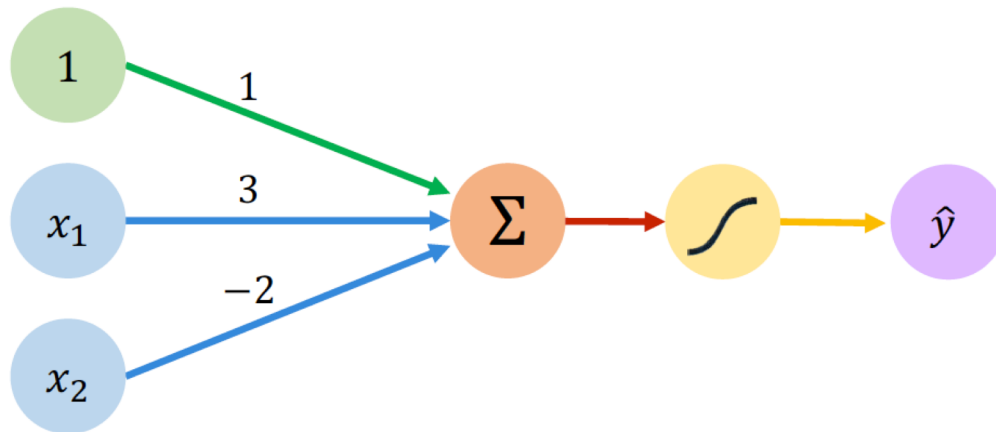


Linear Activation functions produce linear decisions no matter the network size



Non-linearities allow us to approximate arbitrarily complex functions

# The Perceptron: Example



We have:  $w_0 = 1$  and  $\mathbf{W} = \begin{bmatrix} 3 \\ -2 \end{bmatrix}$

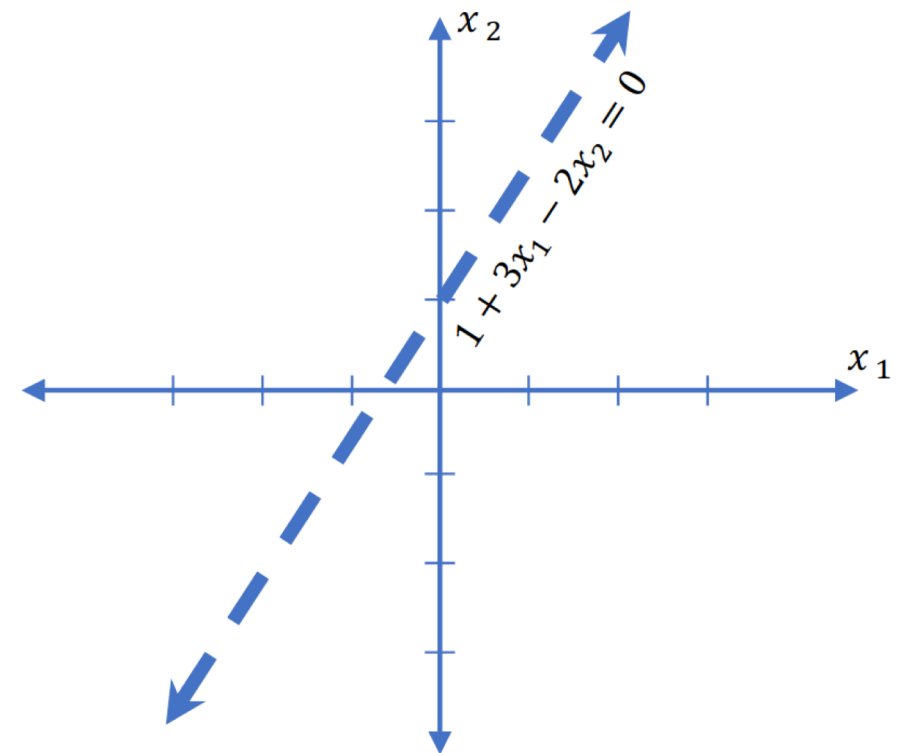
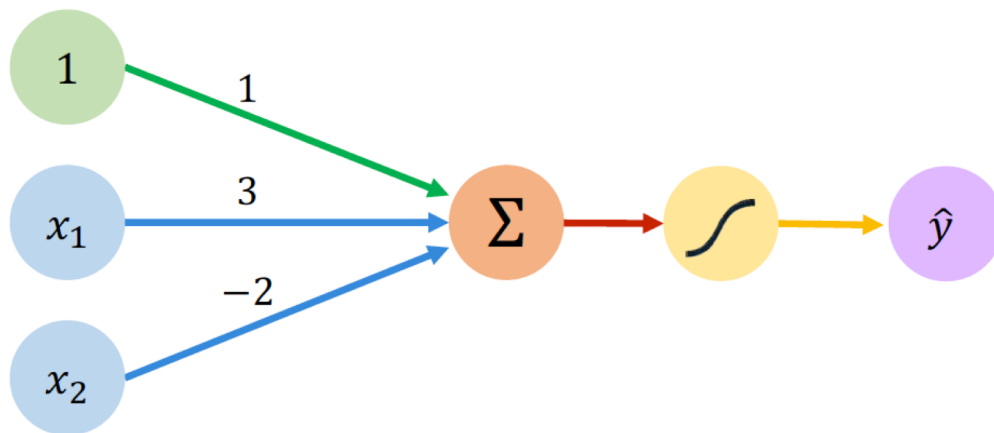
$$\begin{aligned}\hat{y} &= g(w_0 + \mathbf{X}^T \mathbf{W}) \\ &= g\left(1 + \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}^T \begin{bmatrix} 3 \\ -2 \end{bmatrix}\right) \\ \hat{y} &= g(1 + 3x_1 - 2x_2)\end{aligned}$$

This is just a line in 2D!

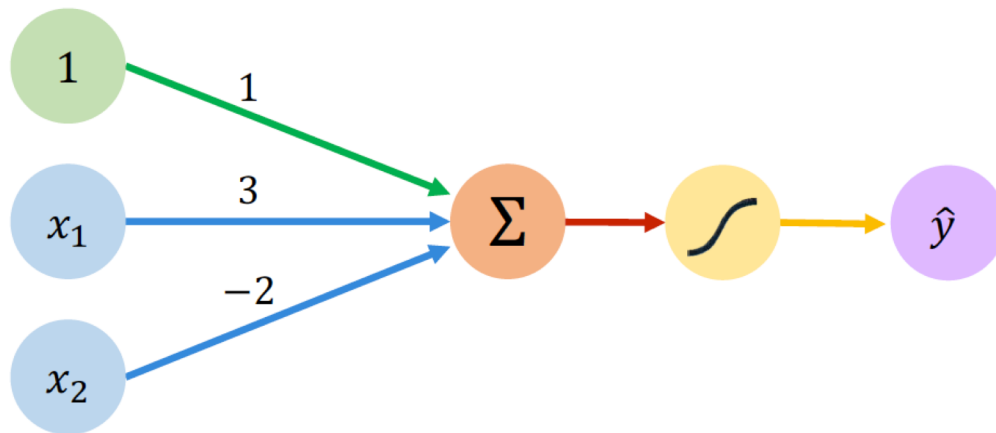
# The Perceptron: Example

Plot this line equal to 0 in the **feature space**:

$$\hat{y} = g(1 + 3x_1 - 2x_2)$$



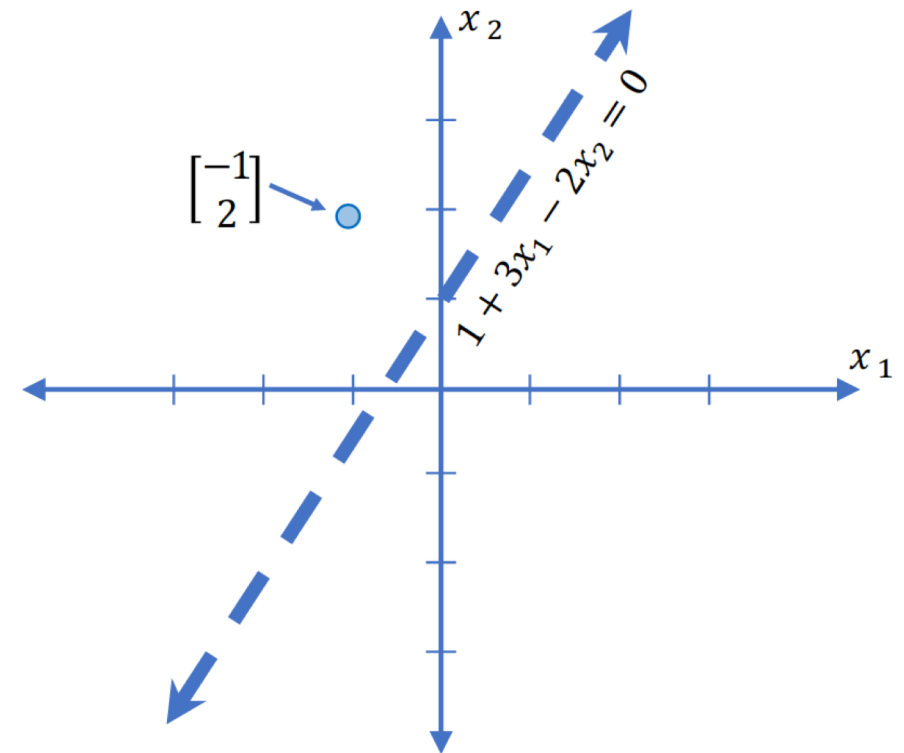
# The Perceptron: Example



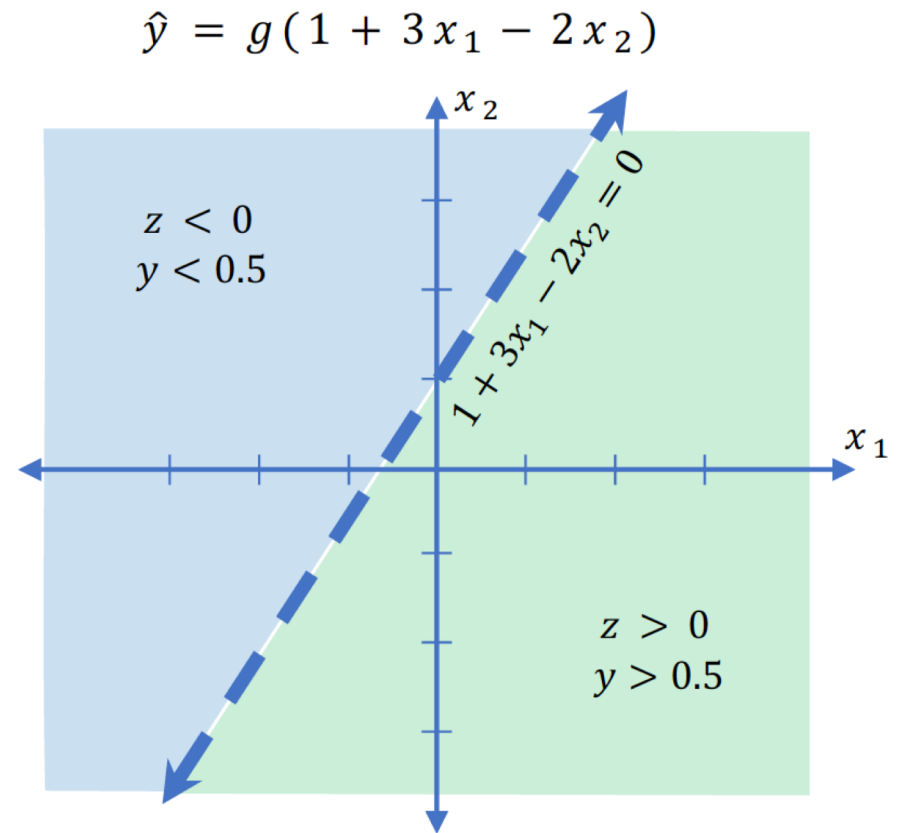
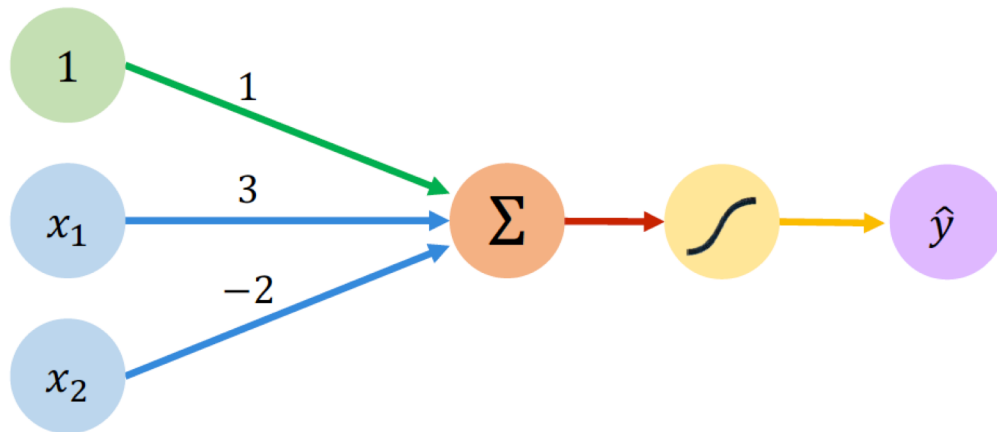
Assume we have input:  $\mathbf{X} = \begin{bmatrix} -1 \\ 2 \end{bmatrix}$

$$\begin{aligned}\hat{y} &= g(1 + (3 * -1) - (2 * 2)) \\ &= g(-6) \approx 0.002\end{aligned}$$

$$\hat{y} = g(1 + 3x_1 - 2x_2)$$

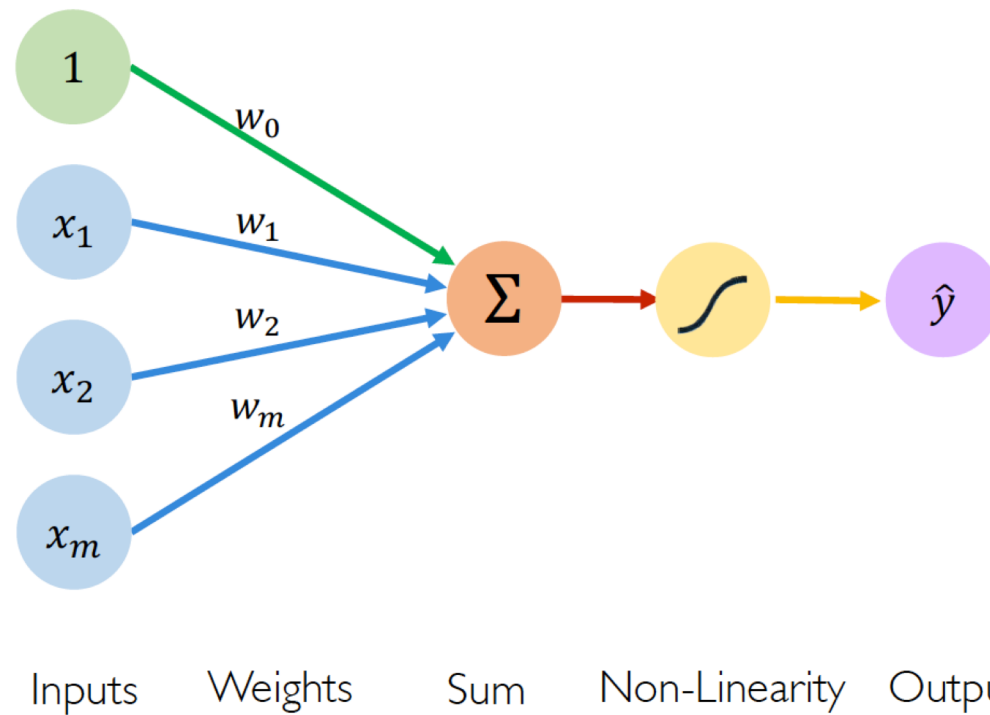


# The Perceptron: Example



# Building Neural Networks with Perceptrons

# The Perceptron: Simplified



# The Perceptron: Simplified

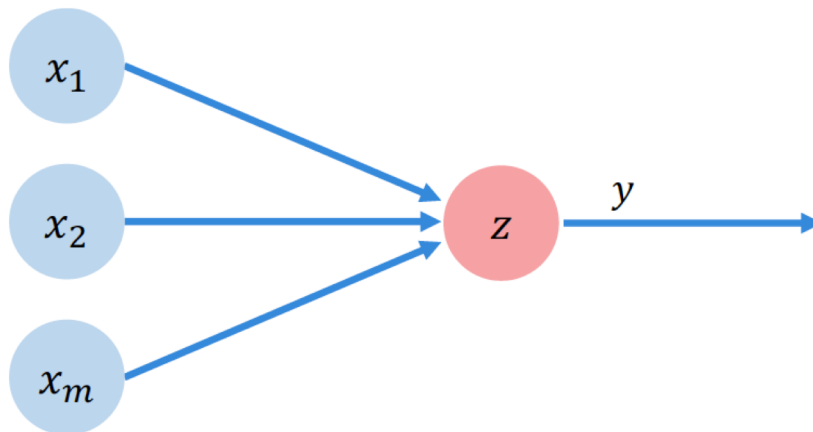


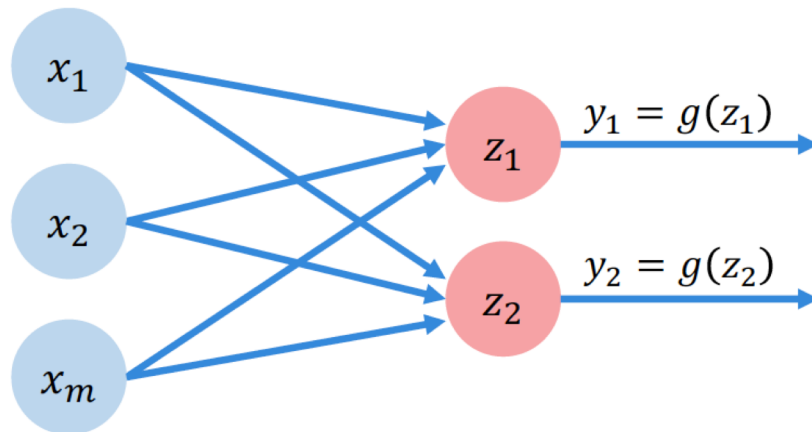
Diagram simplification

- No Bias
- No weights
- $z$ : input to the a.f.
- $y$ : output of the a.f.

$$z = w_0 + \sum_{j=1}^m x_j w_j$$



# Multi Output Perceptron



- Simply add a perceptron
- Same input
- Same process
- What changes are the weights

$$z_i = w_{0,i} + \sum_{j=1}^m x_j w_{j,i}$$

# Dense layer from scratch



```
class MyDenseLayer(tf.keras.layers.Layer):
    def __init__(self, input_dim, output_dim):
        super(MyDenseLayer, self).__init__()

        # Initialize weights and bias
        self.W = self.add_weight([input_dim, output_dim])
        self.b = self.add_weight([1, output_dim])

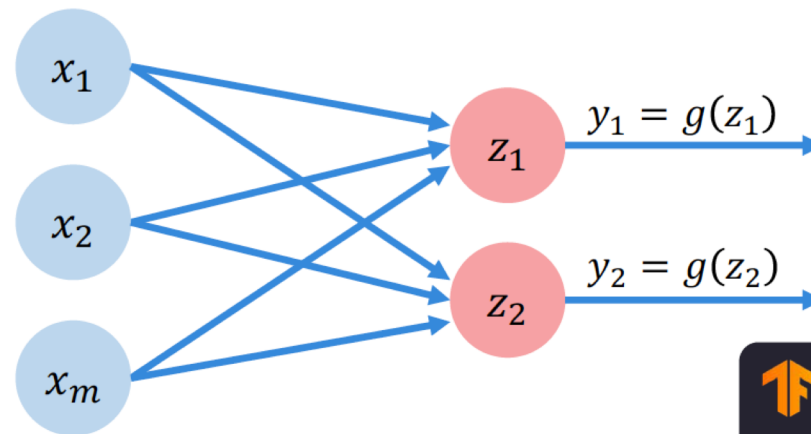
    def call(self, inputs):
        # Forward propagate the inputs
        z = tf.matmul(inputs, self.W) + self.b

        # Feed through a non-linear activation
        output = tf.math.sigmoid(z)

        return output
```

