



上海大学未来技术学院  
SCHOOL OF FUTURE TECHNOLOGY, SHANGHAI UNIVERSITY

上海大学人工智能研究院  
INSTITUTE OF ARTIFICIAL INTELLIGENCE, SHANGHAI UNIVERSITY

# 人工智能导论

## ——第2课：人工智能基本理论

叶林奇

未来技术学院（人工智能研究院）

2023冬季学期



# 提纲

---

一、前馈网络

二、循环网络

三、卷积网络



上海大学  
SHANGHAI UNIVERSITY



## 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

Teaching computers how to **learn a task** directly from **raw data**



Why Deep Learning and Why Now?

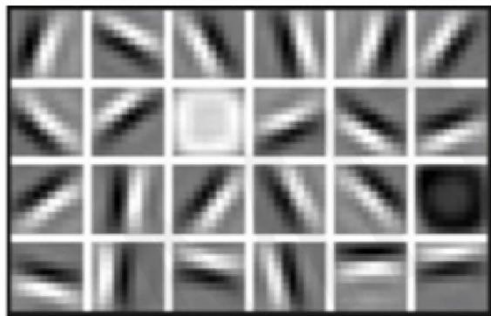


## Why Deep Learning?

Hand engineered features are time consuming, brittle, and not scalable in practice

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

Low Level Features



Lines & Edges

Mid Level Features



Eyes & Nose & Ears

High Level Features



Facial Structure

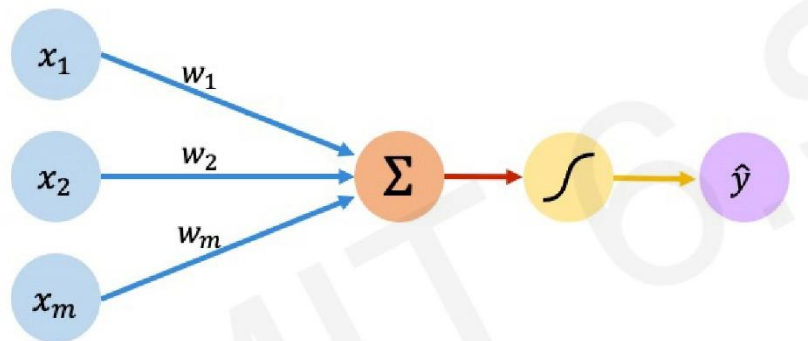


## The Perceptron

The structural building block of deep learning



## The Perceptron: Forward Propagation



Inputs    Weights    Sum    Non-Linearity    Output

Output

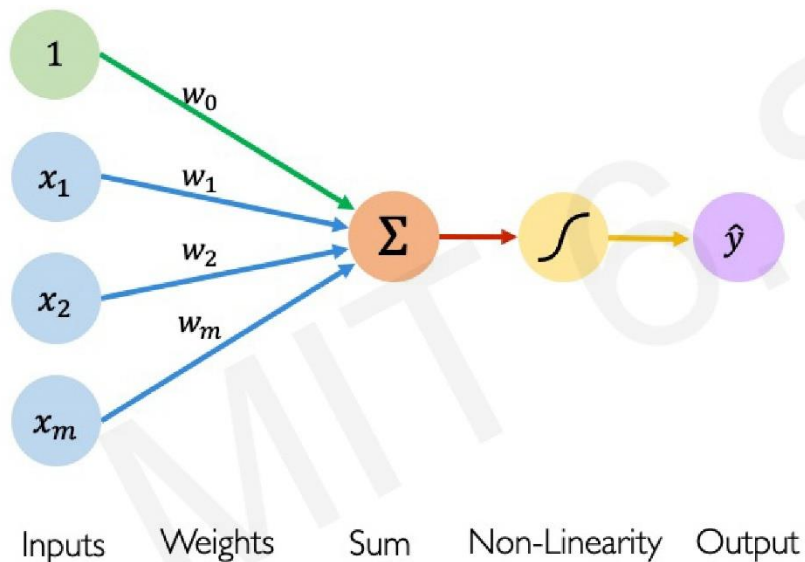
Linear combination of inputs

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

Non-linear activation function



## The Perceptron: Forward Propagation



Output

Linear combination of inputs

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

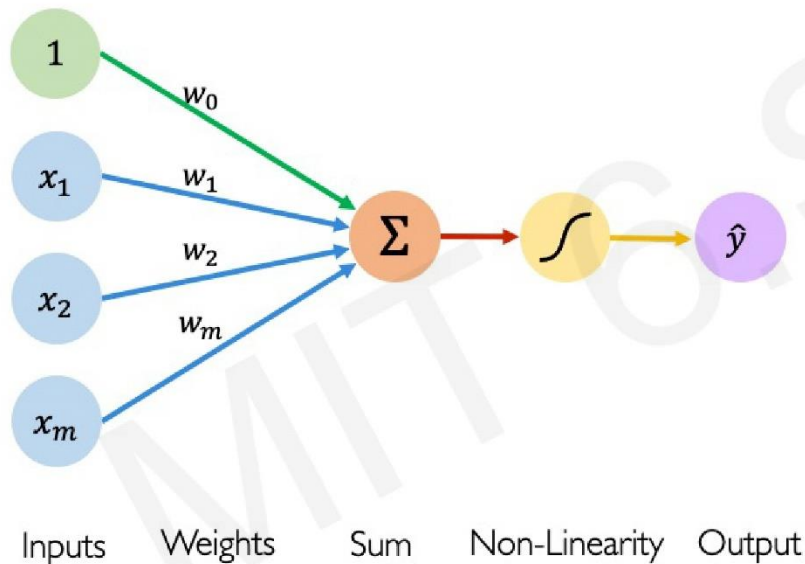
Non-linear activation function

Bias





## The Perceptron: Forward Propagation



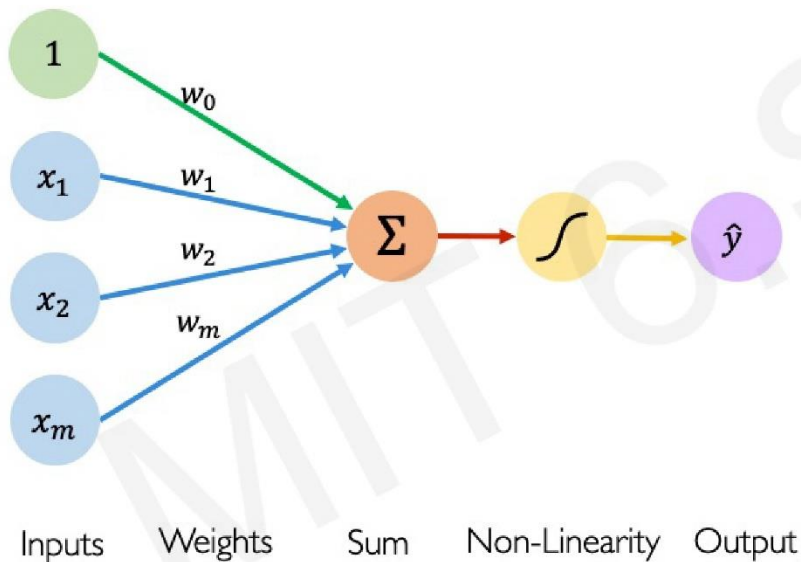
$$\hat{y} = g \left( w_0 + \sum_{i=1}^m x_i w_i \right)$$

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

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



## The Perceptron: Forward Propagation

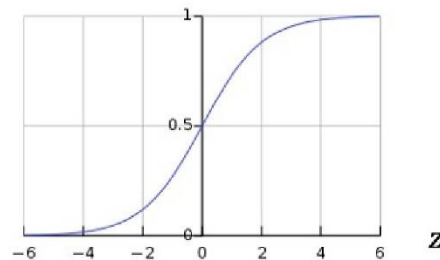


### Activation Functions

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

- Example: sigmoid function

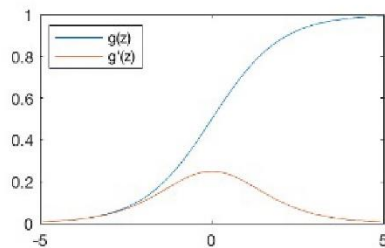
$$g(z) = \sigma(z) = \frac{1}{1 + e^{-z}}$$





## Common Activation Functions

### Sigmoid Function

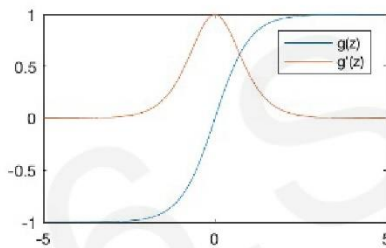


$$g(z) = \frac{1}{1 + e^{-z}}$$

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


 `tf.math.sigmoid(z)`

### Hyperbolic Tangent

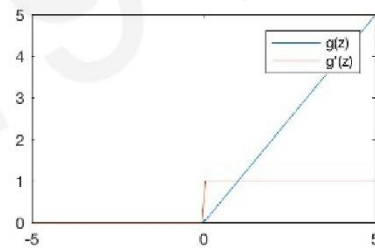


$$g(z) = \frac{e^z - e^{-z}}{e^z + e^{-z}}$$

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

 `tf.math.tanh(z)`

### Rectified Linear Unit (ReLU)



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

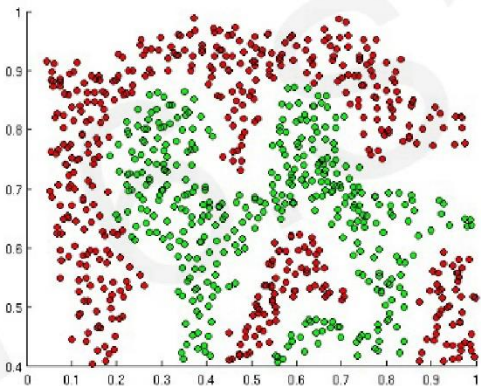
$$g'(z) = \begin{cases} 1, & z > 0 \\ 0, & \text{otherwise} \end{cases}$$

 `tf.nn.relu(z)`



## Importance of Activation Functions

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

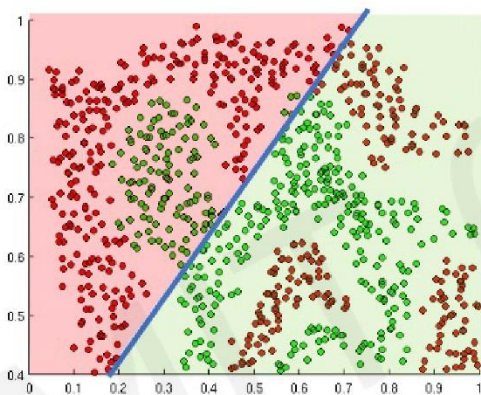


What if we wanted to build a neural network to distinguish green vs red points?



## Importance of Activation Functions

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

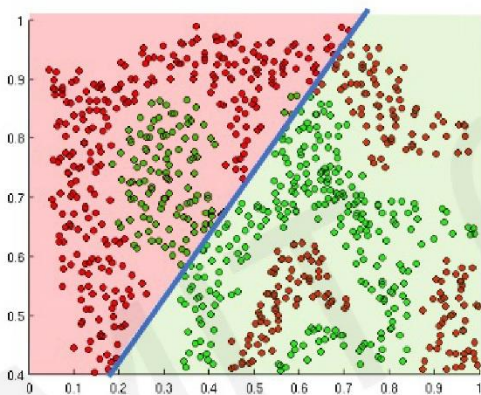


Linear activation functions produce linear decisions no matter the network size

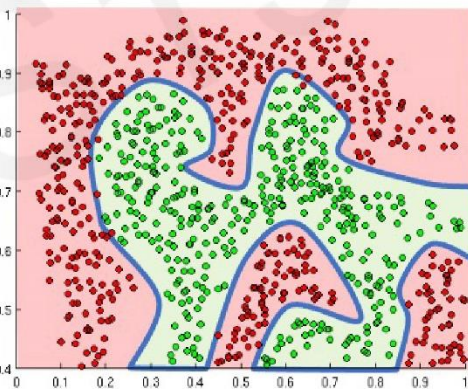


## Importance of Activation Functions

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



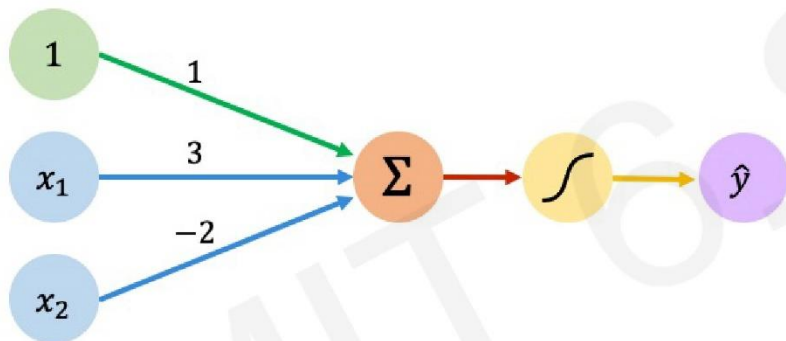
Linear activation functions produce linear decisions no matter the network size



Non-linearities allow us to approximate arbitrarily complex functions



## The Perceptron: Example



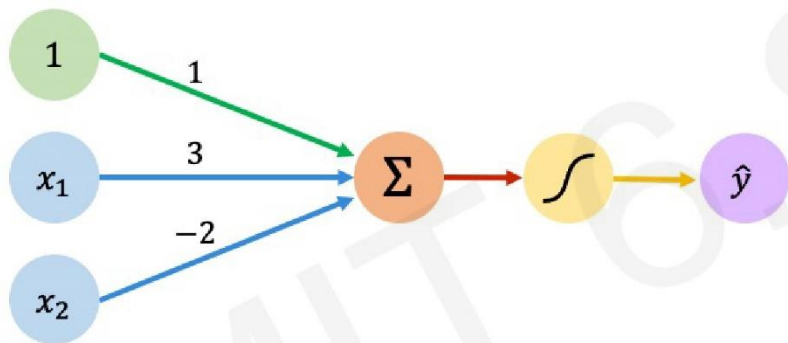
We have:  $w_0 = 1$  and  $\mathbf{w} = \begin{bmatrix} 3 \\ -2 \end{bmatrix}$

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

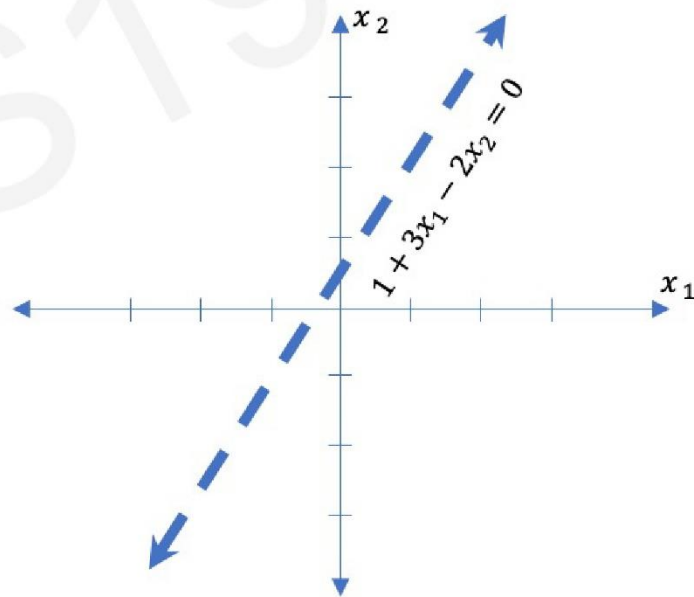
This is just a line in 2D!



## The Perceptron: Example



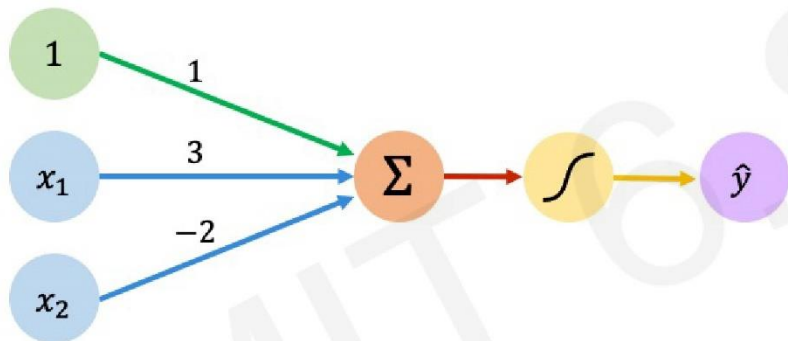
$$\hat{y} = g(1 + 3x_1 - 2x_2)$$





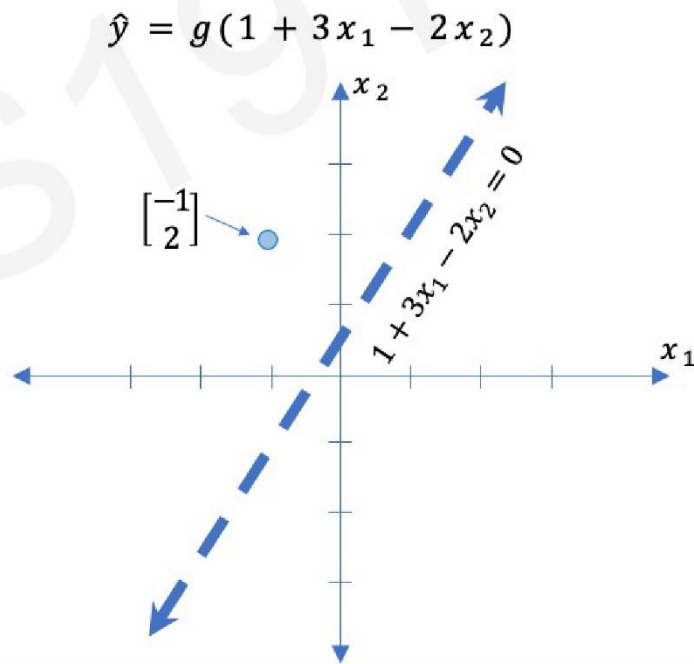


## The Perceptron: Example



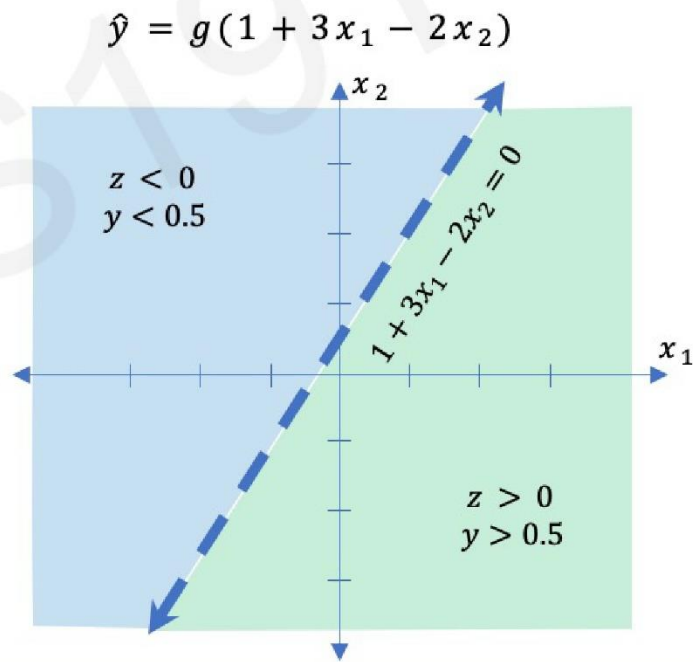
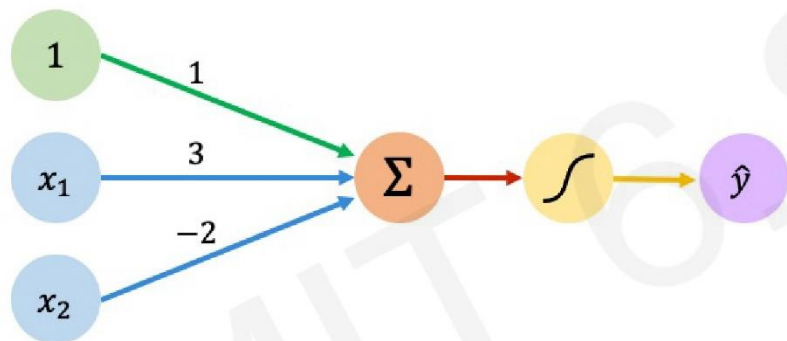
Assume we have input:  $\mathbf{X} = \begin{bmatrix} -1 \\ 2 \end{bmatrix}$

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





## The Perceptron: Example





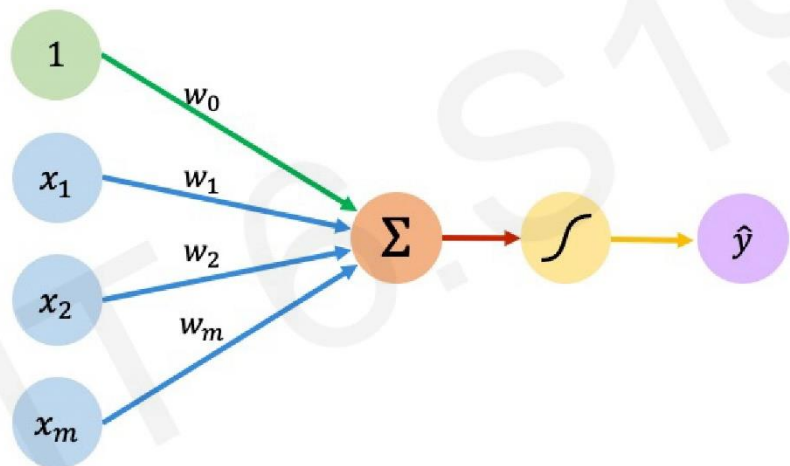
# 人工智能基本理论

Building Neural Networks with Perceptrons



## The Perceptron: Simplified

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



Inputs

Weights

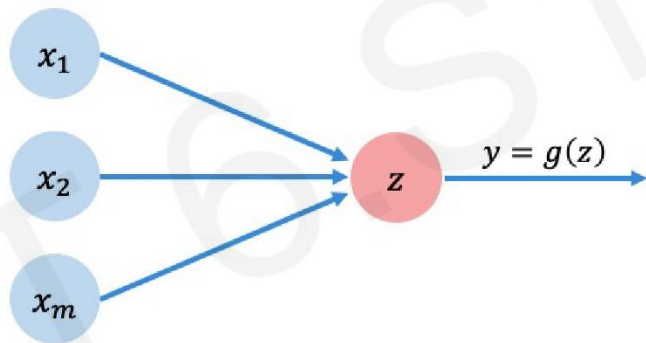
Sum

Non-Linearity

Output



## The Perceptron: Simplified

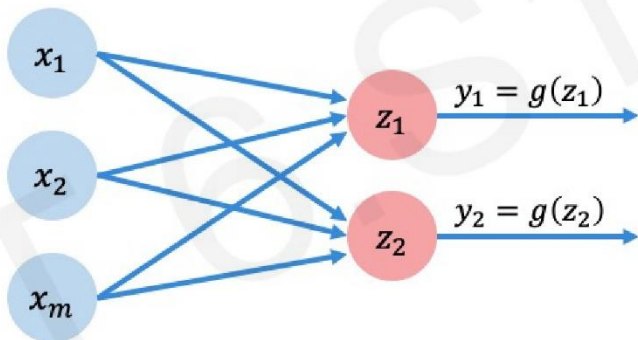


$$z = w_0 + \sum_{j=1}^m x_j w_j$$



## Multi Output Perceptron

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

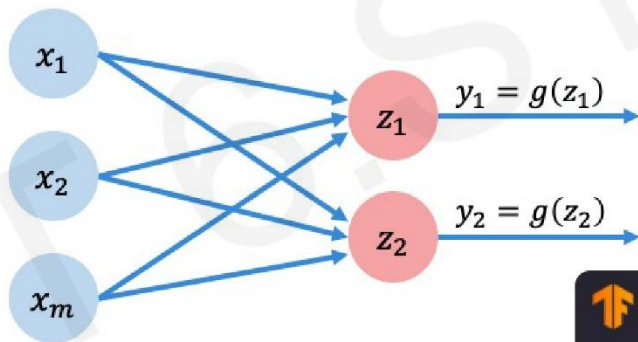


$$z_i = w_{0,i} + \sum_{j=1}^m x_j w_{j,i}$$



## Multi Output Perceptron

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

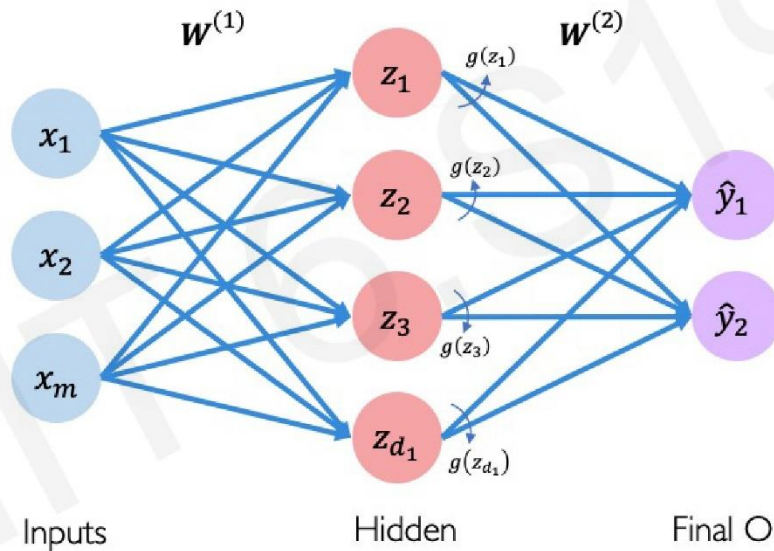


```
import tensorflow as tf
layer = tf.keras.layers.Dense(
    units=2)
```

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



## Single Layer Neural Network



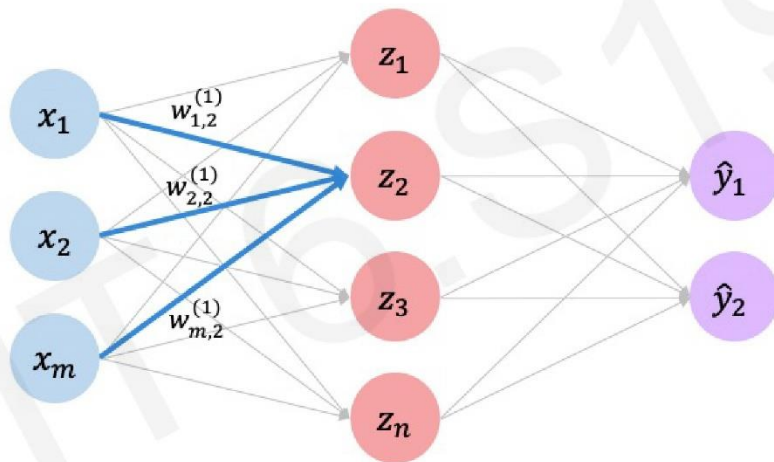
$$z_i = w_{0,i}^{(1)} + \sum_{j=1}^m x_j w_{j,i}^{(1)}$$

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





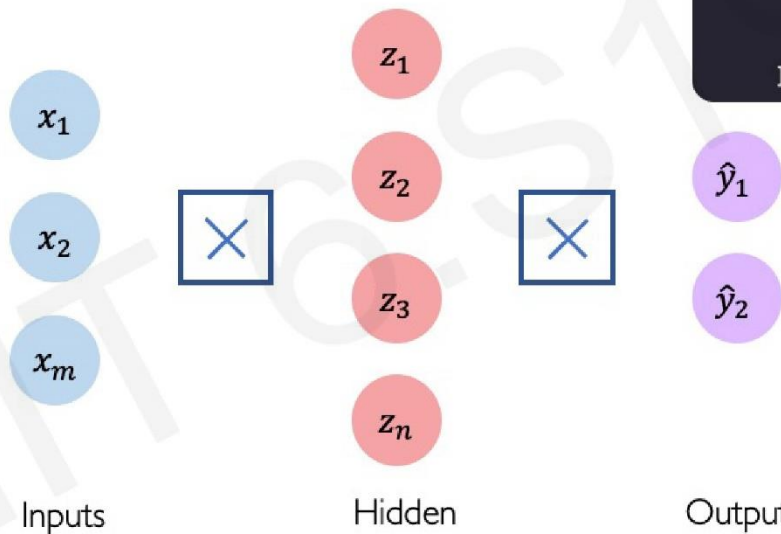
## Single Layer Neural Network



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



## Multi Output Perceptron

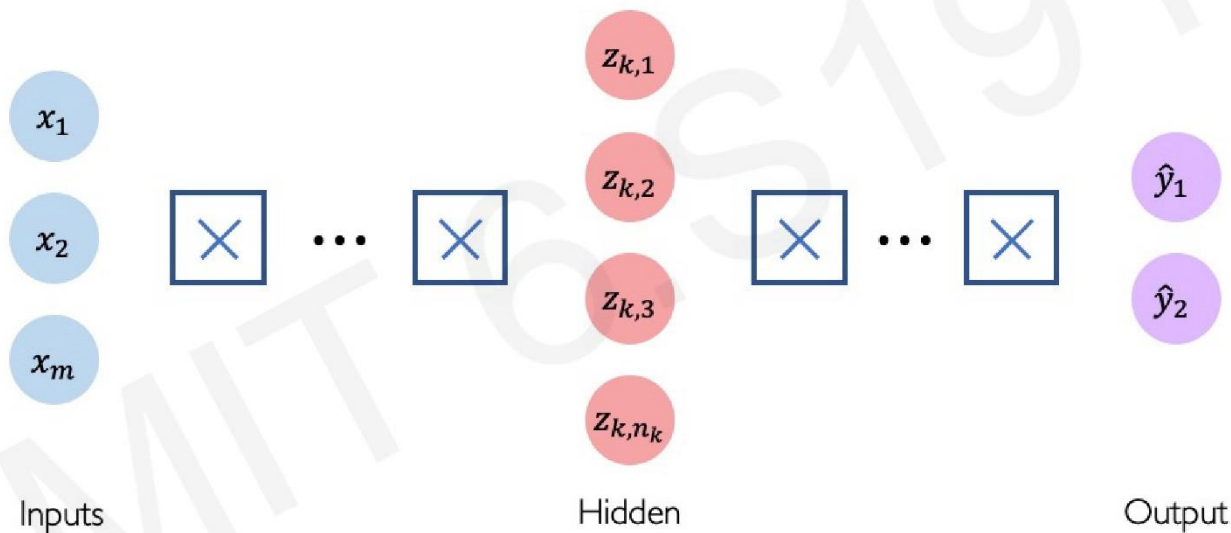


```
import tensorflow as tf

model = tf.keras.Sequential([
    tf.keras.layers.Dense(n),
    tf.keras.layers.Dense(2)
])
```



