



15. Deep Learning

Artificial Intelligence




■ Alan Turing

- Turing test, a method to assess a machine's ability to exhibit human-like intelligent behavior

■ Artificial intelligence, or AI

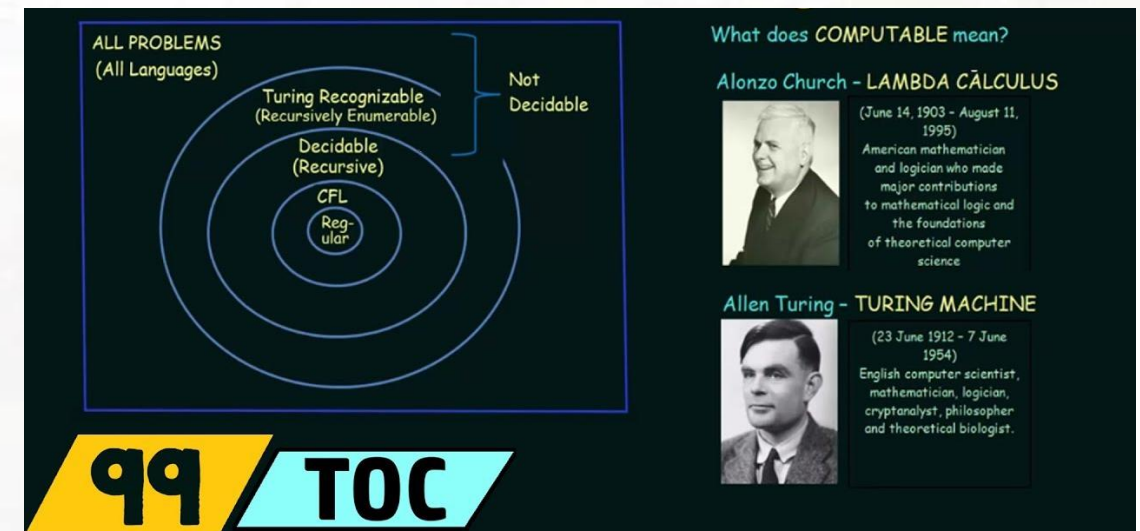


- technology that enables **computers** and machines to simulate human intelligence and problem-solving capabilities.



Turing Test

■ Turing Machine



- a mathematical model of computation describing an abstract machine that manipulates symbols on a strip of tape according to a table of rules
- Despite the model's simplicity, it is capable of implementing **any computer algorithm**.