# Multi Output Perceptron

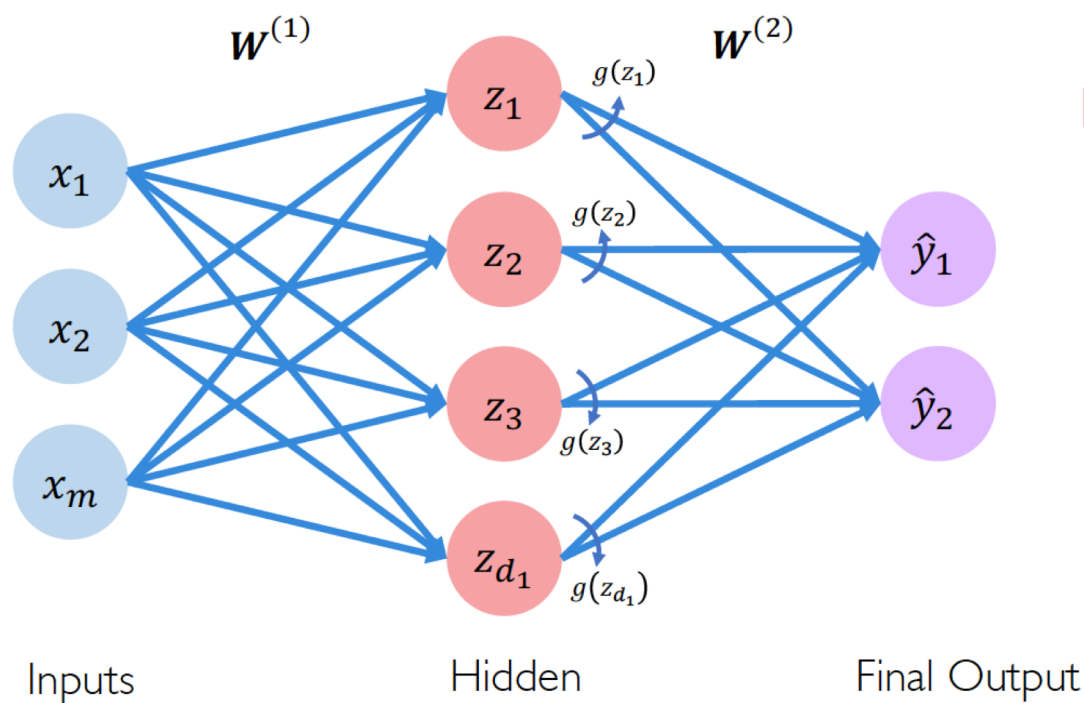
Because all inputs are densely connected to all outputs, these layers are called **Dense** layers



```
import tensorflow as tf  
  
layer = tf.keras.layers.Dense(  
    units=2)
```

$$z_i = w_{0,i} + \sum_{j=1}^m x_j w_{j,i}$$

# Single Layer Neural Network

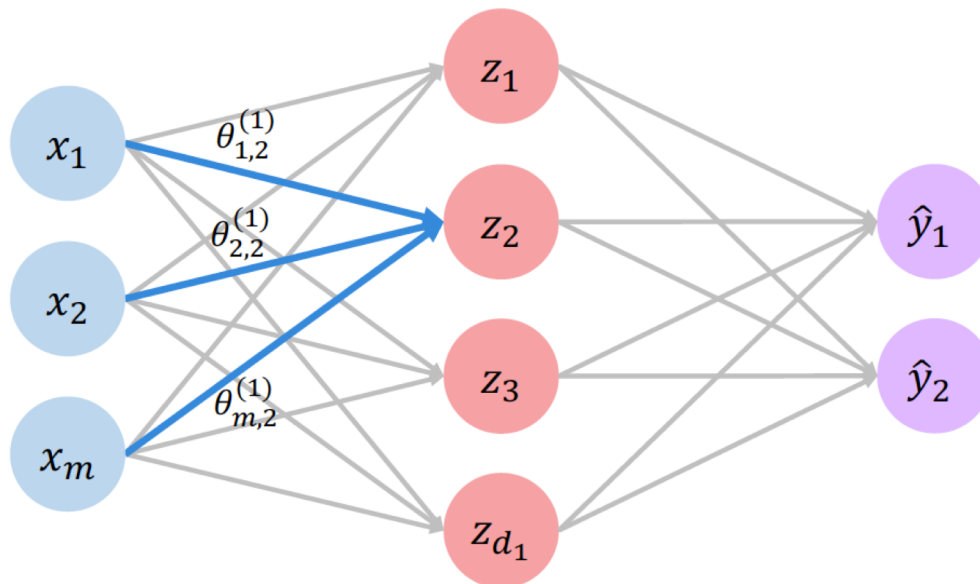


## Hidden layer(s):

- Not observable
- To be learned
- No specific behaviour enforced
- 2 weight matrices
- Same operation as before (dot product, bias, a.f.)

$$z_i = w_{0,i}^{(1)} + \sum_{j=1}^m x_j w_{j,i}^{(1)} \quad \hat{y}_i = g \left( w_{0,i}^{(2)} + \sum_{j=1}^{d_1} z_j w_{j,i}^{(2)} \right)$$

# Single Layer Neural Network

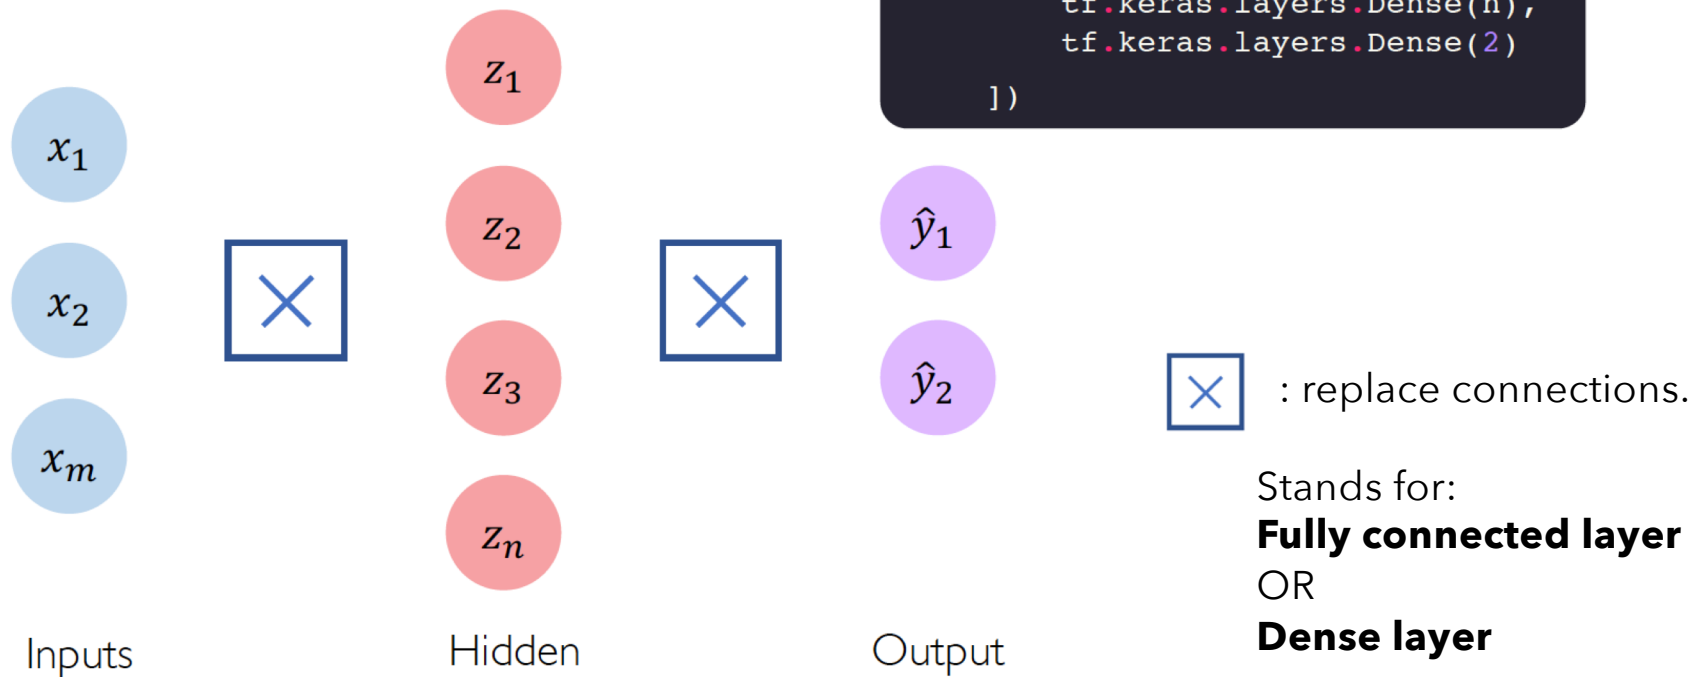


## Zoom in into a single hidden layer, say $z_2$ :

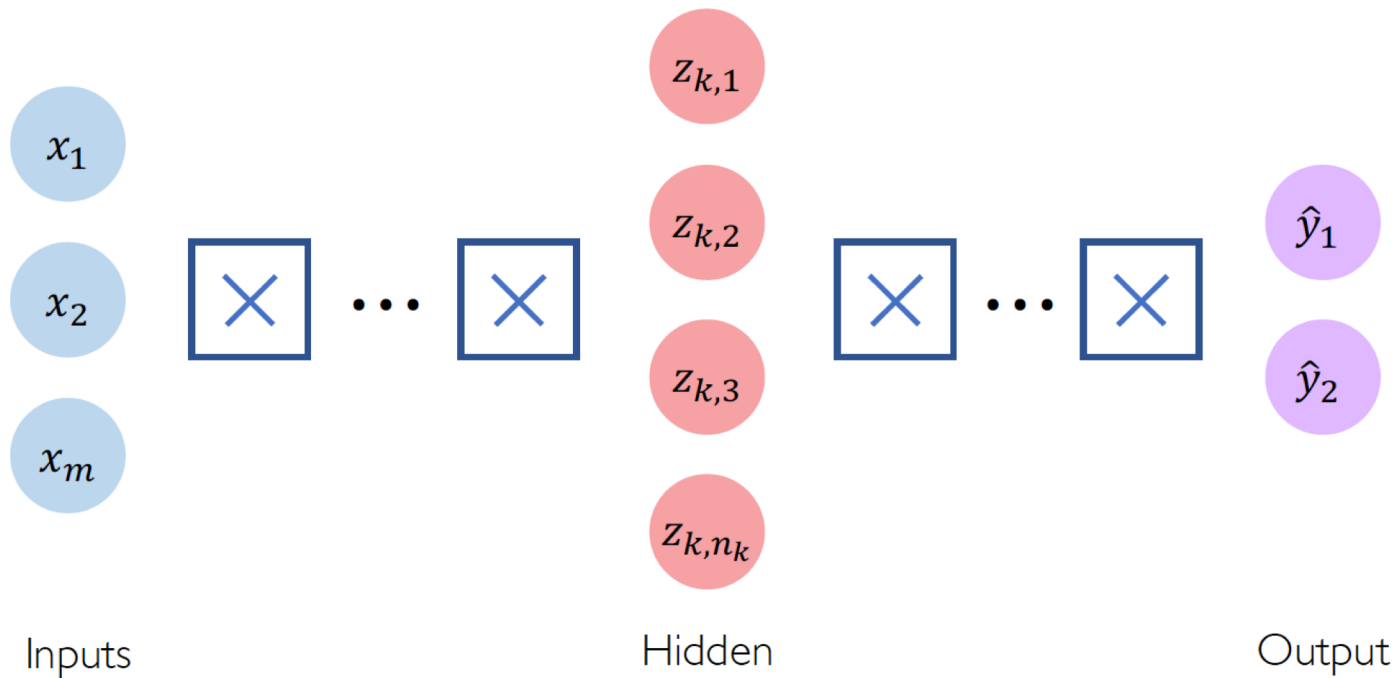
- Same operation as before (dot product, bias, a.f.)
- Same for  $z_3$ , what changes are the **weights**


$$\begin{aligned} z_2 &= w_{0,2}^{(1)} + \sum_{j=1}^m x_j w_{j,2}^{(1)} \\ &= w_{0,2}^{(1)} + x_1 w_{1,2}^{(1)} + x_2 w_{2,2}^{(1)} + x_m w_{m,2}^{(1)} \end{aligned}$$

# Multi Output Perceptron



# Deep Neural Network



```
  
import tensorflow as tf  
  
model = tf.keras.Sequential([  
    tf.keras.layers.Dense(n1),  
    tf.keras.layers.Dense(n2),  
    :  
    tf.keras.layers.Dense(2)  
])
```

$$z_{k,i} = w_{0,i}^{(k)} + \sum_{j=1}^{n_{k-1}} g(z_{k-1,j}) w_{j,i}^{(k)}$$

# Deep Neural Network

- Stack hidden layer back to back to back to create increasingly deeper and deeper models.
- Output computed going deeper into the NN and computing these weighted sums over and over and over again with these a.f. repeatedly applied



# Applying Neural Networks

# Example Problem

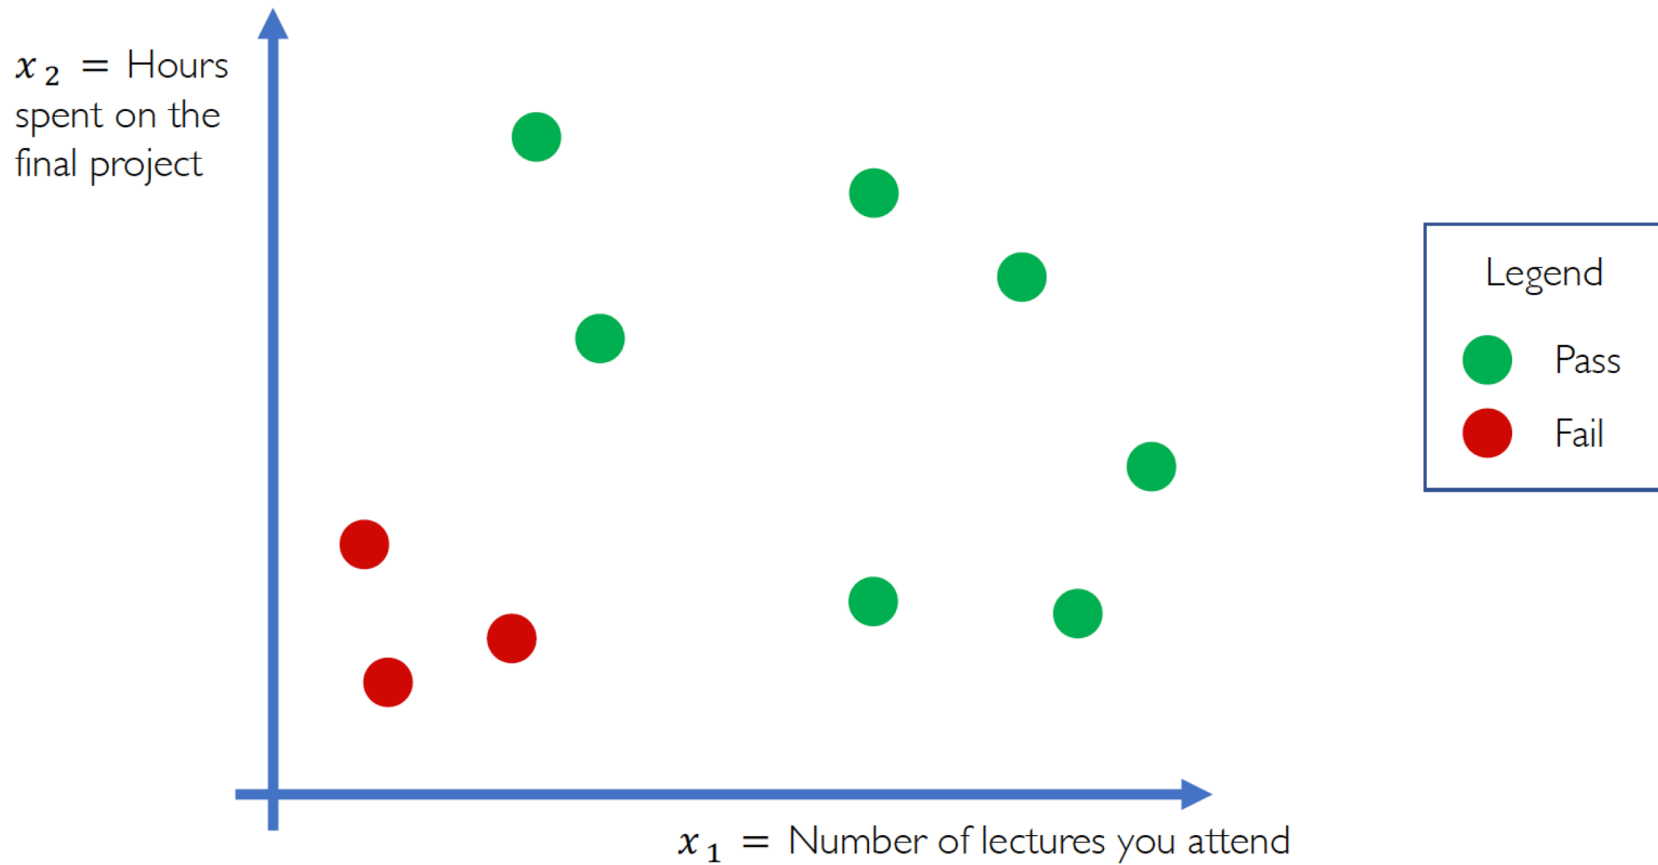
Will I pass this class?