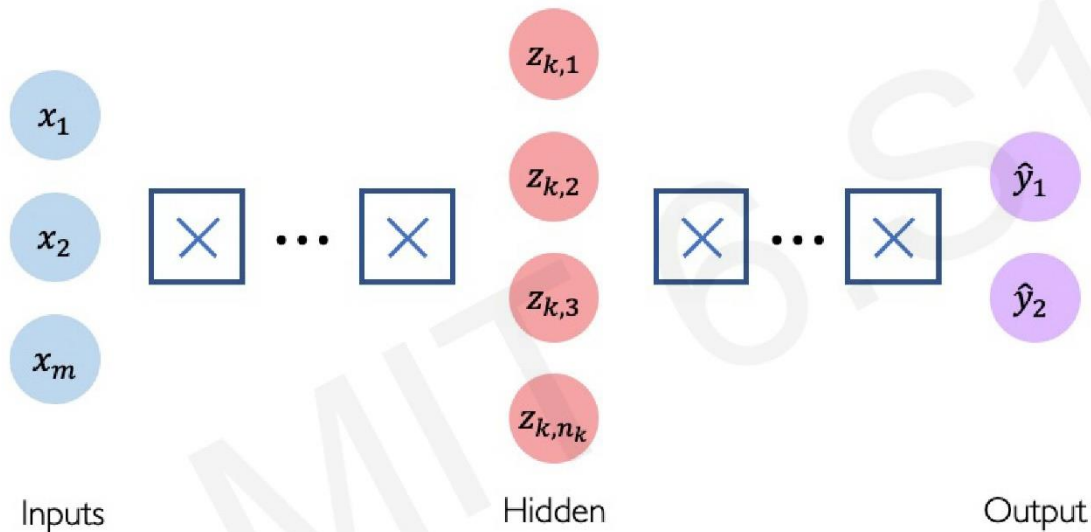
## Deep Neural Network



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



## Deep Neural Network



```
import tensorflow as tf

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

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



# 人工智能基本理论

Applying Neural Networks



## Example Problem

Will I pass this class?

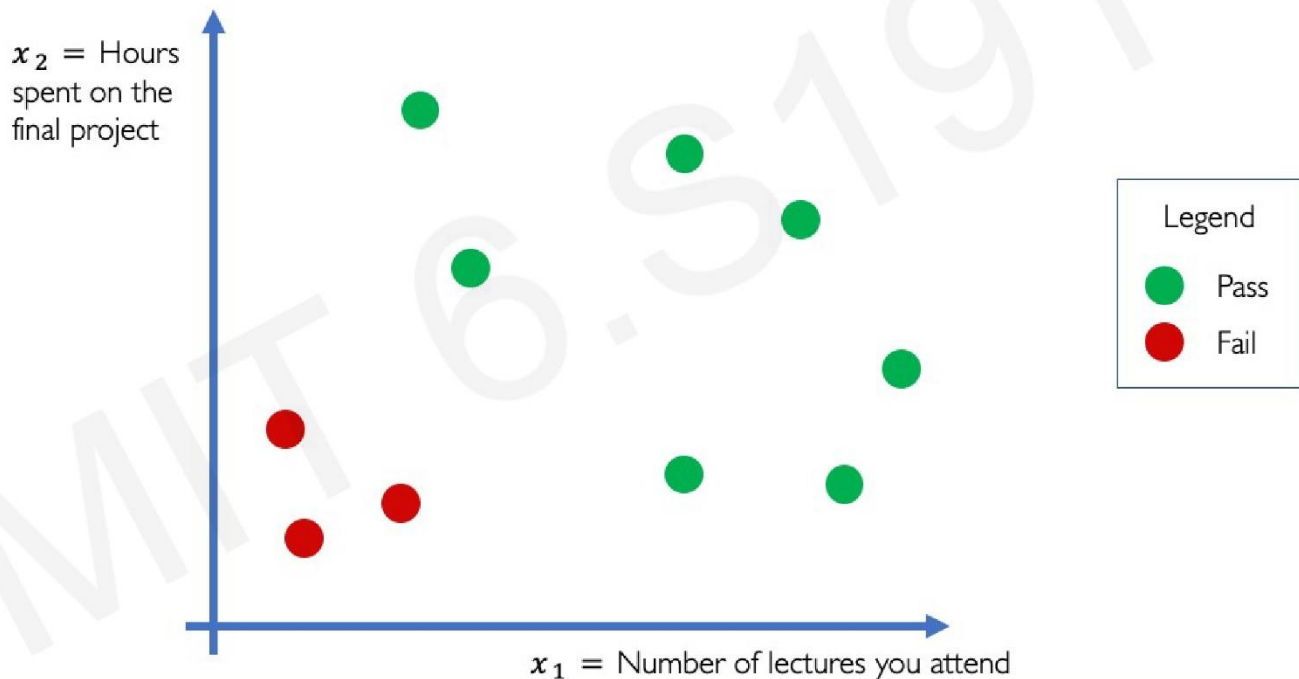
Let's start with a simple two feature model

$x_1$  = Number of lectures you attend

$x_2$  = Hours spent on the final project

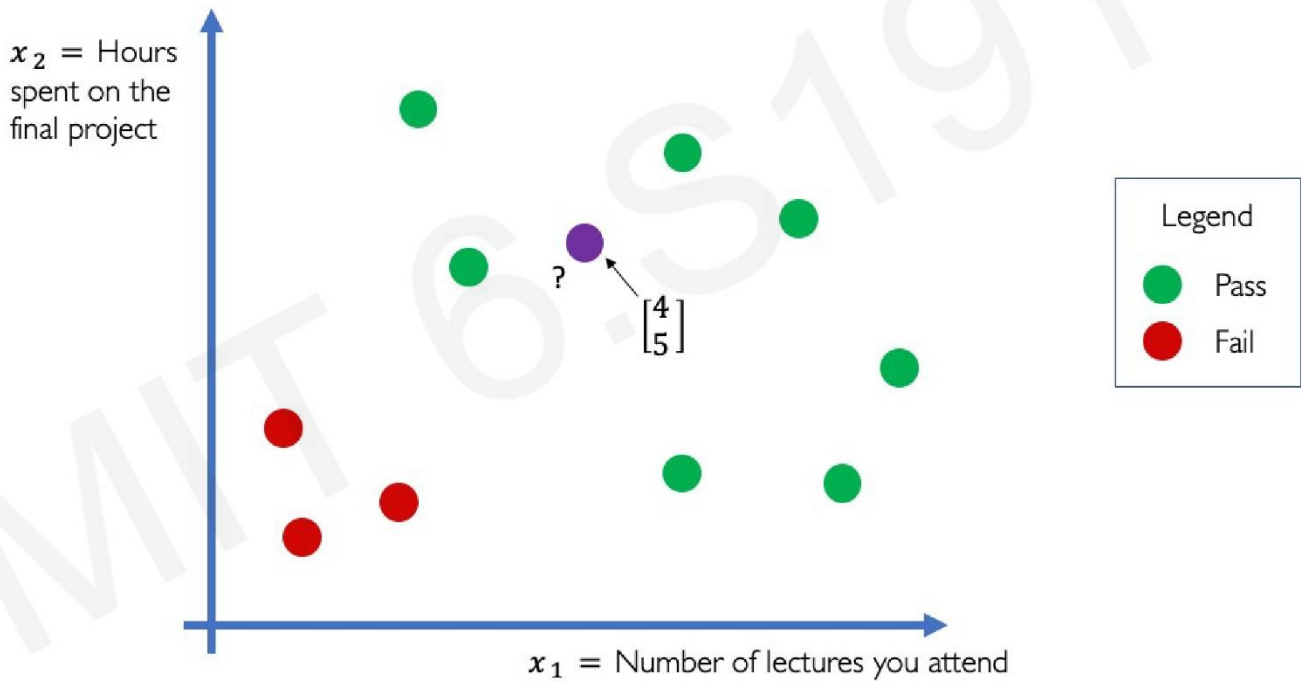


## Example Problem: Will I pass this class?





## Example Problem: Will I pass this class?

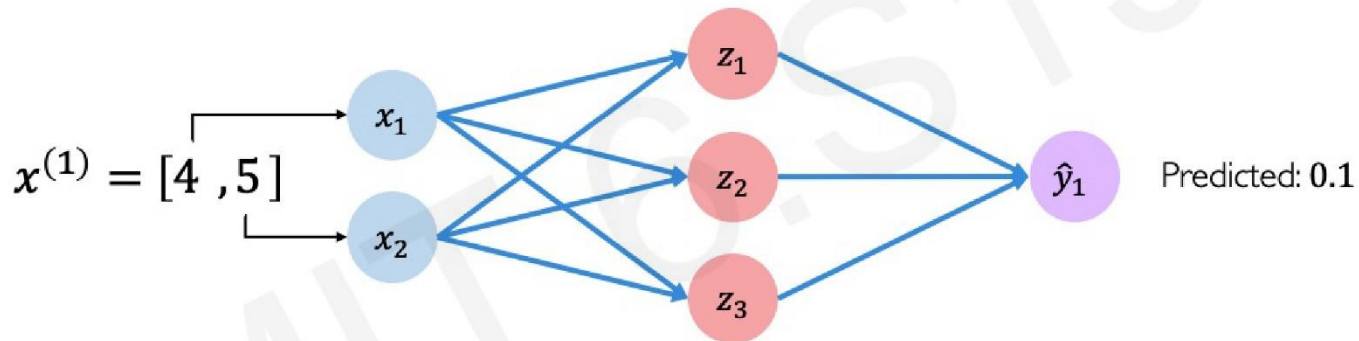






# 人工智能基本理论

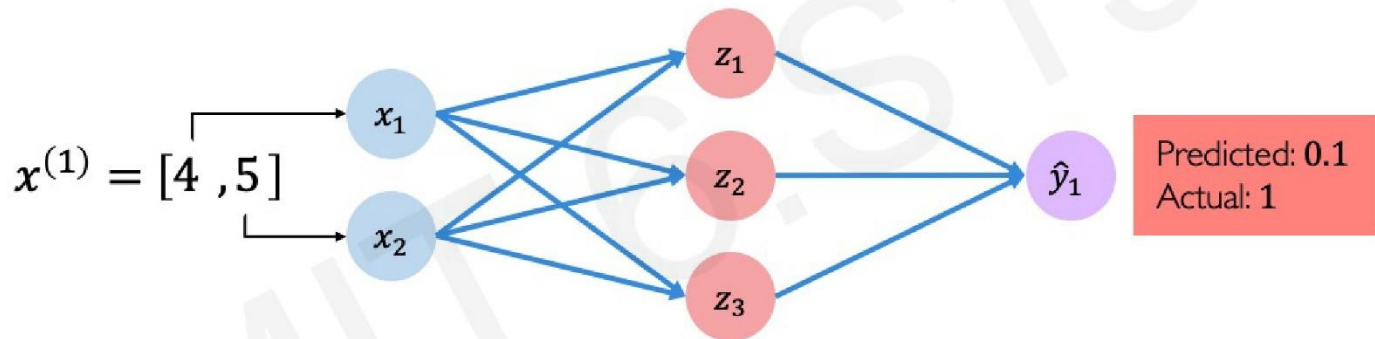
Example Problem: Will I pass this class?





# 人工智能基本理论

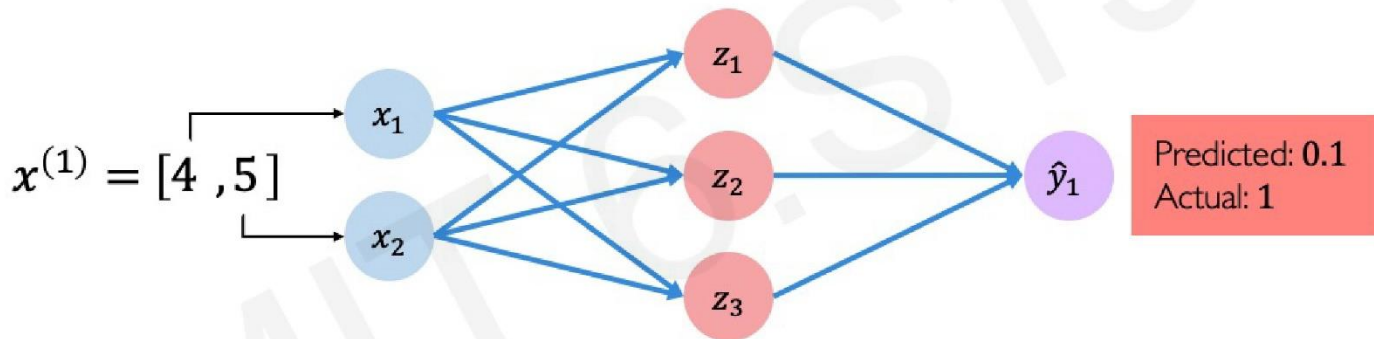
Example Problem: Will I pass this class?





## Quantifying Loss

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

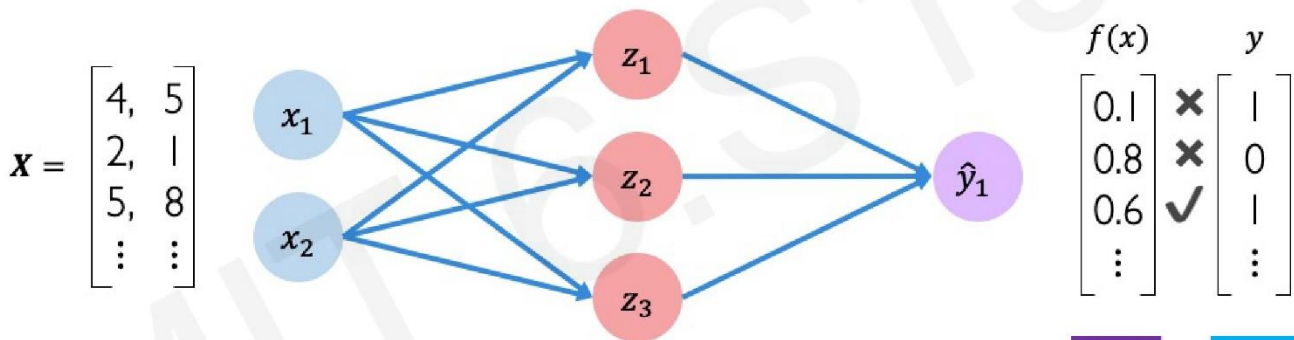


$$\mathcal{L}(\underbrace{f(x^{(i)}; \mathbf{W})}_{\text{Predicted}}, \underbrace{y^{(i)}}_{\text{Actual}})$$



## Empirical Loss

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



Also known as:

- Objective function
- Cost function
- Empirical Risk

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

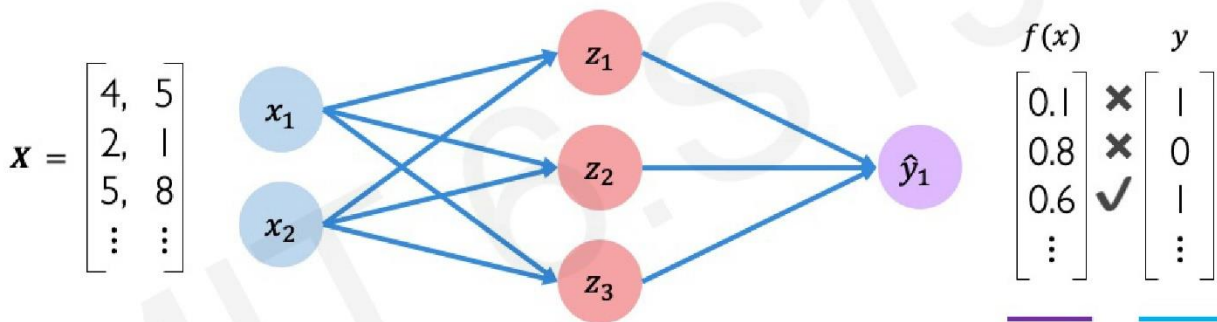
Predicted

Actual



## Binary Cross Entropy Loss

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



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

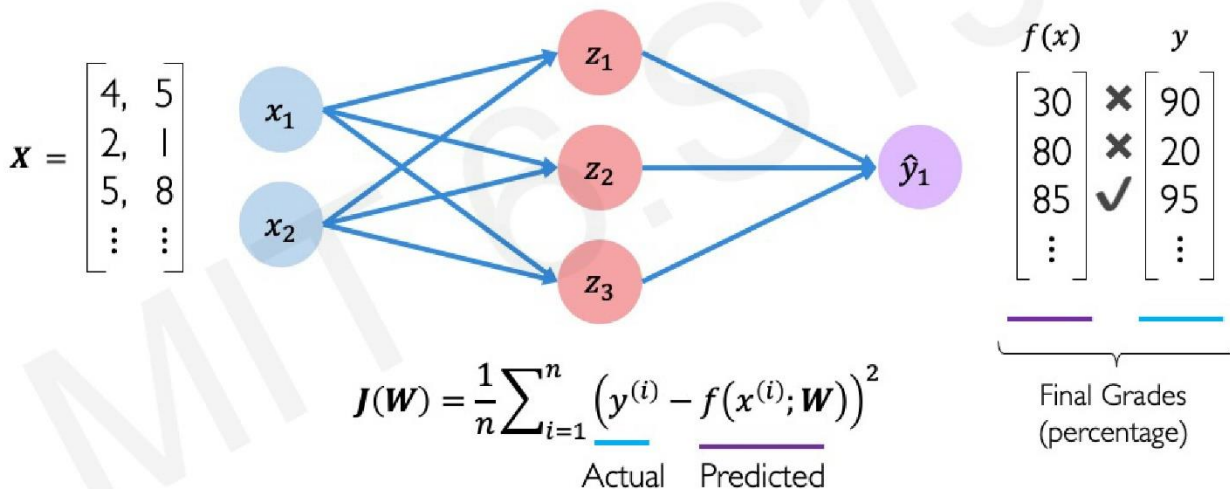


```
loss = tf.reduce_mean( tf.nn.softmax_cross_entropy_with_logits(y, predicted) )
```



## 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)) )  
loss = tf.keras.losses.MSE( y, predicted )
```



# 人工智能基本理论

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:

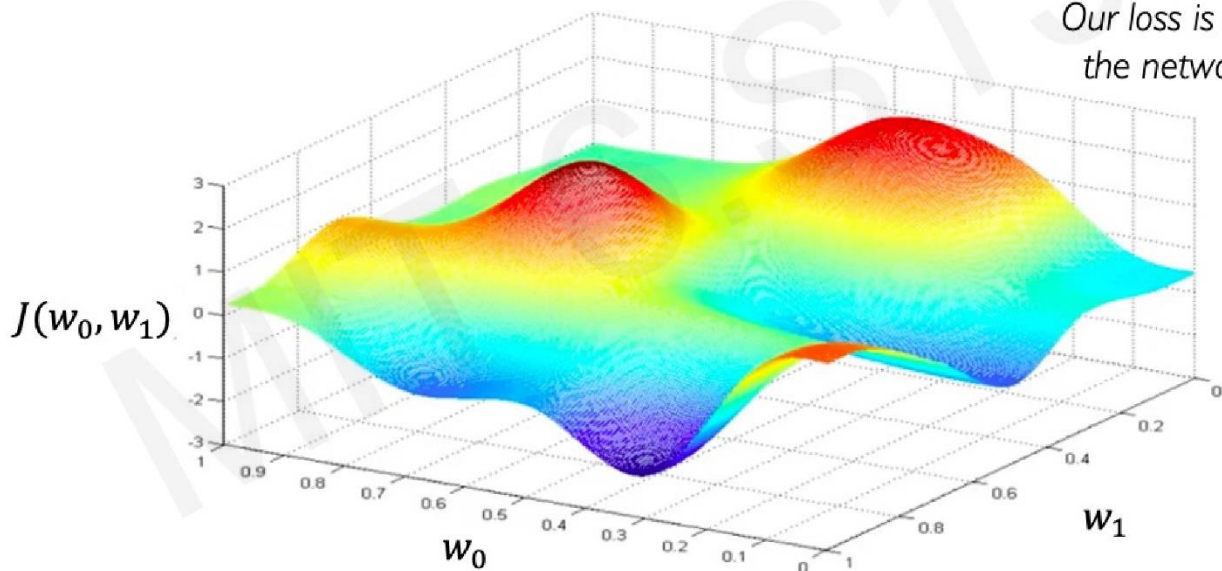
$$\mathbf{W} = \{\mathbf{W}^{(0)}, \mathbf{W}^{(1)}, \dots\}$$



## Loss Optimization

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

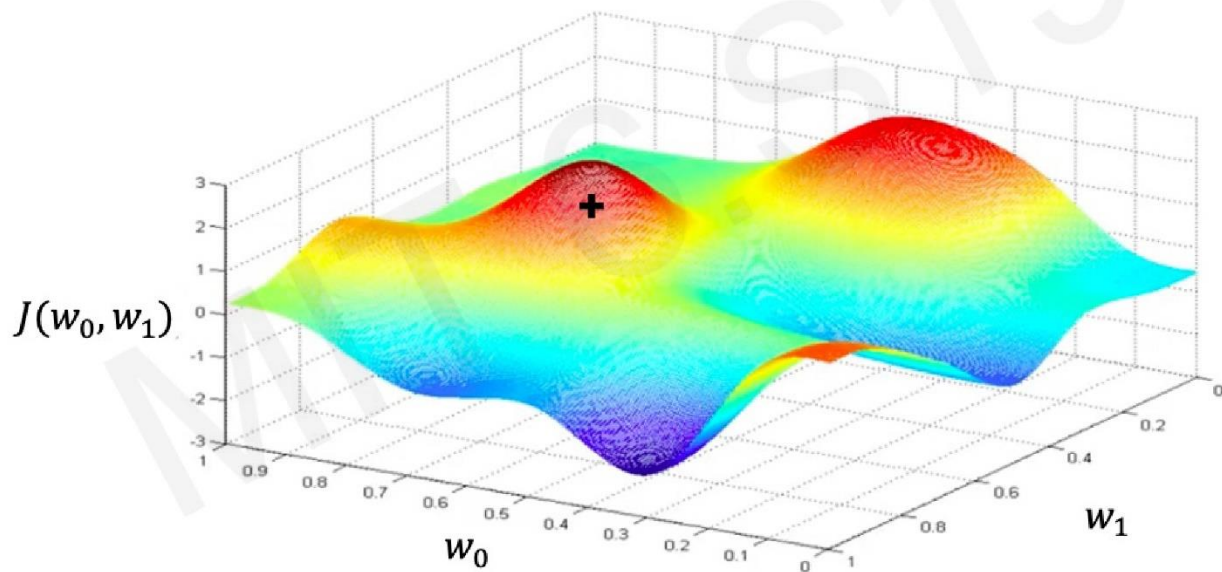
Remember:  
*Our loss is a function of  
the network weights!*





## Loss Optimization

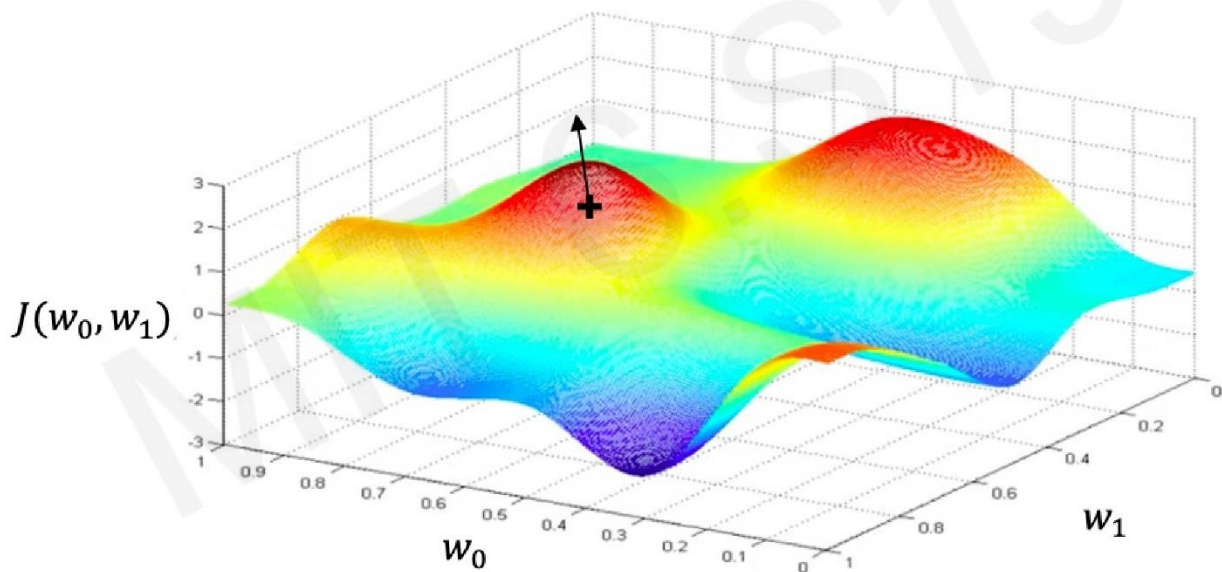
Randomly pick an initial  $(w_0, w_1)$





## Loss Optimization

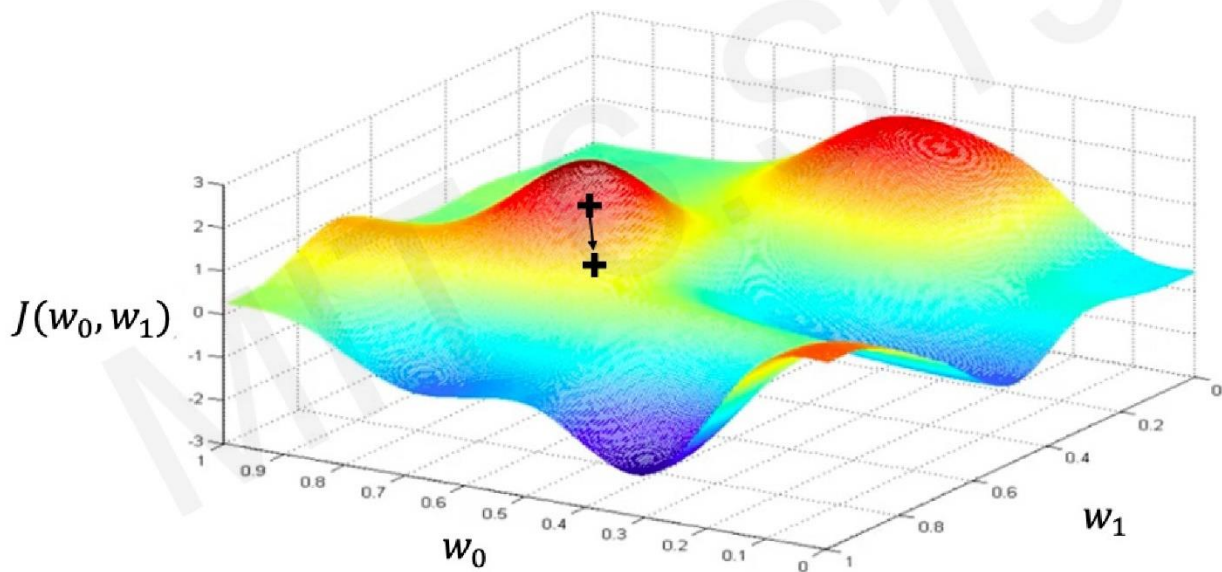
Compute gradient,  $\frac{\partial J(w)}{\partial w}$





## Loss Optimization

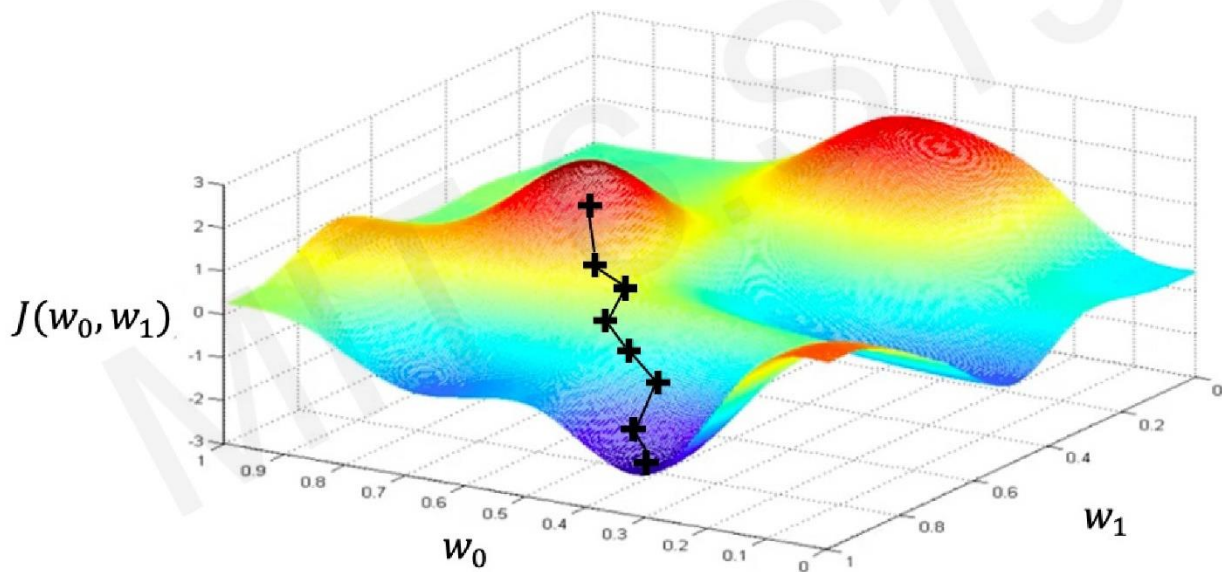
Take small step in opposite direction of gradient





## 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,  $\frac{\partial J(W)}{\partial W}$
4. Update weights,  $W \leftarrow W - \eta \frac{\partial J(W)}{\partial W}$
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

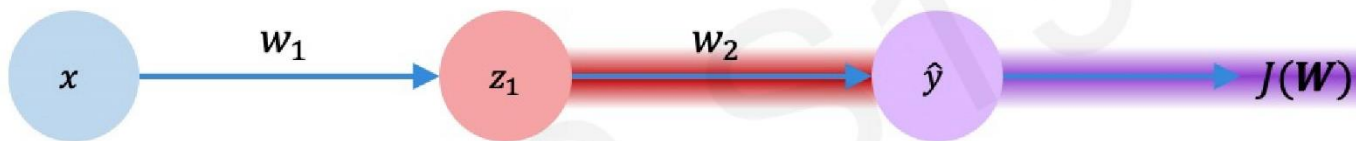


$$\frac{\partial J(W)}{\partial w_2} =$$

Let's use the chain rule!