Artificial Intelligence



浙江大學
ZHEJIANG UNIVERSITY



Turing Machine




Foundation model




Text to Image




ChatGPT



SAM



SORA



MLLM



DriveGPT




Genie



The Development of AI



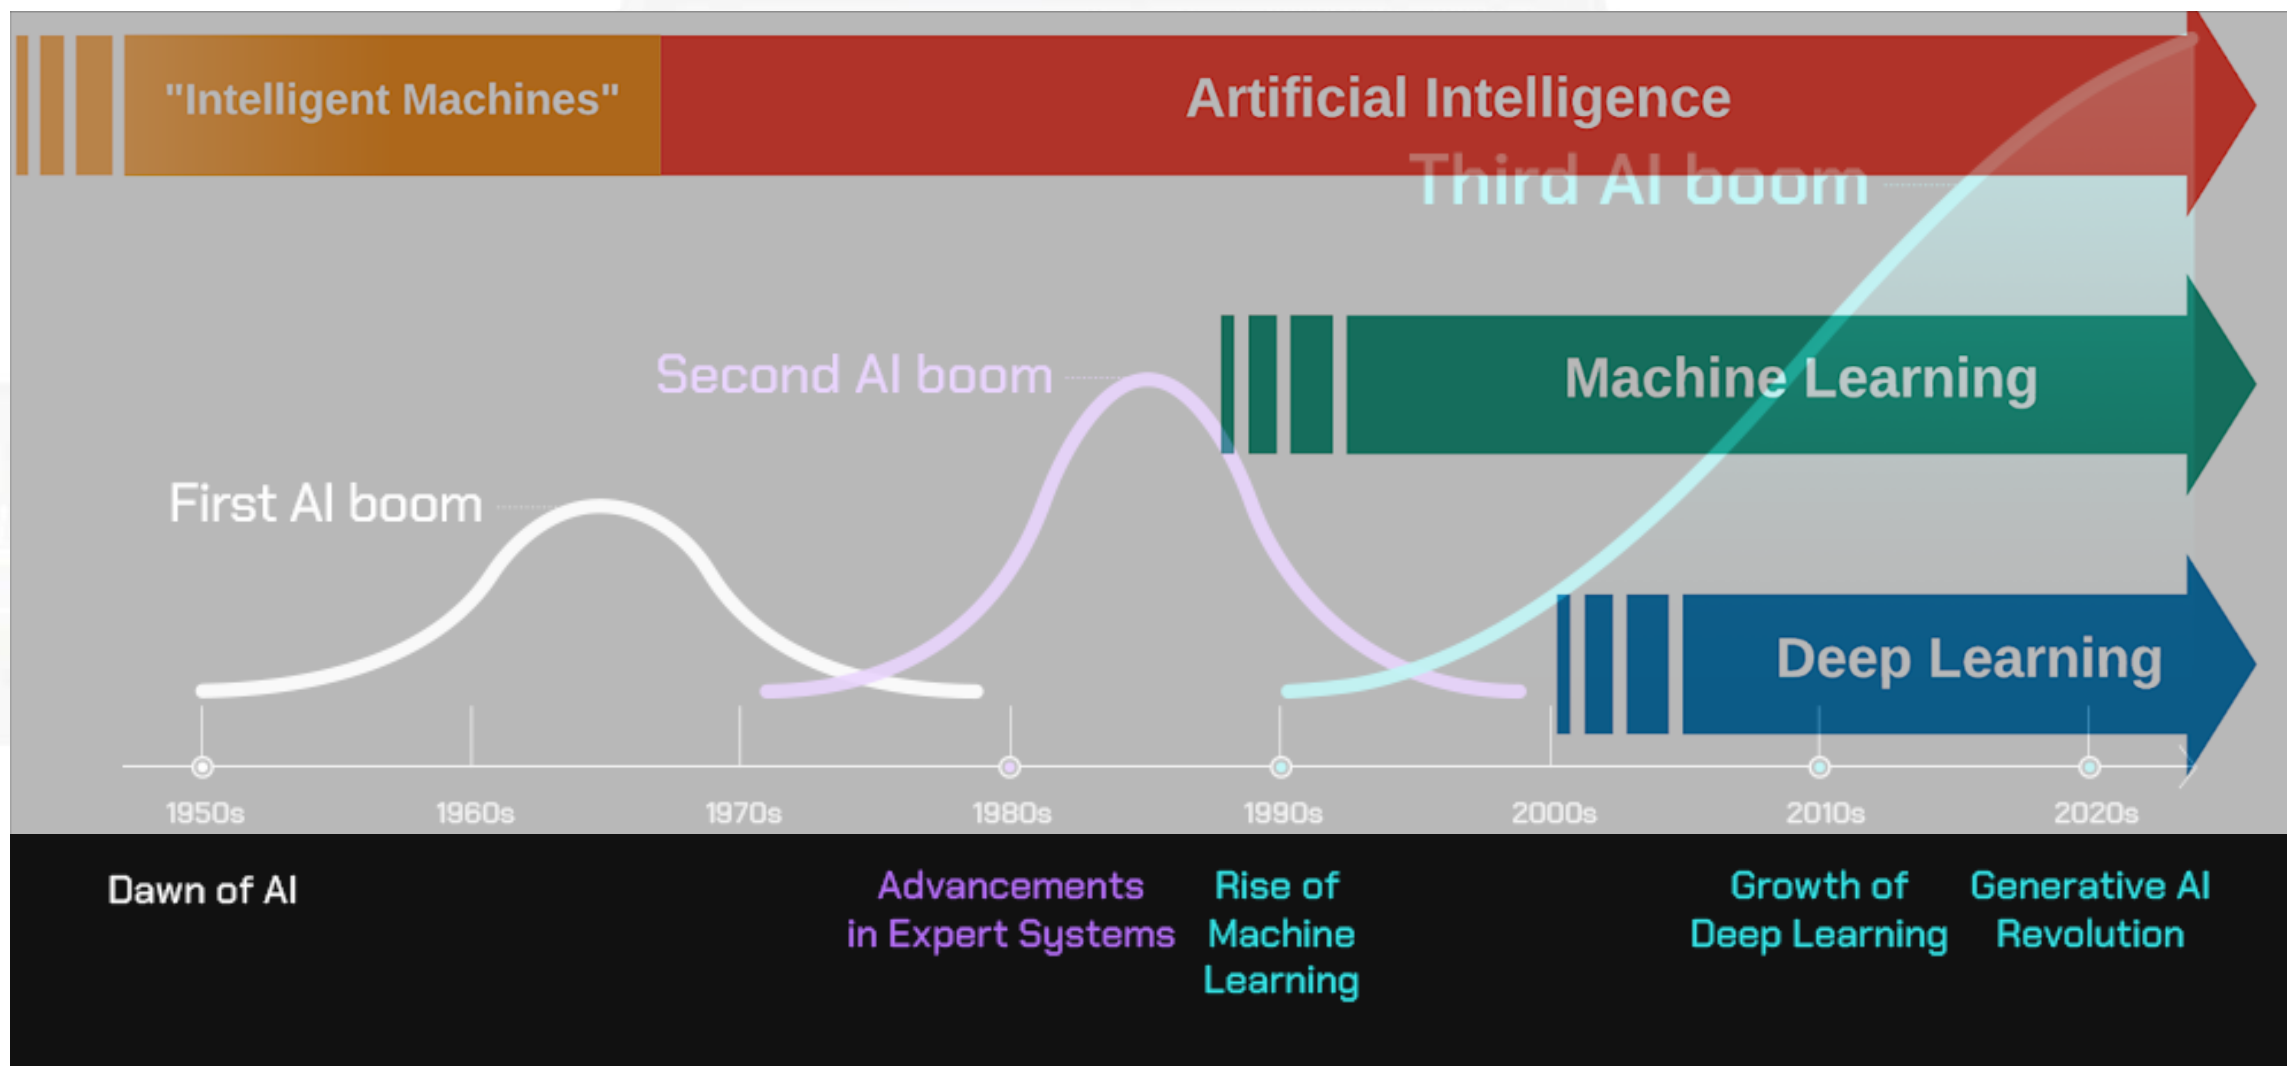
浙江大學