Let's start with a simple two feature model

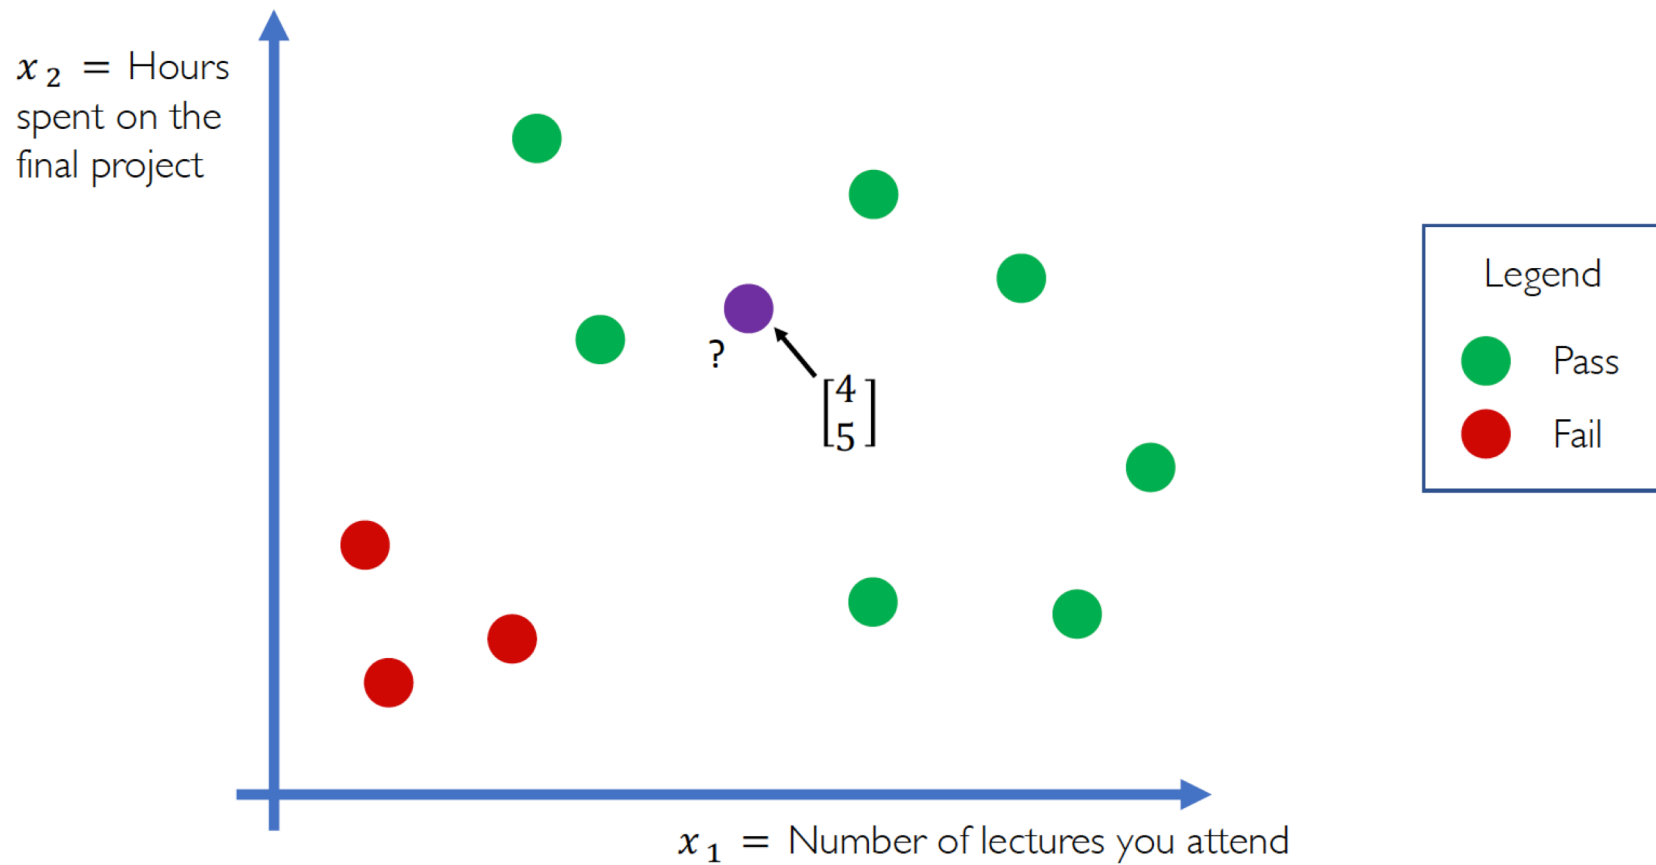
$x_1$  = Number of lectures you attend

$x_2$  = Hours spent on the final project

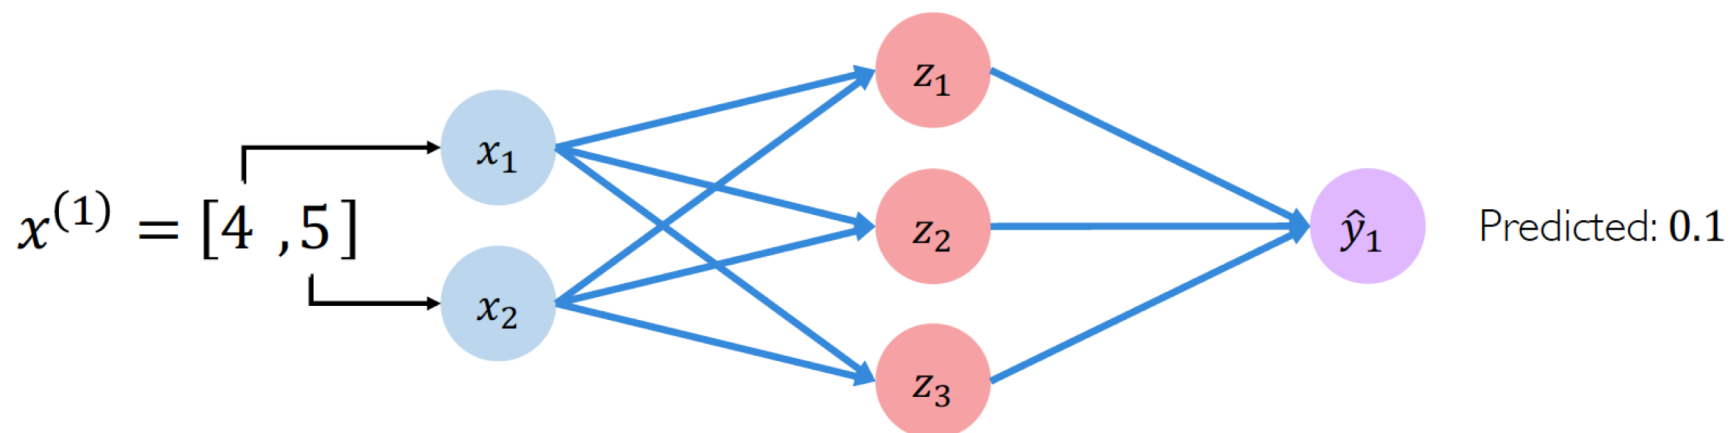
# Example Problem: Will I pass this class?



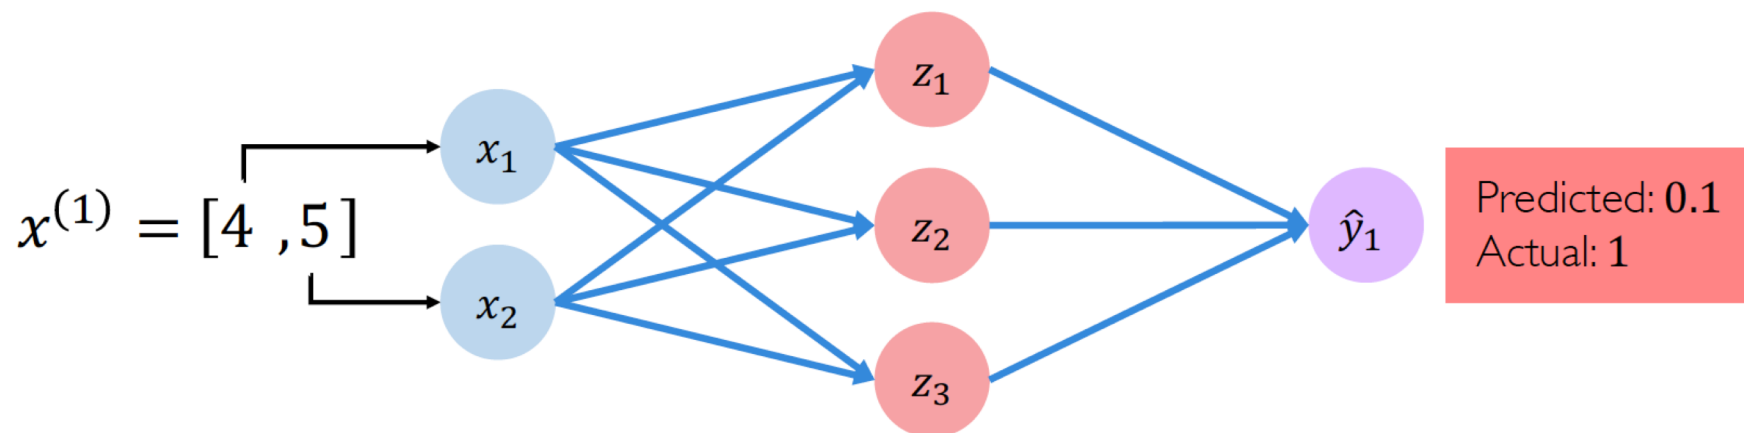
# Example Problem: Will I pass this class?



# Example Problem: Will I pass this class?



# Example Problem: Will I pass this class?



# Example Problem: Will I pass the exam?

Why **Wrong prediction**?

➤ Because the network is not trained

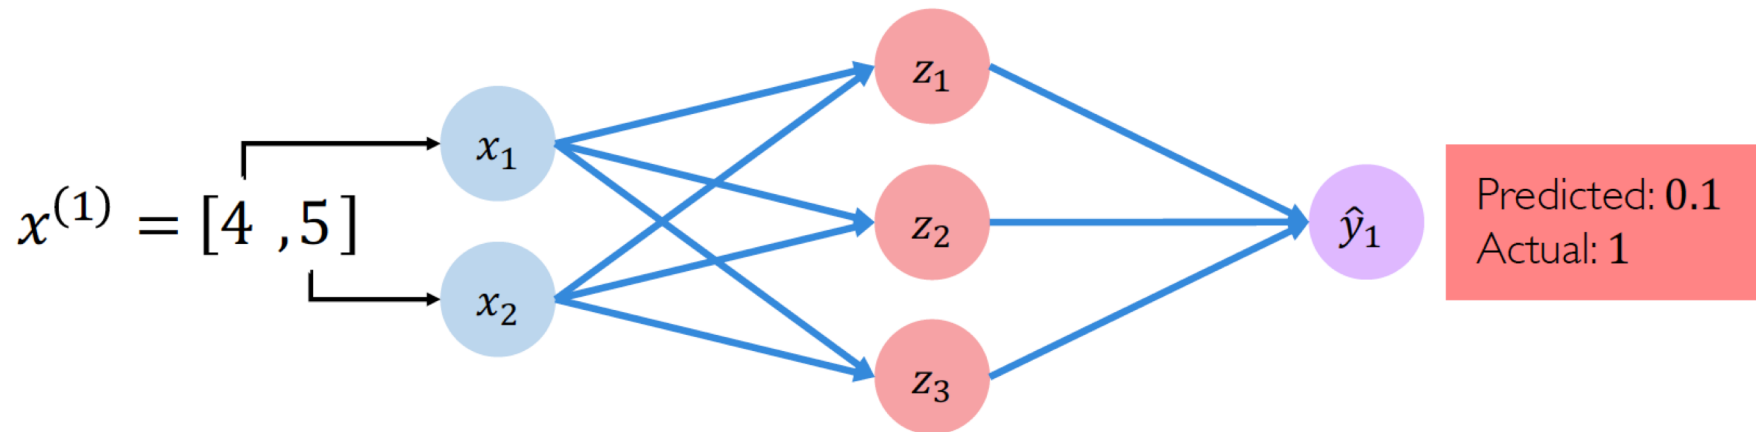
**Train a network:** teach it to get the right answer. How?

➤ tell it when it makes a mistake, so to correct it in the future

The **Loss** of a network is what quantify the wrong prediction

# Quantifying Loss

The **loss** of our network measures the cost incurred from incorrect predictions



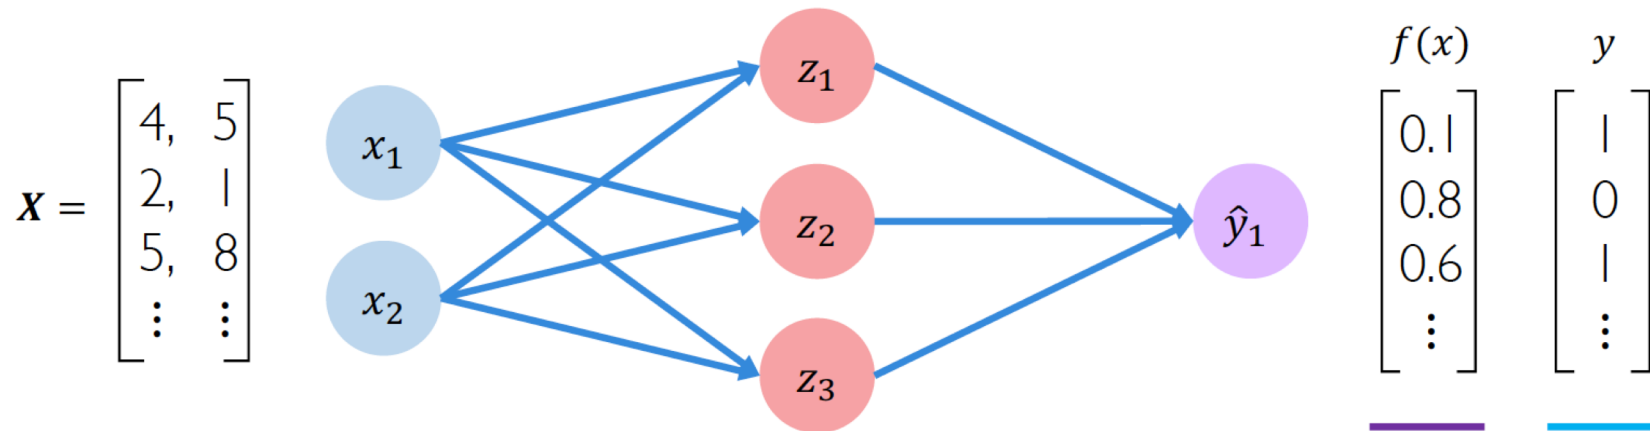
$$\mathcal{L}(\underbrace{f(x^{(i)}; \mathbf{W})}_{\text{Predicted}}, \underbrace{y^{(i)}}_{\text{Actual}})$$



# Empirical Loss

When we train a network, we do not want to minimize the loss for a particular student, but **the loss across the entire training set**

The **empirical loss** measures the total loss over our entire dataset



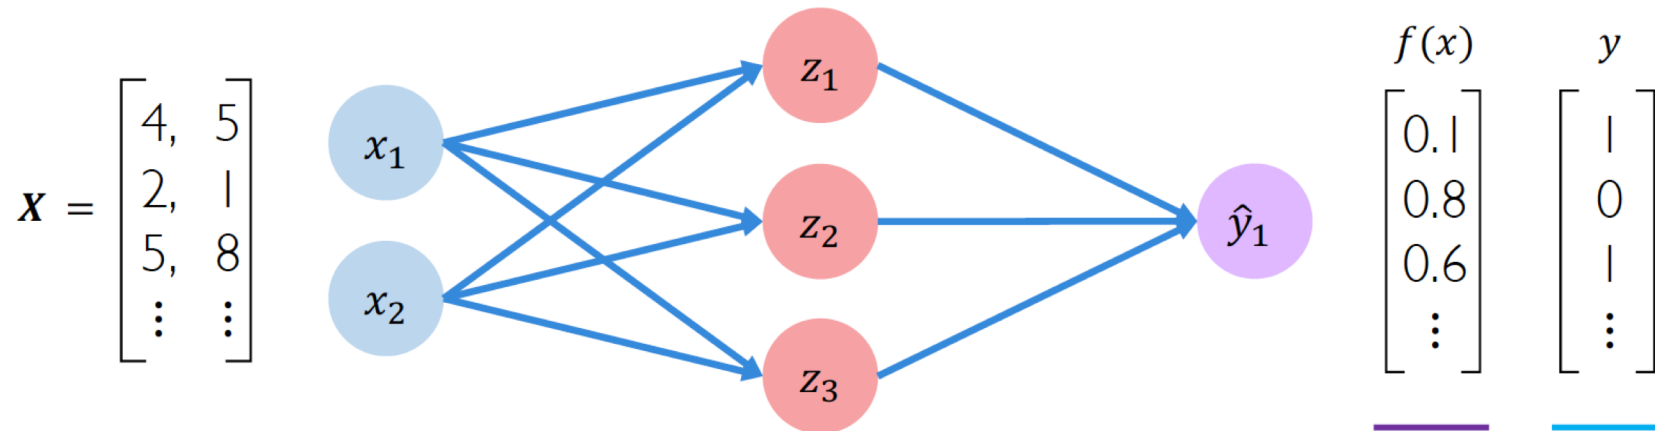
Also known as:

- Objective function
- Cost function
- Empirical Risk

$$J(W) = \frac{1}{n} \sum_{i=1}^n \mathcal{L}(\underbrace{f(x^{(i)}; W)}_{\text{Predicted}}, \underbrace{y^{(i)}}_{\text{Actual}})$$

# Binary Cross Entropy Loss

*Cross entropy loss can be used with models that output a probability between 0 and 1*



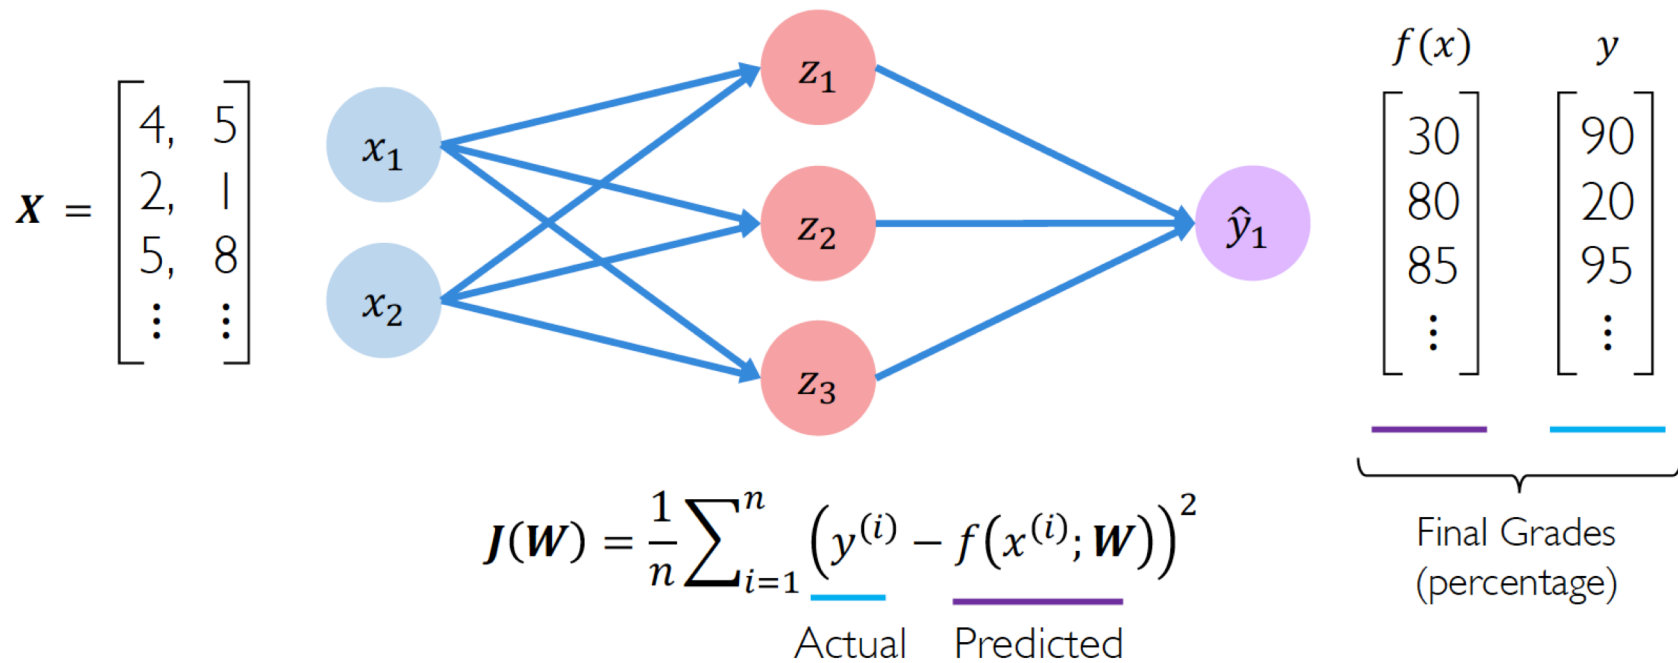
$$J(\mathbf{W}) = \frac{1}{n} \sum_{i=1}^n \underbrace{y^{(i)}}_{\text{Actual}} \log \left( \underbrace{f(x^{(i)}; \mathbf{W})}_{\text{Predicted}} \right) + (1 - \underbrace{y^{(i)}}_{\text{Actual}}) \log \left( 1 - \underbrace{f(x^{(i)}; \mathbf{W})}_{\text{Predicted}} \right)$$



```
loss = tf.reduce_mean( tf.nn.softmax_cross_entropy_with_logits(y, predicted) )
```