## Computing Gradients: Backpropagation

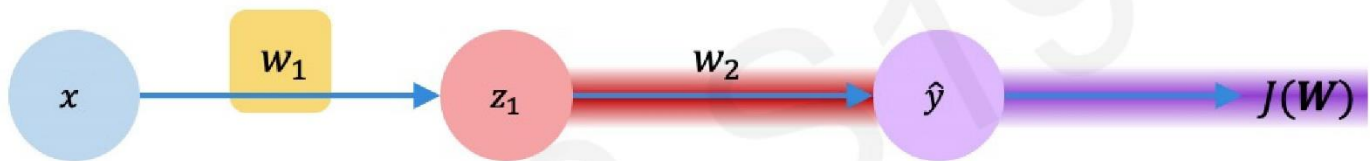


$$\frac{\partial J(W)}{\partial w_2} = \frac{\partial J(W)}{\partial \hat{y}} * \frac{\partial \hat{y}}{\partial w_2}$$

—      —



## Computing Gradients: Backpropagation



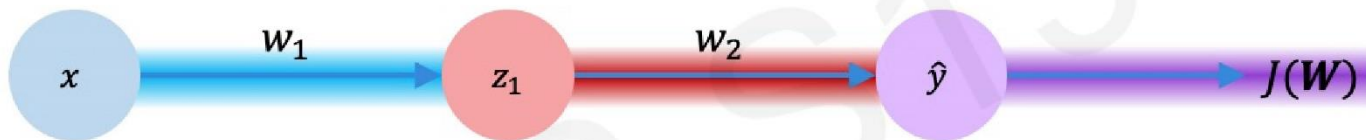
$$\frac{\partial J(W)}{\partial w_1} = \frac{\partial J(W)}{\partial \hat{y}} * \frac{\partial \hat{y}}{\partial w_1}$$

Apply chain rule!

Apply chain rule!



## Computing Gradients: Backpropagation

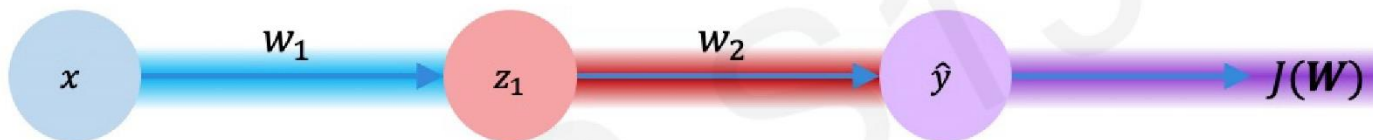


$$\frac{\partial J(\mathbf{W})}{\partial w_1} = \frac{\partial J(\mathbf{W})}{\partial \hat{y}} * \frac{\partial \hat{y}}{\partial z_1} * \frac{\partial z_1}{\partial w_1}$$

(Purple bar under  $\frac{\partial J(\mathbf{W})}{\partial \hat{y}}$ , Red bar under  $\frac{\partial \hat{y}}{\partial z_1}$ , Blue bar under  $\frac{\partial z_1}{\partial w_1}$ )



## 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

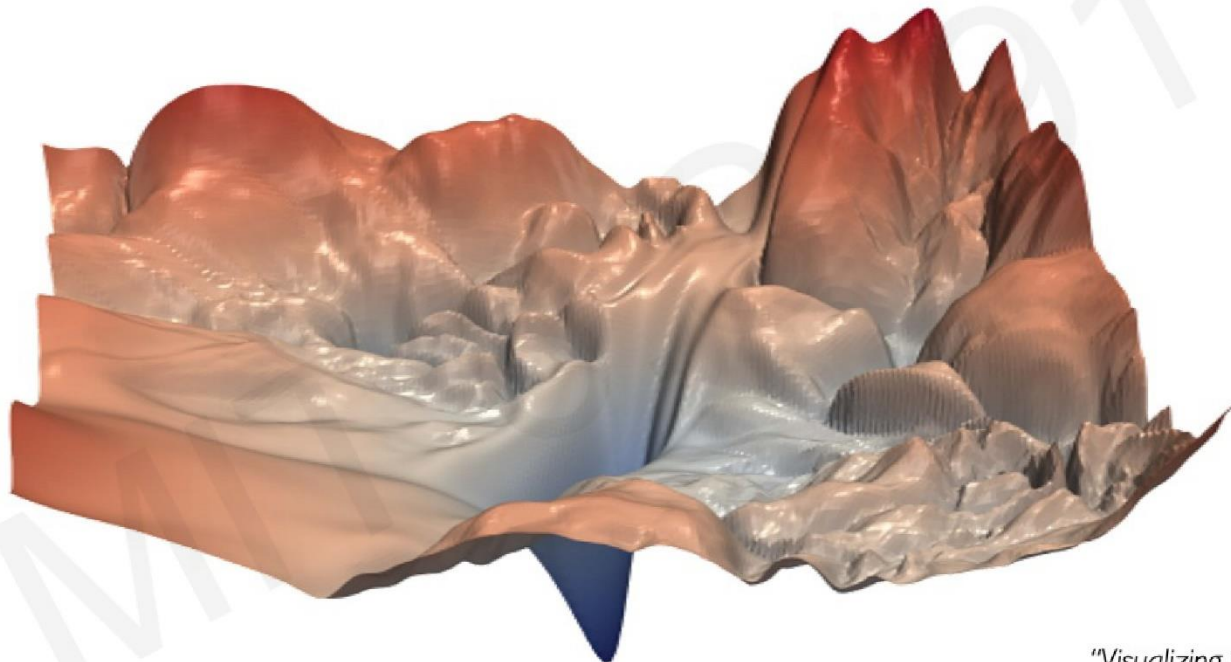


## Neural Networks in Practice: Optimization





## 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

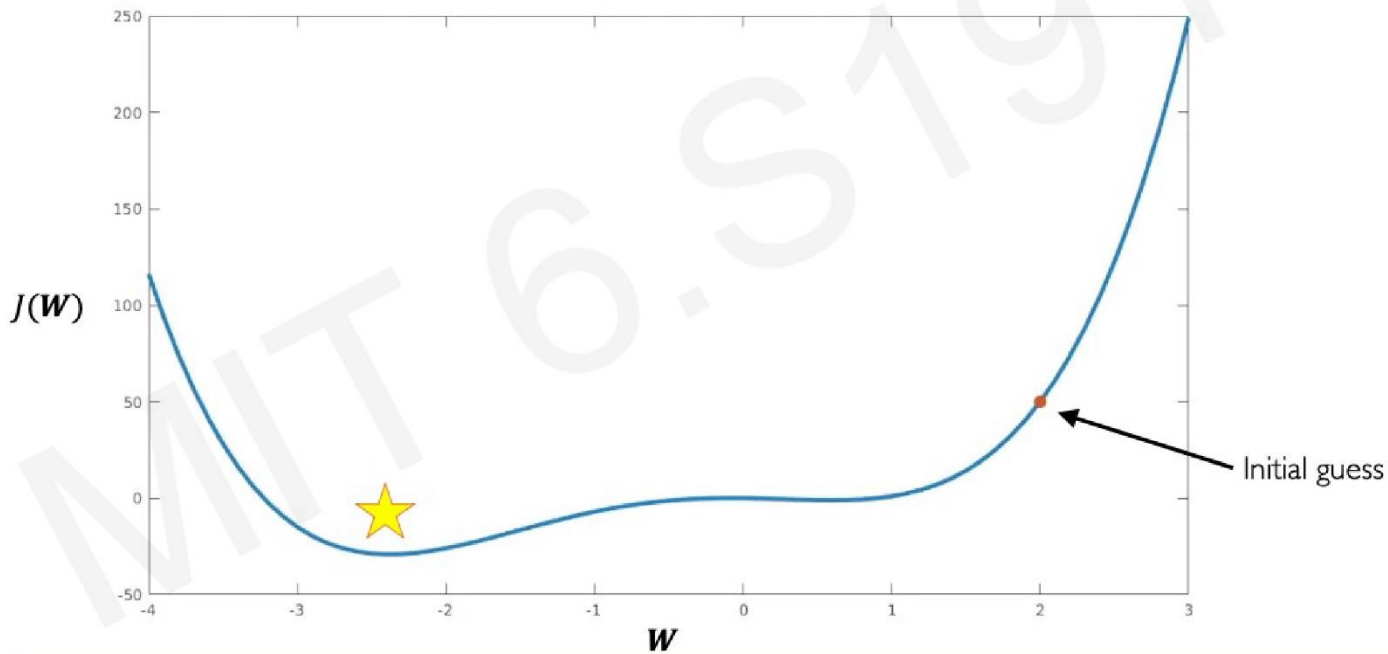
$$W \leftarrow W - \eta \frac{\partial J(W)}{\partial W}$$

How can we set the learning rate?



## Setting the Learning Rate

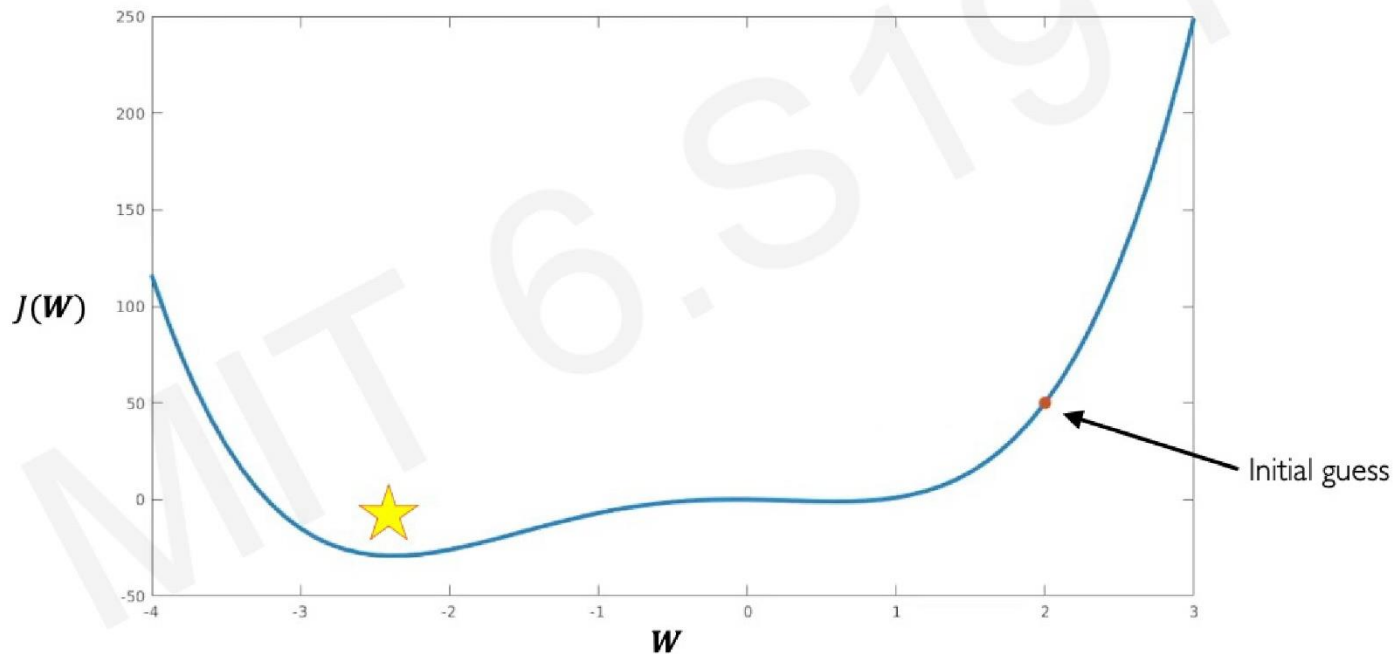
*Small learning rate converges slowly and gets stuck in false local minima*





## Setting the Learning Rate

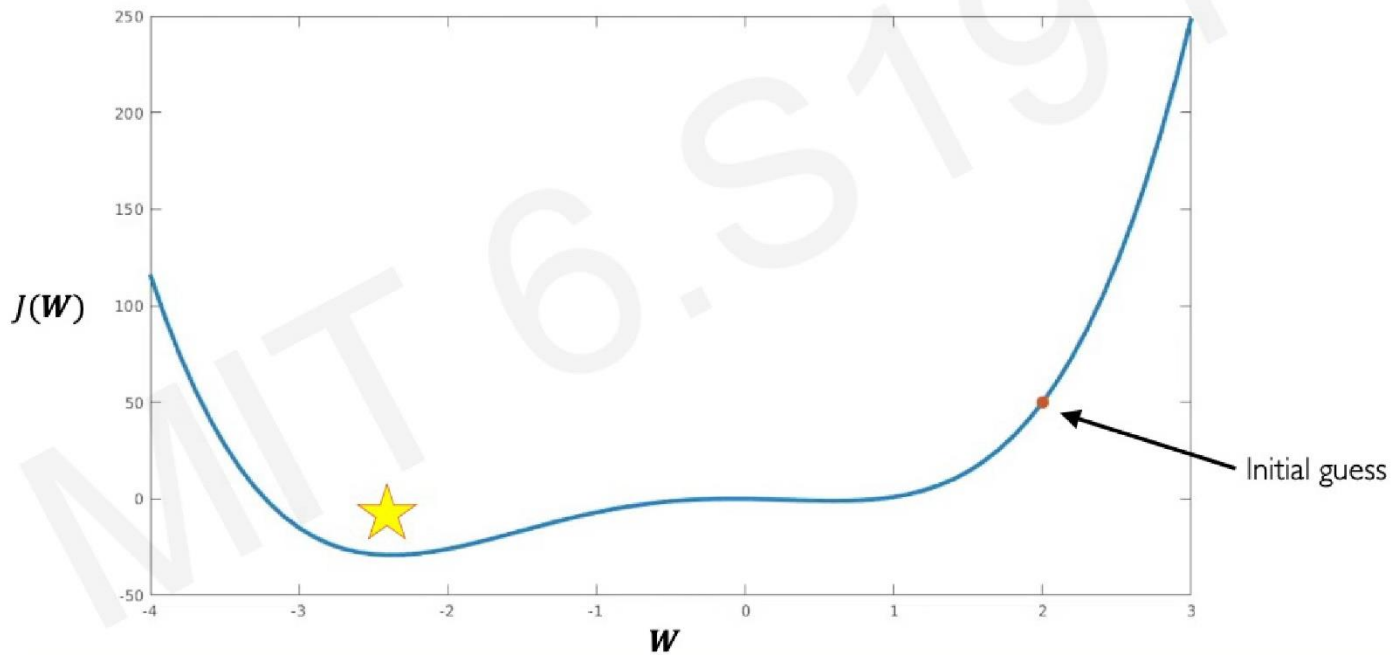
*Large learning rates overshoot, become unstable and diverge*





## 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”

MIT 6.S191



## 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 no longer fixed
- Can be made larger or smaller depending on:
  - how large gradient is
  - how fast learning is happening
  - size of particular weights
  - etc...



## 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/>



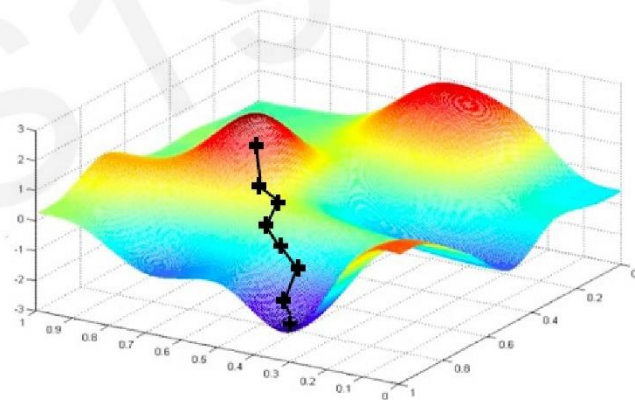
## Neural Networks in Practice: Mini-batches



## 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

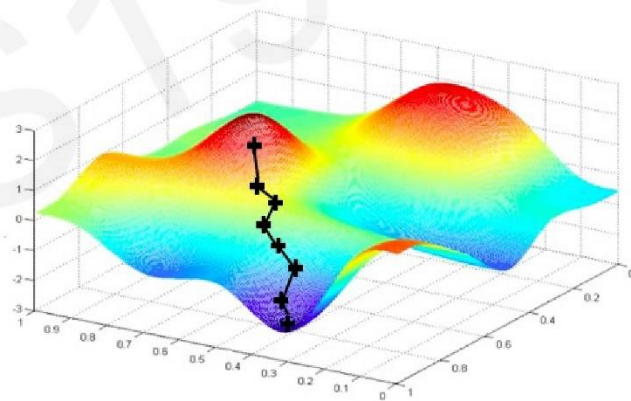




## Gradient Descent

### Algorithm

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



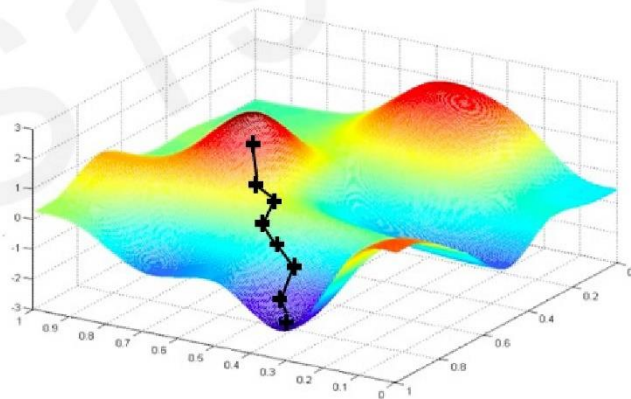
Can be very  
computationally  
intensive to compute!



## 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

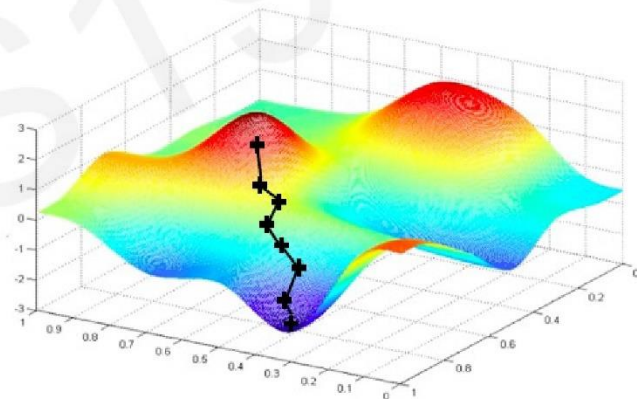




## Stochastic Gradient Descent

### 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



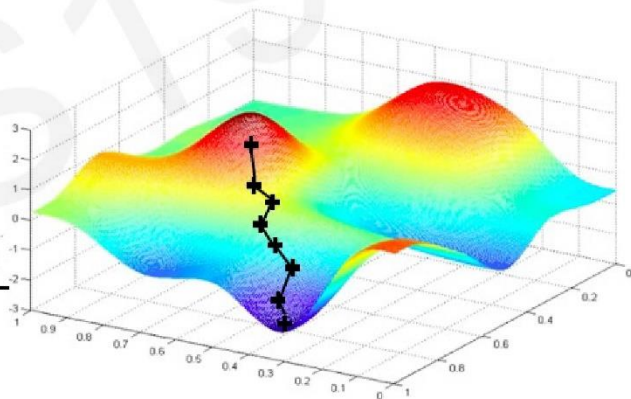
Easy to compute but  
very noisy (stochastic)!



## 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



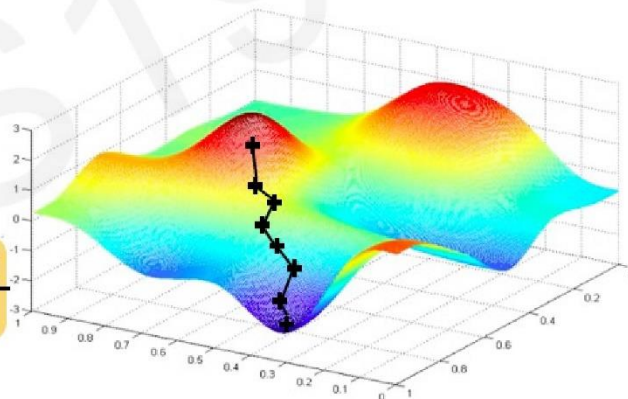




## Stochastic Gradient Descent

### 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**

Smoother convergence

Allows for larger learning rates



## Mini-batches while training

**More accurate estimation of gradient**

Smoother convergence

Allows for larger learning rates

**Mini-batches lead to fast training!**

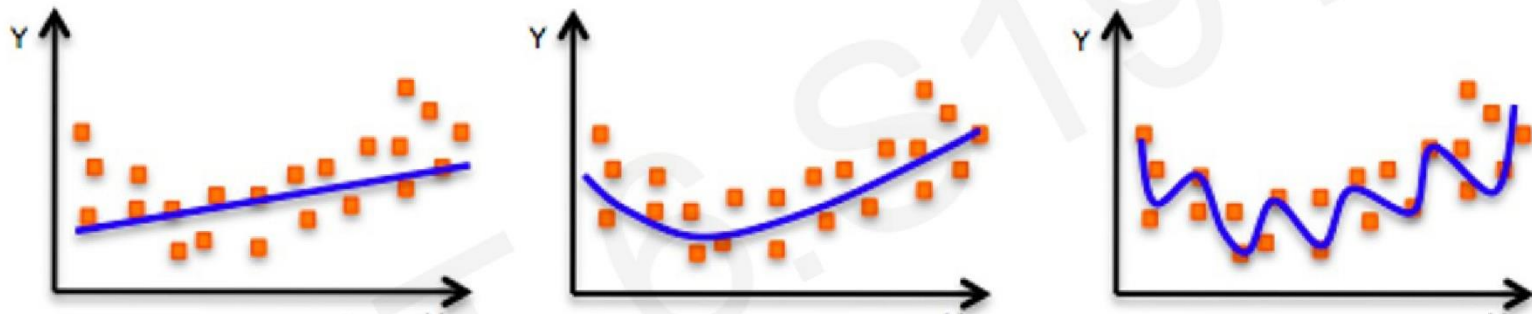
Can parallelize computation + achieve significant speed increases on GPU's



## 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*

MIT 6.S191



## Regularization

*What is it?*

*Technique that constrains our optimization problem to discourage complex models*

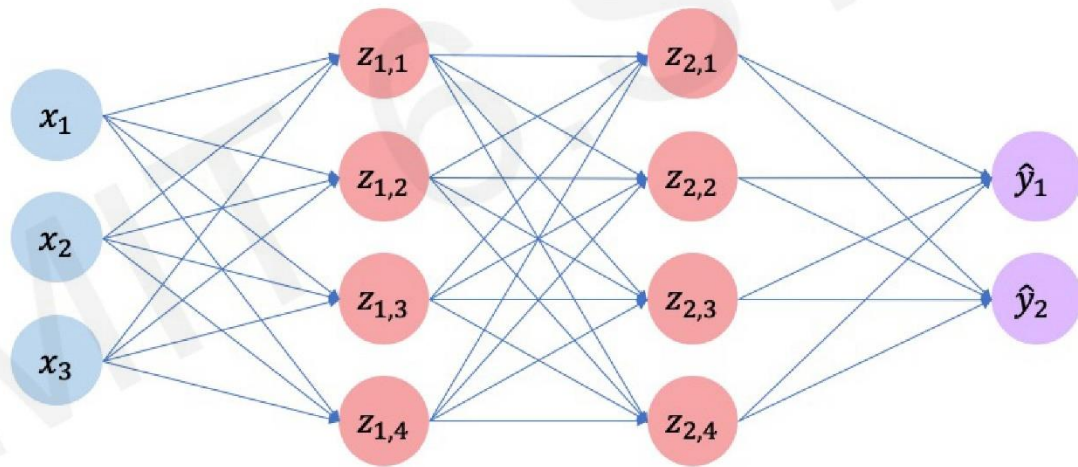
***Why do we need it?***

*Improve generalization of our model on unseen data*



## Regularization I: Dropout

- During training, randomly set some activations to 0



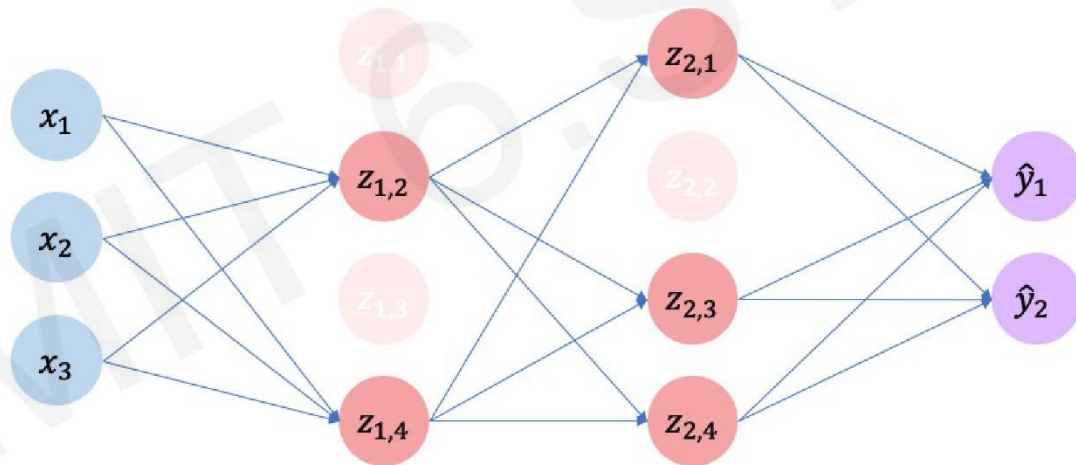




## Regularization I: Dropout

- During training, randomly set some activations to 0
  - Typically 'drop' 50% of activations in layer
  - Forces network to not rely on any 1 node

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

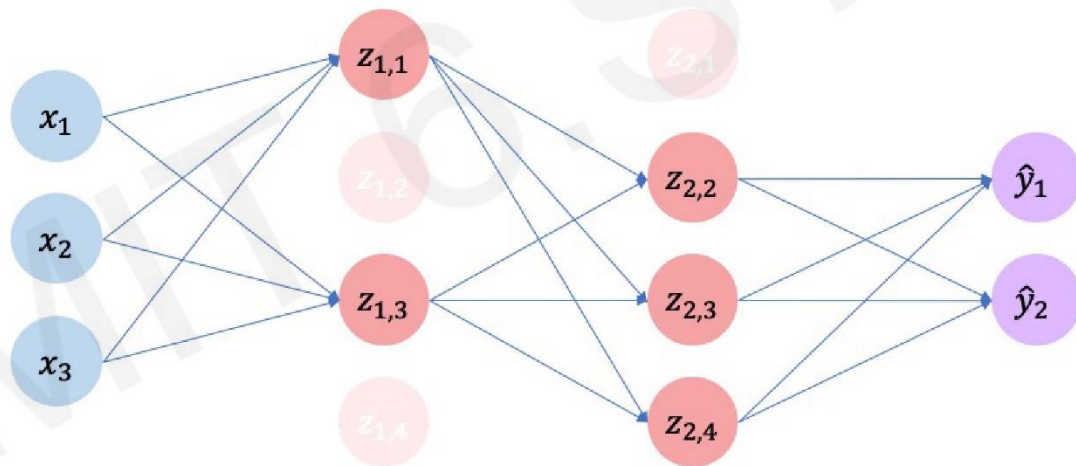




## Regularization I: Dropout

- During training, randomly set some activations to 0
  - Typically 'drop' 50% of activations in layer
  - Forces network to not rely on any 1 node

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





## Regularization 2: Early Stopping

- Stop training before we have a chance to overfit





## Regularization 2: Early Stopping

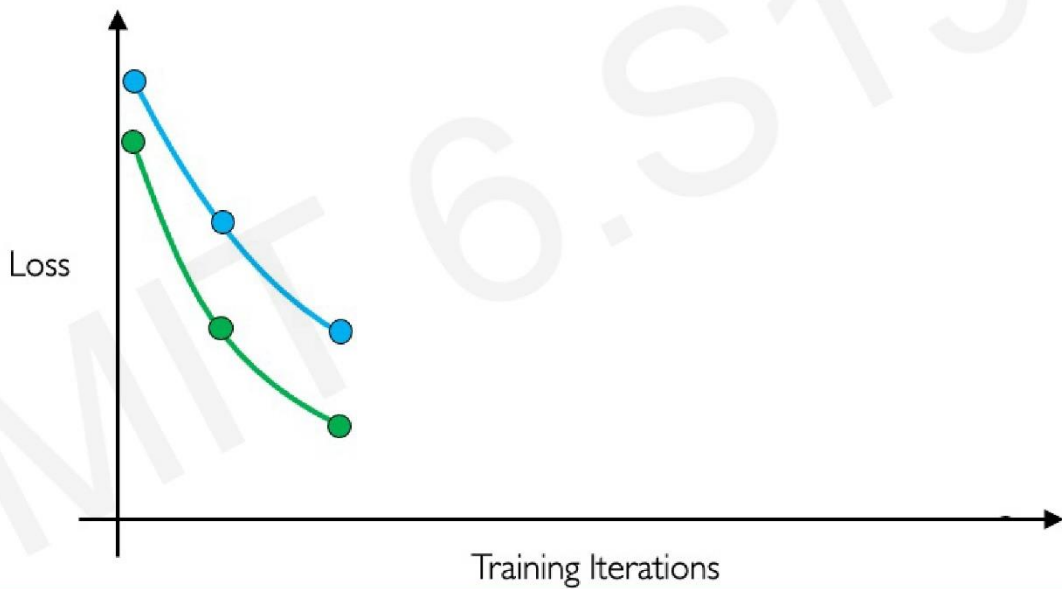
- Stop training before we have a chance to overfit





## Regularization 2: Early Stopping

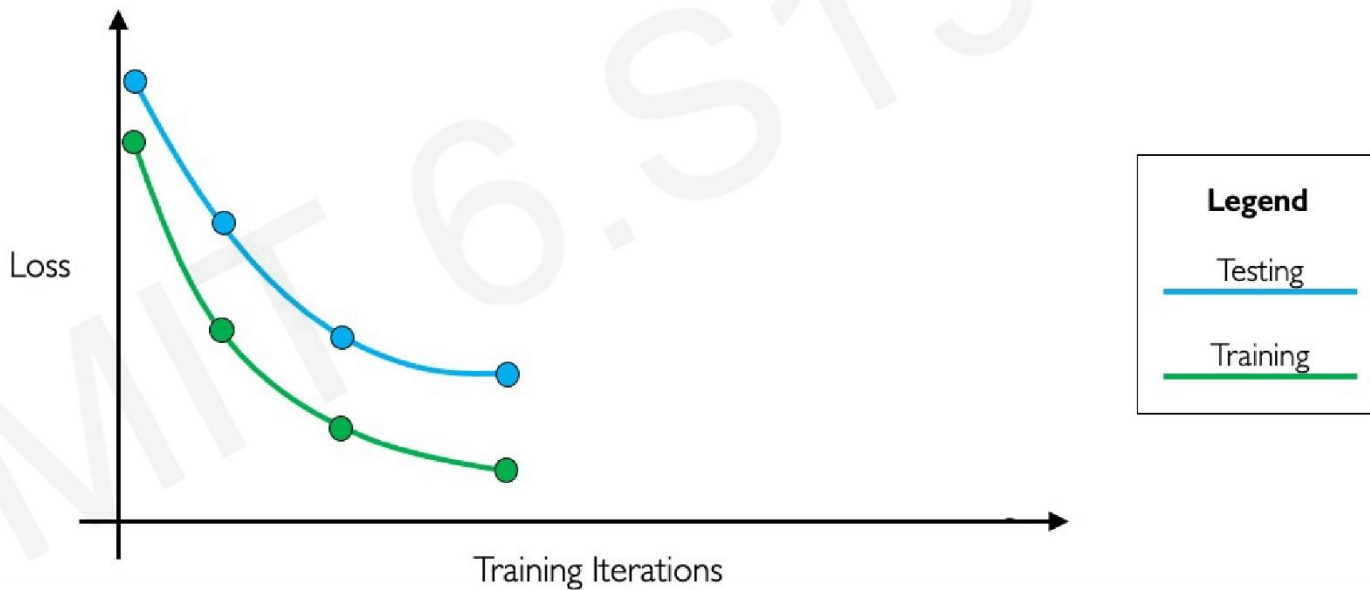
- Stop training before we have a chance to overfit