ZHEJIANG UNIVERSITY



The Development of AI



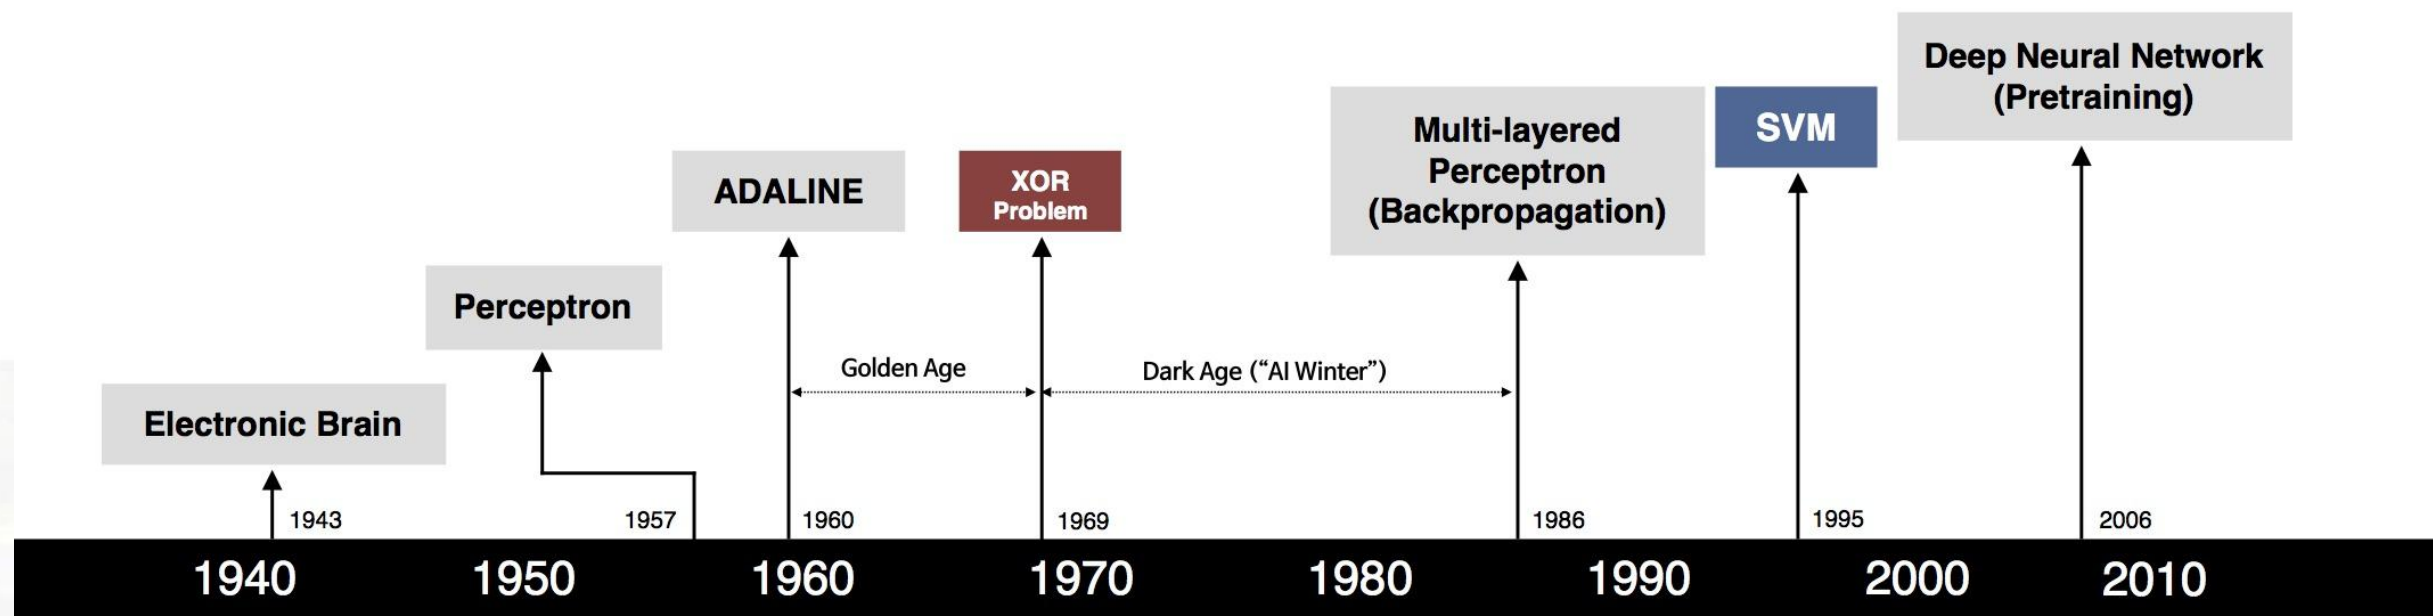
浙江大學
ZHEJIANG UNIVERSITY




The Development of DL



浙江大學
ZHEJIANG UNIVERSITY




S. McCulloch - W. Pitts



- Adjustable Weights
- Weights are not Learned



F. Rosenblatt




- Learnable Weights and Threshold



B. Widrow - M. Hoff




M. Minsky - S. Papert



- XOR Problem




D. Rumelhart - G. Hinton - R. Williams



- Solution to nonlinearly separable problems
- Big computation, local optima and overfitting