# Mean Squared Error Loss

*Mean squared error loss can be used with regression models that output continuous real numbers*



```
loss = tf.reduce_mean( tf.square(tf.subtract(y, predicted)) )
```

# Training Neural Networks

# Loss Optimization

We want to find the network weights that **achieve the lowest loss**

$$W^* = \operatorname{argmin}_W \frac{1}{n} \sum_{i=1}^n \mathcal{L}(f(x^{(i)}; W), y^{(i)})$$

$$W^* = \operatorname{argmin}_W J(W)$$

# Loss Optimization

We want to find the network weights that **achieve the lowest loss**

$$\mathbf{W}^* = \operatorname{argmin}_{\mathbf{W}} \frac{1}{n} \sum_{i=1}^n \mathcal{L}(f(x^{(i)}; \mathbf{W}), y^{(i)})$$

$$\mathbf{W}^* = \operatorname{argmin}_{\mathbf{W}} J(\mathbf{W})$$



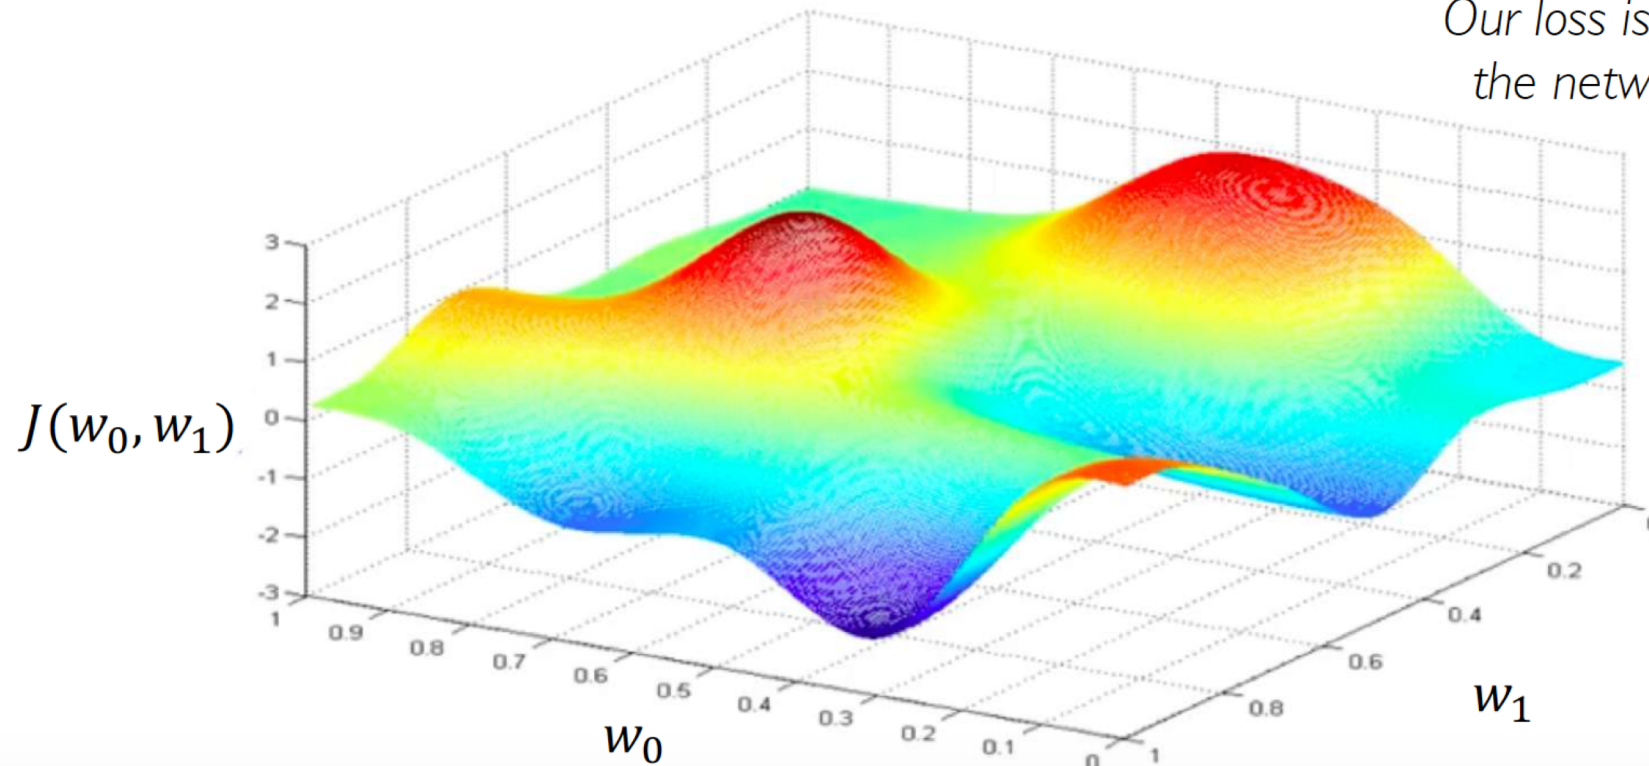
Remember:

$$\mathbf{W} = \{\mathbf{W}^{(0)}, \mathbf{W}^{(1)}, \dots\}$$

# Loss Optimization

$$W^* = \operatorname{argmin}_W J(W)$$

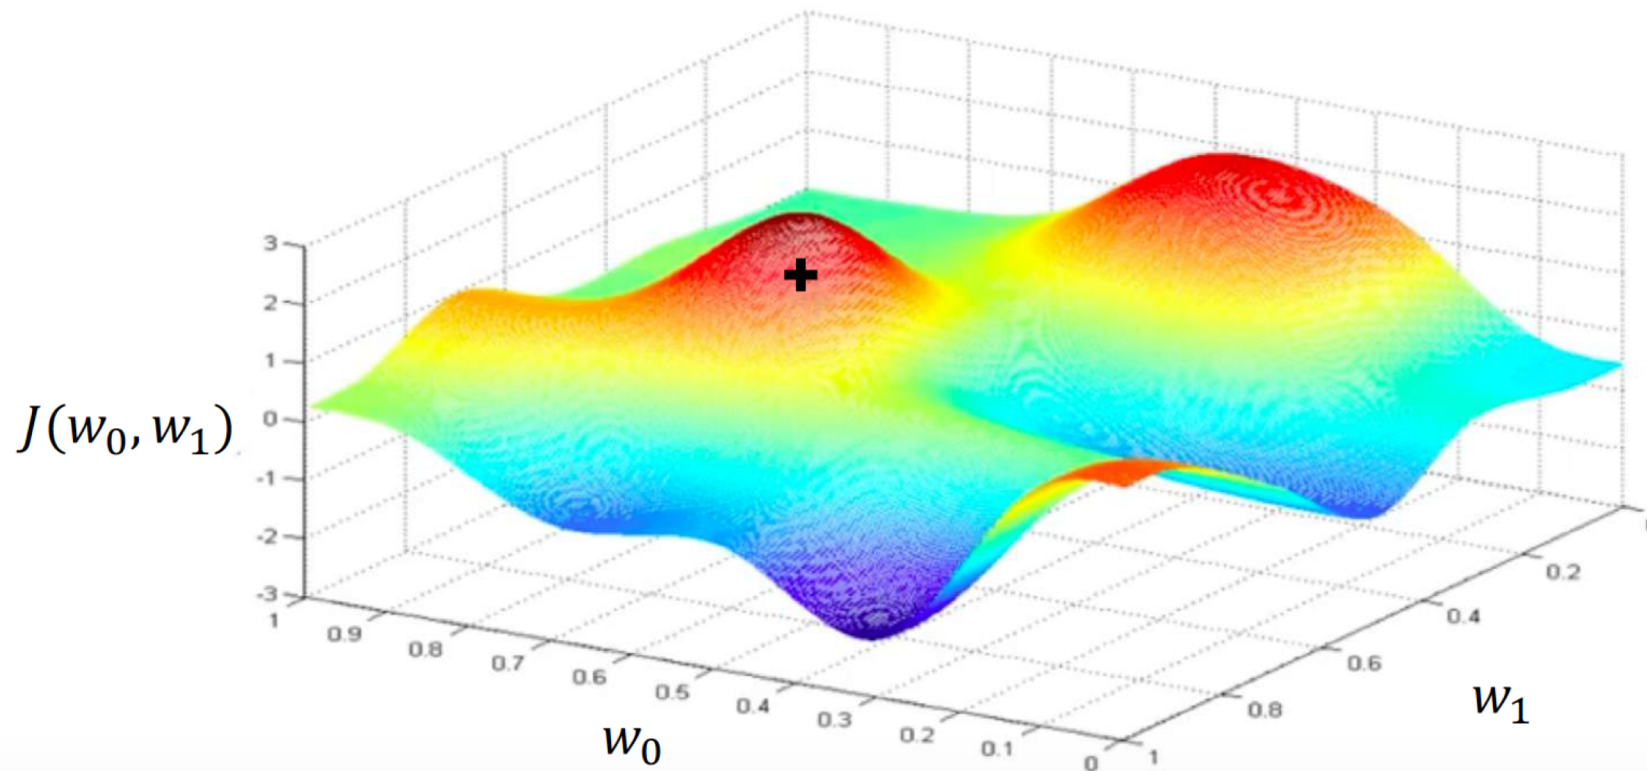
Remember:  
*Our loss is a function of  
the network weights!*