## Regularization 2: Early Stopping

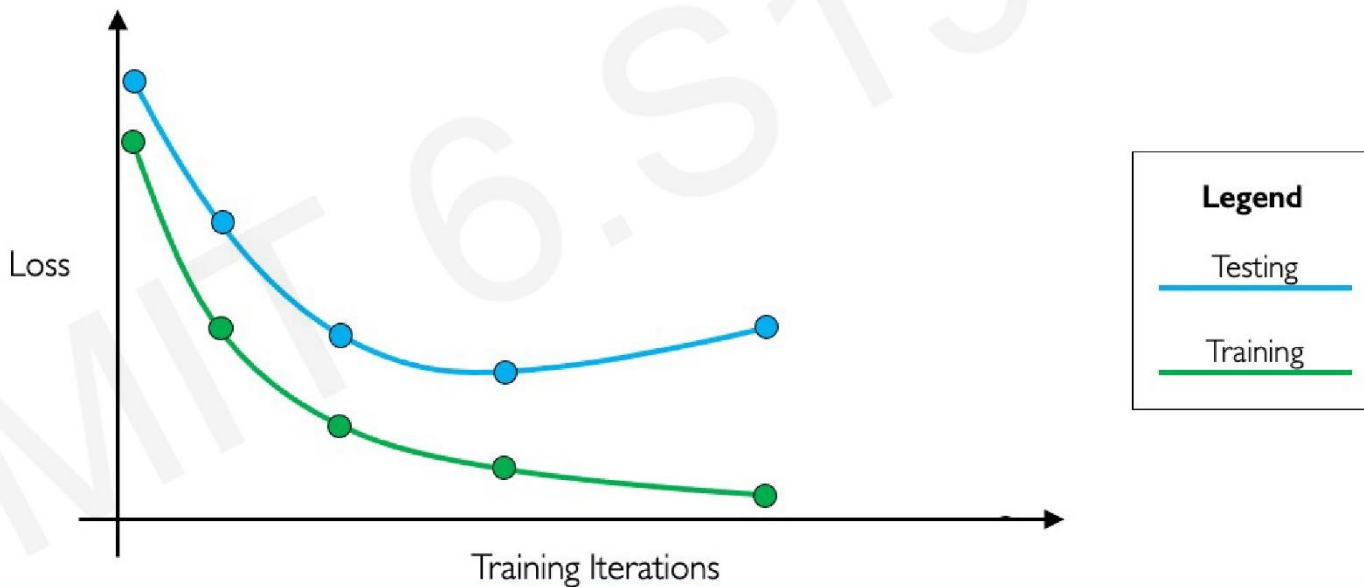
- Stop training before we have a chance to overfit





## Regularization 2: Early Stopping

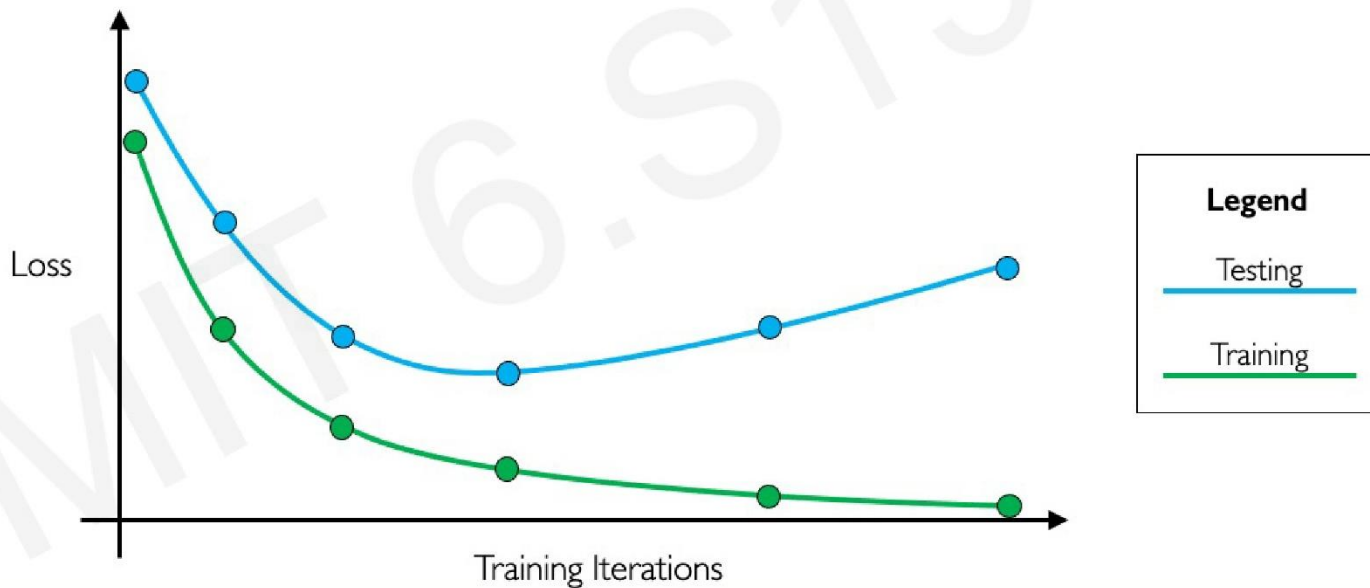
- Stop training before we have a chance to overfit





## Regularization 2: Early Stopping

- Stop training before we have a chance to overfit







## Regularization 2: Early Stopping

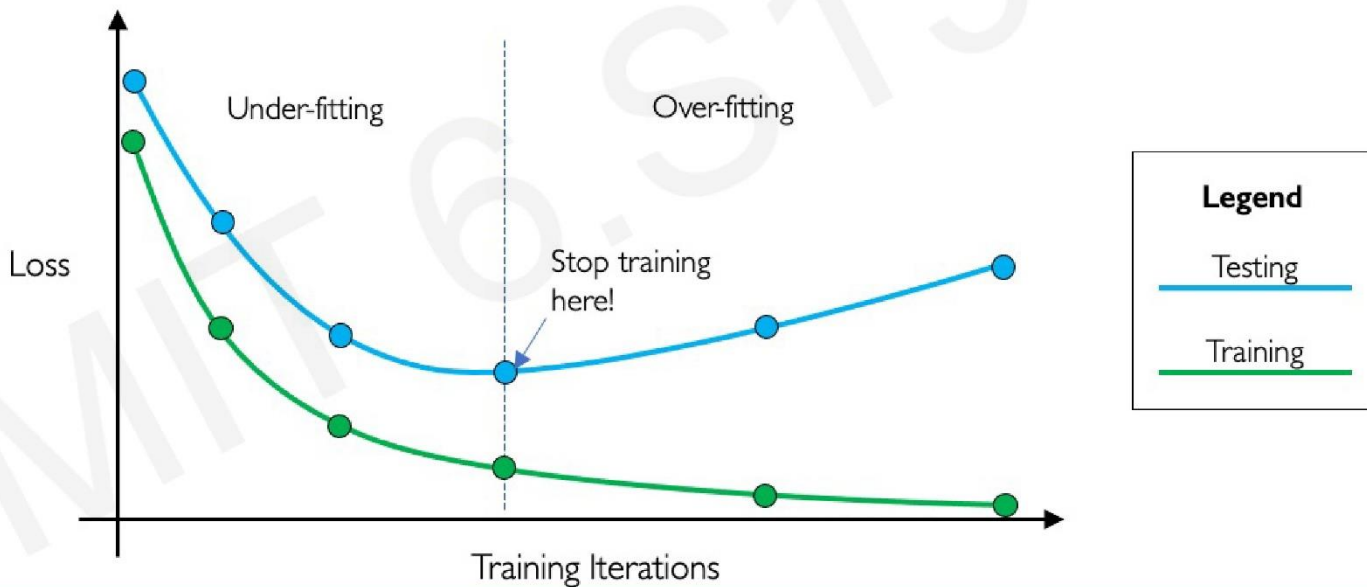
- Stop training before we have a chance to overfit





## Regularization 2: Early Stopping

- Stop training before we have a chance to overfit

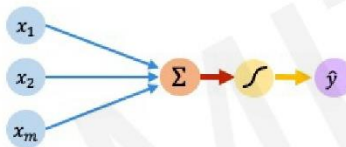




## Core Foundation Review

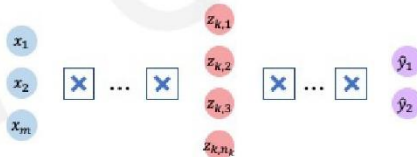
### The Perceptron

- Structural building blocks
- Nonlinear activation functions



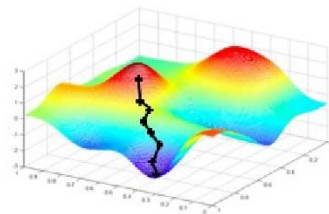
### Neural Networks

- Stacking Perceptrons to form neural networks
- Optimization through backpropagation



### Training in Practice

- Adaptive learning
- Batching
- Regularization



# 提纲

---

一、前馈网络

二、循环网络

三、卷积网络



上海大学  
SHANGHAI UNIVERSITY



# 人工智能基本理论

## Sequences in the Wild





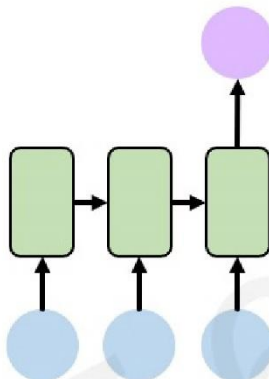
## Sequence Modeling Applications



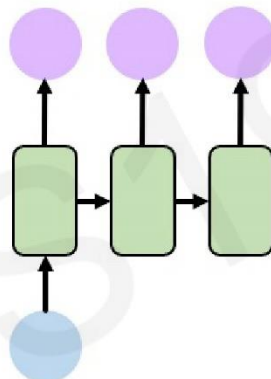
One to One  
**Binary Classification**



"Will I pass this class?"  
Student → Pass?



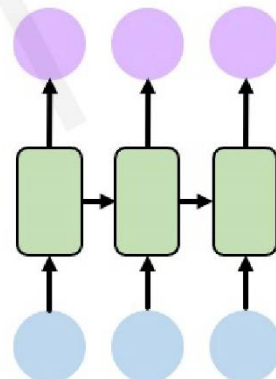
Many to One  
**Sentiment Classification**



One to Many  
**Image Captioning**



"A baseball player throws a ball!"



Many to Many  
**Machine Translation**



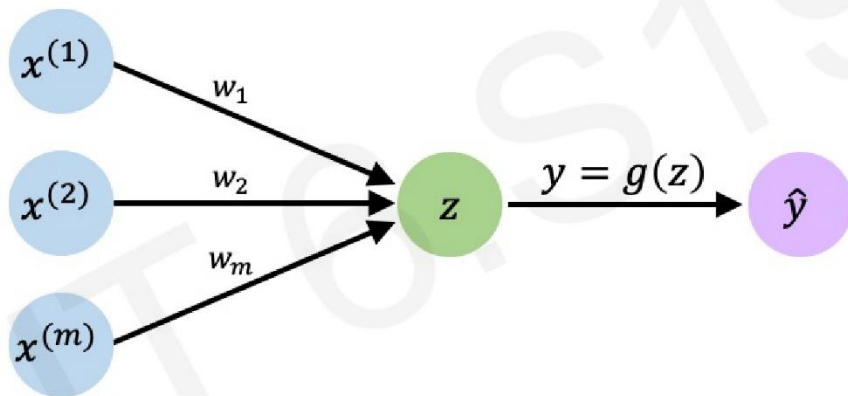


## Neurons with Recurrence

MIT 6.S191



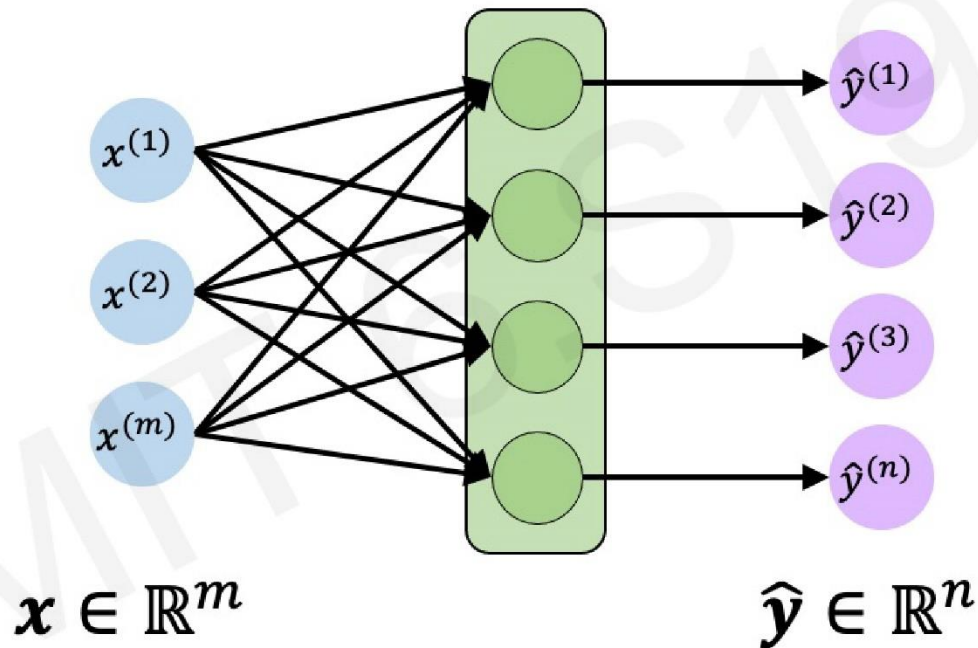
## The Perceptron Revisited





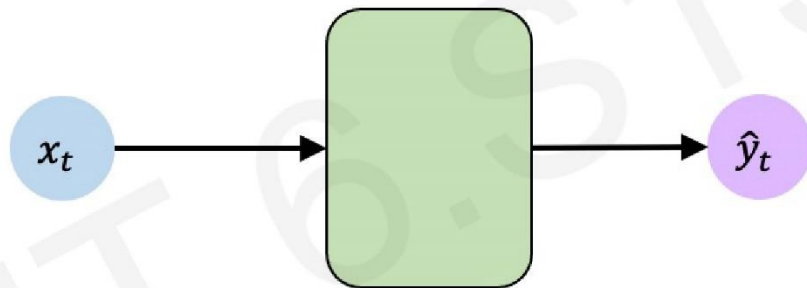


## Feed-Forward Networks Revisited





# Feed-Forward Networks Revisited

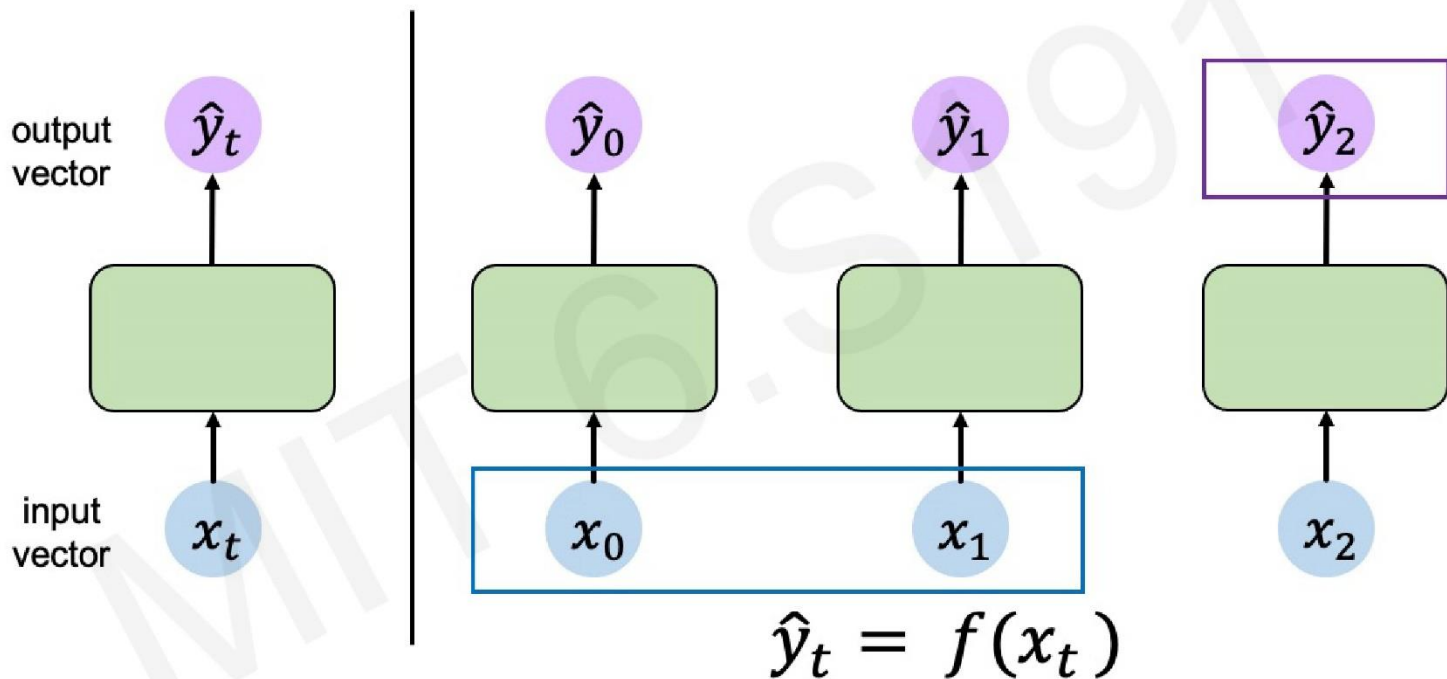


$$x_t \in \mathbb{R}^m$$

$$\hat{y}_t \in \mathbb{R}^n$$

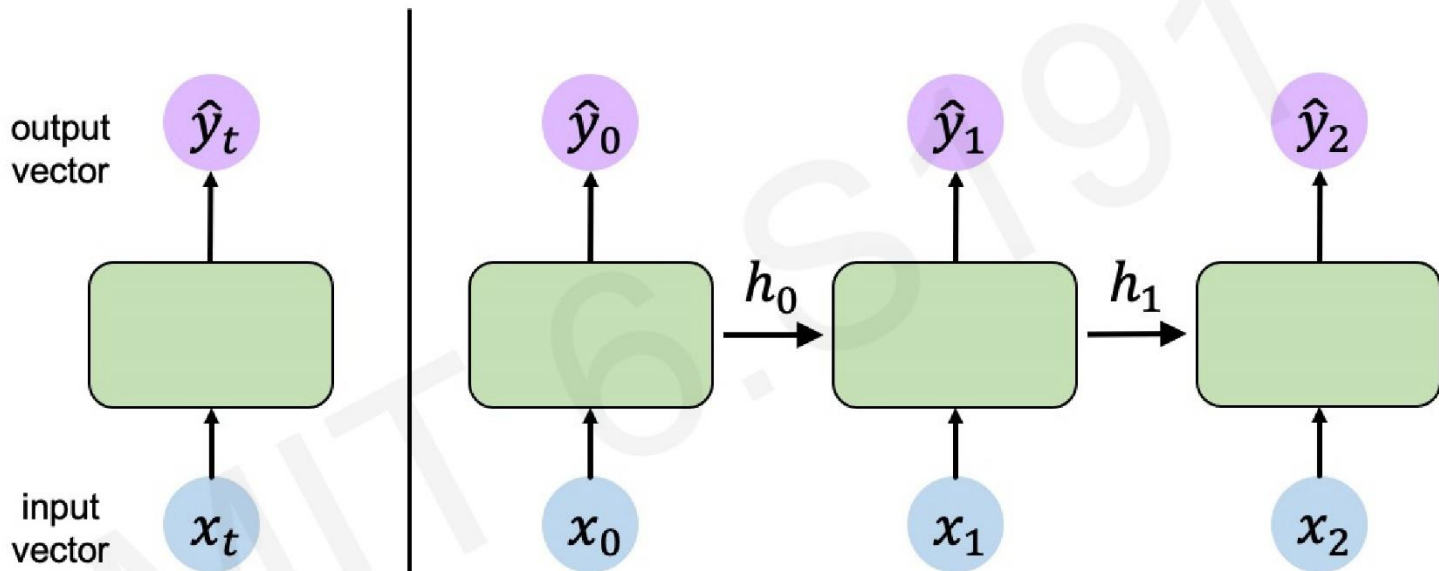


## Handling Individual Time Steps





## Neurons with Recurrence

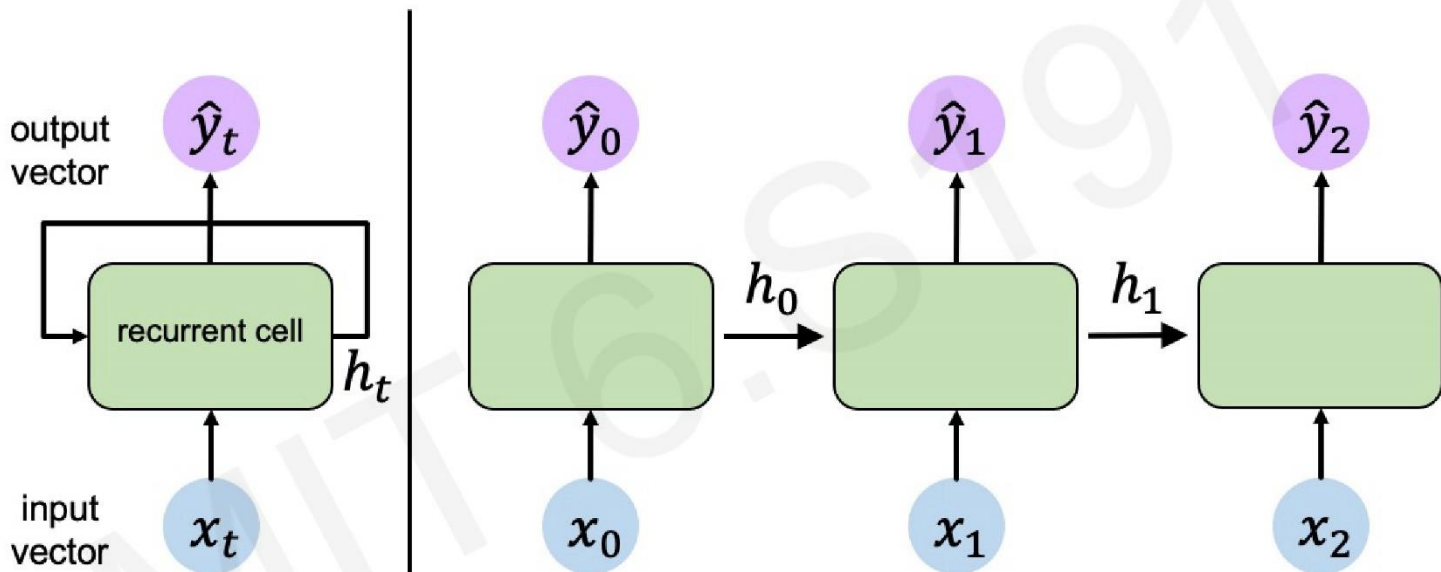


$$\hat{y}_t = f(x_t, h_{t-1})$$

output      input      past memory



## Neurons with Recurrence



$$\hat{y}_t = f(x_t, h_{t-1})$$

output      input      past memory

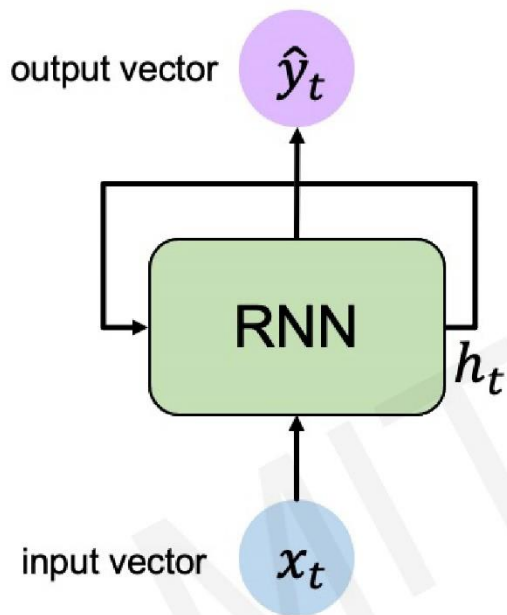


## Recurrent Neural Networks (RNNs)

MIT 6.S191



## Recurrent Neural Networks (RNNs)



Apply a **recurrence relation** at every time step to process a sequence:

$$h_t = f_W(x_t, h_{t-1})$$

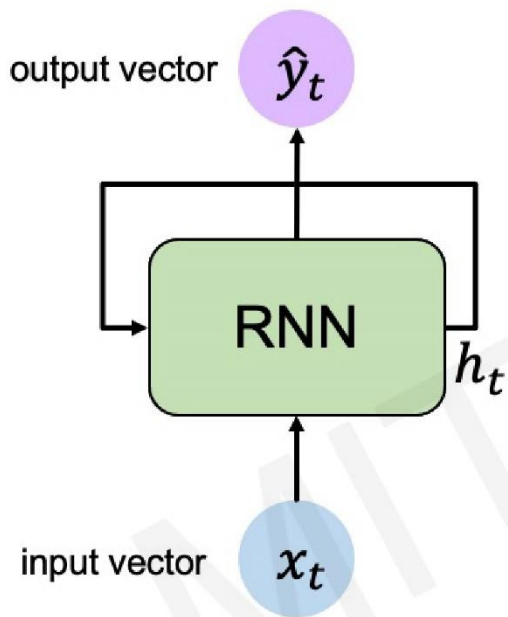
cell state      function with weights  $W$       input      old state

Note: the same function and set of parameters are used at every time step

RNNs have a **state**,  $h_t$ , that is updated **at each time step** as a sequence is processed



## RNN State Update and Output



Update Hidden State

$$h_t = \tanh(W_{hh}^T h_{t-1} + W_{xh}^T x_t)$$

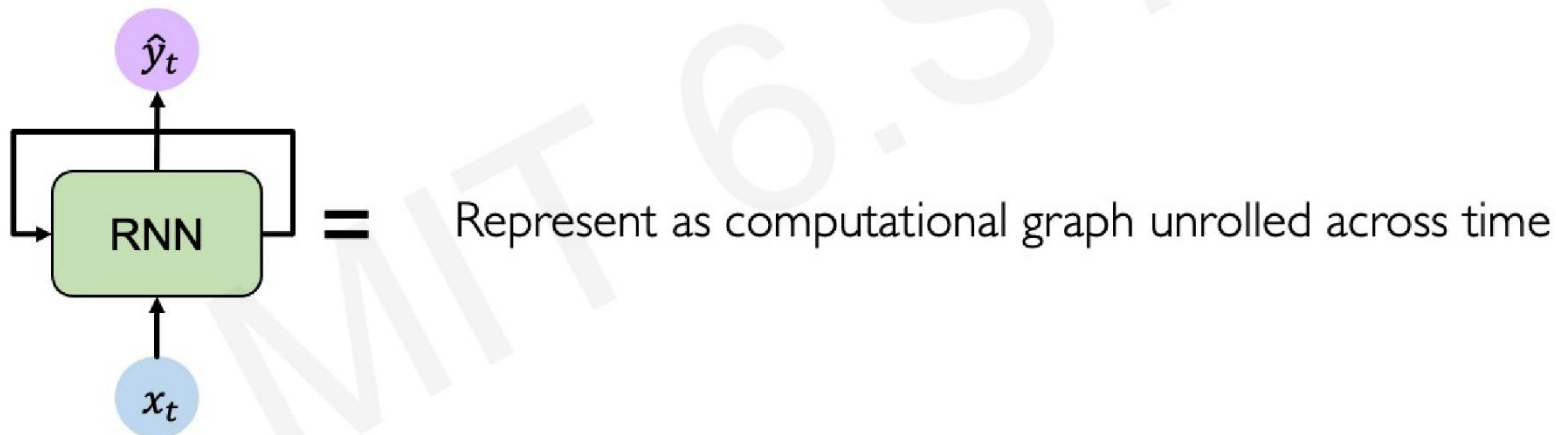
Input Vector

$x_t$



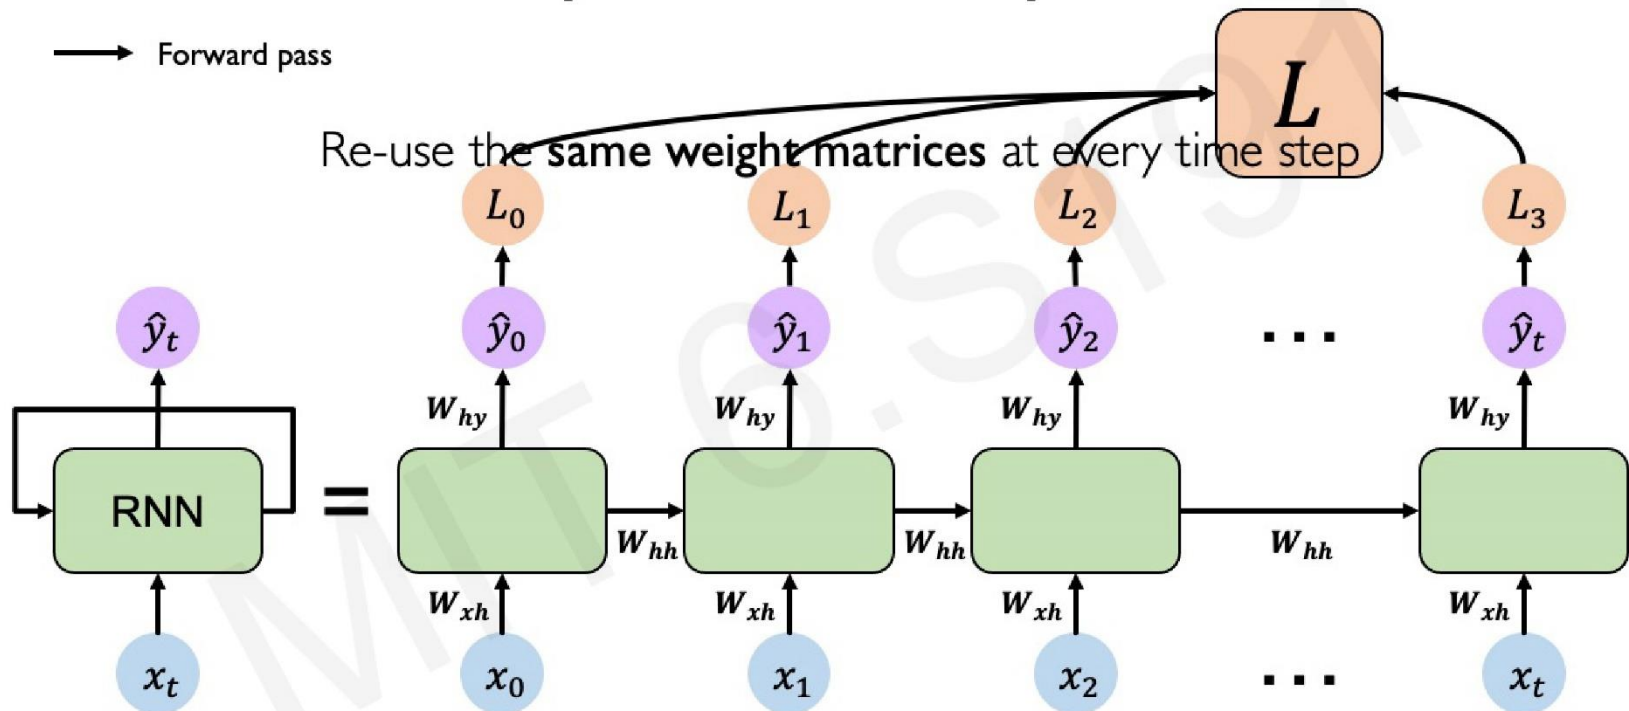


## RNNs: Computational Graph Across Time





## RNNs: Computational Graph Across Time

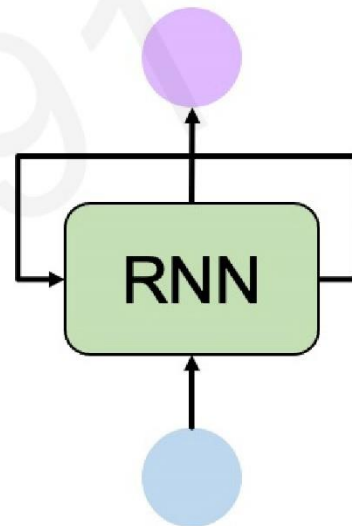




## Sequence Modeling: Design Criteria

To model sequences, we need to:

1. Handle **variable-length** sequences
2. Track **long-term** dependencies
3. Maintain information about **order**
4. **Share parameters** across the sequence



**Recurrent Neural Networks (RNNs)** meet these sequence modeling design criteria



A Sequence Modeling Problem:  
Predict the Next Word



## A Sequence Modeling Problem: Predict the Next Word

“This morning I took my cat for a walk.”