V. Vapnik - C. Cortes



- Limitations of learning prior knowledge
- Kernel function: Human Intervention



G. Hinton - S. Ruslan

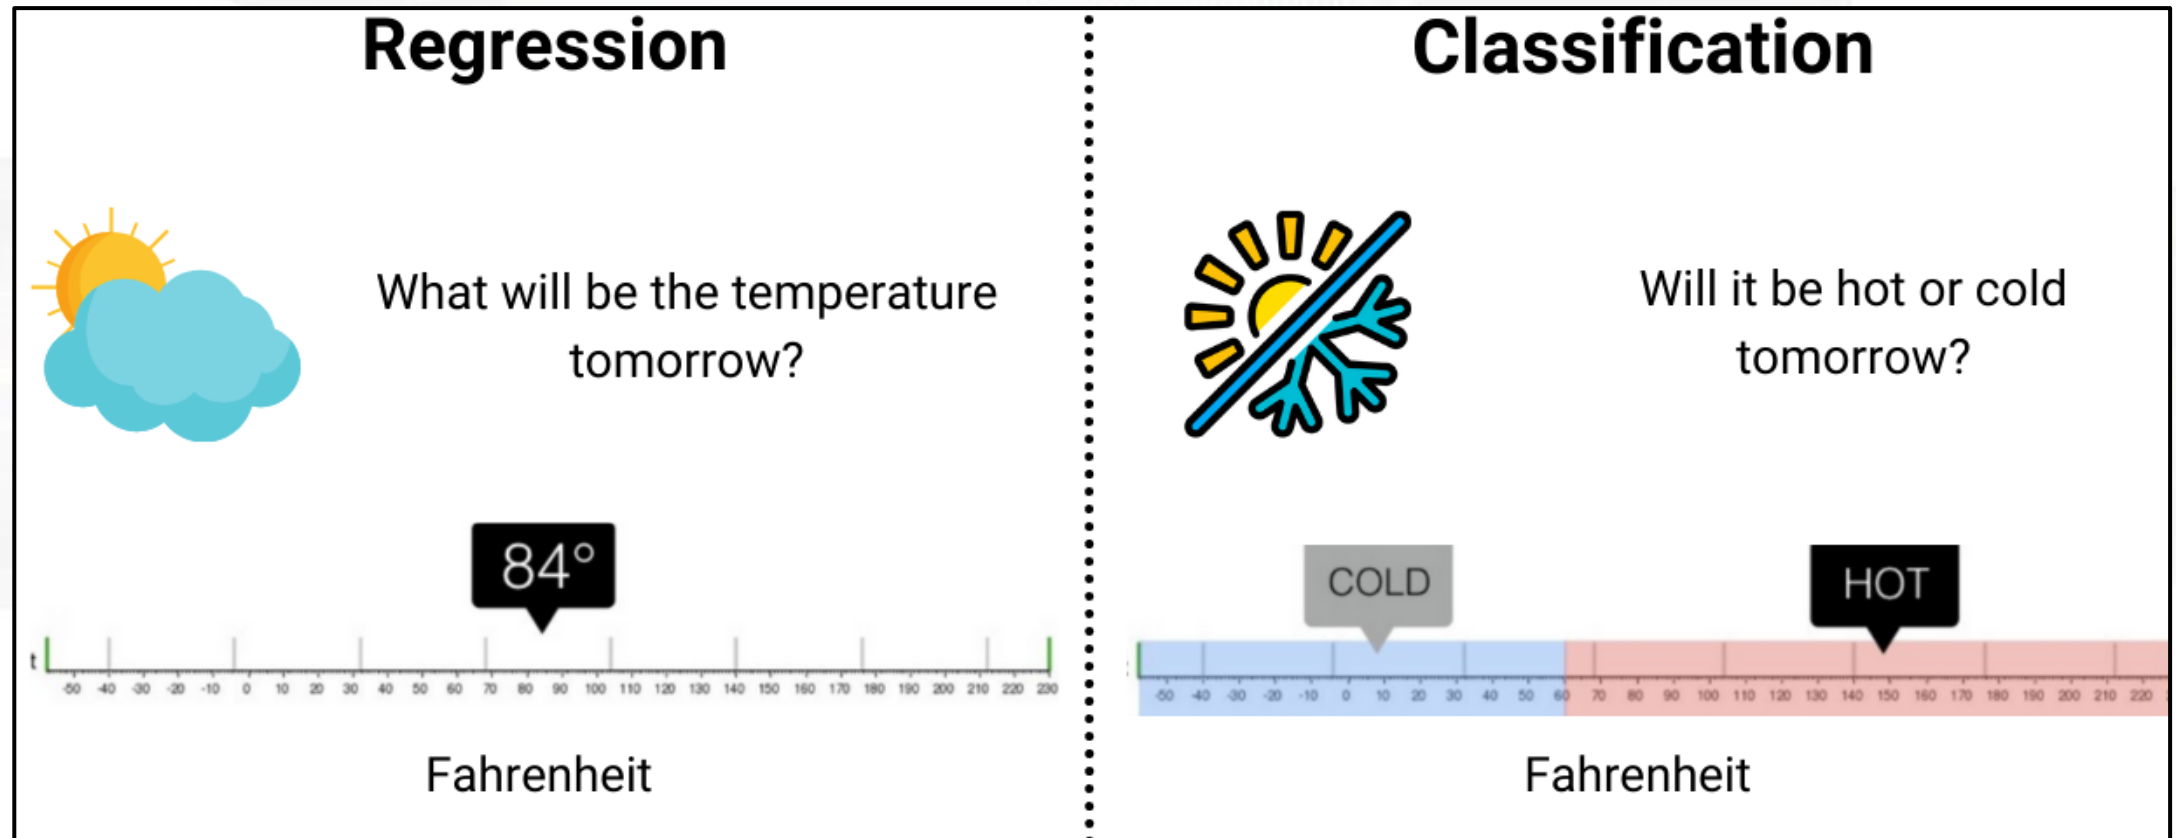