# Loss Optimization

Loss optimization through **gradient descent**

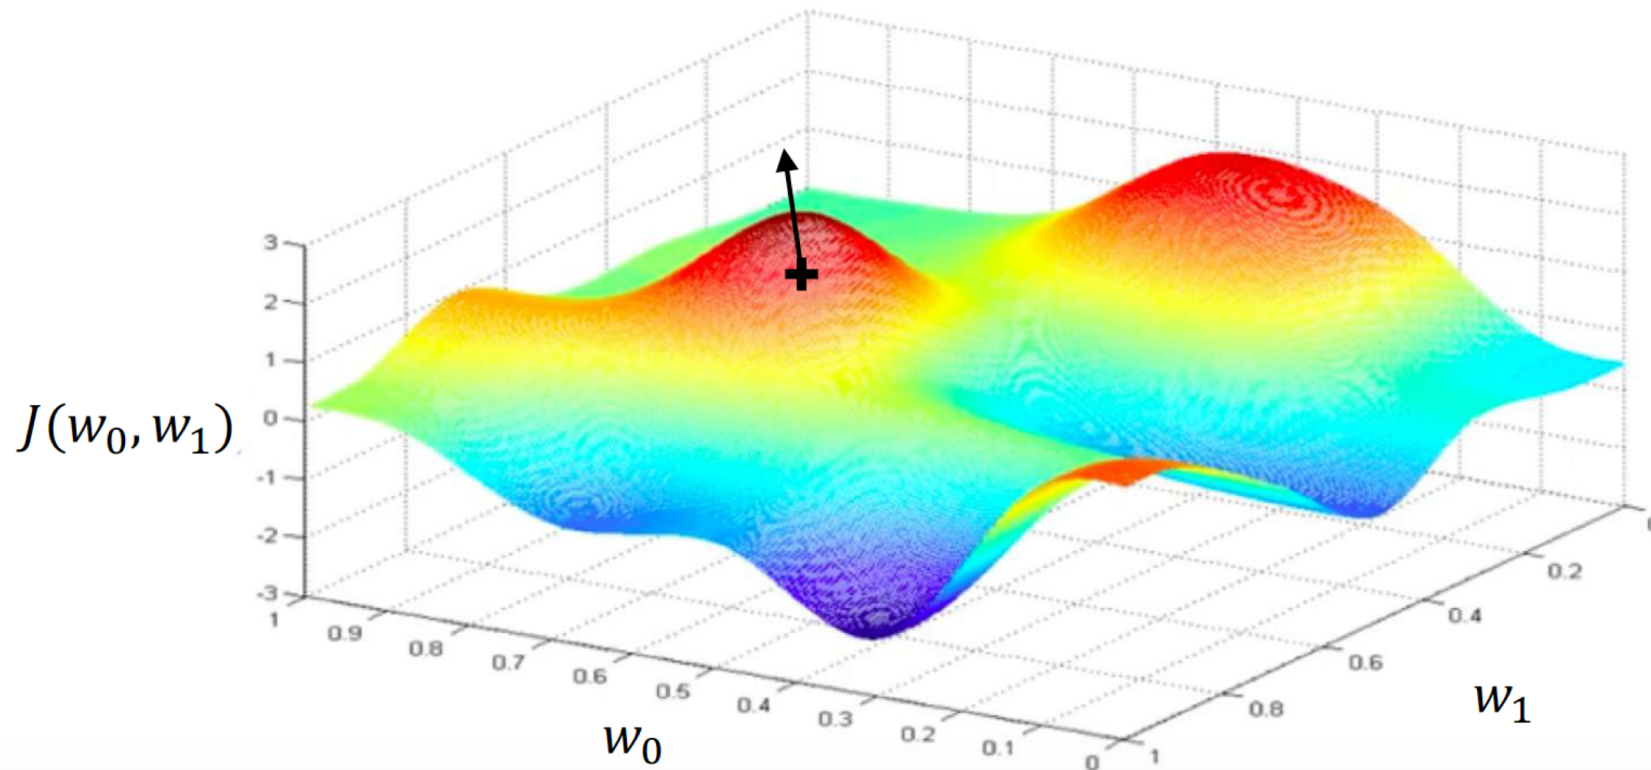
Randomly pick an initial  $(w_0, w_1)$





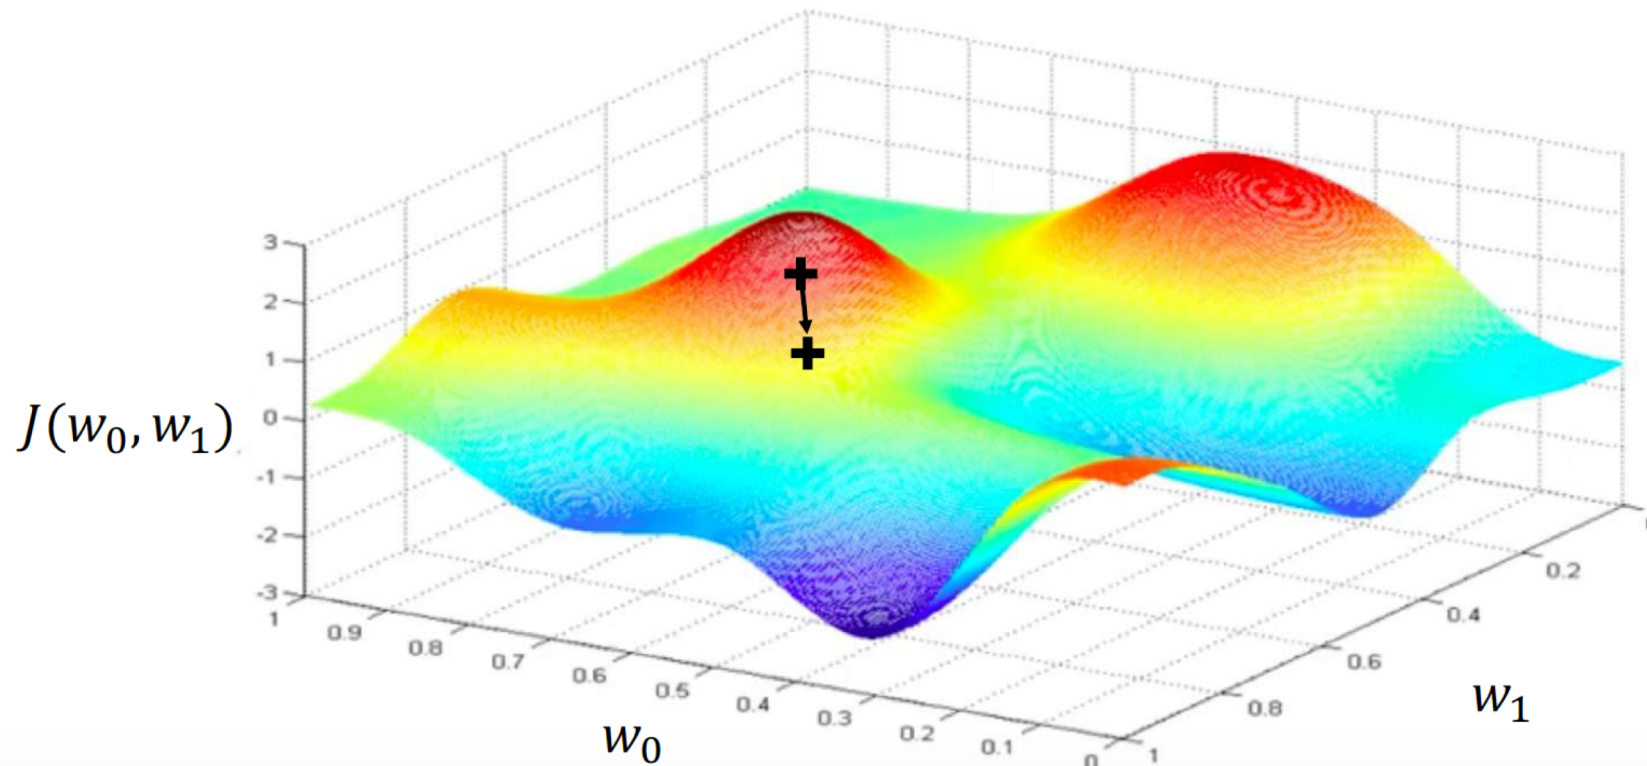
# Loss Optimization

Compute gradient,  $\frac{\partial J(W)}{\partial W}$



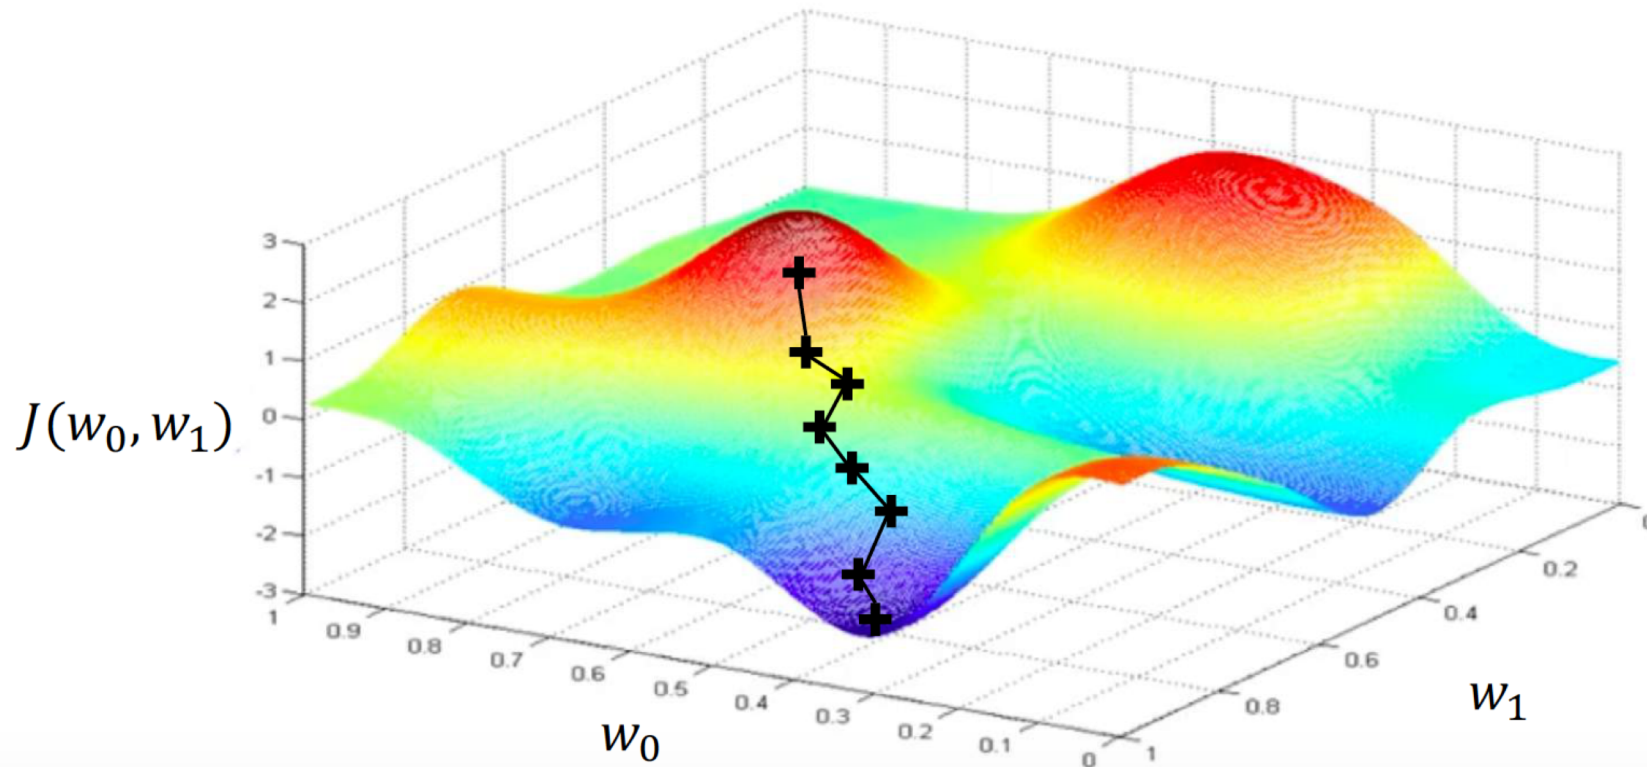
# Loss Optimization

Take small step in opposite direction of gradient



# Gradient Descent

Repeat until convergence



# Gradient Descent

## Algorithm

1. Initialize weights randomly  $\sim \mathcal{N}(0, \sigma^2)$
2. Loop until convergence:
3. Compute gradient,  $\frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
4. Update weights,  $\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
5. Return weights

# Gradient Descent

## Algorithm

1. Initialize weights randomly  $\sim \mathcal{N}(0, \sigma^2)$
2. Loop until convergence:
3. Compute gradient,  $\frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
4. Update weights,  $\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
5. Return weights



```
import tensorflow as tf

weights = tf.Variable([tf.random.normal()])

while True:    # loop forever
    with tf.GradientTape() as g:
        loss = compute_loss(weights)
        gradient = g.gradient(loss, weights)

    weights = weights - lr * gradient
```

# Gradient Descent

## Algorithm

1. Initialize weights randomly  $\sim \mathcal{N}(0, \sigma^2)$
2. Loop until convergence:
3. Compute gradient,  $\frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
4. Update weights,  $\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
5. Return weights



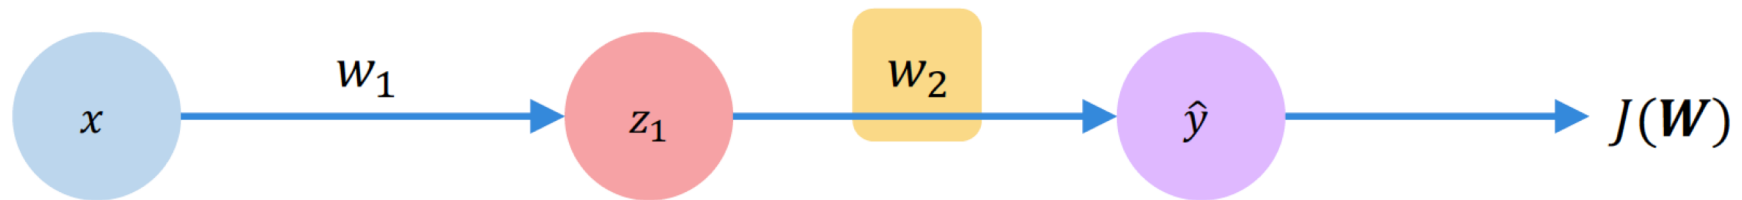
```
import tensorflow as tf

weights = tf.Variable([tf.random.normal()])

while True:    # loop forever
    with tf.GradientTape() as g:
        loss = compute_loss(weights)
        gradient = g.gradient(loss, weights)

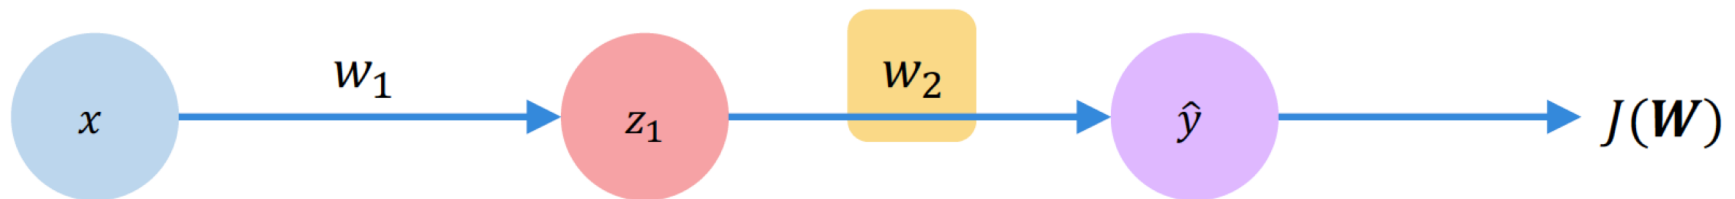
    weights = weights - lr * gradient
```

# Computing Gradients: Backpropagation



*How does a small change in one weight (ex.  $w_2$ ) affect the final loss  $J(\mathbf{W})$ ?*

# Computing Gradients: Backpropagation

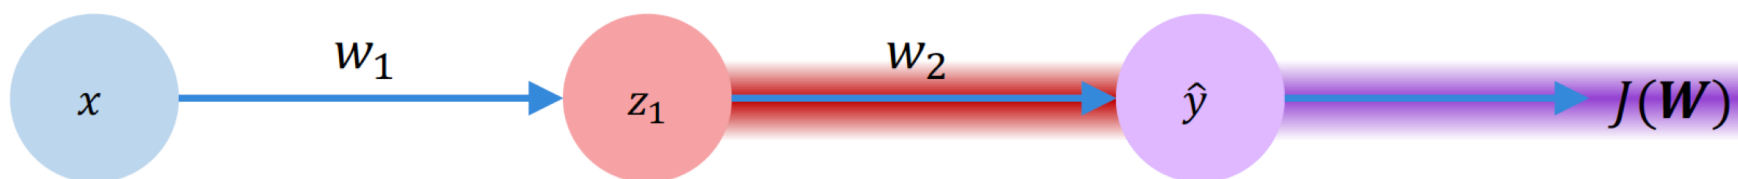


$$\frac{\partial J(\mathbf{W})}{\partial w_2} =$$

Let's use the chain rule!

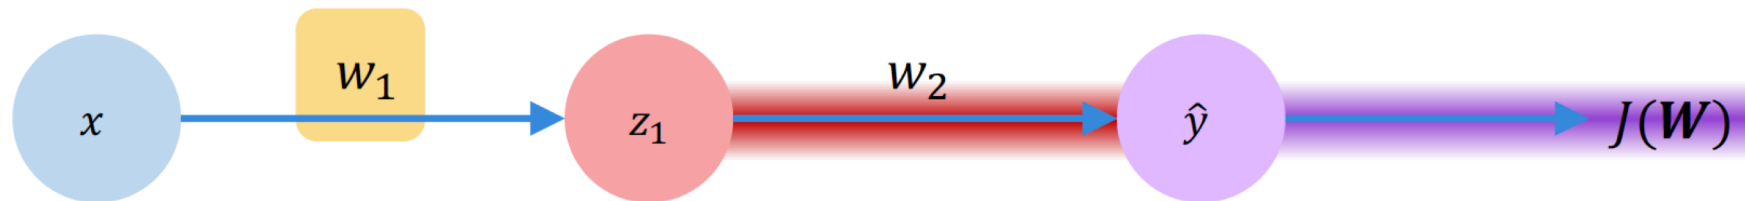


# Computing Gradients: Backpropagation



$$\frac{\partial J(\mathbf{W})}{\partial w_2} = \underbrace{\frac{\partial J(\mathbf{W})}{\partial \hat{y}}}_{\text{purple}} * \underbrace{\frac{\partial \hat{y}}{\partial w_2}}_{\text{red}}$$

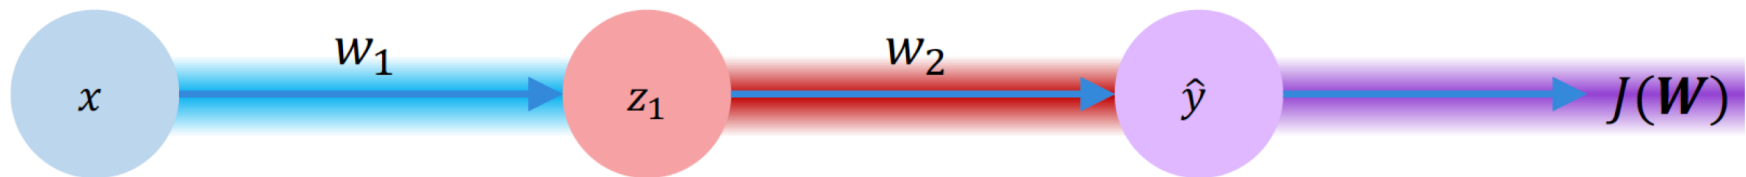
# Computing Gradients: Backpropagation



$$\frac{\partial J(W)}{\partial w_1} = \frac{\partial J(W)}{\partial \hat{y}} * \frac{\partial \hat{y}}{\partial w_1}$$

Apply chain rule!                      Apply chain rule!

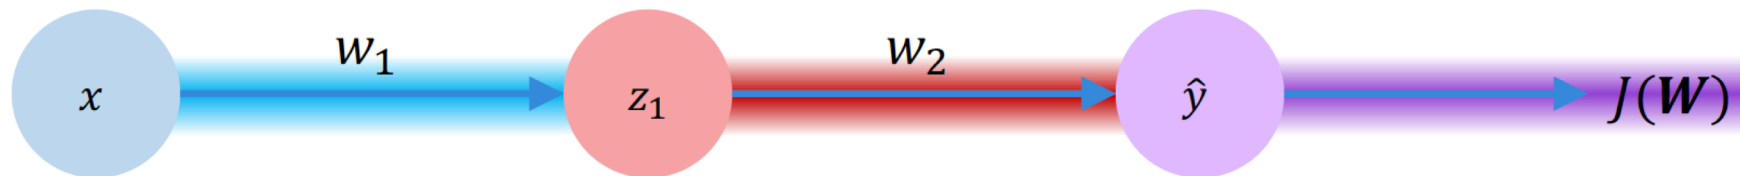
# Computing Gradients: Backpropagation



$$\frac{\partial J(W)}{\partial w_1} = \frac{\partial J(W)}{\partial \hat{y}} * \frac{\partial \hat{y}}{\partial z_1} * \frac{\partial z_1}{\partial w_1}$$



# Computing Gradients: Backpropagation

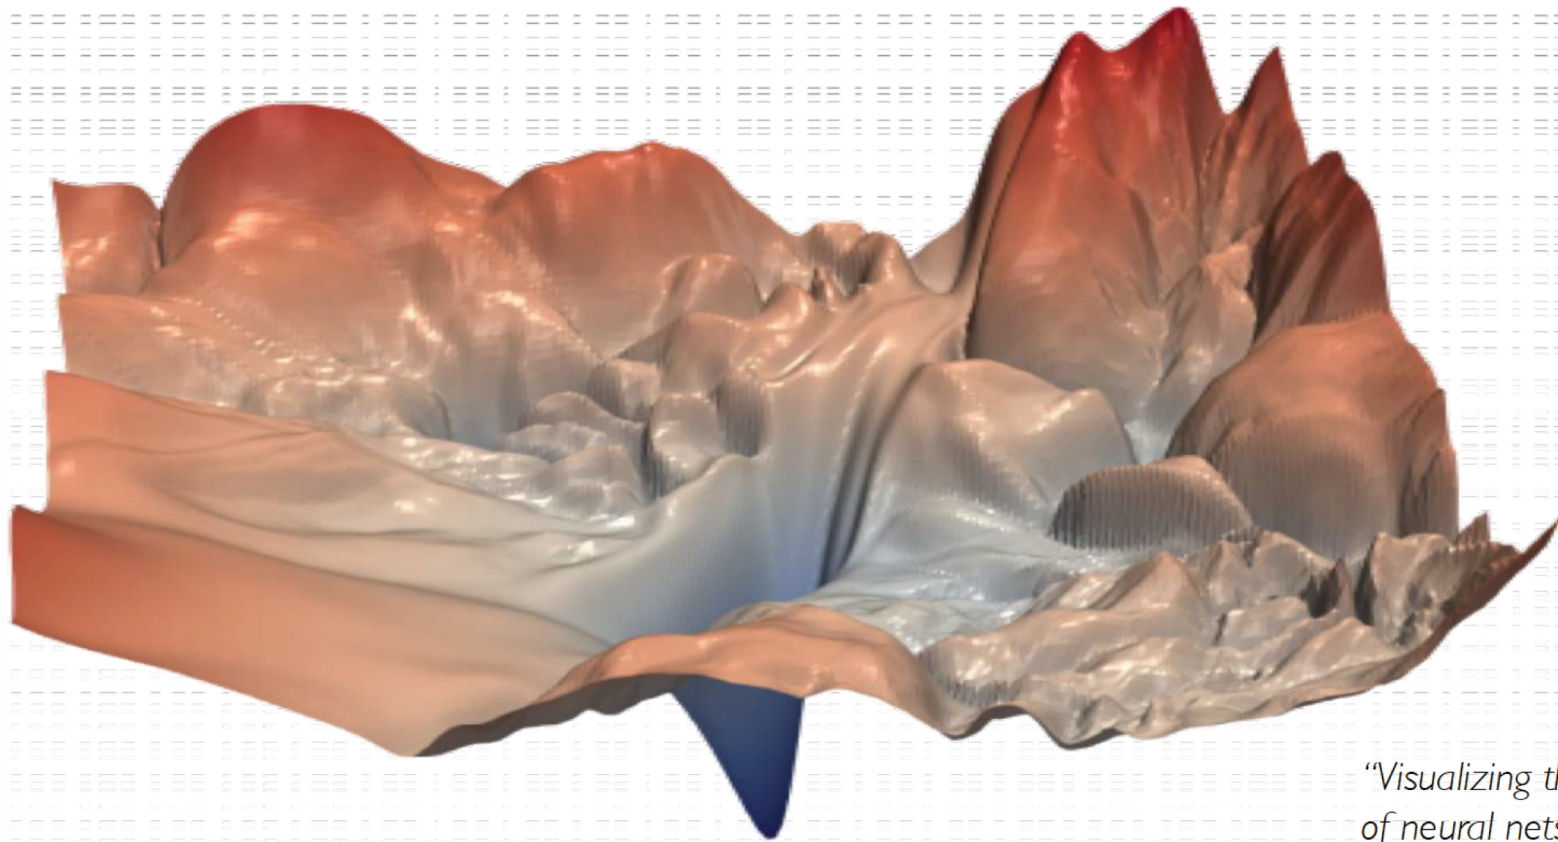


$$\frac{\partial J(W)}{\partial w_1} = \frac{\partial J(W)}{\partial \hat{y}} * \frac{\partial \hat{y}}{\partial z_1} * \frac{\partial z_1}{\partial w_1}$$

Repeat this for **every weight in the network** using gradients from later layers

# Neural Networks in practice: Optimization

# Training Neural Networks is Difficult



*“Visualizing the loss landscape of neural nets”. Dec 2017.*

# Loss Functions Can Be Difficult to Optimize

## Remember:

Optimization through gradient descent

$$\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$$

# Loss Functions Can Be Difficult to Optimize

## Remember:

Optimization through gradient descent

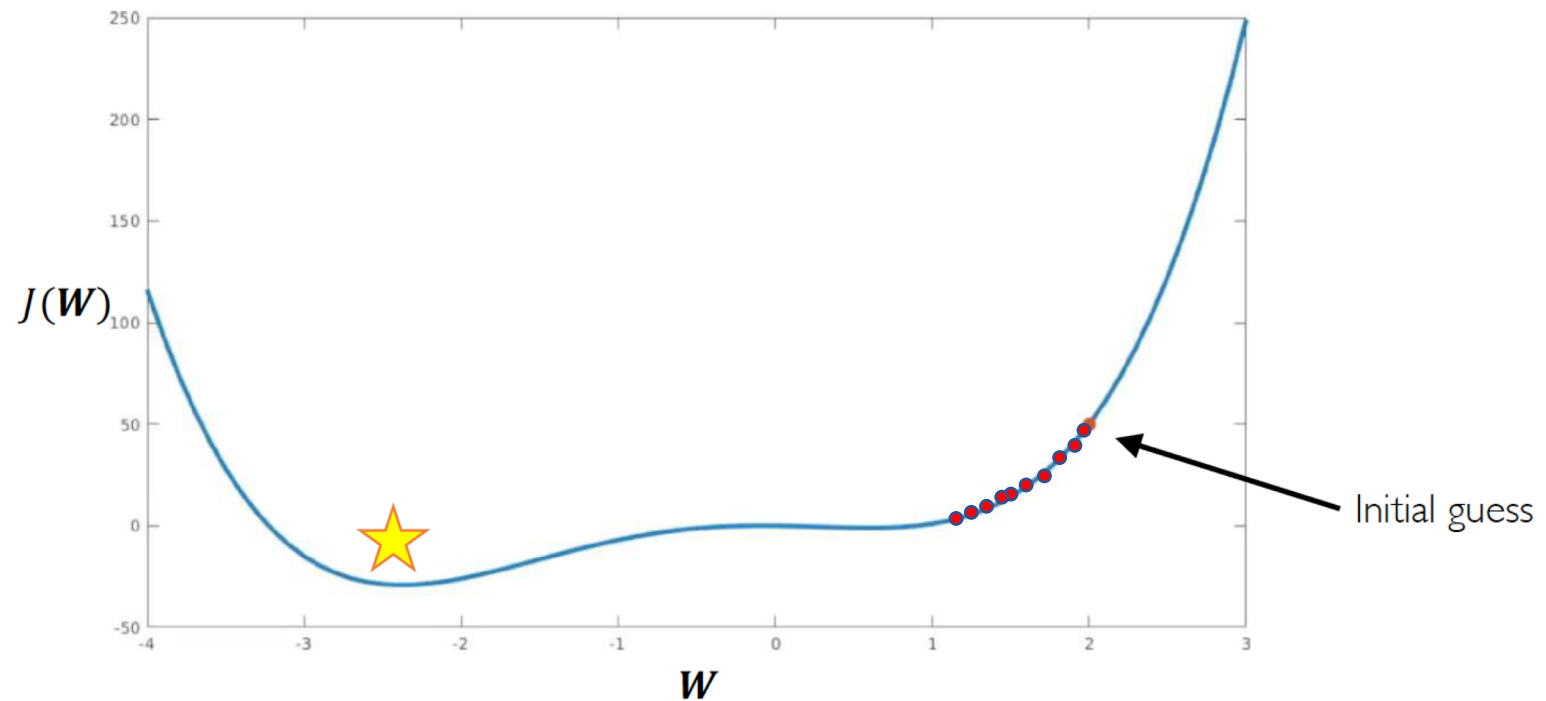
$$W \leftarrow W - \eta \frac{\partial J(W)}{\partial W}$$

How can we set the  
learning rate?