MIT 6.S191



## A Sequence Modeling Problem: Predict the Next Word

“This morning I took my cat for a walk.”

given these words

MIT 6.S191



## A Sequence Modeling Problem: Predict the Next Word

“This morning I took my cat for a walk.”

given these words

predict the  
next word

MIT 6.S191



## A Sequence Modeling Problem: Predict the Next Word

“This morning I took my cat for a walk.”

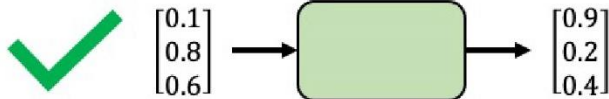
given these words

predict the  
next word

### Representing Language to a Neural Network



*Neural networks cannot interpret words*

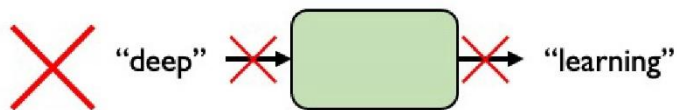


*Neural networks require numerical inputs*

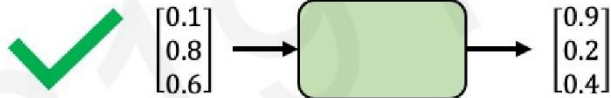




## Encoding Language for a Neural Network

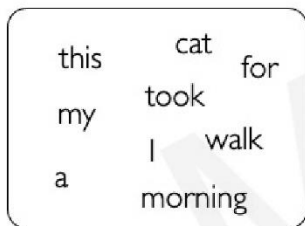


Neural networks cannot interpret words

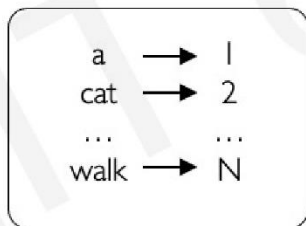


Neural networks require numerical inputs

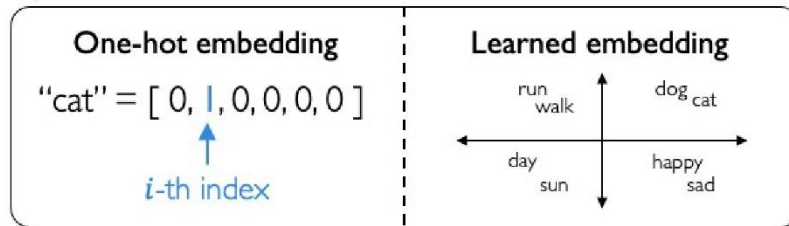
Embedding: transform indexes into a vector of fixed size.



**1. Vocabulary:**  
Corpus of words



**2. Indexing:**  
Word to index



**3. Embedding:**  
Index to fixed-sized vector



## Handle Variable Sequence Lengths

The food was great

vs.

We visited a restaurant for lunch

vs.

We were hungry but cleaned the house before eating



## Model Long-Term Dependencies

“**France** is where I grew up, but I now live in Boston. I speak fluent \_\_\_\_.”



We need information from **the distant past** to accurately predict the correct word.



## Capture Differences in Sequence Order



The food was good, not bad at all.

vs.

The food was bad, not good at all.

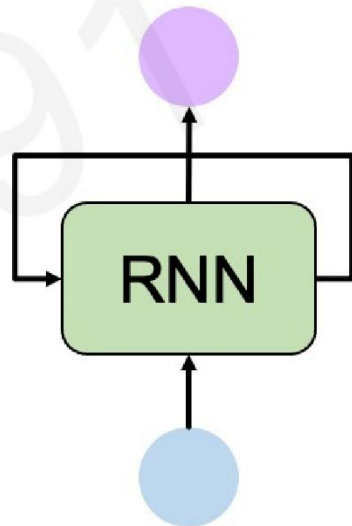




## Sequence Modeling: Design Criteria

To model sequences, we need to:

1. Handle **variable-length** sequences
2. Track **long-term** dependencies
3. Maintain information about **order**
4. **Share parameters** across the sequence



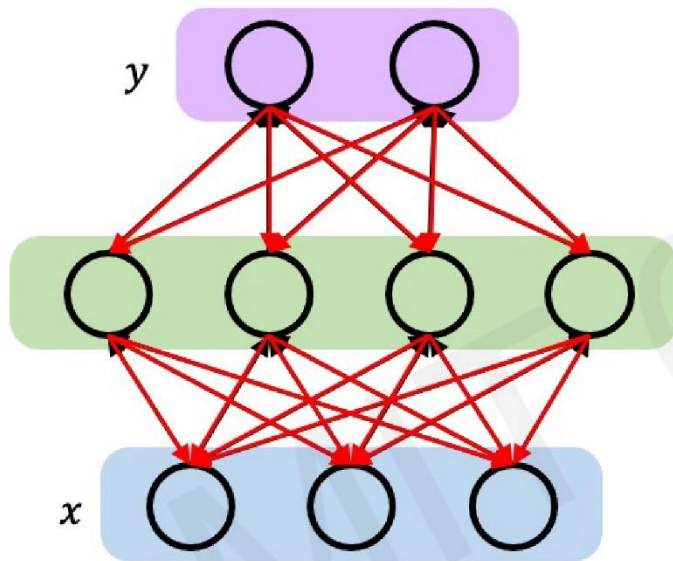
**Recurrent Neural Networks (RNNs)** meet these sequence modeling design criteria



# Backpropagation Through Time (BPTT)



## Recall: Backpropagation in Feed Forward Models

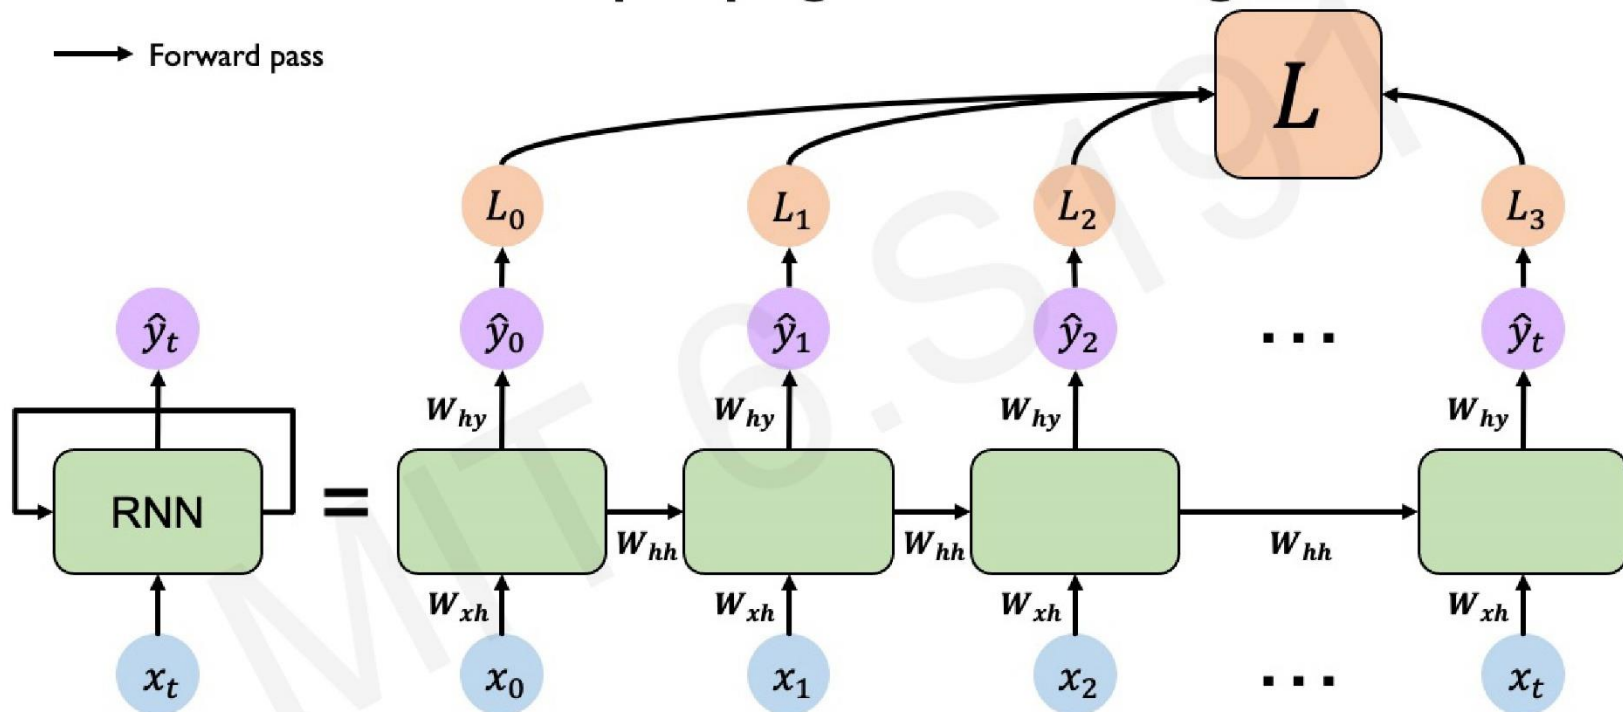


### Backpropagation algorithm:

1. Take the derivative (gradient) of the loss with respect to each parameter
2. Shift parameters in order to minimize loss



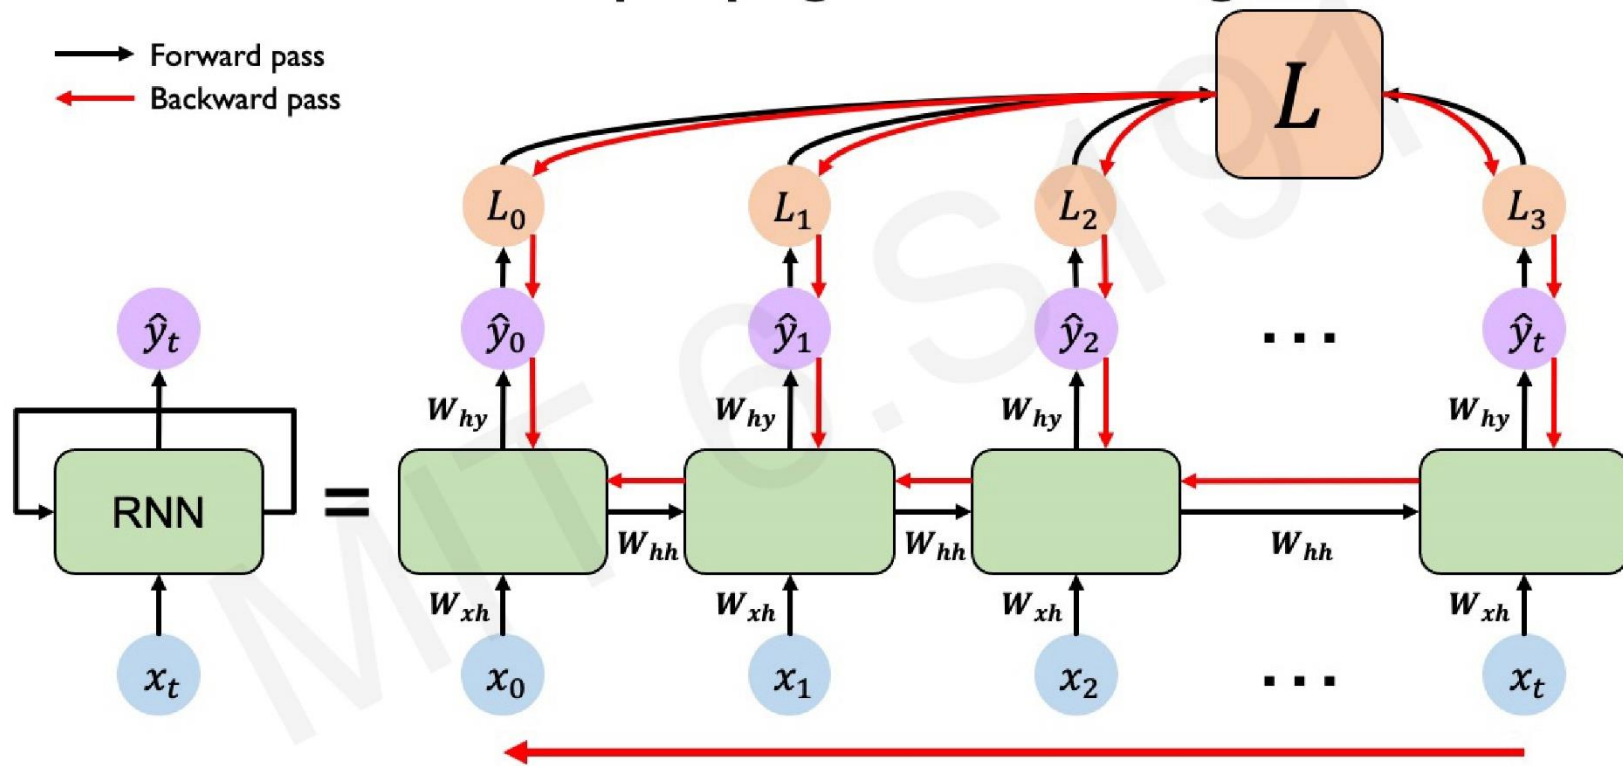
## RNNs: Backpropagation Through Time





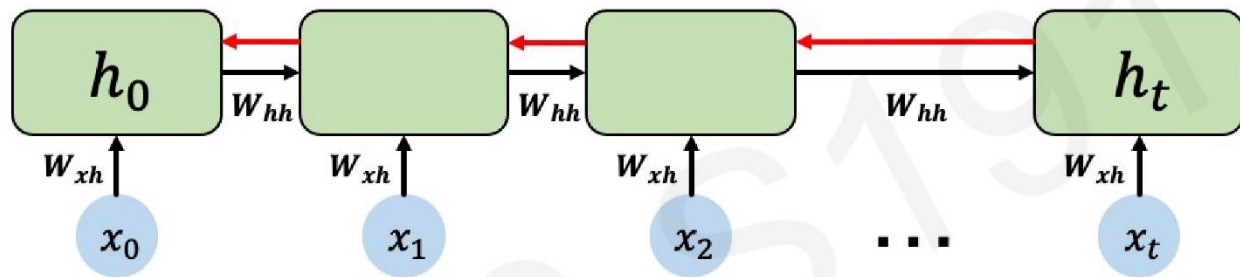


## RNNs: Backpropagation Through Time



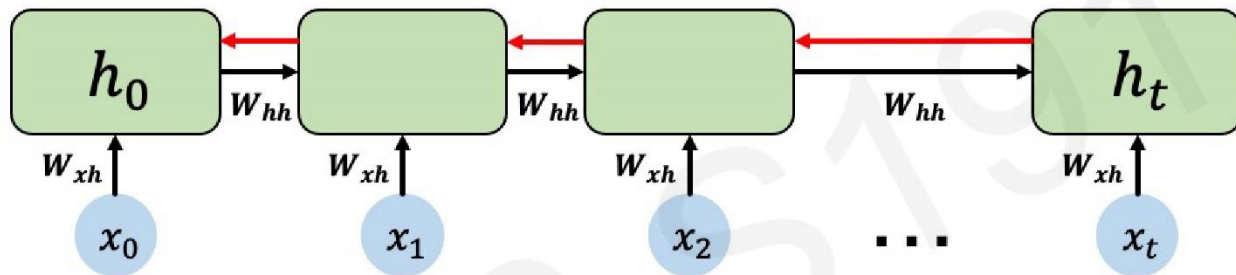


## Standard RNN Gradient Flow





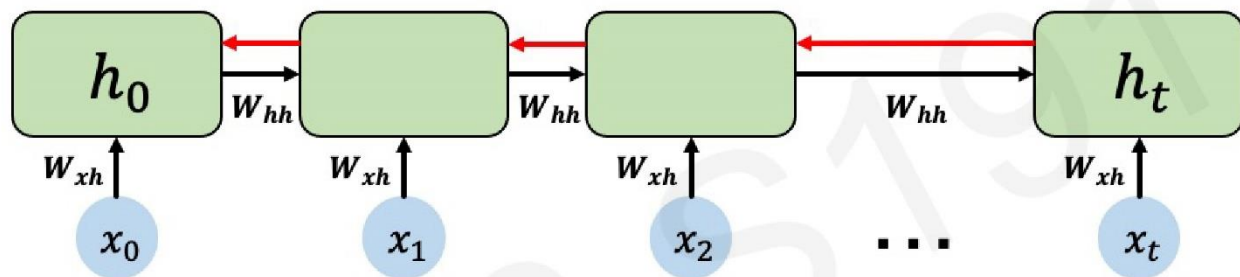
## Standard RNN Gradient Flow



Computing the gradient wrt  $h_0$  involves many factors of  $W_{hh}$  + repeated gradient computation!



## Standard RNN Gradient Flow: Exploding Gradients



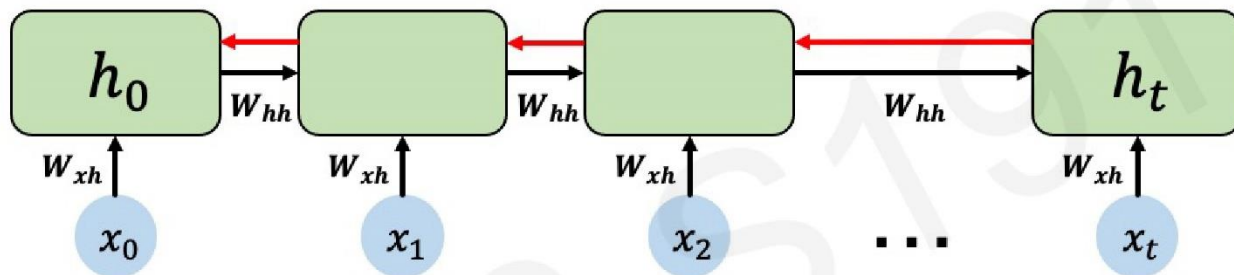
Computing the gradient wrt  $h_0$  involves many factors of  $W_{hh}$  + repeated gradient computation!

Many values  $> 1$ :  
**exploding gradients**

Gradient clipping to  
scale big gradients



## Standard RNN Gradient Flow: Vanishing Gradients



Computing the gradient wrt  $h_0$  involves **many factors** of  $W_{hh}$  + repeated gradient computation!

Many values  $> 1$ :  
exploding gradients

Gradient clipping to  
scale big gradients

Many values  $< 1$ :  
vanishing gradients

1. Activation function
2. Weight initialization
3. Network architecture



## The Problem of Long-Term Dependencies

Why are vanishing gradients a problem?

Multiply many **small numbers** together



Errors due to further back time steps  
have smaller and smaller gradients



Bias parameters to capture short-term  
dependencies



## The Problem of Long-Term Dependencies

“The clouds are in the \_\_\_\_”

Why are vanishing gradients a problem?

Multiply many **small numbers** together



Errors due to further back time steps  
have smaller and smaller gradients



Bias parameters to capture short-term  
dependencies



## The Problem of Long-Term Dependencies

Why are vanishing gradients a problem?

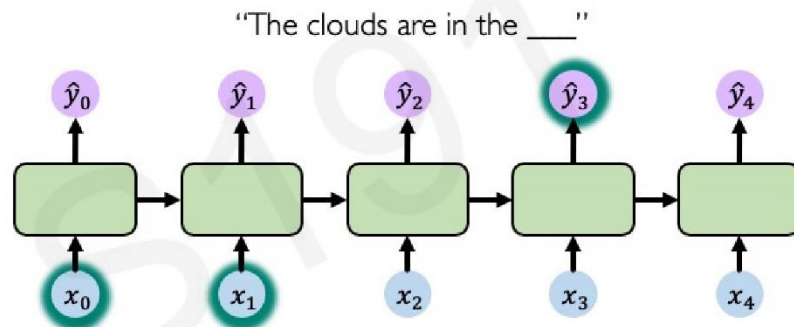
Multiply many **small numbers** together



Errors due to further back time steps  
have smaller and smaller gradients



Bias parameters to capture short-term  
dependencies







## The Problem of Long-Term Dependencies

Why are vanishing gradients a problem?

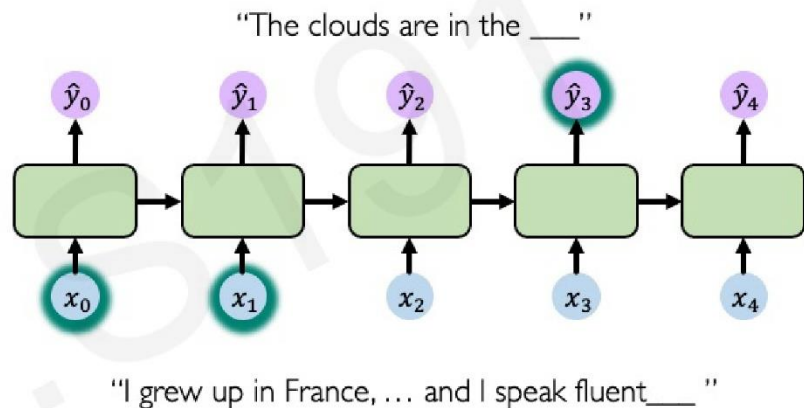
Multiply many **small numbers** together



Errors due to further back time steps  
have smaller and smaller gradients



Bias parameters to capture short-term  
dependencies





## The Problem of Long-Term Dependencies

Why are vanishing gradients a problem?

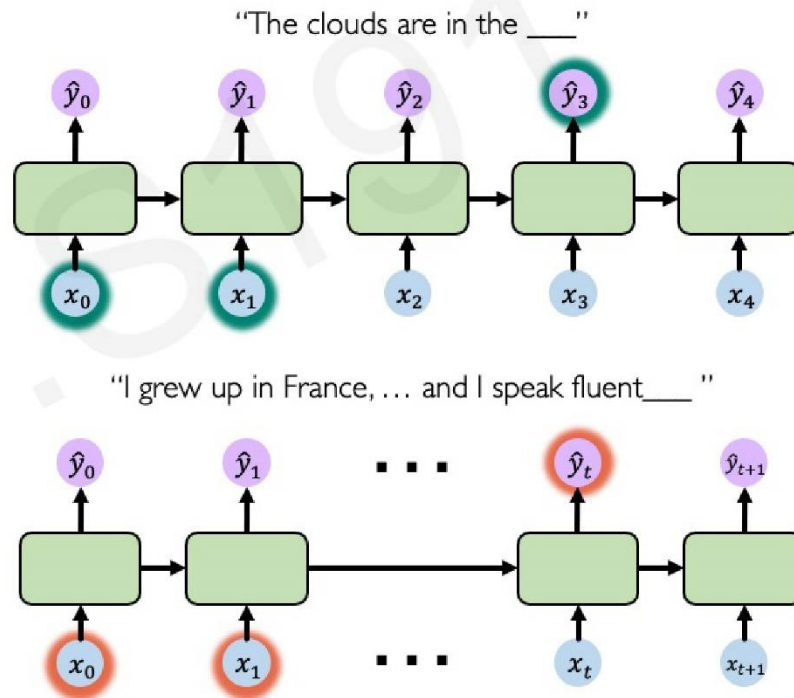
Multiply many **small numbers** together



Errors due to further back time steps  
have smaller and smaller gradients

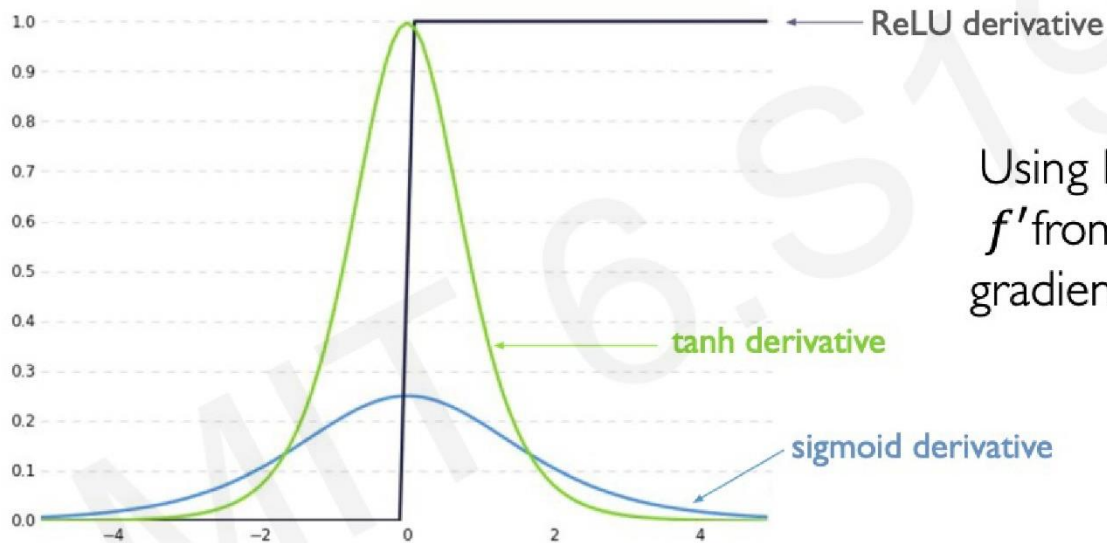


Bias parameters to capture short-term  
dependencies





## Trick #1: Activation Functions



Using ReLU prevents  $f'$  from shrinking the gradients when  $x > 0$



## Trick #2: Parameter Initialization

Initialize **weights** to identity matrix

Initialize **biases** to zero

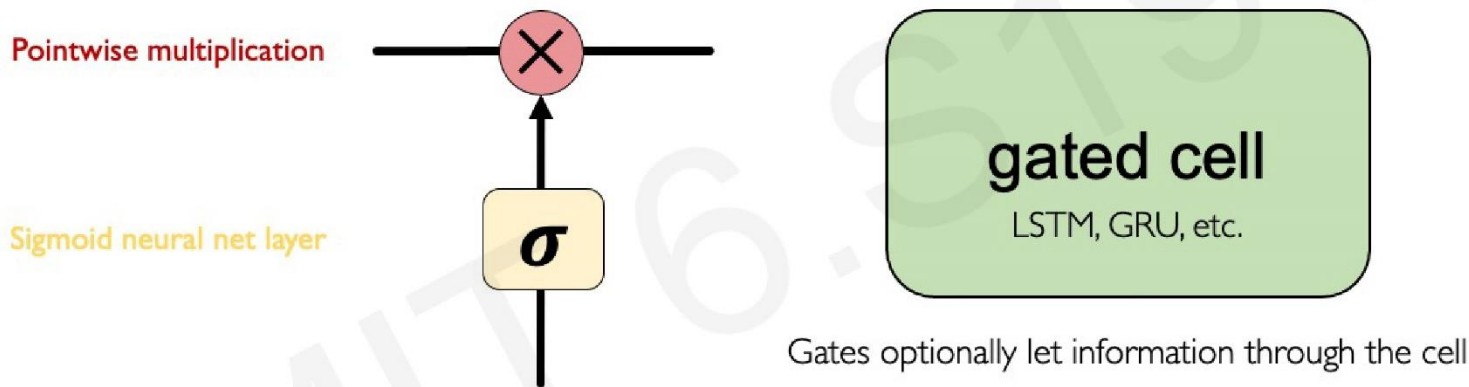
$$I_n = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{pmatrix}$$

This helps prevent the weights from shrinking to zero.



## Trick #3: Gated Cells

Idea: use **gates** to selectively **add** or **remove** information within **each recurrent unit** with



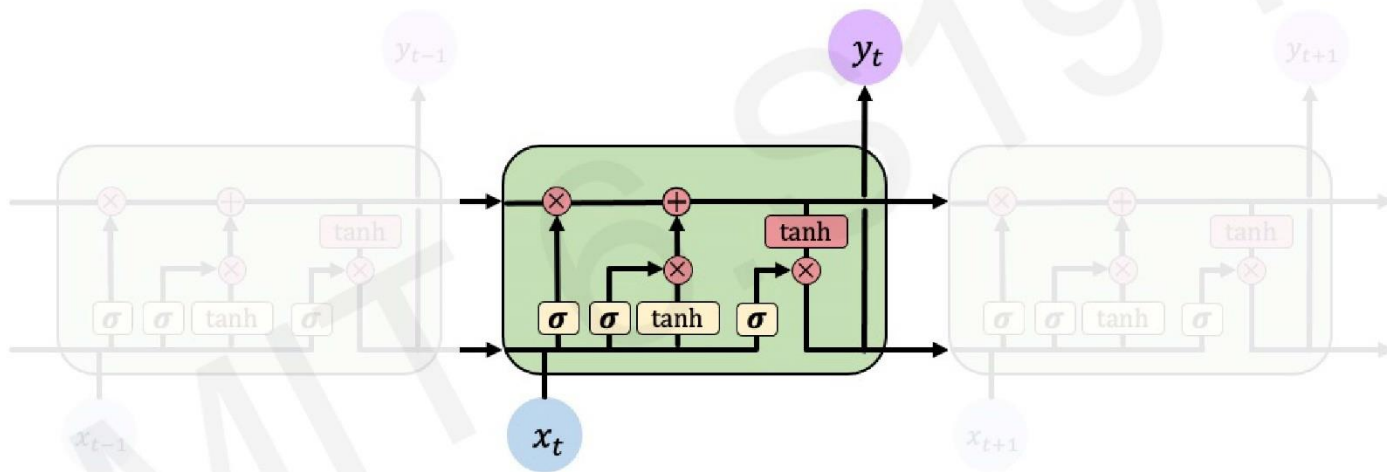
**Long Short Term Memory (LSTMs)** networks rely on a gated cell to track information throughout many time steps.