- Hierarchical feature Learning

Data View: Fundamental Problems



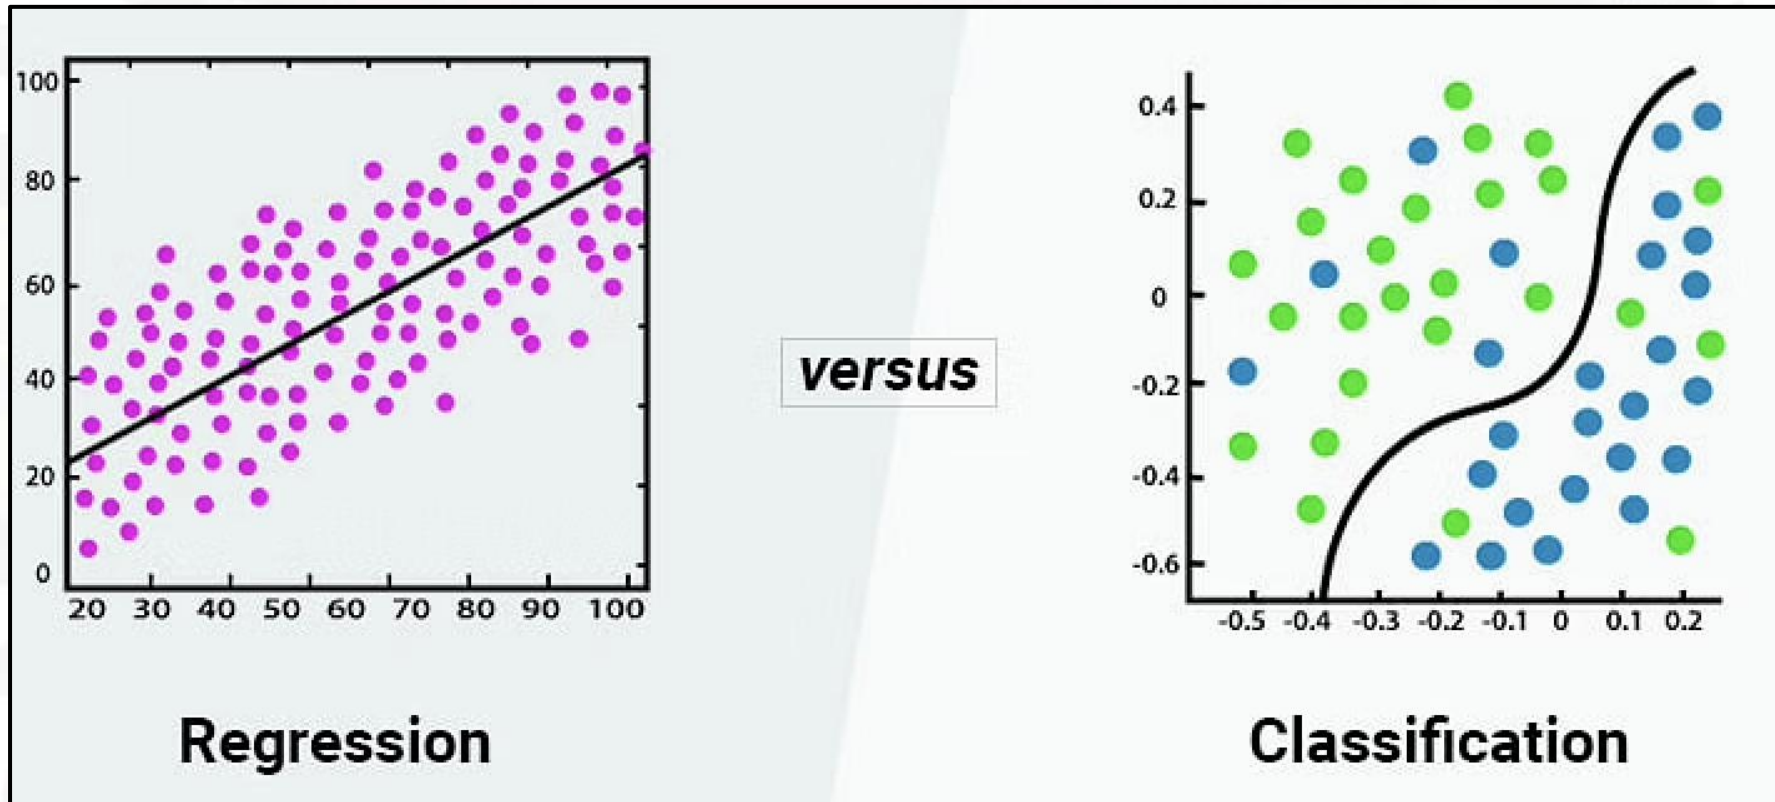
■ Output: Continuous description or discrete state