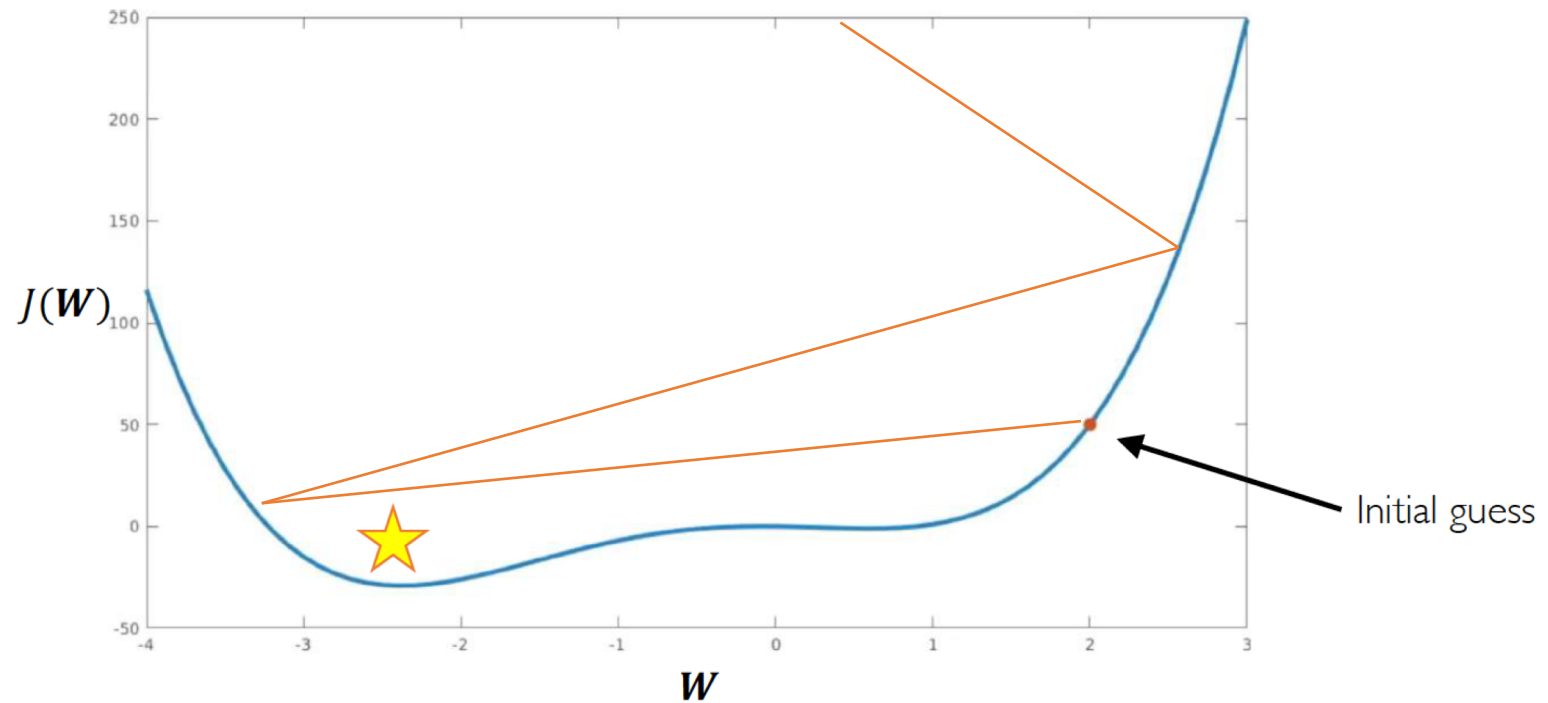
# Setting the Learning Rate

*Small learning rate converges slowly and gets stuck in false local minima*



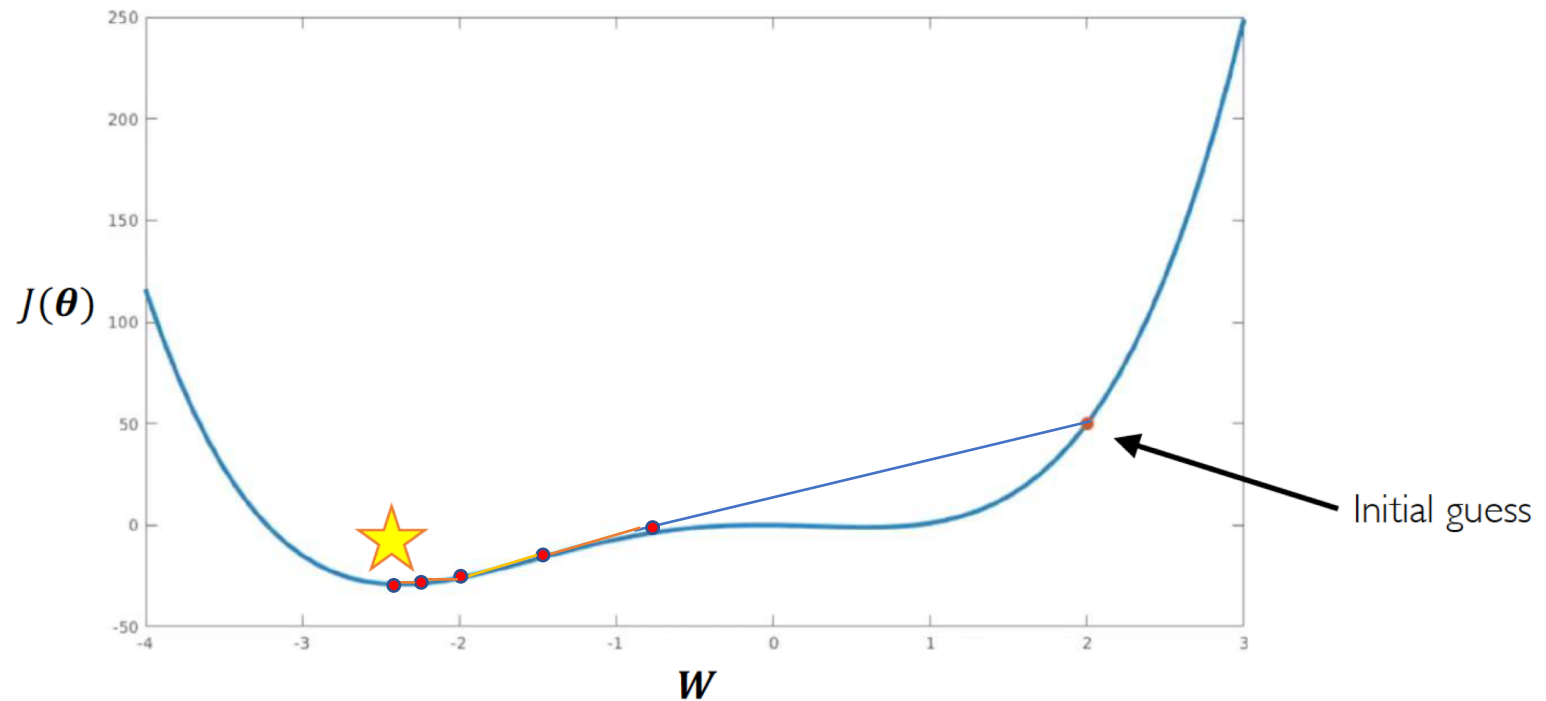
# Setting the Learning Rate

*Large learning rates overshoot, become unstable and diverge*



# Setting the Learning Rate

*Stable learning rates converge smoothly and avoid local minima*



# How to deal with this?

## Idea 1:

Try lots of different learning rates and see what works “just right”

---

# How to deal with this?

## Idea 1:

Try lots of different learning rates and see what works “just right”

## Idea 2:






Do something smarter!

Design an adaptive learning rate that “adapts” to the landscape

# Adaptive Learning Rates

- Learning rates are non longer fixed
- Can be made larger or smaller depending on:
  - How large gradient is
  - How fast learning is happening
  - Size of particular weights
  - ...

# Gradient Descent Algorithms

Algorithm	TF Implementation	Reference
• SGD	 <code>tf.keras.optimizers.SGD</code>	Kiefer & Wolfowitz. "Stochastic Estimation of the Maximum of a Regression Function." 1952.
• Adam	 <code>tf.keras.optimizers.Adam</code>	Kingma et al. "Adam: A Method for Stochastic Optimization." 2014.
• Adadelta	 <code>tf.keras.optimizers.Adadelta</code>	Zeiler et al. "ADADELTA: An Adaptive Learning Rate Method." 2012.
• Adagrad	 <code>tf.keras.optimizers.Adagrad</code>	Duchi et al. "Adaptive Subgradient Methods for Online Learning and Stochastic Optimization." 2011.
• RMSProp	 <code>tf.keras.optimizers.RMSProp</code>	

Additional details: <http://ruder.io/optimizing-gradient-descent/>

# Putting it all together



```
import tensorflow as tf

model = tf.keras.Sequential([...])

# pick your favorite optimizer
optimizer = tf.keras.optimizer.SGD()

while True: # loop forever

    # forward pass through the network
    prediction = model(x)

    with tf.GradientTape() as tape:
        # compute the loss
        loss = compute_loss(y, prediction)

    # update the weights using the gradient
    grads = tape.gradient(loss, model.trainable_variables)
    optimizer.apply_gradients(zip(grads, model.trainable_variables))
```



Can replace with any TensorFlow optimizer!

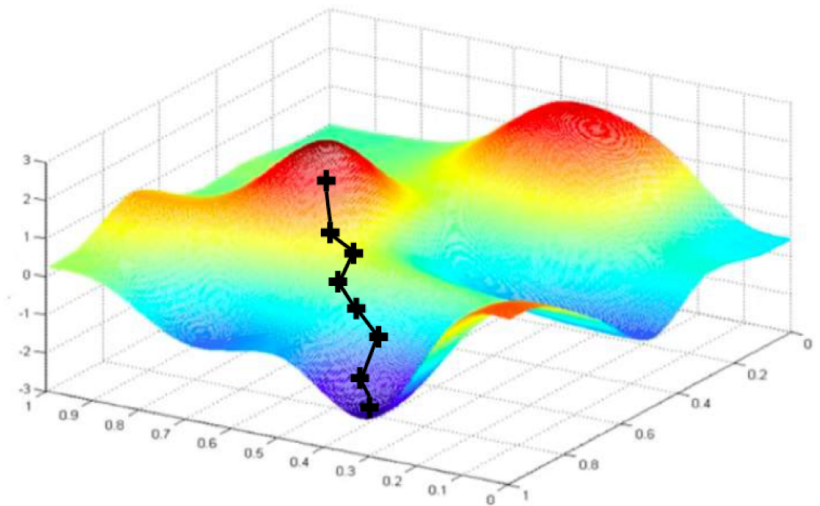


# Neural Networks in practice: Mini-batches

# Gradient Descent

## Algorithm

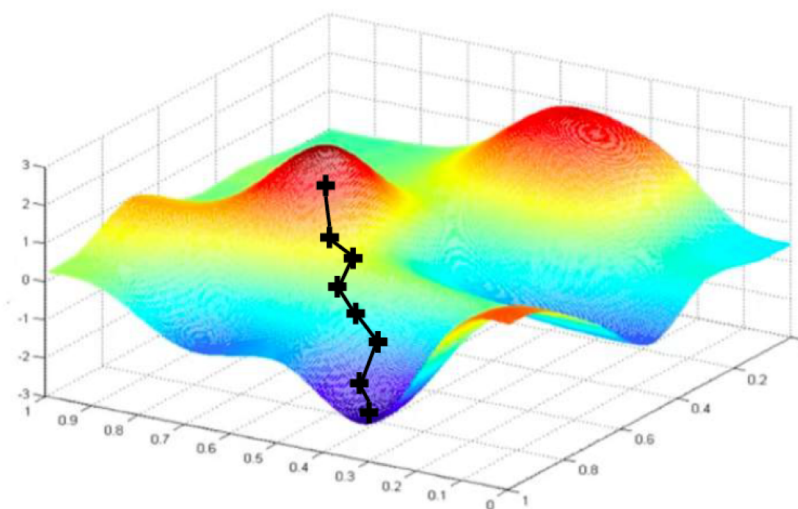
1. Initialize weights randomly  $\sim \mathcal{N}(0, \sigma^2)$
2. Loop until convergence:
3.     Compute gradient,  $\frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
4.     Update weights,  $\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
5. Return weights



# Gradient Descent

## Algorithm

1. Initialize weights randomly  $\sim \mathcal{N}(0, \sigma^2)$
2. Loop until convergence:
3. Compute gradient,  $\frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
4. Update weights,  $\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
5. Return weights

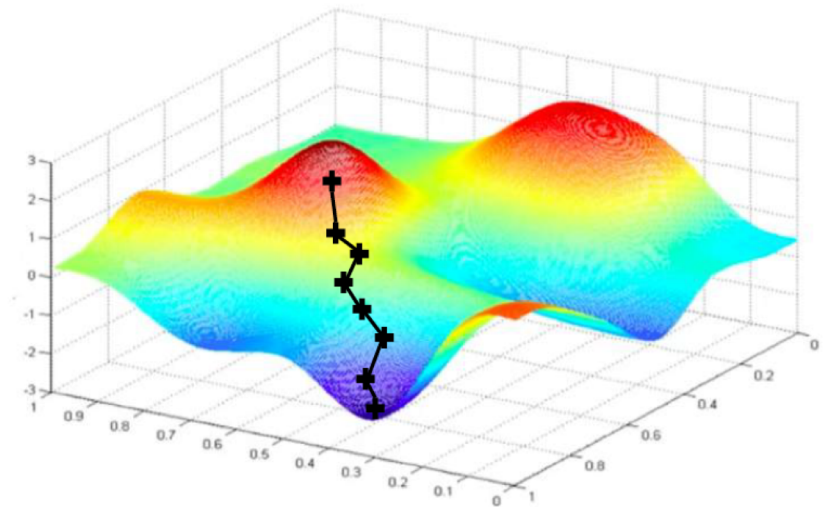


Can be very  
computational to  
compute!

# Stochastic Gradient Descent

## Algorithm

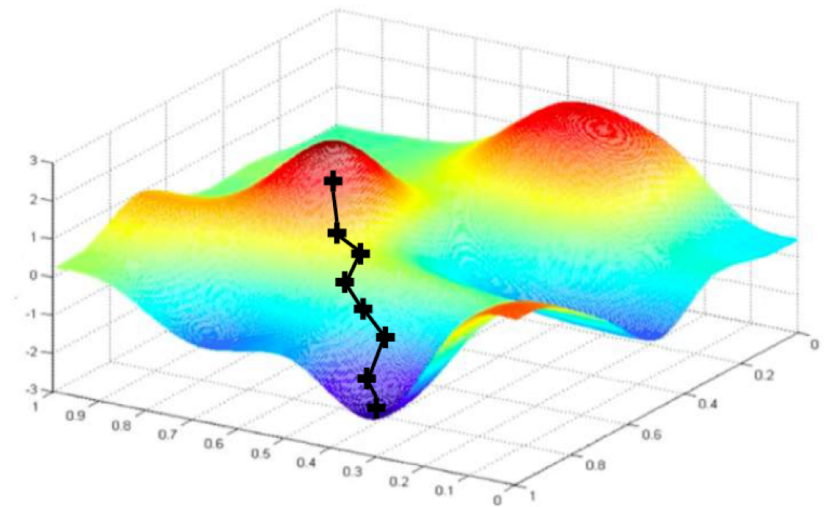
1. Initialize weights randomly  $\sim \mathcal{N}(0, \sigma^2)$
2. Loop until convergence:
  3. Pick single data point  $i$
  4. Compute gradient,  $\frac{\partial J_i(\mathbf{W})}{\partial \mathbf{W}}$
  5. Update weights,  $\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
6. Return weights



# Stochastic Gradient Descent

## Algorithm

1. Initialize weights randomly  $\sim \mathcal{N}(0, \sigma^2)$
2. Loop until convergence:
3. Pick single data point  $i$
4. Compute gradient,  $\frac{\partial J_i(\mathbf{W})}{\partial \mathbf{W}}$
5. Update weights,  $\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
6. Return weights

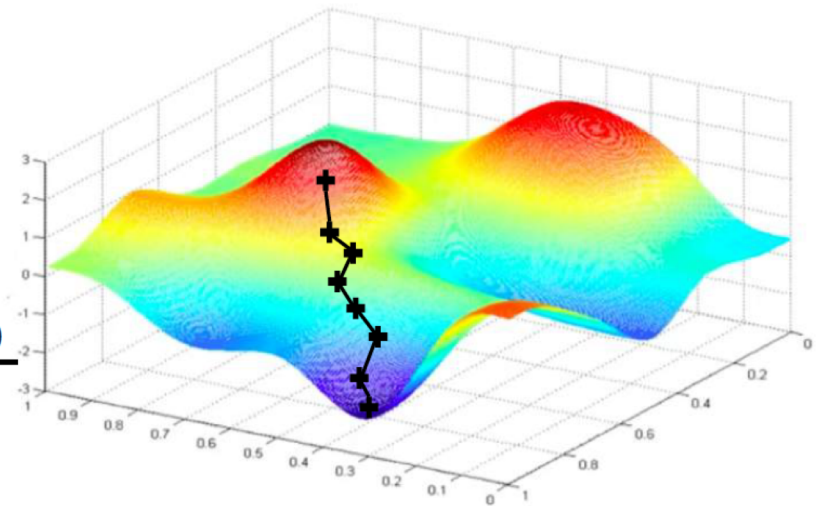


Easy to compute but  
**very noisy**  
(stochastic)!