## Long Short Term Memory (LSTMs)

Gated LSTM cells control information flow:

- 1) Forget
- 2) Store
- 3) Update
- 4) Output



LSTM cells are able to track information throughout many timesteps



```
tf.keras.layers.LSTM(num_units)
```



## LSTMs: Key Concepts

1. Maintain a **cell state**
2. Use **gates** to control the **flow of information**
  - **Forget** gate gets rid of irrelevant information
  - **Store** relevant information from current input
  - Selectively **update** cell state
  - **Output** gate returns a filtered version of the cell state
3. Backpropagation through time with partially **uninterrupted gradient flow**



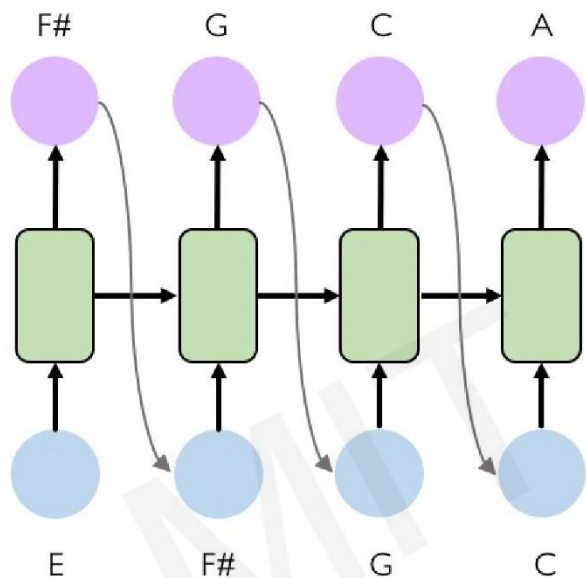
## RNN Applications & Limitations

MIT 6.S191





## Example Task: Music Generation



**Input:** sheet music

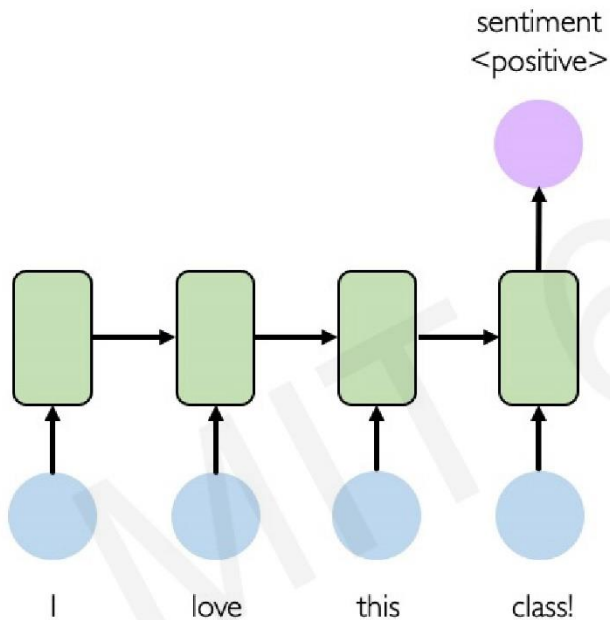
**Output:** next character in sheet music

Listening to  
3rd movement





## Example Task: Sentiment Classification



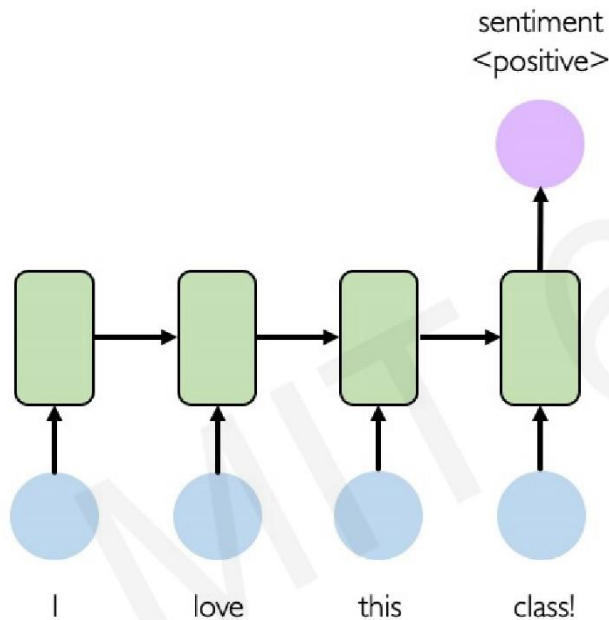
**Input:** sequence of words

**Output:** probability of having positive sentiment

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



## Example Task: Sentiment Classification



### Tweet sentiment classification



Ivar Hagendoorn  
@IvarHagendoorn

Follow



The @MIT Introduction to #DeepLearning is definitely one of the best courses of its kind currently available online [introtodeeplearning.com](http://introtodeeplearning.com)

12:45 PM - 12 Feb 2018



Angels-Cave  
@AngelsCave

Follow



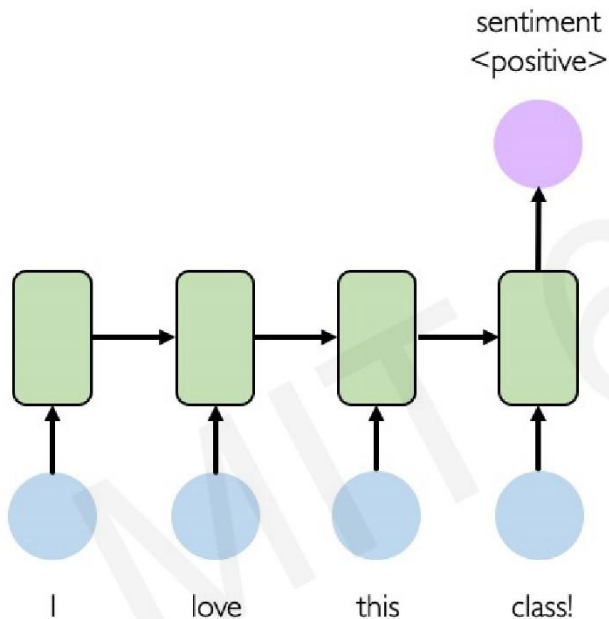
Replying to @Kazuki2048

I wouldn't mind a bit of snow right now. We haven't had any in my bit of the Midlands this winter! :(




2:19 AM - 25 Jan 2019



## Limitations of Recurrent Models



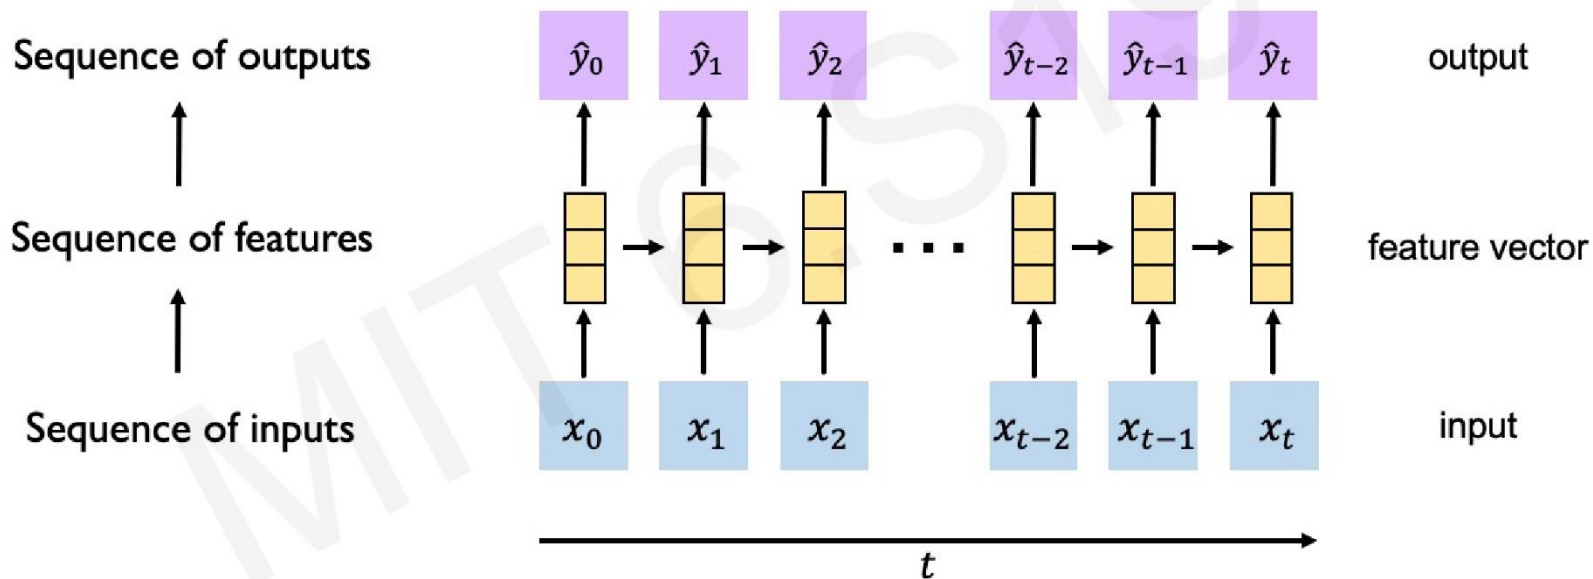
### Limitations of RNNs

-  Encoding bottleneck
-  Slow, no parallelization
-  Not long memory



## Goal of Sequence Modeling

RNNs: recurrence to model sequence dependencies






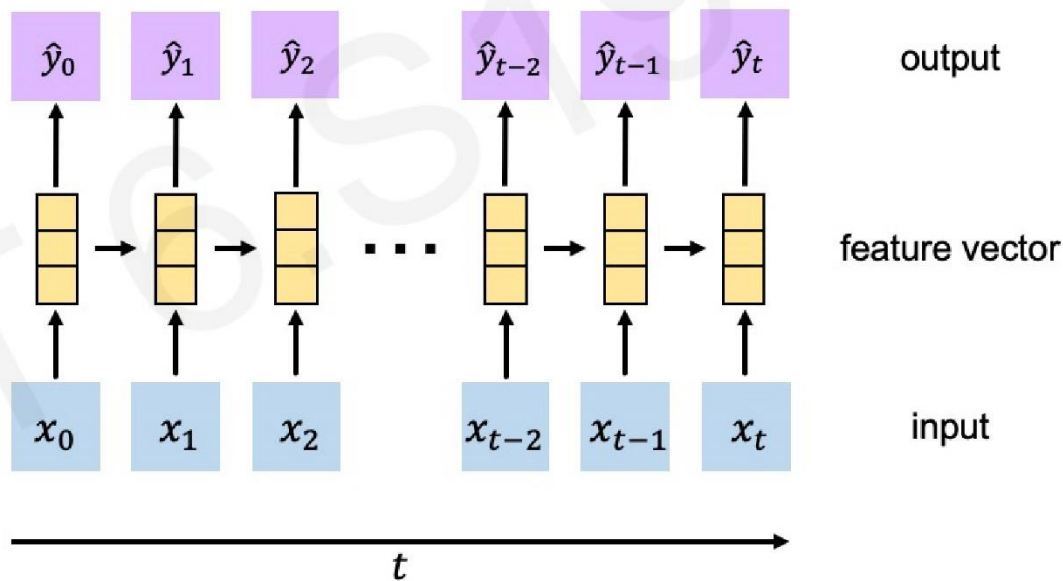


## Goal of Sequence Modeling

RNNs: recurrence to model sequence dependencies

### Limitations of RNNs

-  Encoding bottleneck
-  Slow, no parallelization
-  Not long memory





## Goal of Sequence Modeling

Can we eliminate the need for recurrence entirely?

### Desired Capabilities



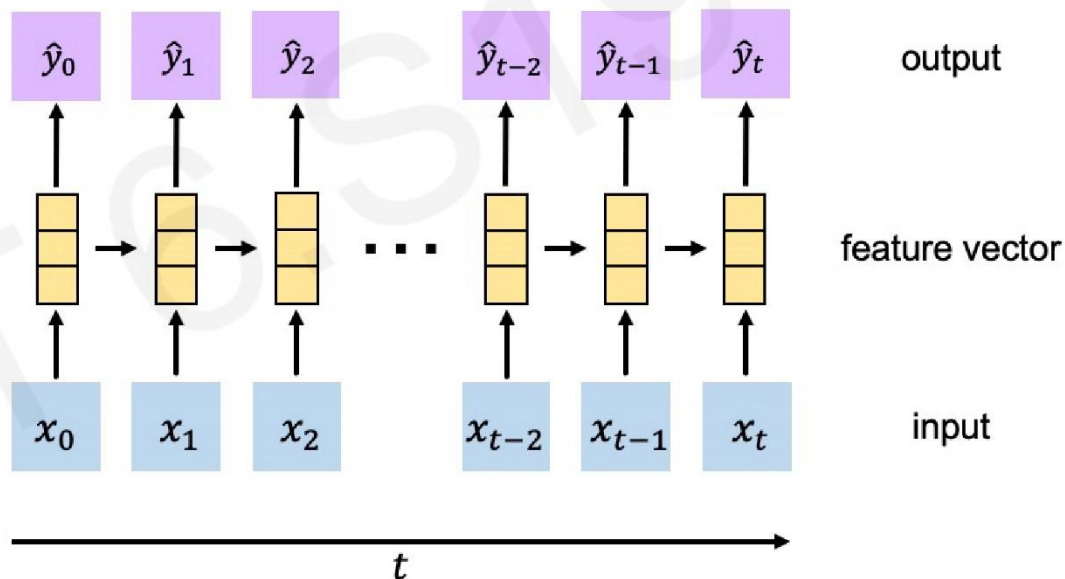
Continuous stream



Parallelization



Long memory





## Goal of Sequence Modeling

Can we eliminate the need for recurrence entirely?

### Desired Capabilities



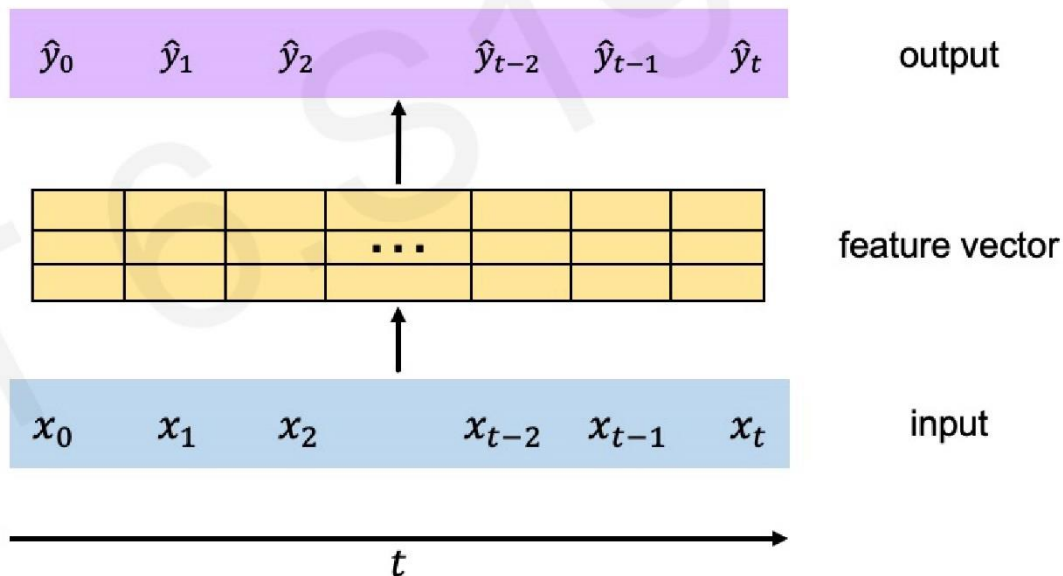
Continuous stream



Parallelization



Long memory







## Goal of Sequence Modeling

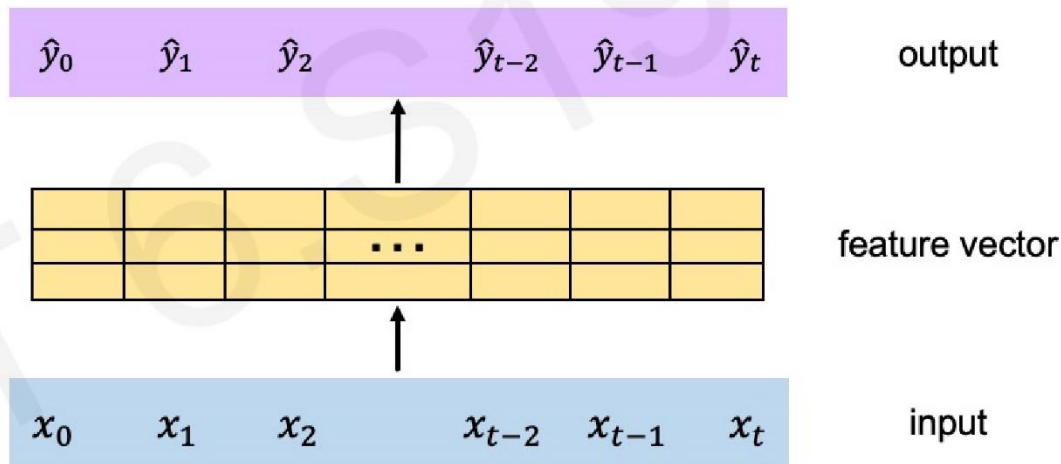
Idea 1: Feed everything into dense network

- ✓ No recurrence
- ✗ Not scalable
- ✗ No order
- ✗ No long memory



Idea: Identify and attend to what's important

Can we eliminate the need for recurrence entirely?





# 人工智能基本理论

Attention Is All You Need

MIT 6.S191



## Intuition Behind Self-Attention

Attending to the most important parts of an input.



1. Identify which parts to attend to
2. Extract the features with high attention

Similar to a search problem!



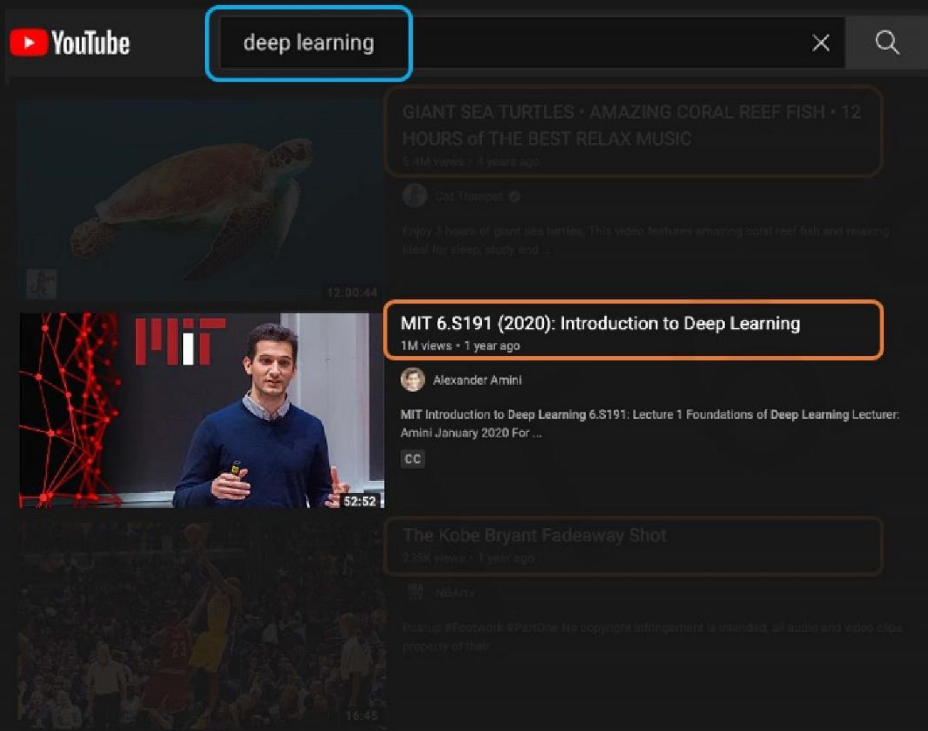
## A Simple Example: Search

How can I learn  
more about  
neural networks?





## Understanding Attention with Search



Query (Q)

Key ( $K_1$ )

Key ( $K_2$ )

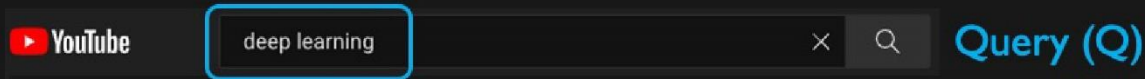
Key ( $K_3$ )

How similar is the key to the query?

1. **Compute attention mask:** how similar is each key to the desired query?



## Understanding Attention with Search



Query (Q)

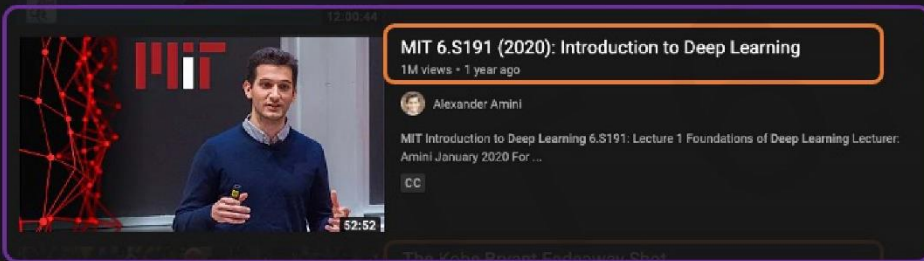
GIANT SEA TURTLES • AMAZING CORAL REEF FISH • 12 HOURS of THE BEST RELAX MUSIC

Key ( $K_1$ )

5.4M views • 4 years ago

Dot Thomas

Enjoy 3 hours of giant sea turtles. This video features amazing coral reef fish and relaxing music for sleep, study and ...



Key ( $K_2$ )

Value (V)

The Kobe Bryant Fadeaway Shot

Key ( $K_3$ )

228K views • 1 year ago

NBA.com

Podcast #Feedback #Parody No copyright infringement is intended, all audio and video clips property of their ...

2. Extract values based on attention:  
Return the values highest attention



## Learning Self-Attention with Neural Networks

Goal: identify and attend to most important features in input.

1. Encode **position** information
2. Extract **query, key, value** for search
3. Compute **attention weighting**
4. Extract **features with high attention**



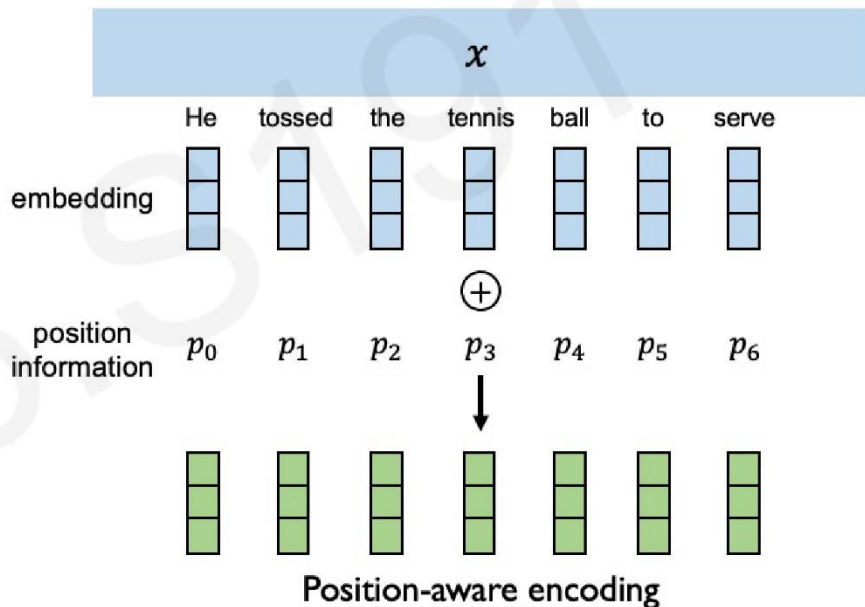
Data is fed in all at once! Need to encode position information to understand order.



## Learning Self-Attention with Neural Networks

Goal: identify and attend to most important features in input.

1. Encode **position** information
2. Extract **query, key, value** for search
3. Compute attention weighting
4. Extract features with high attention



Data is fed in all at once! Need to encode position information to understand order.

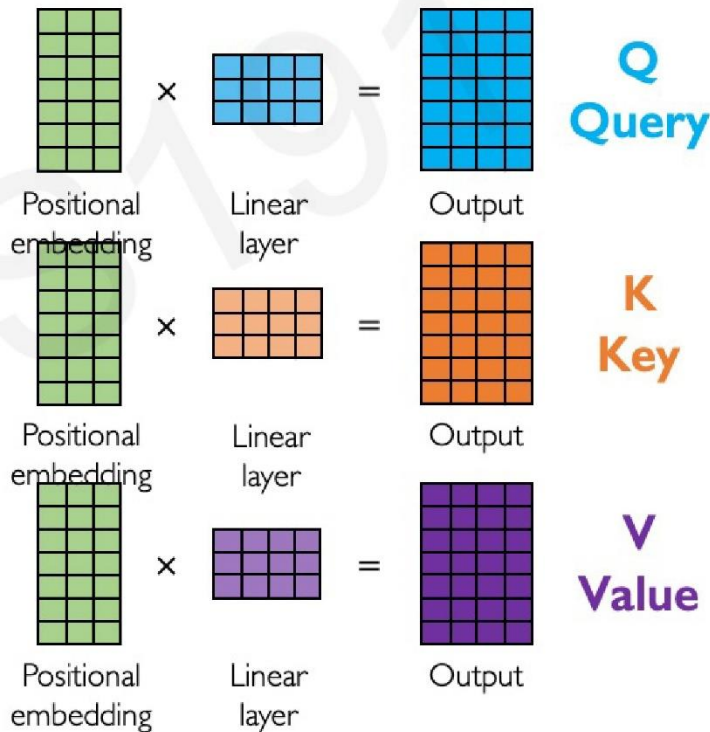




## Learning Self-Attention with Neural Networks

Goal: identify and attend to most important features in input.

1. Encode **position** information
2. Extract **query, key, value** for search
3. Compute attention weighting
4. Extract features with high attention





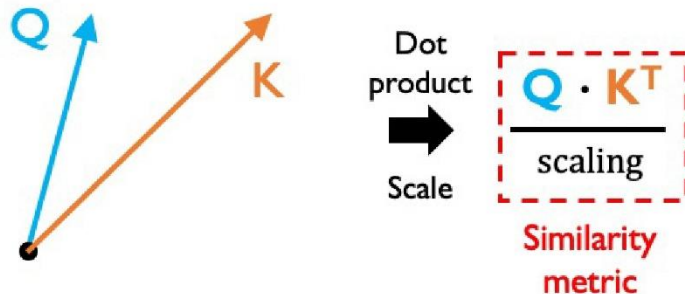
## Learning Self-Attention with Neural Networks

Goal: identify and attend to most important features in input.

1. Encode **position** information
2. Extract **query, key, value** for search
3. Compute **attention weighting**
4. Extract features with high attention

Attention score: compute pairwise similarity between each **query** and **key**

How to compute similarity between two sets of features?



Also known as the "cosine similarity"



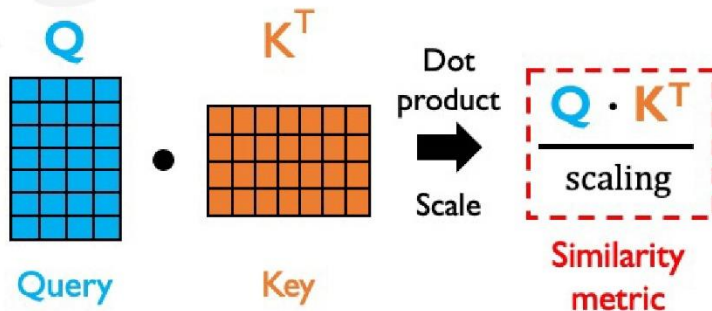
## Learning Self-Attention with Neural Networks

Goal: identify and attend to most important features in input.

1. Encode **position** information
2. Extract **query, key, value** for search
3. Compute **attention weighting**
4. Extract features with high attention

Attention score: compute pairwise similarity between each **query** and **key**

How to compute similarity between two sets of features?



Also known as the “cosine similarity”



## Learning Self-Attention with Neural Networks

Goal: identify and attend to most important features in input.

1. Encode **position** information
2. Extract **query, key, value** for search
3. Compute **attention weighting**
4. Extract features with high attention

Attention weighting: where to attend to!  
How similar is the key to the query?

	He	tossed	the	tennis	ball	to	serve
He	1	0	0	0	0	0	0
tossed	0	1	0	0	0	0	0
the	0	0	1	0	0	0	0
tennis	0	0	0	1	0	0	0
ball	0	0	0	0	1	0	0
to	0	0	0	0	0	1	0
serve	0	0	0	0	0	0	1

$$\text{softmax} \left( \frac{Q \cdot K^T}{\text{scaling}} \right)$$

Attention weighting

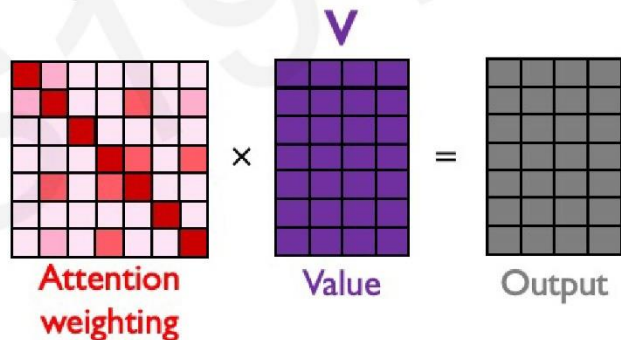


## Learning Self-Attention with Neural Networks

Goal: identify and attend to most important features in input.