Data View: Fundamental Problems



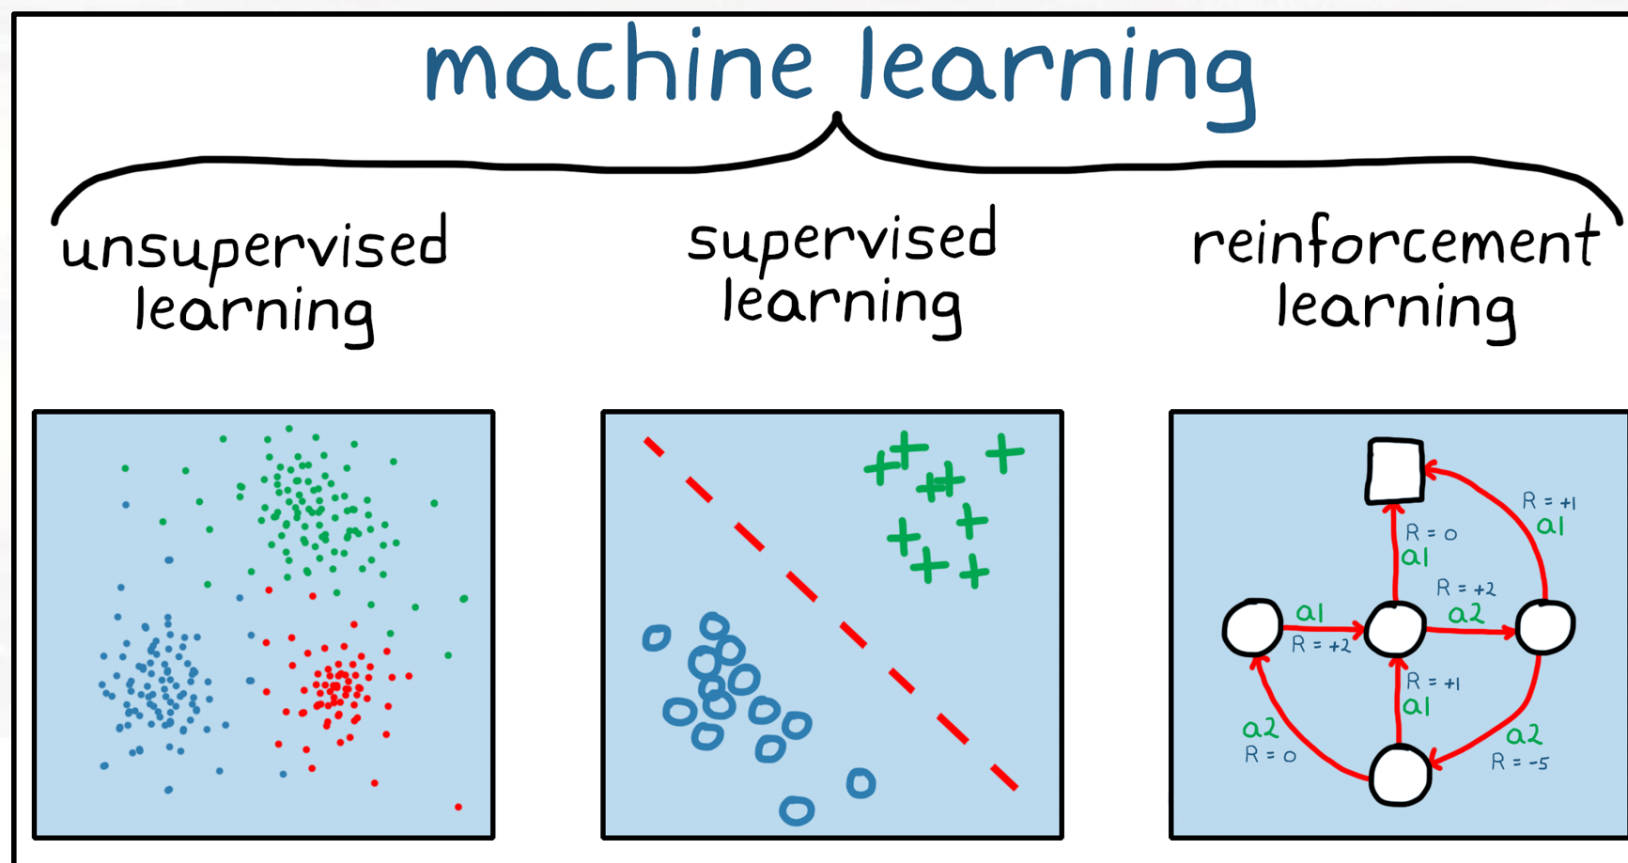
- Output: Continuous description or discrete state