# Stochastic Gradient Descent

## Algorithm

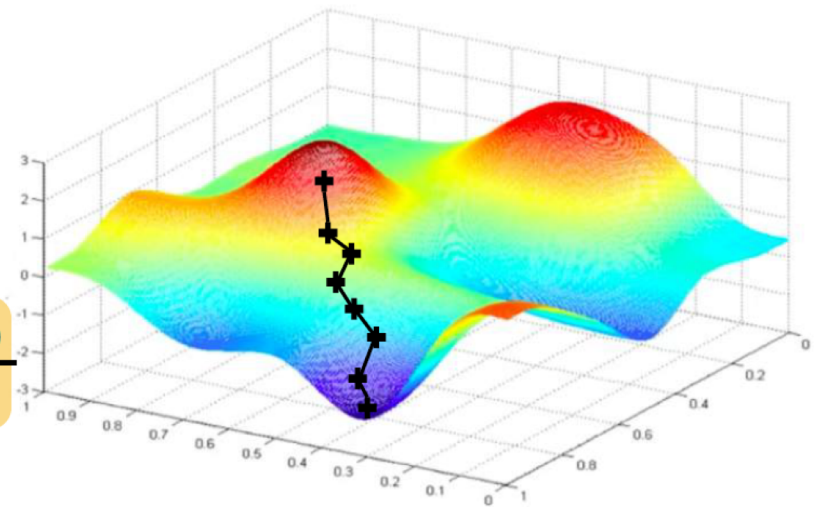
1. Initialize weights randomly  $\sim \mathcal{N}(0, \sigma^2)$
2. Loop until convergence:
3. Pick batch of  $B$  data points
4. Compute gradient,  $\frac{\partial J(\mathbf{W})}{\partial \mathbf{W}} = \frac{1}{B} \sum_{k=1}^B \frac{\partial J_k(\mathbf{W})}{\partial \mathbf{W}}$
5. Update weights,  $\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
6. Return weights



# Stochastic Gradient Descent

## Algorithm

1. Initialize weights randomly  $\sim \mathcal{N}(0, \sigma^2)$
2. Loop until convergence:
3. Pick batch of  $B$  data points
4. Compute gradient,  $\frac{\partial J(\mathbf{W})}{\partial \mathbf{W}} = \frac{1}{B} \sum_{k=1}^B \frac{\partial J_k(\mathbf{W})}{\partial \mathbf{W}}$
5. Update weights,  $\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
6. Return weights



Fast to compute and a much better estimate of the true gradient!

# Mini-batches while training

- More accurate estimation of gradient

Smother convergence

Allows for larger learning rates

- Mini-batches lead to fast training!

Can parallelize computation + achieve significant speed increases on GPU's

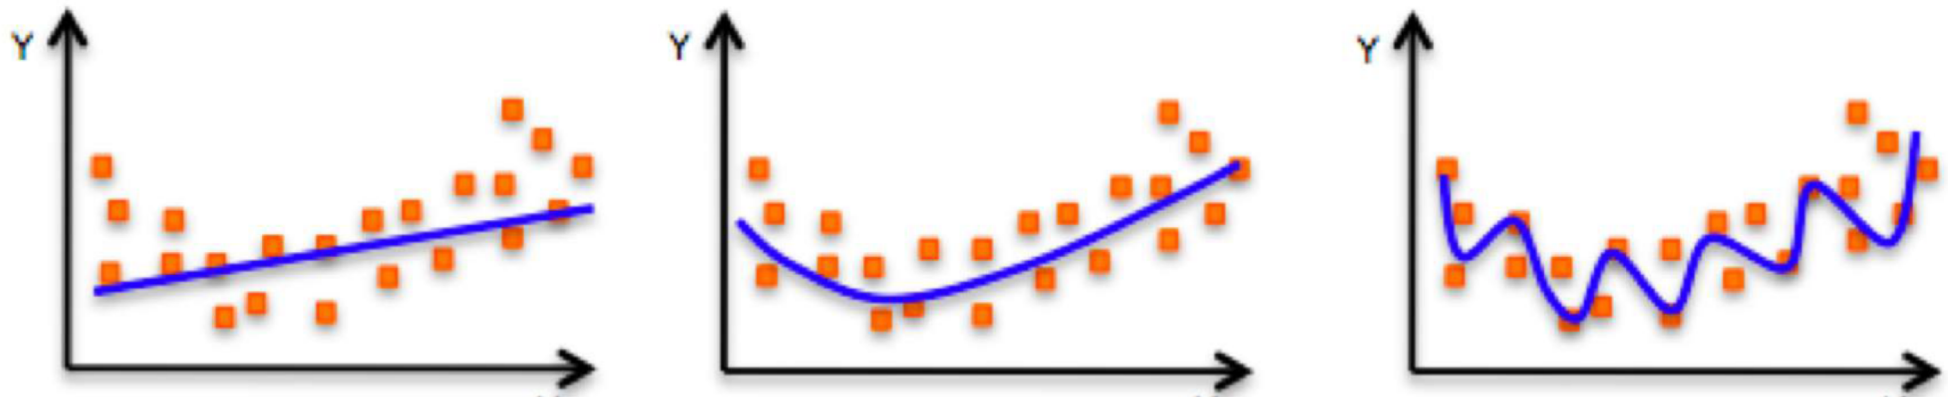


# Some terminology

- One **epoch** is when the entire dataset is passed forward and backward through the neural network only once (multiple times are usually needed)
- The **batch** size is the number of training examples in a mini-batch
- An **iteration** is the number of batches needed to complete one epoch
- Ex. For a dataset of 10000 sample with mini-batch size 1000, 10 iterations will complete 1 epoch

# Neural Networks in Practice: Overfitting

# The Problem of Overfitting



## Underfitting

Model does not have capacity to fully learn the data

## Ideal fit

## Overfitting

Too complex, extra parameters, does not generalize well

# Regularization

- What is it?

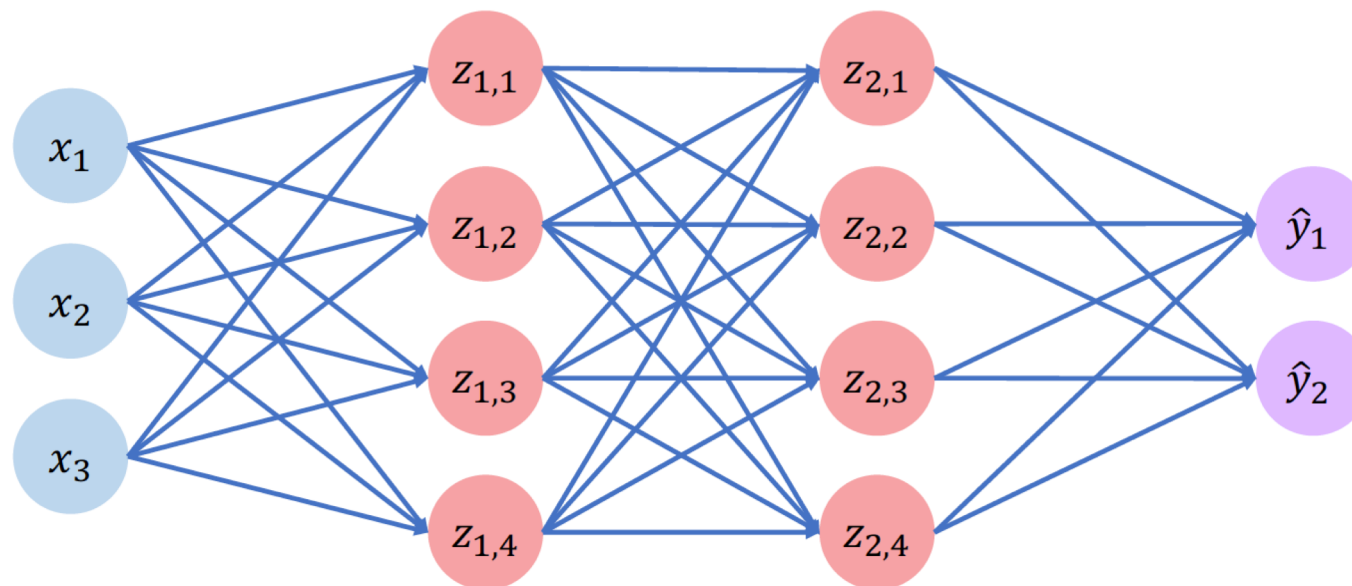
Technique that constrains our optimization problem to discourage complex models

- Why do we need it?

Improve generalization of our model on unseen data

# Regularization I: Dropout

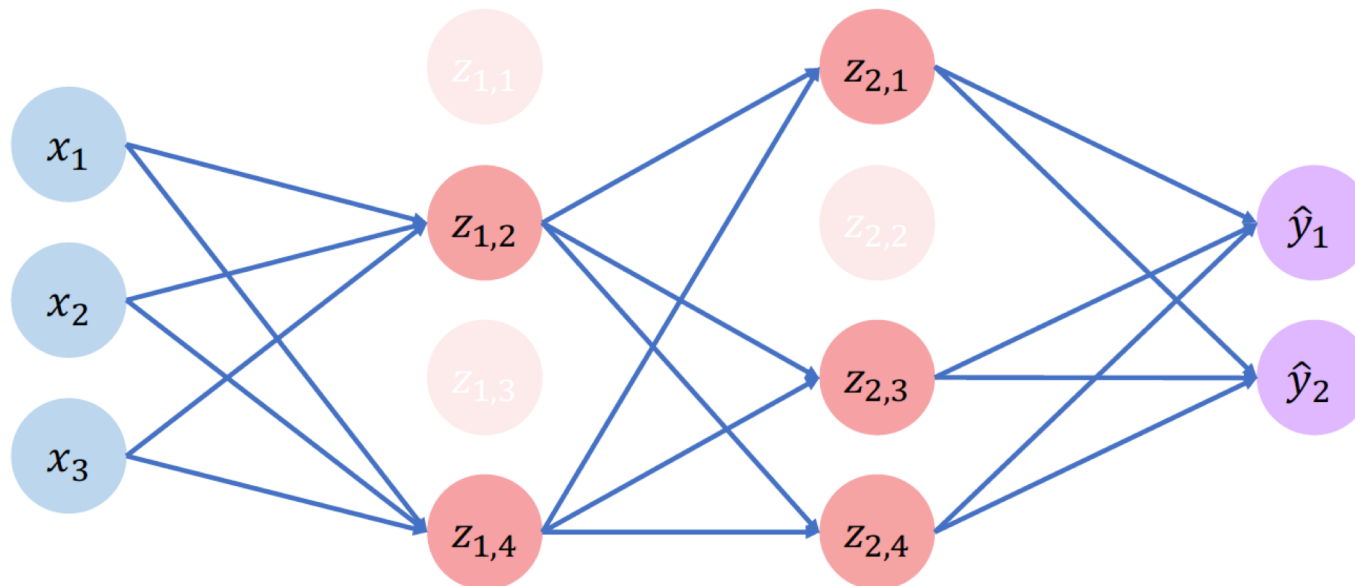
- During training, randomly set some activations to 0



# Regularization I: Dropout

- During training, randomly set some activations to 0
  - Typically 'drop' 50% of activations in layer

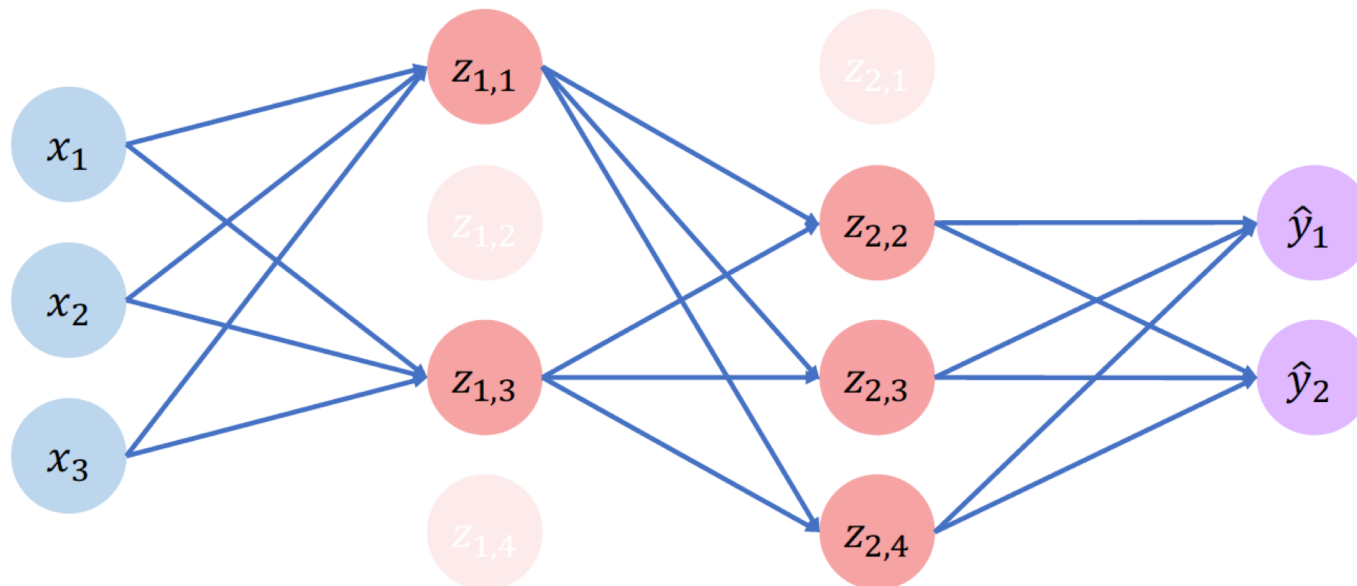
```
tf.keras.layers.Dropout(p=0.5)
```



# Regularization I: Dropout

- During training, randomly set some activations to 0
  - Typically 'drop' 50% of activations in layer

```
tf.keras.layers.Dropout(p=0.5)
```



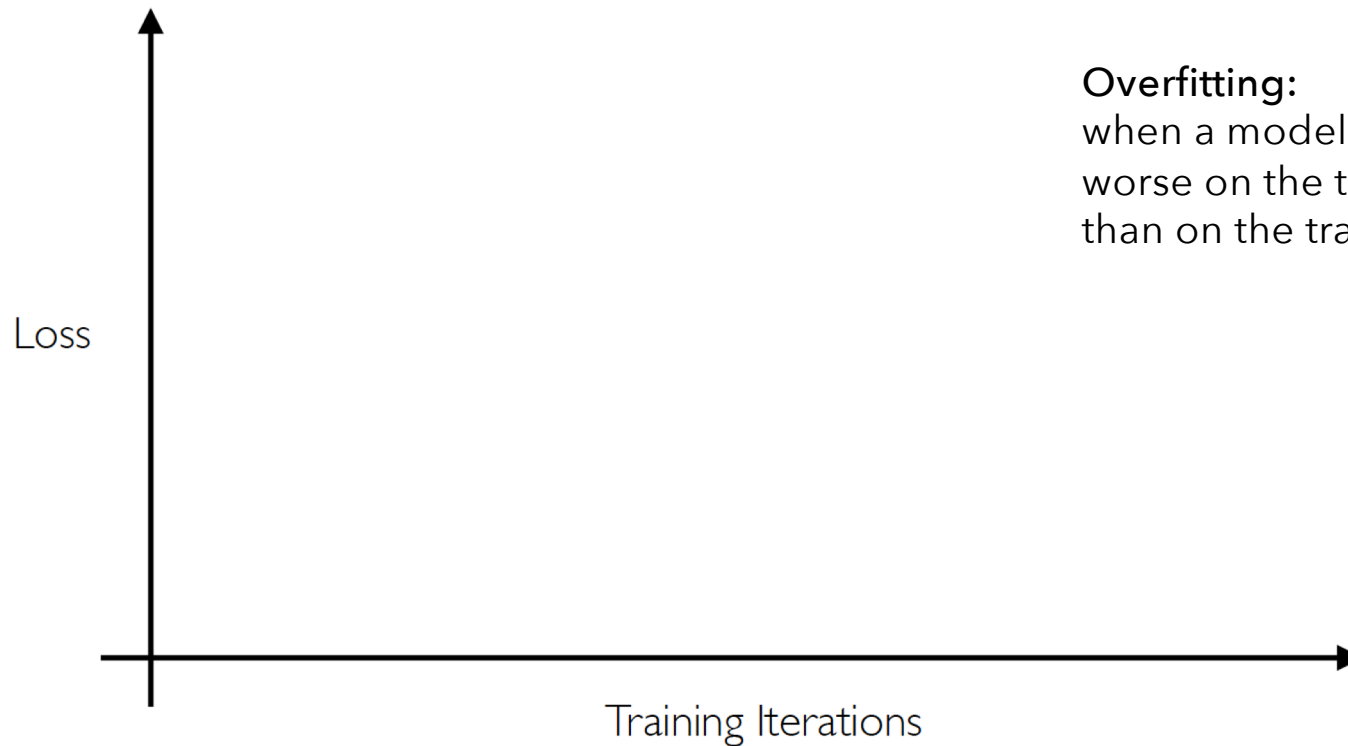
# Regularization I: Dropout

- the network is not going to rely too heavily on any particular path through the network
- instead it's going to **find a whole ensemble of different paths**, because it doesn't know which path is going to be dropped out at any given time