1. Encode **position** information
2. Extract **query, key, value** for search
3. Compute **attention weighting**
4. Extract **features with high attention**

Last step: self-attend to extract features



$$\text{softmax} \left( \frac{Q \cdot K^T}{\text{scaling}} \right) \cdot V = A(Q, K, V)$$

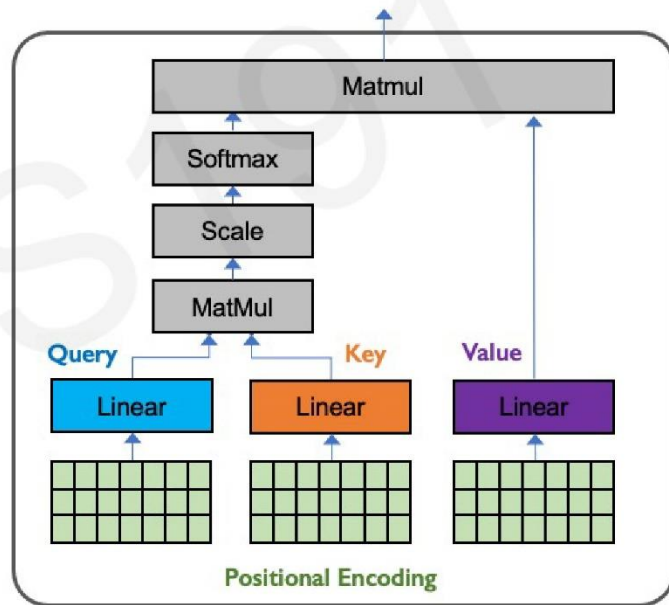


## Learning Self-Attention with Neural Networks

Goal: identify and attend to most important features in input.

1. Encode **position** information
2. Extract **query, key, value** for search
3. Compute **attention weighting**
4. Extract **features with high attention**

These operations form a self-attention head that can plug into a larger network. Each head attends to a different part of input.



$$\text{softmax} \left( \frac{Q \cdot K^T}{\text{scaling}} \right) \cdot V$$



## Applying Multiple Self-Attention Heads



Attention weighting

×



Value

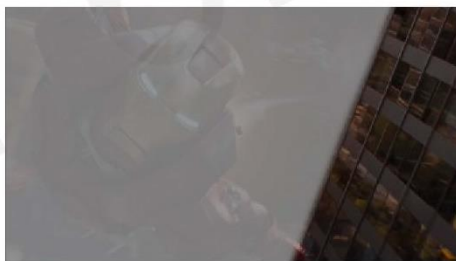
=



Output



Output of attention head 1



Output of attention head 2



Output of attention head 3



## Self-Attention Applied

### Language Processing

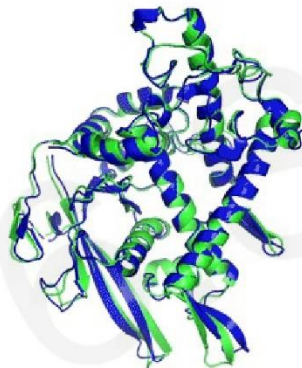


An armchair in the shape of an avocado

BERT, GPT-3

Devlin et al., *NAACL* 2019  
Brown et al., *NeurIPS* 2020

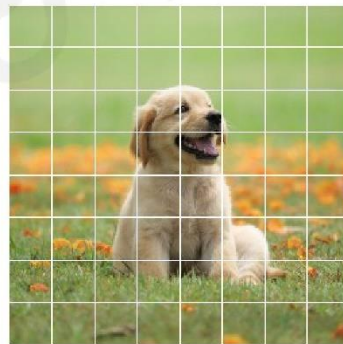
### Biological Sequences



AlphaFold2

Jumper et al., *Nature* 2021

### Computer Vision



Vision Transformers

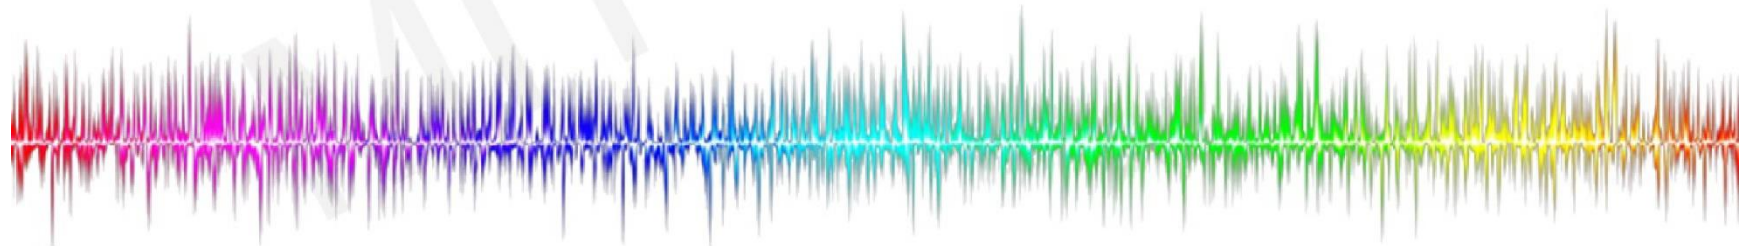
Dosovitskiy et al., *ICLR* 2020





## Deep Learning for Sequence Modeling: Summary

1. RNNs are well suited for **sequence modeling** tasks
2. Model sequences via a **recurrence relation**
3. Training RNNs with **backpropagation through time**
4. Models for **music generation**, classification, machine translation, and more
5. Self-attention to model **sequences without recurrence**



# 提纲

---

一、前馈网络

二、循环网络

三、卷积网络



上海大学  
SHANGHAI UNIVERSITY



# | 人工智能基本理论

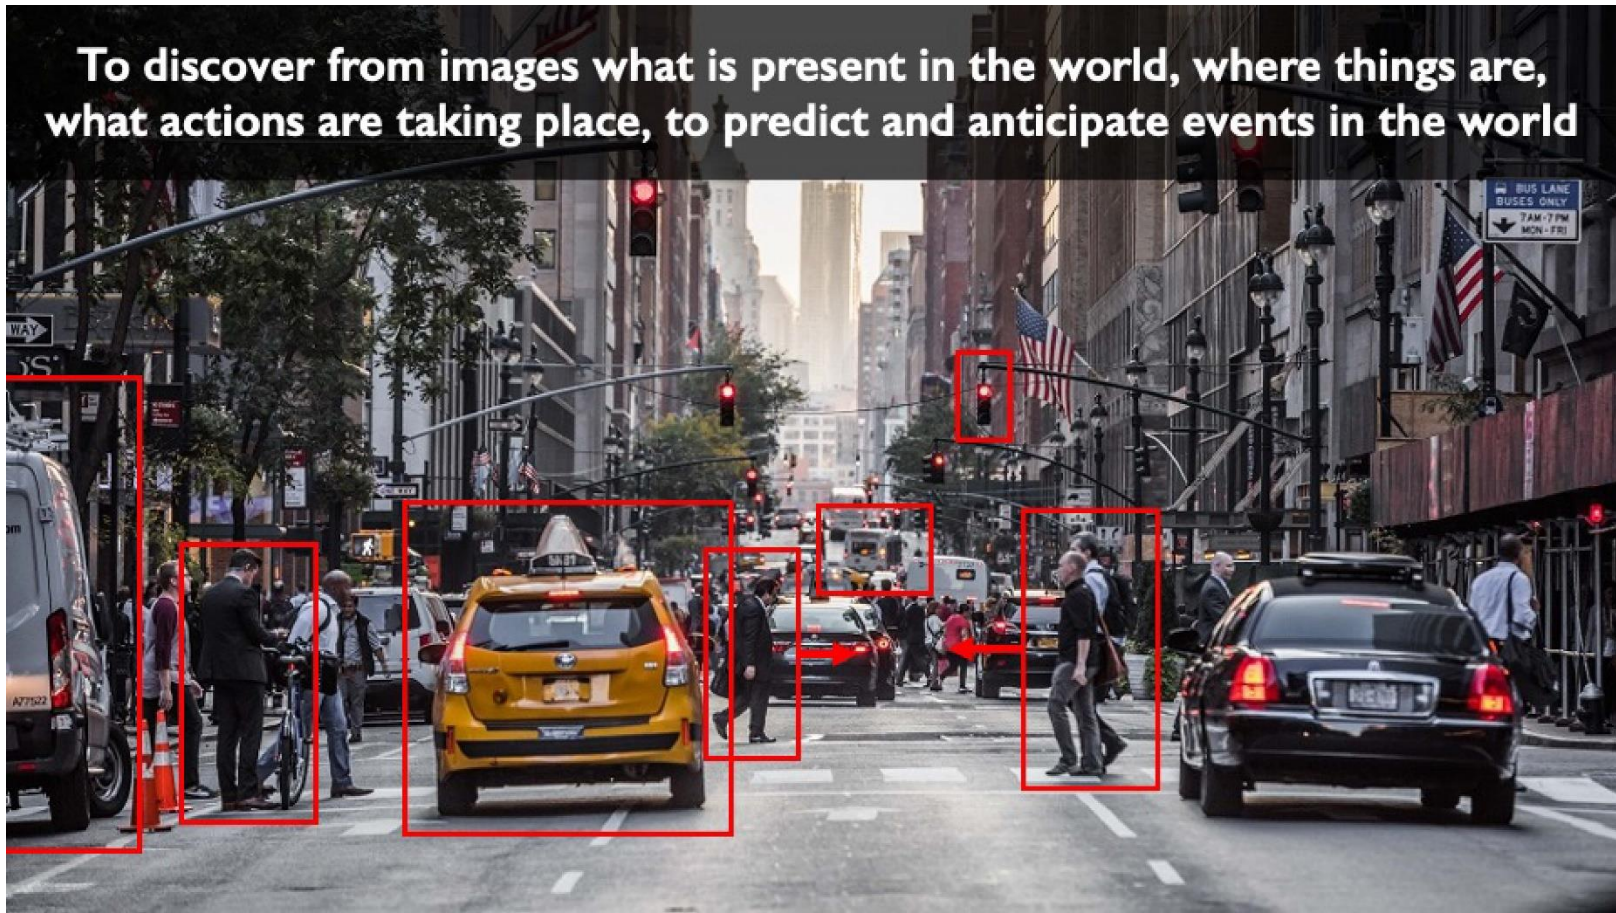
**“To know what is  
where by looking.”**





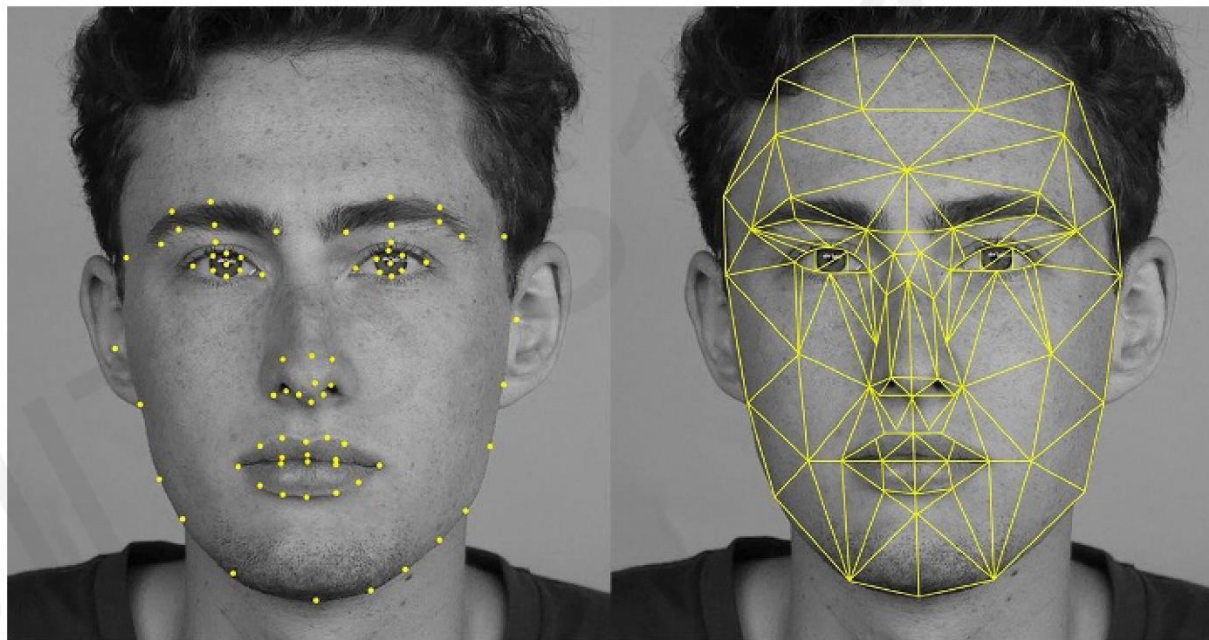
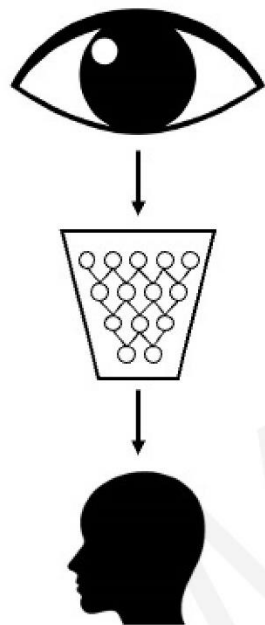
# 人工智能基本理论

To discover from images what is present in the world, where things are, what actions are taking place, to predict and anticipate events in the world



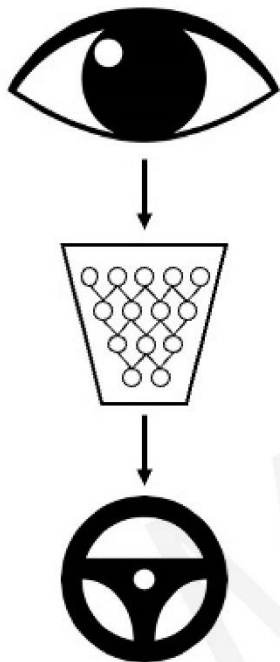


## Impact: Facial Detection & Recognition





## Impact: Self-Driving Cars



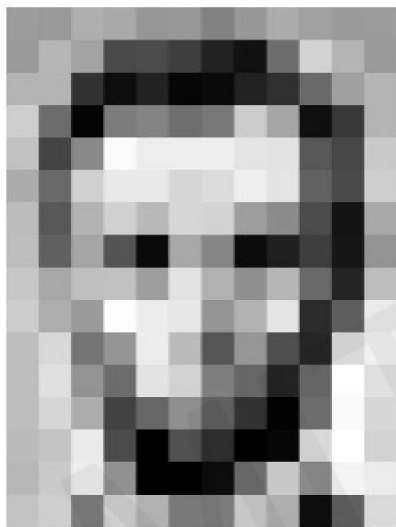


What Computers “See”

MIT 6.S191



## Images are Numbers



157	153	174	168	150	152	129	151	172	161	155	156
155	182	163	74	75	62	83	17	110	210	180	154
180	180	50	14	84	6	10	83	48	105	159	181
206	109	5	124	131	111	120	204	166	15	56	180
194	68	137	251	237	239	239	228	227	87	71	201
172	105	207	233	233	214	220	239	228	98	74	206
188	88	179	209	185	215	211	158	139	75	20	169
189	97	165	84	10	168	134	11	31	62	22	148
199	168	191	193	158	227	178	143	182	105	36	190
205	174	155	252	236	231	149	178	228	43	95	234
190	216	116	149	236	187	85	150	79	38	218	241
190	224	147	108	227	210	127	102	35	101	255	224
190	214	173	66	103	143	95	50	2	109	249	215
187	196	235	75	1	81	47	0	6	217	255	211
183	202	237	145	0	0	12	108	209	138	243	236
195	206	123	207	177	121	123	200	175	13	96	218

What the computer sees

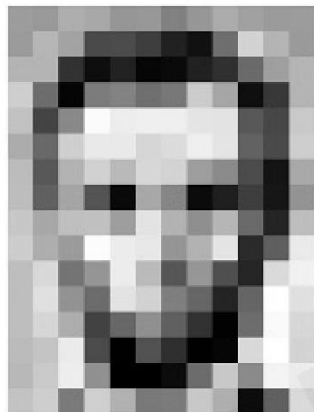
157	153	174	168	150	152	129	151	172	161	155	156
155	182	163	74	75	62	83	17	110	210	180	154
180	180	50	14	84	6	10	83	48	105	159	181
206	109	5	124	131	111	120	204	166	15	56	180
194	68	137	251	237	239	239	228	227	87	71	201
172	105	207	233	233	214	220	239	228	98	74	206
188	88	179	209	185	215	211	158	139	75	20	169
189	97	165	84	10	168	134	11	31	62	22	148
199	168	191	193	158	227	178	143	182	105	36	190
205	174	155	252	236	231	149	178	228	43	95	234
190	216	116	149	236	187	85	150	79	38	218	241
190	224	147	108	227	210	127	102	35	101	255	224
190	214	173	66	103	143	95	50	2	109	249	215
187	196	235	75	1	81	47	0	6	217	255	211
183	202	237	145	0	0	12	108	209	138	243	236
195	206	123	207	177	121	123	200	175	13	96	218

An image is just a matrix of numbers  $[0,255]$ !  
i.e.,  $1080 \times 1080 \times 3$  for an RGB image





## Tasks in Computer Vision



Input Image



187	153	174	168	156	182	129	151	172	161	186	156
195	182	163	74	75	82	33	17	113	216	180	154
180	180	50	14	34	6	10	33	48	106	159	181
206	109	5	124	131	111	120	204	166	15	56	180
194	68	127	261	237	230	230	228	227	87	71	201
172	105	207	233	233	214	230	239	239	88	74	206
188	88	179	208	186	216	211	168	139	75	10	149
189	97	165	84	16	168	134	11	31	62	22	148
199	168	191	193	198	227	178	143	182	106	36	190
205	174	155	252	236	231	143	178	228	43	95	234
190	210	116	149	236	187	85	150	79	38	210	241
190	234	147	108	227	210	127	102	35	101	266	234
190	214	173	65	103	143	96	63	3	108	240	215
187	196	295	75	1	81	47	0	6	217	245	211
183	202	237	145	0	0	12	108	200	138	243	236
195	200	123	207	177	121	123	200	175	13	96	218

Pixel Representation

classification

Lincoln

0.8

Washington

0.1

Jefferson

0.05

Obama

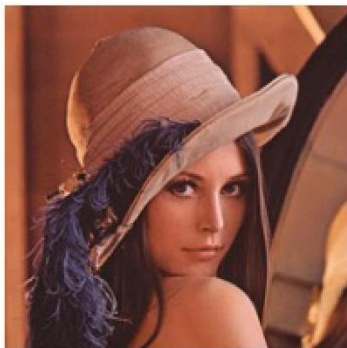
0.05

- **Regression:** output variable takes continuous value
- **Classification:** output variable takes class label. Can produce probability of belonging to a particular class



## High Level Feature Detection

Let's identify key features in each image category



Nose,  
Eyes,  
Mouth



Wheels,  
License Plate,  
Headlights



Door,  
Windows,  
Steps



## Manual Feature Extraction

Domain knowledge

Define features

Detect features to classify

Viewpoint variation



Scale variation



Deformation



Occlusion



Illumination conditions



Background clutter



Intra-class variation

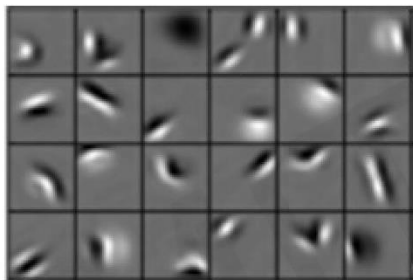




## Learning Feature Representations

Can we learn a **hierarchy of features** directly from the data instead of hand engineering?

Low level features



Edges, dark spots

Mid level features



Eyes, ears, nose

High level features



Facial structure



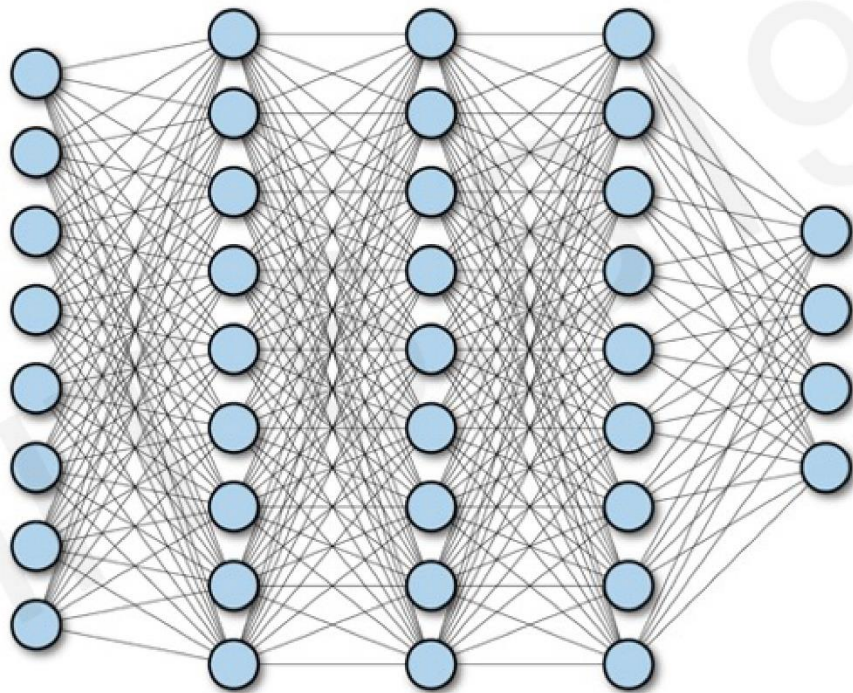
# 人工智能基本理论

Learning Visual Features

MIT 6.S191



## Fully Connected Neural Network

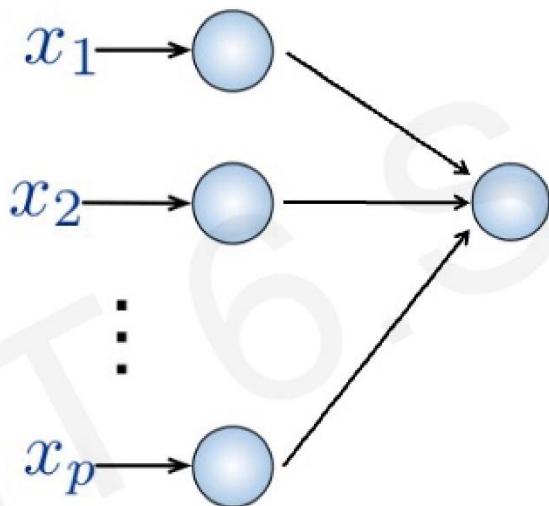




## Fully Connected Neural Network

### Input:

- 2D image
- Vector of pixel values



### Fully Connected:

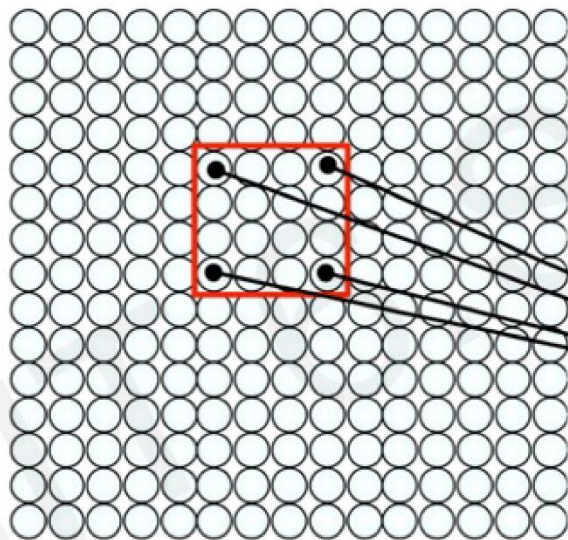
- Connect neuron in hidden layer to all neurons in input layer
- No spatial information!
- And many, many parameters!

How can we use **spatial structure** in the input to inform the architecture of the network?



## Using Spatial Structure

**Input:** 2D image.  
Array of pixel values

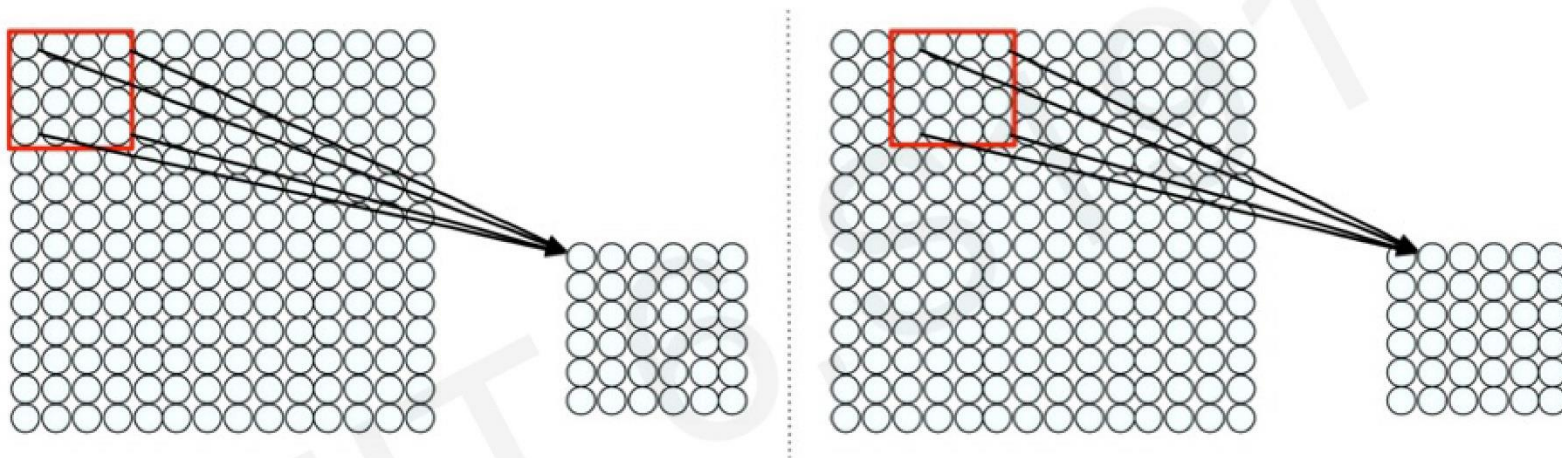


**Idea:** connect patches of input  
to neurons in hidden layer.  
Neuron connected to region of  
input. Only “sees” these values.





## Using Spatial Structure



Connect patch in input layer to a single neuron in subsequent layer.

Use a sliding window to define connections.

How can we **weight** the patch to detect particular features?

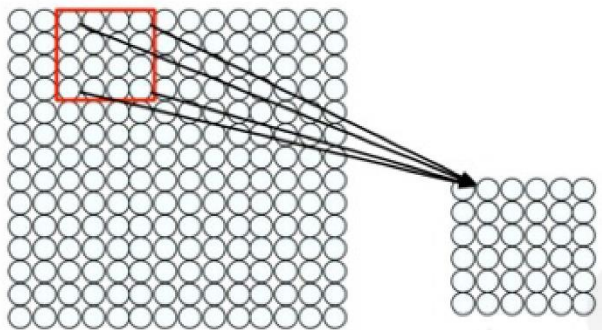


## Applying Filters to Extract Features

- 1) Apply a set of weights – a filter – to extract **local features**
- 2) Use **multiple filters** to extract different features
- 3) Spatially **share** parameters of each filter  
(features that matter in one part of the input should matter elsewhere)



## Feature Extraction with Convolution



- Filter of size 4x4 : 16 different weights
- Apply this same filter to 4x4 patches in input
- Shift by 2 pixels for next patch

This “patchy” operation is **convolution**

- 1) Apply a set of weights – a filter – to extract **local features**
- 2) Use **multiple filters** to extract different features
- 3) **Spatially share** parameters of each filter



## The Convolution Operation

We slide the 3x3 filter over the input image, element-wise multiply, and add the outputs:

1	1	1	0	0
0	1	1	1	0
0	0	1 <sub>x1</sub>	1 <sub>x0</sub>	1 <sub>x1</sub>
0	0	1 <sub>x0</sub>	1 <sub>x1</sub>	0 <sub>x0</sub>
0	1	1 <sub>x1</sub>	0 <sub>x0</sub>	0 <sub>x1</sub>



1	0	1
0	1	0
1	0	1

filter

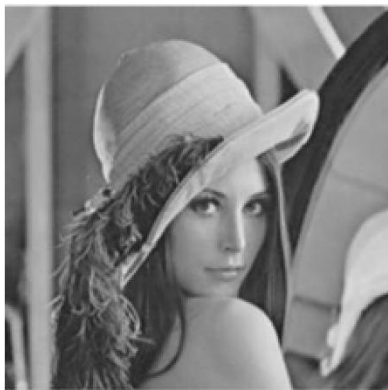


4	3	4
2	4	3
2	3	4

feature map



## Producing Feature Maps



Original



Sharpen



Edge Detect



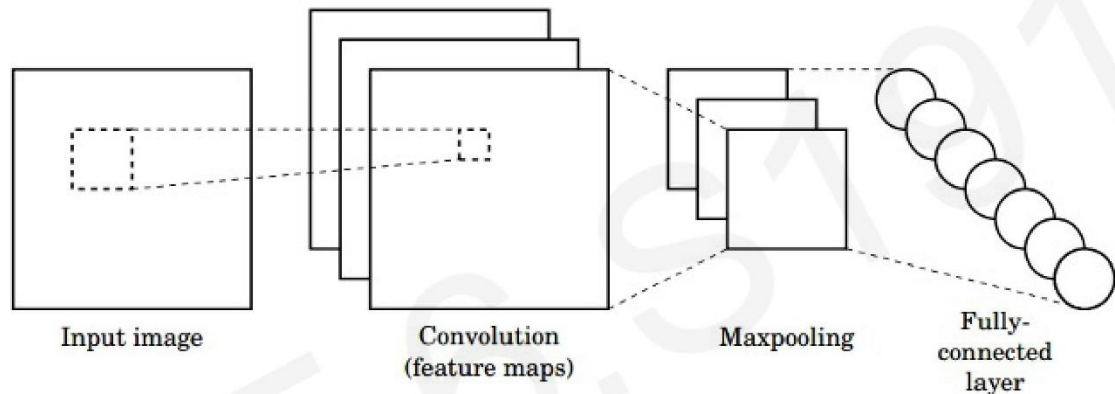
"Strong" Edge Detect



# Convolutional Neural Networks (CNNs)




## CNNs for Classification



1. **Convolution:** Apply filters to generate feature maps.
2. **Non-linearity:** Often ReLU.
3. **Pooling:** Downsampling operation on each feature map.

 `tf.keras.layers.Conv2D`

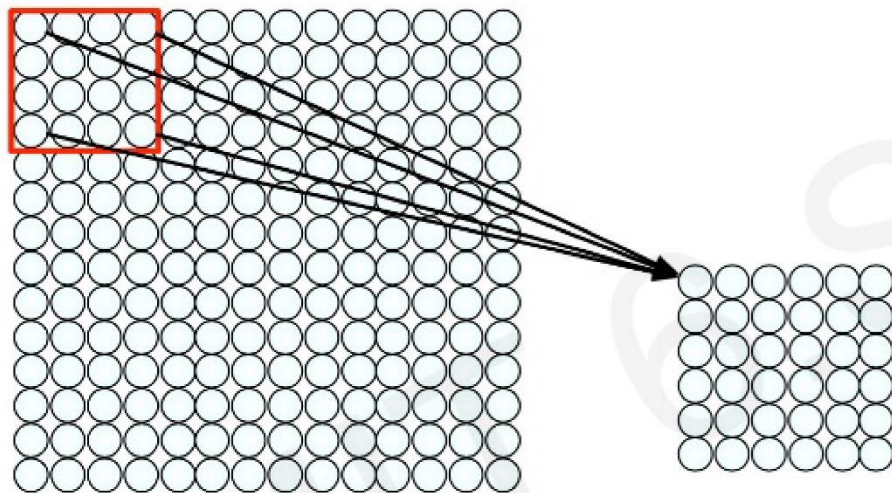
 `tf.keras.activations.*`

 `tf.keras.layers.MaxPool2D`

**Train model with image data.**  
**Learn weights of filters in convolutional layers.**



## Convolutional Layers: Local Connectivity



```
tf.keras.layers.Conv2D
```

### For a neuron in hidden layer:

- Take inputs from patch
- Compute weighted sum
- Apply bias

4x4 filter: matrix of weights  $w_{ij}$

$$\sum_{i=1}^4 \sum_{j=1}^4 w_{ij} x_{i+p,j+q} + b$$

for neuron (p,q) in hidden layer

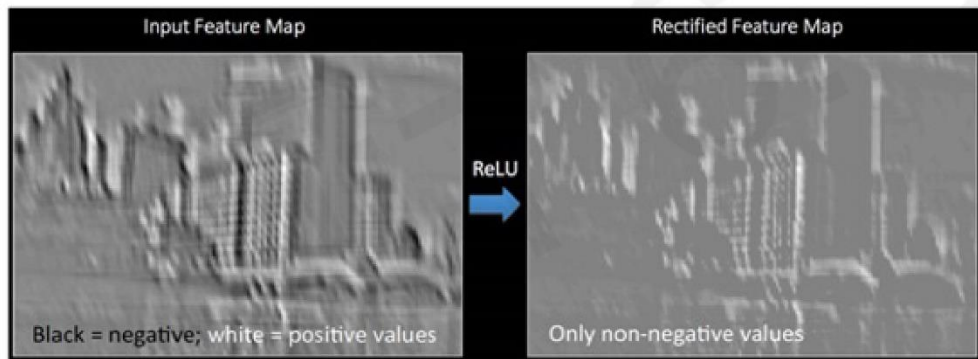
- 1) applying a window of weights
- 2) computing linear combinations
- 3) activating with non-linear function



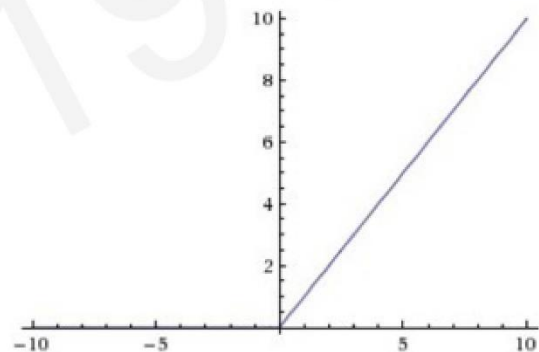


## Introducing Non-Linearity

- Apply after every convolution operation (i.e., after convolutional layers)
- ReLU: pixel-by-pixel operation that replaces all negative values by zero. **Non-linear operation!**



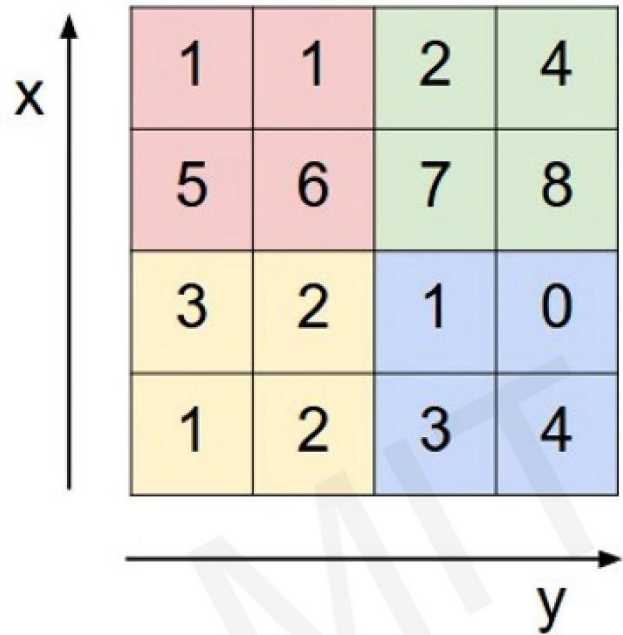
### Rectified Linear Unit (ReLU)



`tf.keras.layers.ReLU`



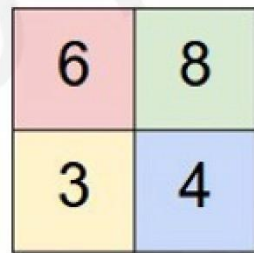
## Pooling



max pool with 2x2 filters  
and stride 2



```
tf.keras.layers.MaxPool2D(  
    pool_size=(2,2),  
    strides=2  
)
```

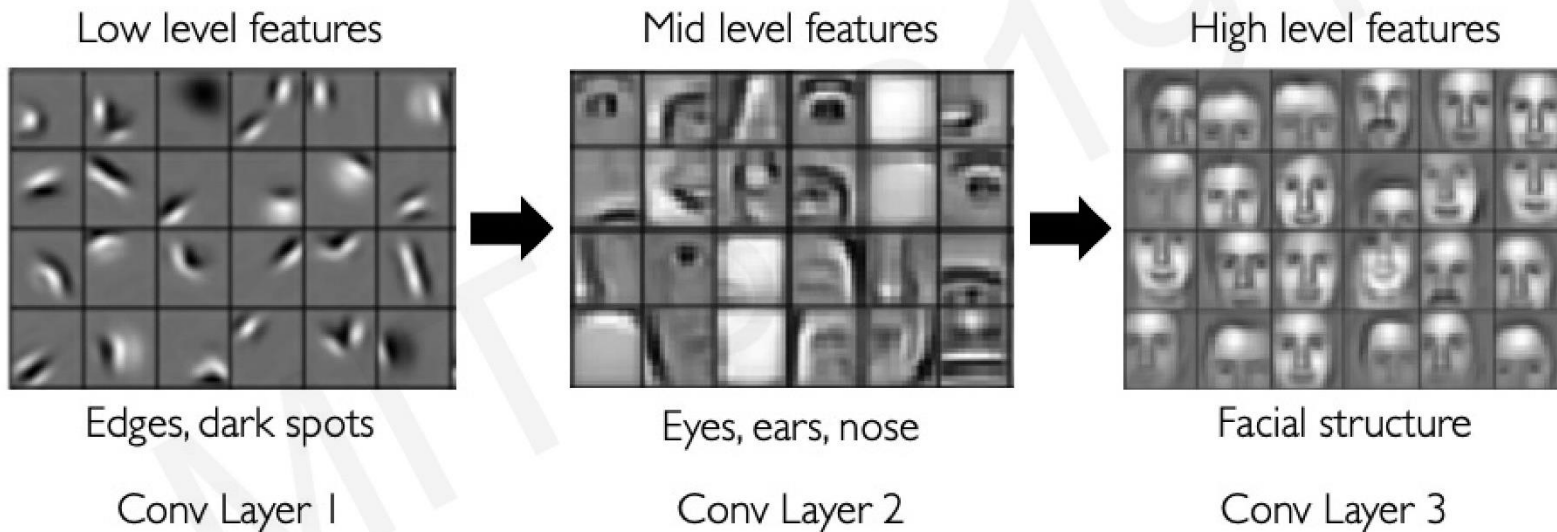


- 1) Reduced dimensionality
- 2) Spatial invariance

How else can we downsample and preserve spatial invariance?

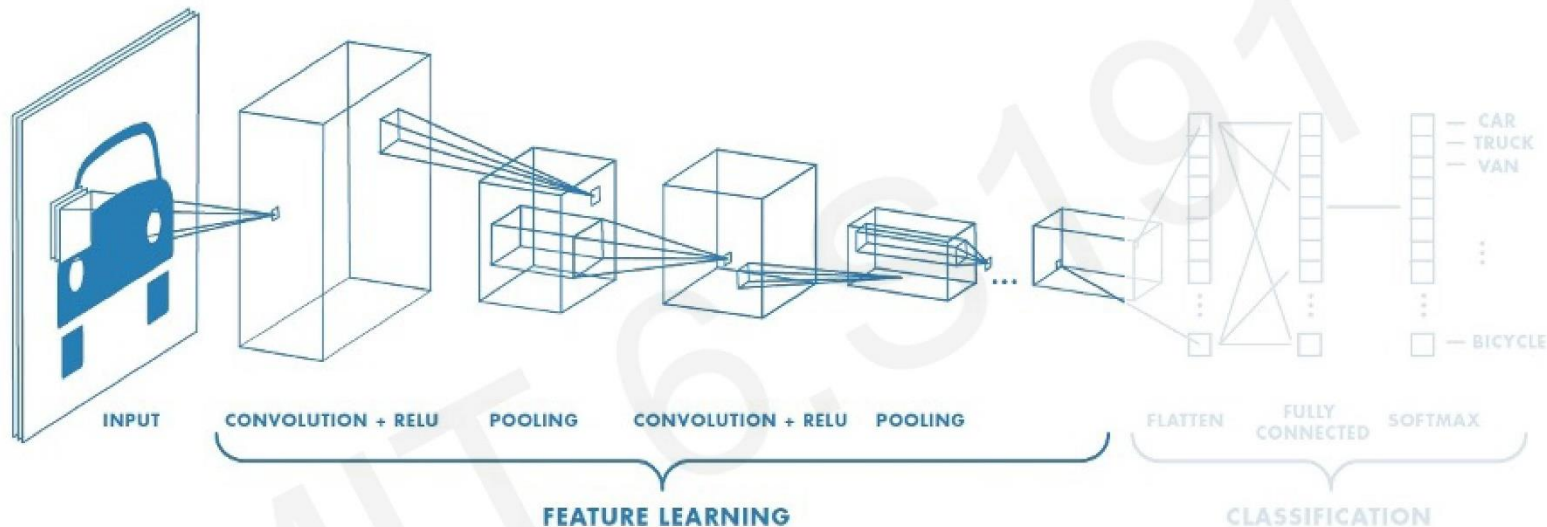


## Representation Learning in Deep CNNs





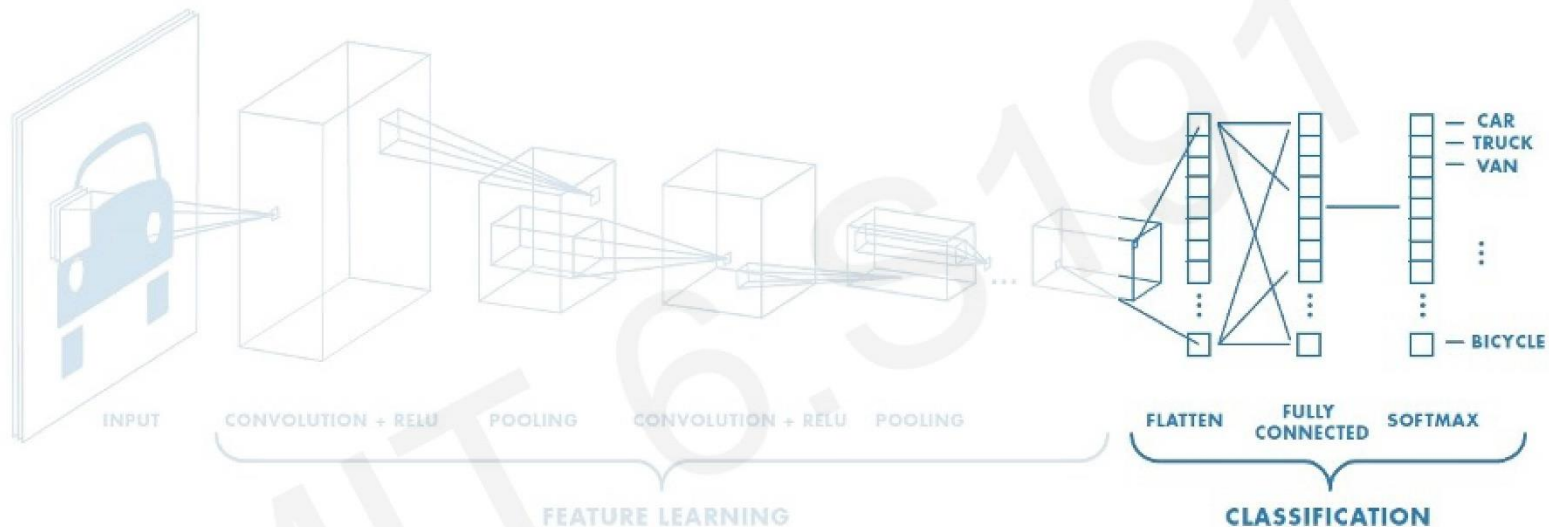
## CNNs for Classification: Feature Learning



1. Learn features in input image through **convolution**
2. Introduce **non-linearity** through activation function (real-world data is non-linear!)
3. Reduce dimensionality and preserve spatial invariance with **pooling**



## CNNs for Classification: Class Probabilities



- CONV and POOL layers output high-level features of input
- Fully connected layer uses these features for classifying input image
- Express output as **probability** of image belonging to a particular class

$$\text{softmax}(y_i) = \frac{e^{y_i}}{\sum_j e^{y_j}}$$

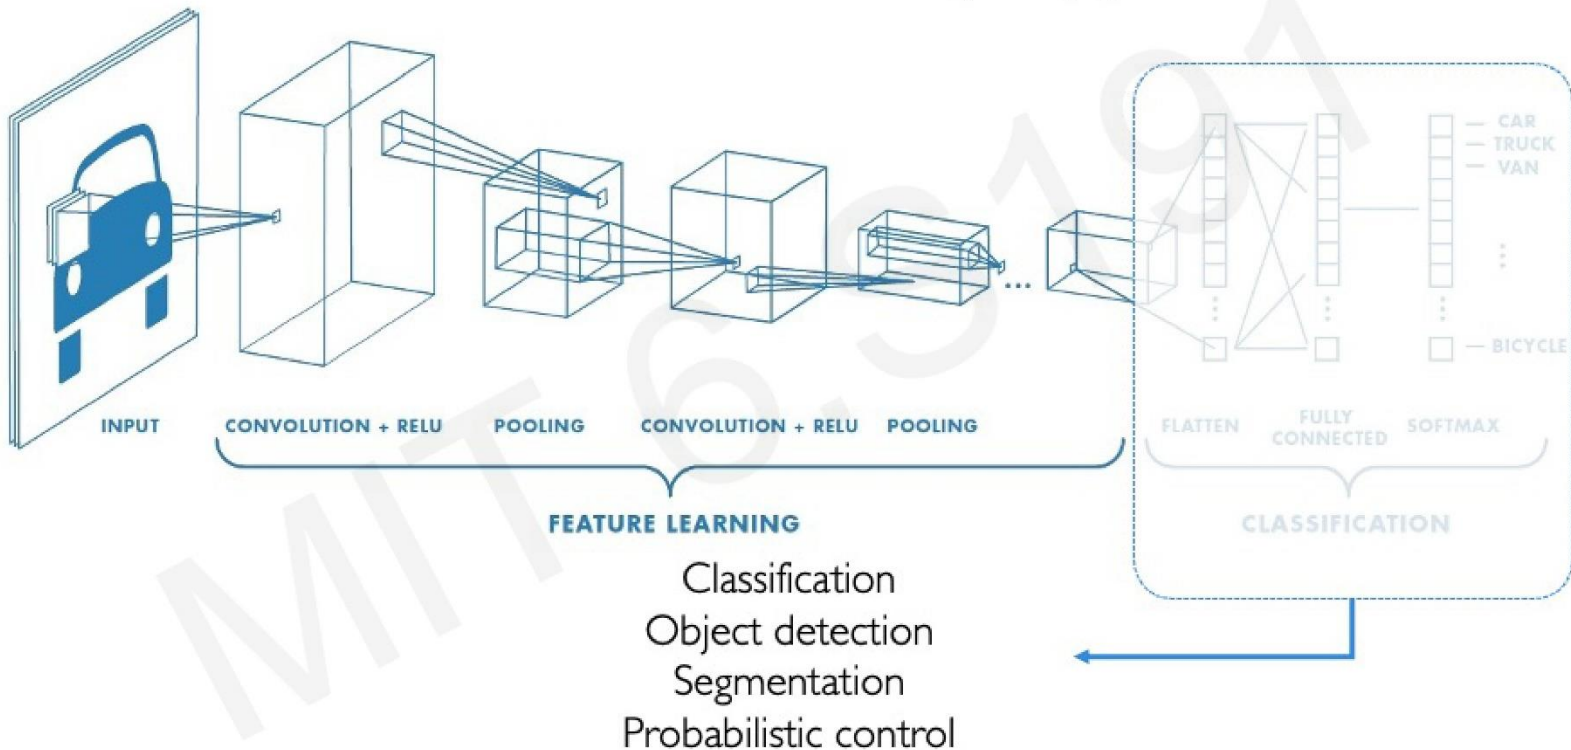


# 人工智能基本理论

An Architecture for Many Applications



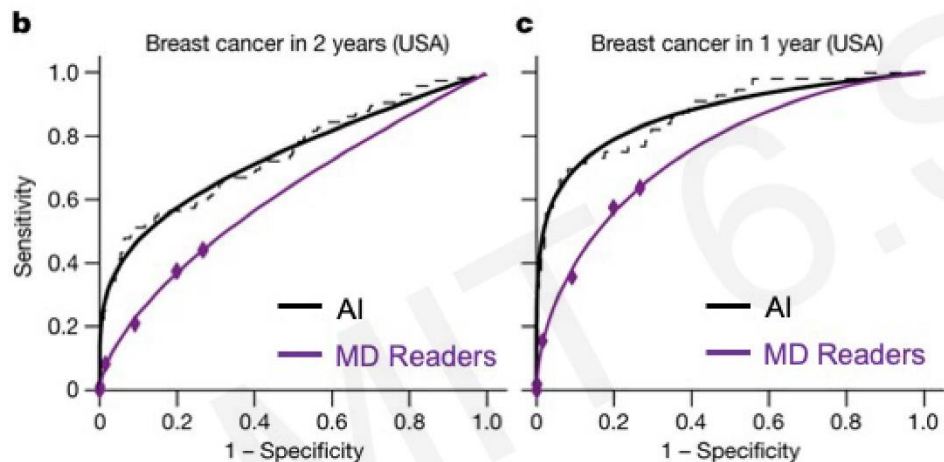
## An Architecture for Many Applications



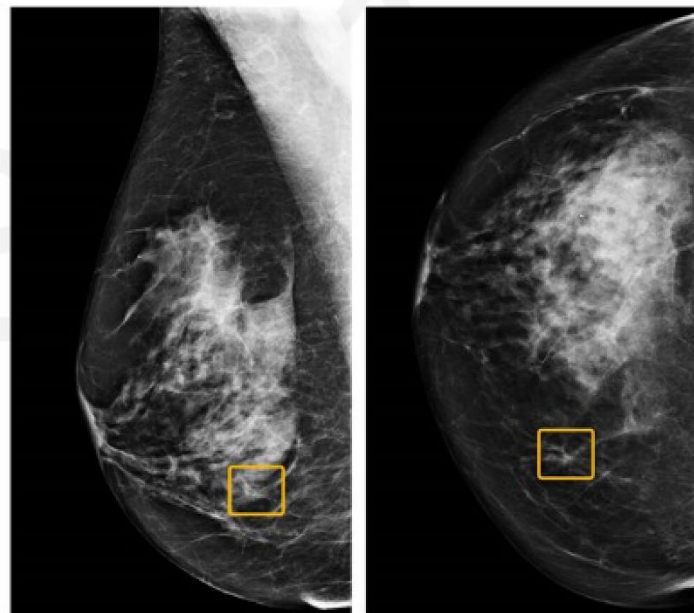


## Classification: Breast Cancer Screening

### International evaluation of an AI system for breast cancer screening



CNN-based system outperformed expert radiologists at detecting breast cancer from mammograms



Breast cancer case missed by radiologist but detected by AI

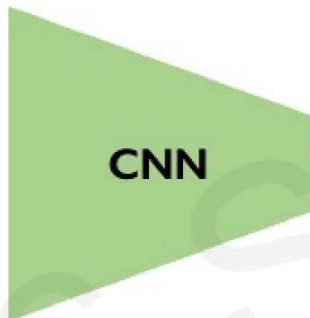




## Object Detection



Image  $x$

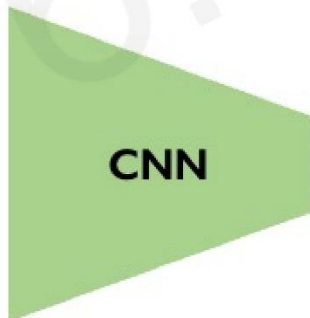


Taxi

Class label  $y$



Image  $x$

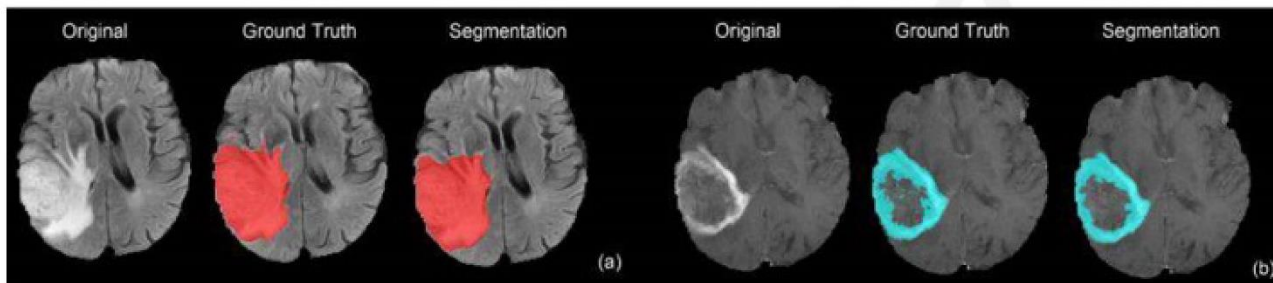


Label  $(x, y, w, h)$

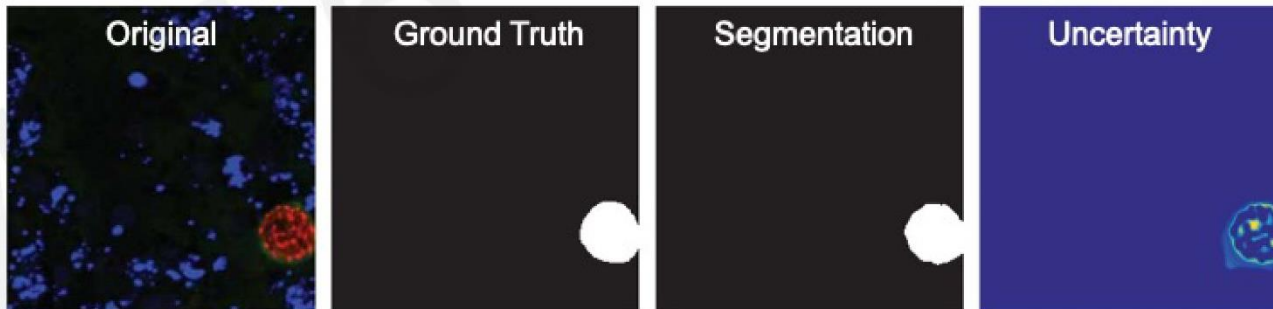


## Semantic Segmentation: Biomedical Image Analysis

Brain Tumors  
Dong+ *MIUA* 2017.



Malaria Infection  
Soleimany+ *arXiv* 2019.



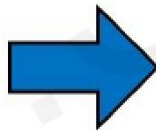
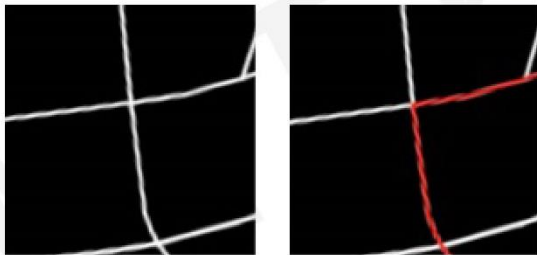


## Continuous Control: Navigation from Vision

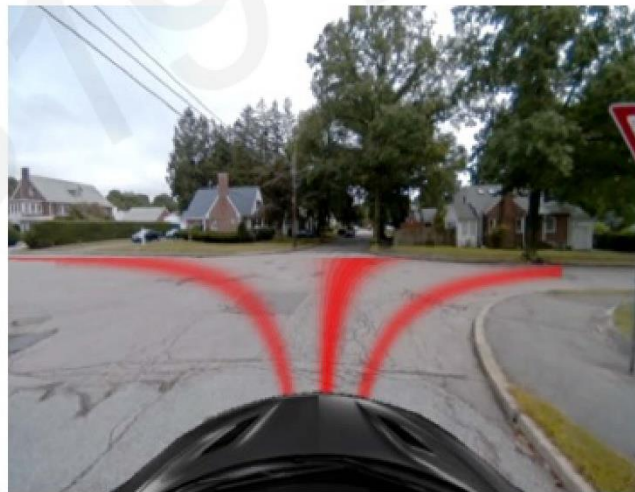
Raw Perception  
 $I$   
(ex. camera)



Coarse Maps  
 $M$   
(ex. GPS)



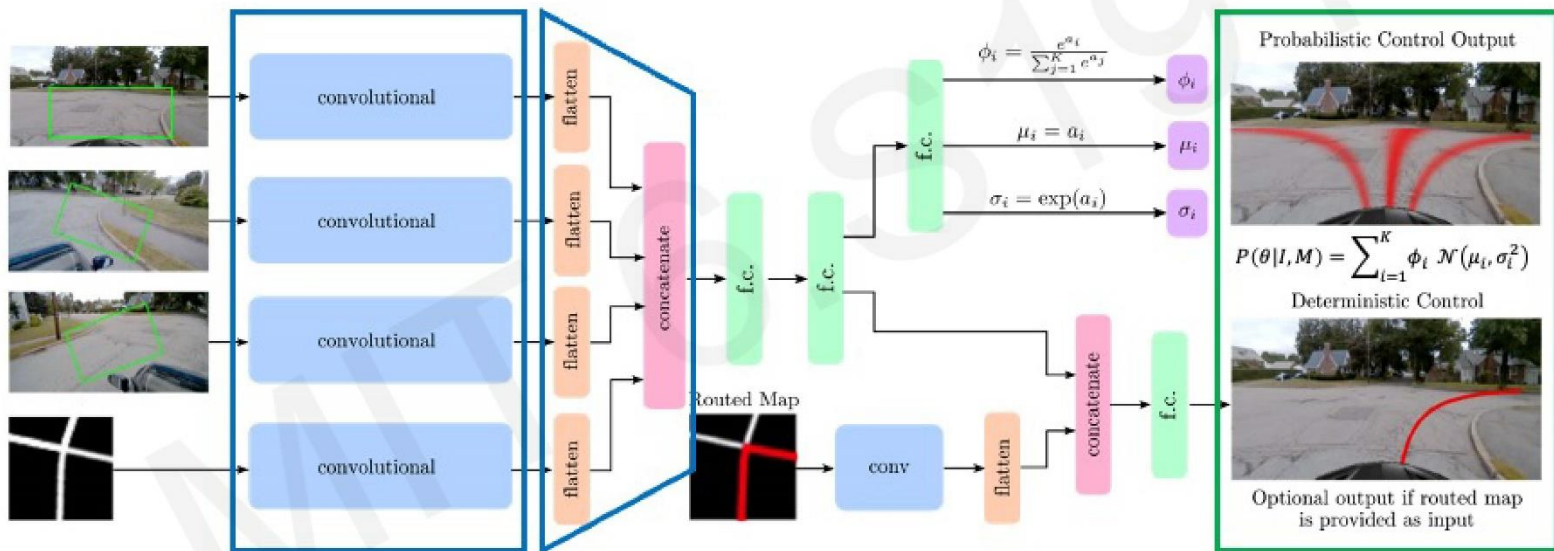
Possible Control Commands





## End-to-End Framework for Autonomous Navigation

Entire model is trained end-to-end **without any human labelling or annotations**



$$L = -\log(P(\theta|I, M))$$



## Deep Learning for Computer Vision: Summary

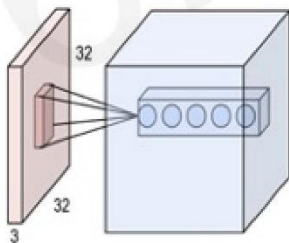
### Foundations

- Why computer vision?
- Representing images
- Convolutions for feature extraction



### CNNs

- CNN architecture
- Application to classification
- ImageNet



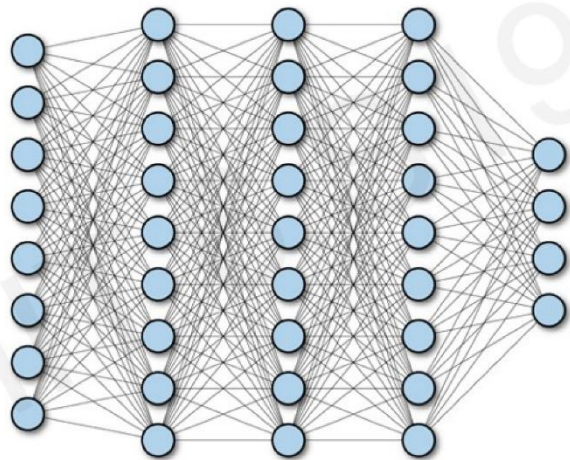
### Applications

- Segmentation, image captioning, control
- Security, medicine, robotics

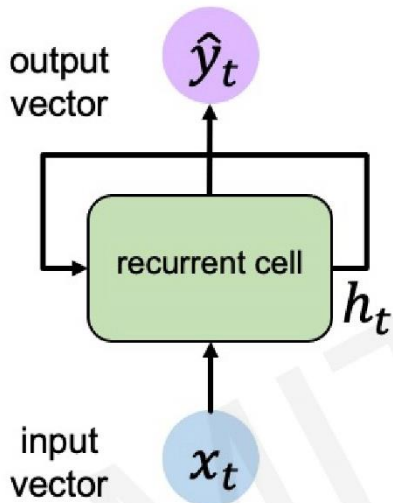




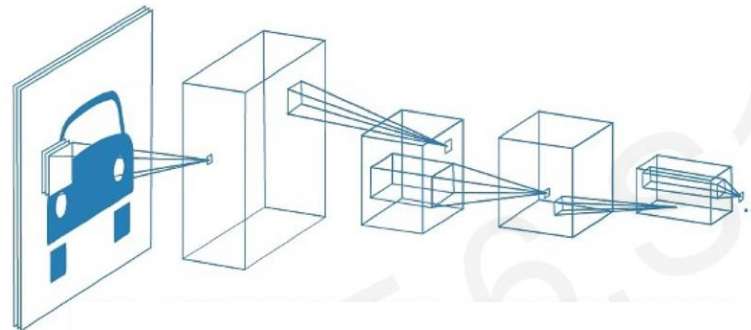
# 人工智能基本理论



前馈神经网络  
MLP



循环神经网络  
RNN



卷积神经网络  
CNN



谢谢大家