# Regularization 2: Early Stopping

- Stop training before we have a chance to overfit



**Overfitting:**  
when a model starts to perform worse on the test (validation) set than on the training set

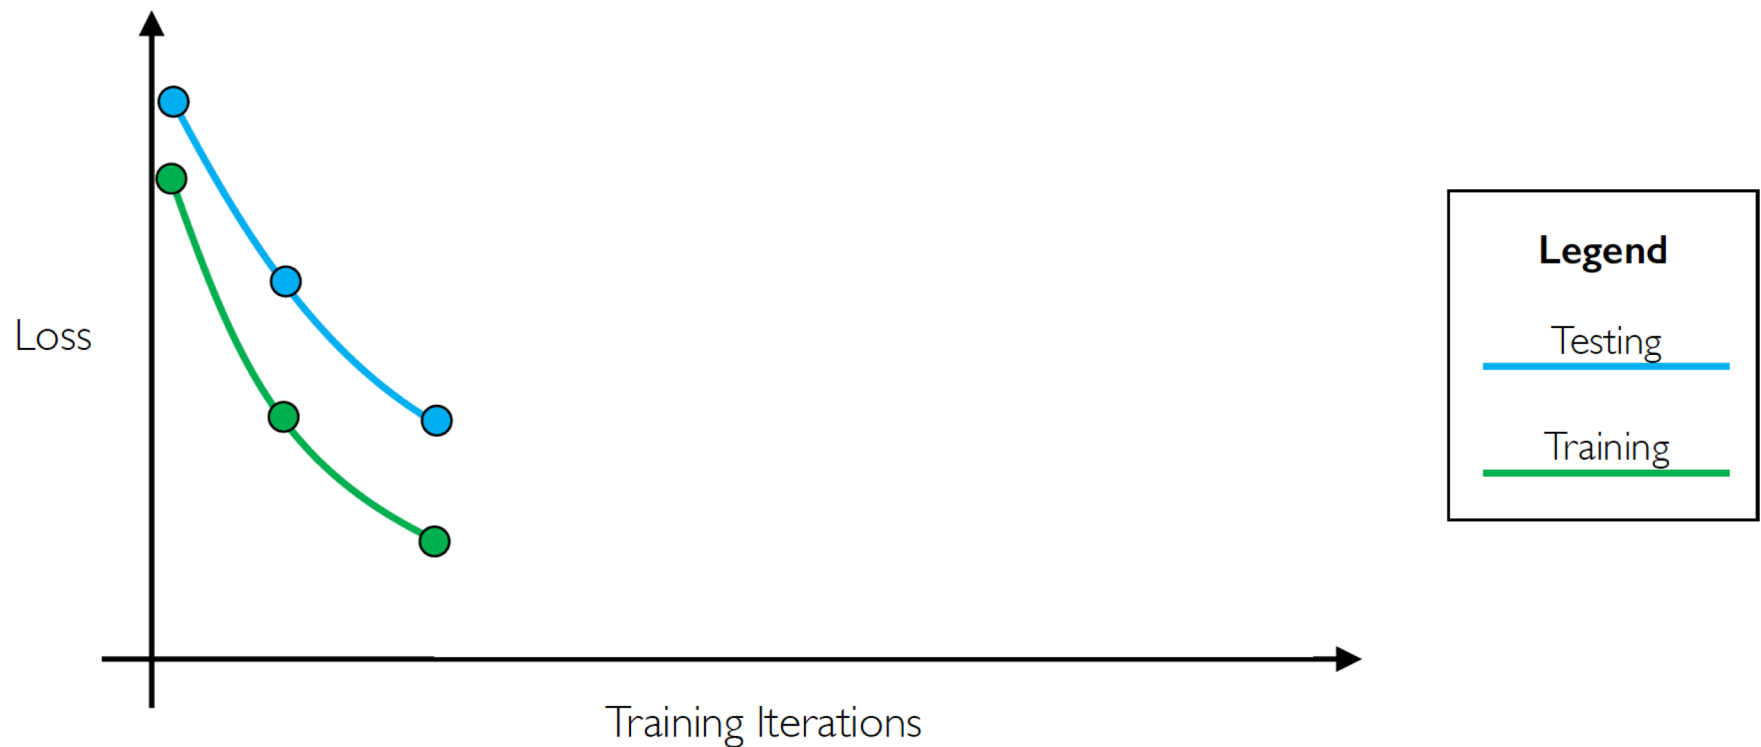
# Regularization 2: Early Stopping

- Stop training before we have a chance to overfit



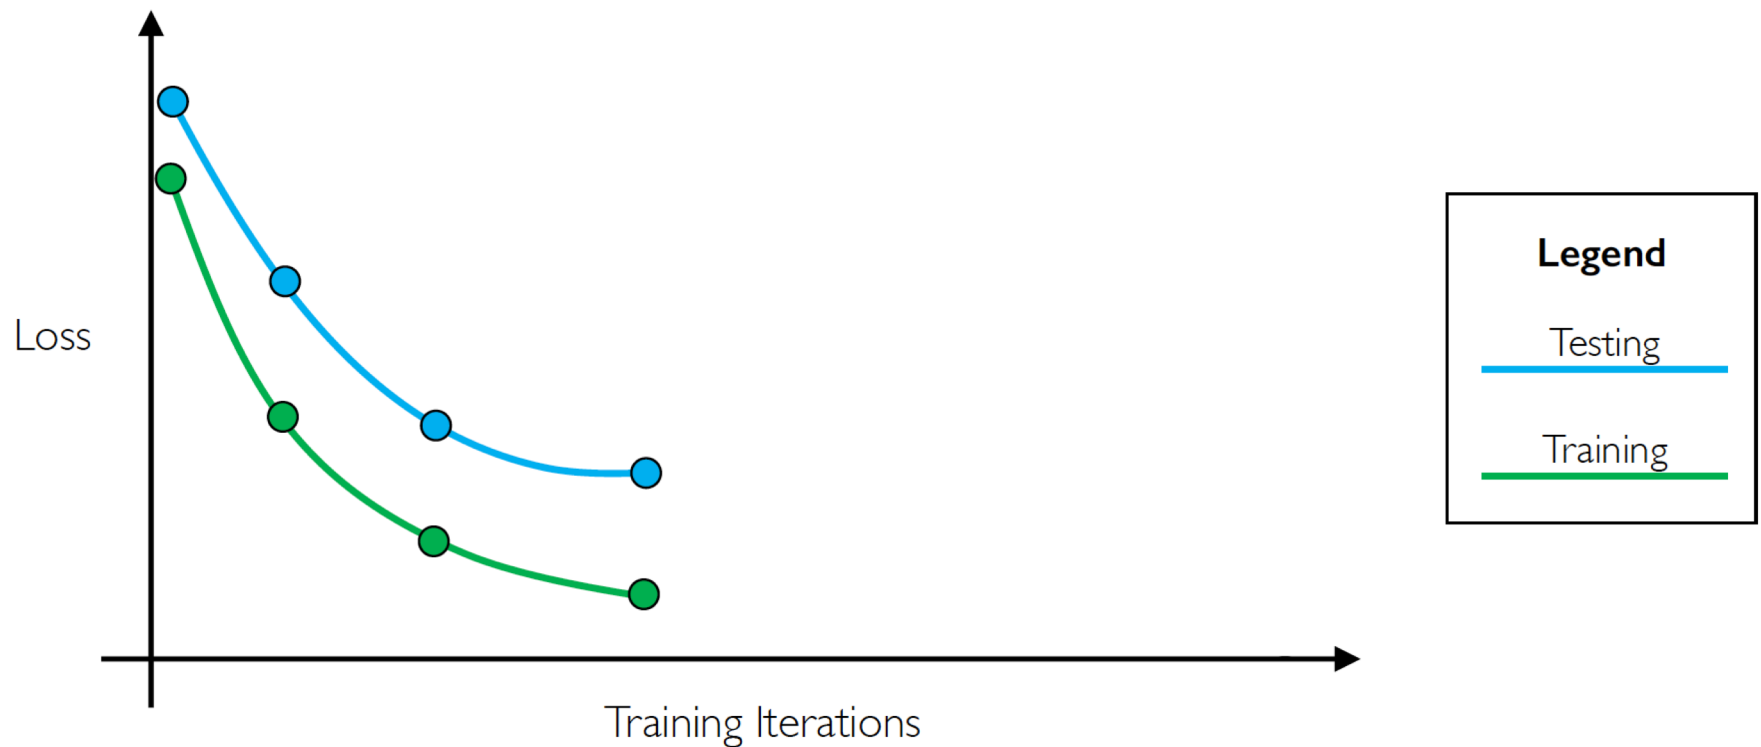
# Regularization 2: Early Stopping

- Stop training before we have a chance to overfit



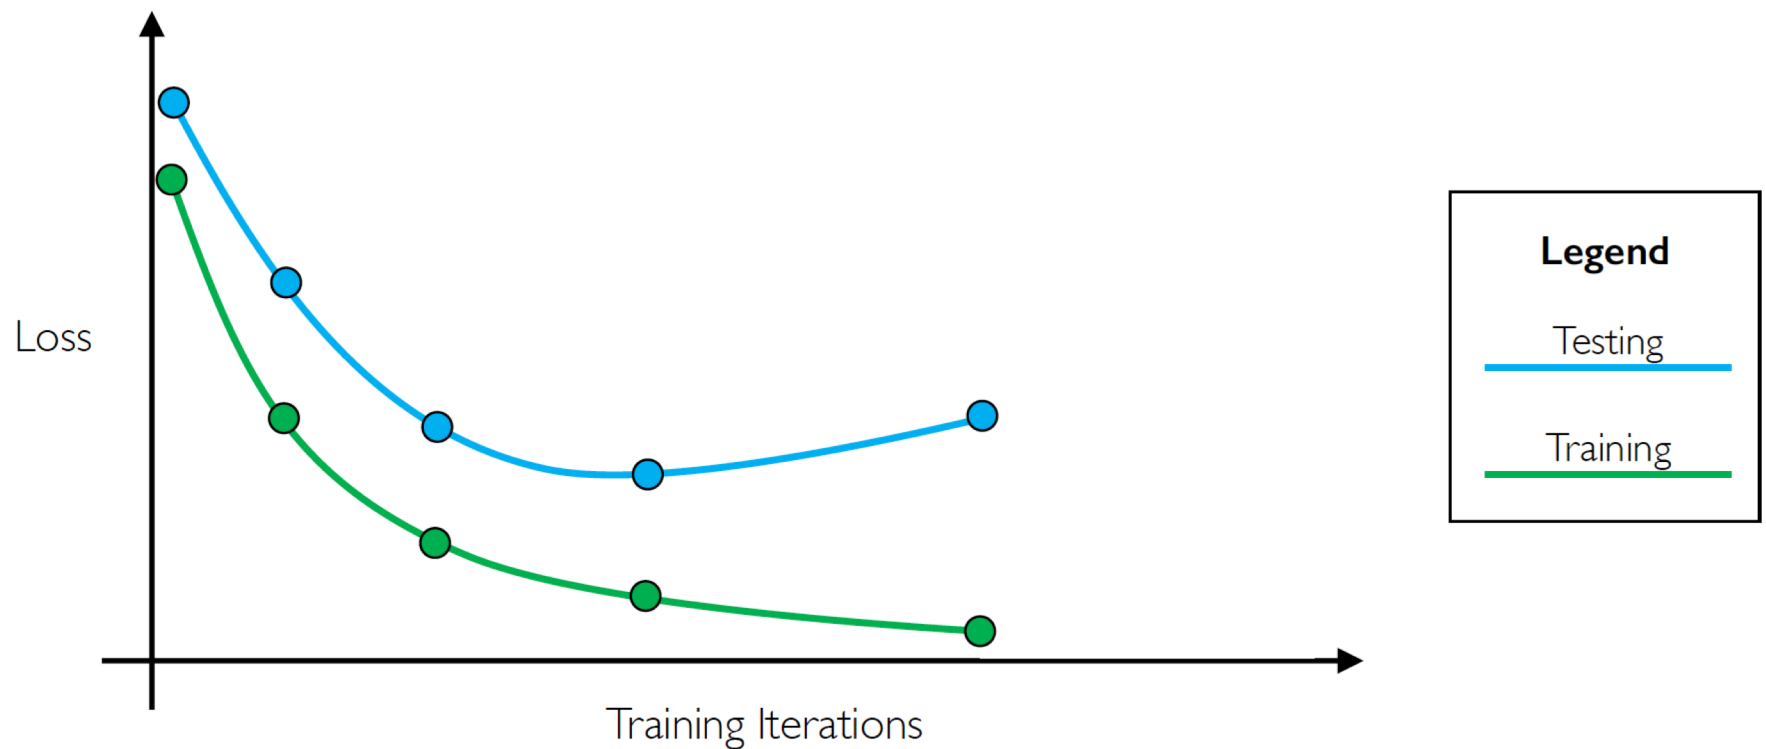
# Regularization 2: Early Stopping

- Stop training before we have a chance to overfit



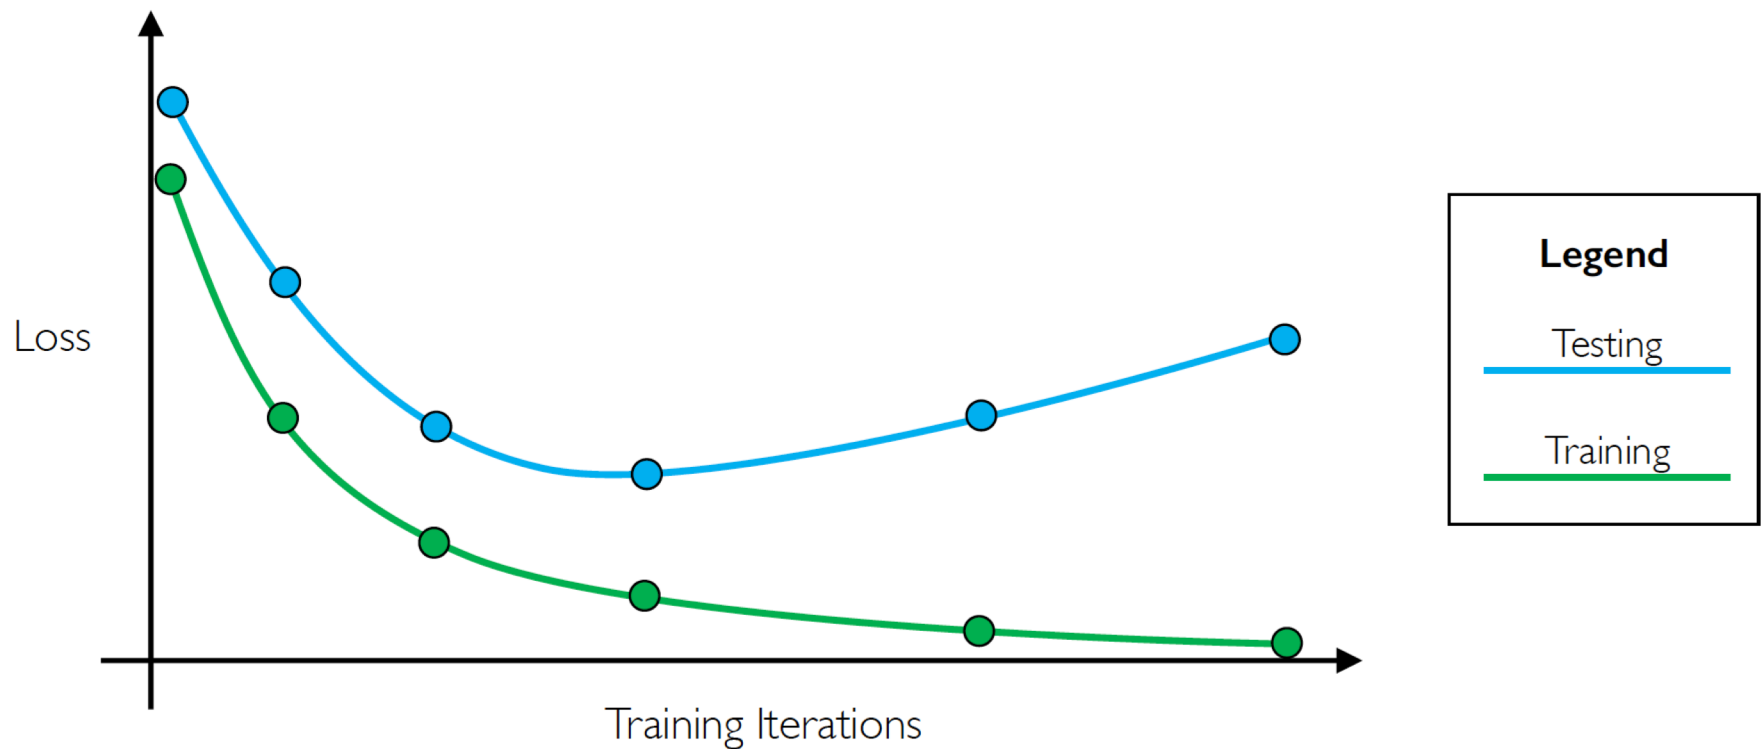
# Regularization 2: Early Stopping

- Stop training before we have a chance to overfit



# Regularization 2: Early Stopping

- Stop training before we have a chance to overfit



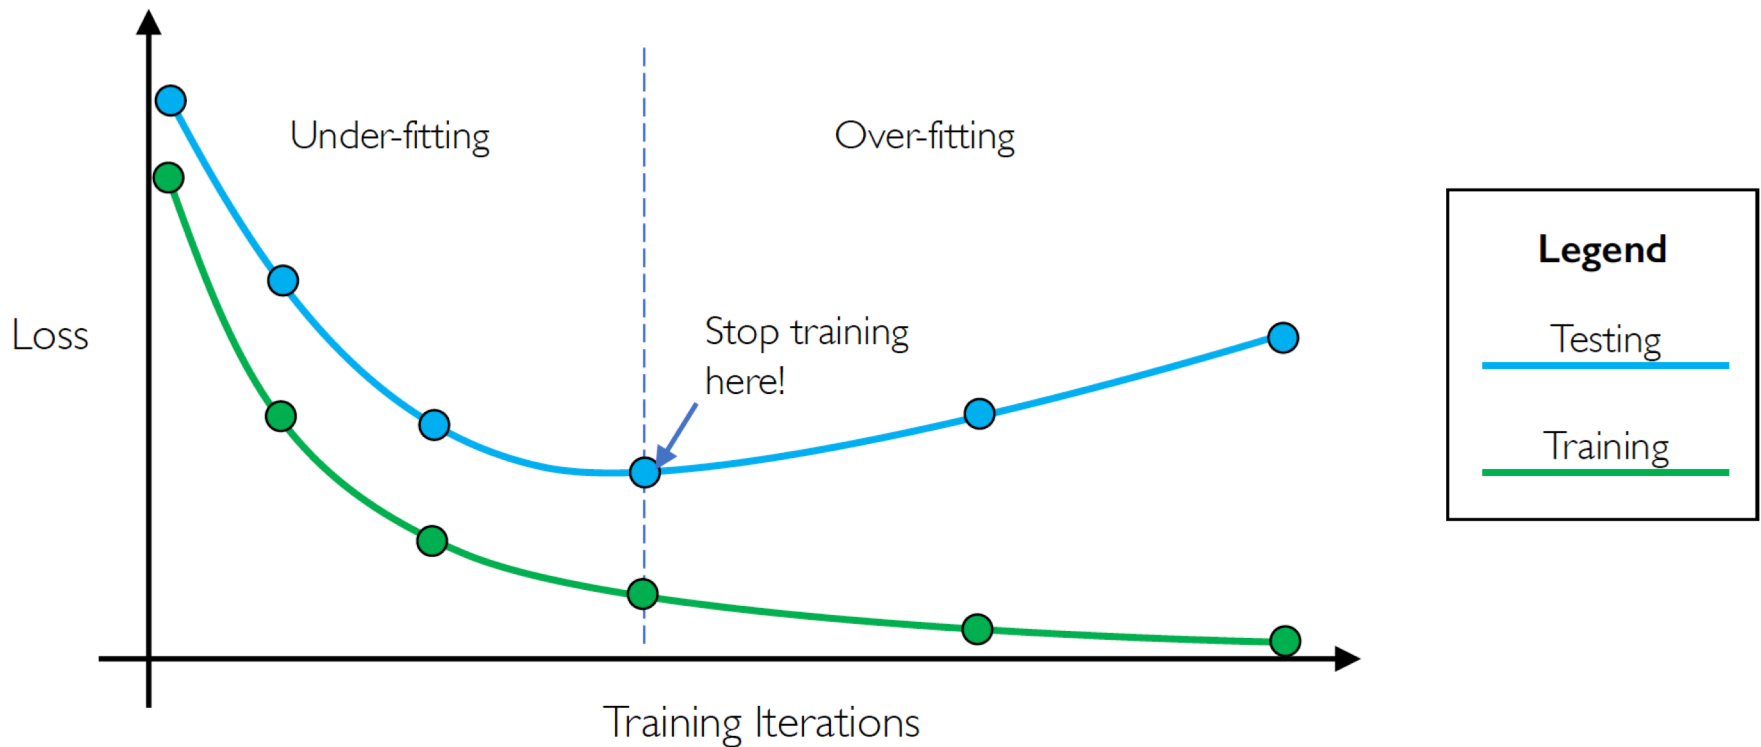
# Regularization 2: Early Stopping

- Stop training before we have a chance to overfit



# Regularization 2: Early Stopping

- Stop training before we have a chance to overfit

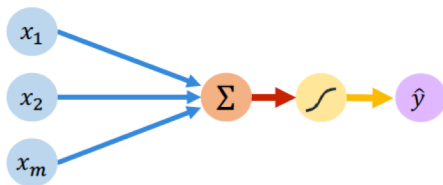




# Core Foundation Review

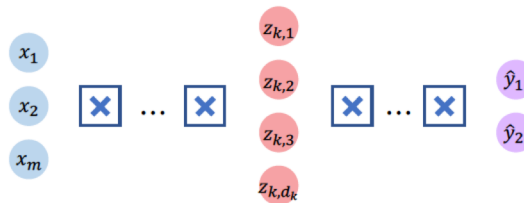
## The Perceptron

- Structural building blocks
- Nonlinear activation functions



## Neural Networks

- Stacking Perceptrons to form neural networks
- Optimization through backpropagation



## Training in Practice

- Adaptive learning
- Batching
- Regularization