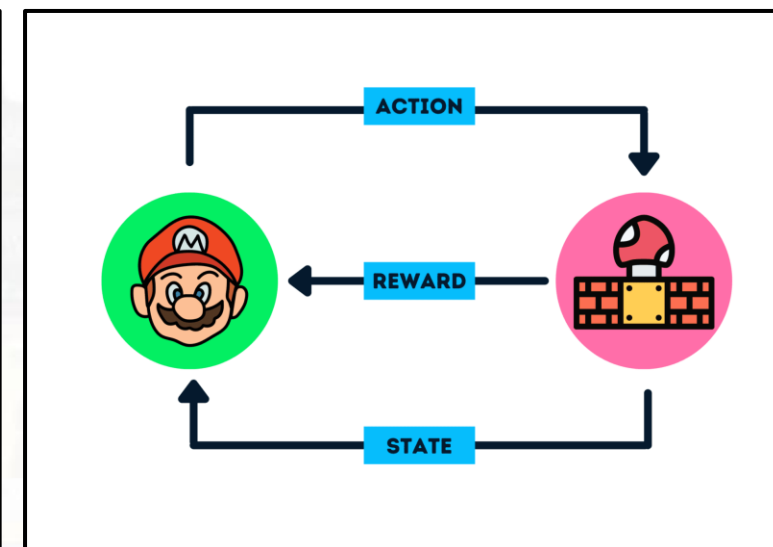
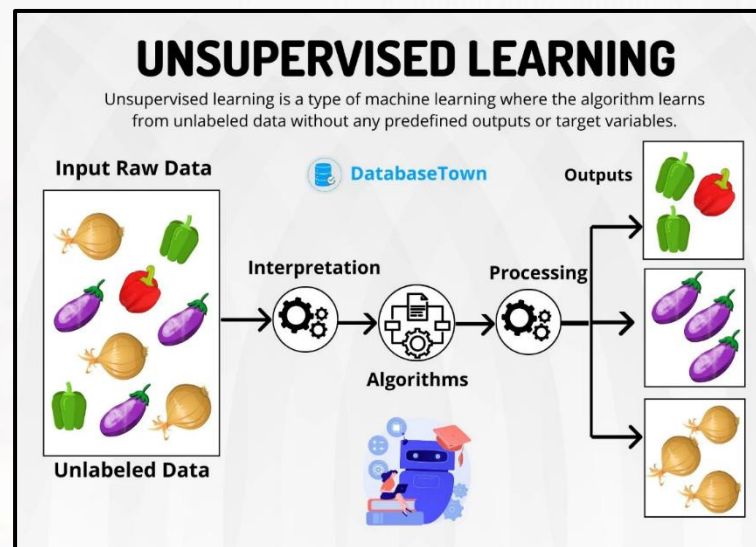
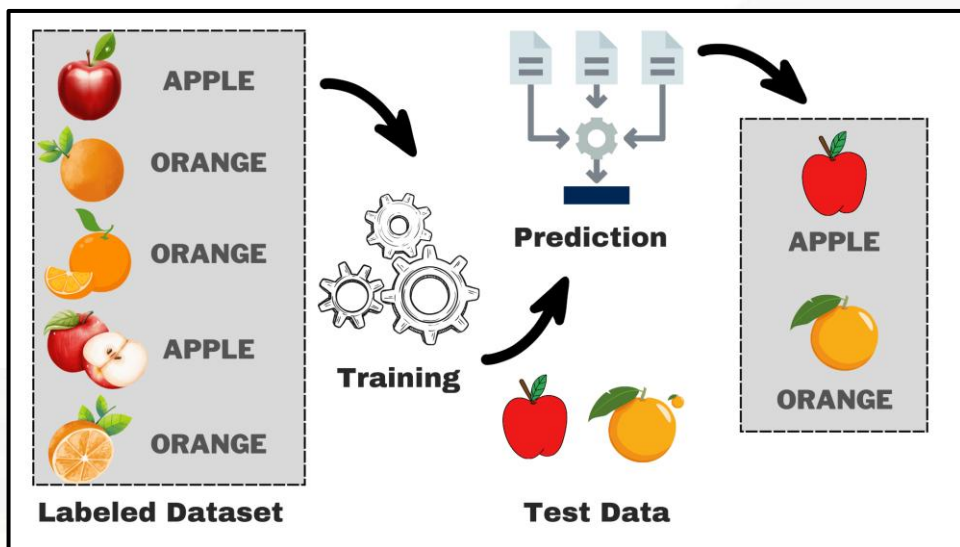
Data view: Learning Paradigm



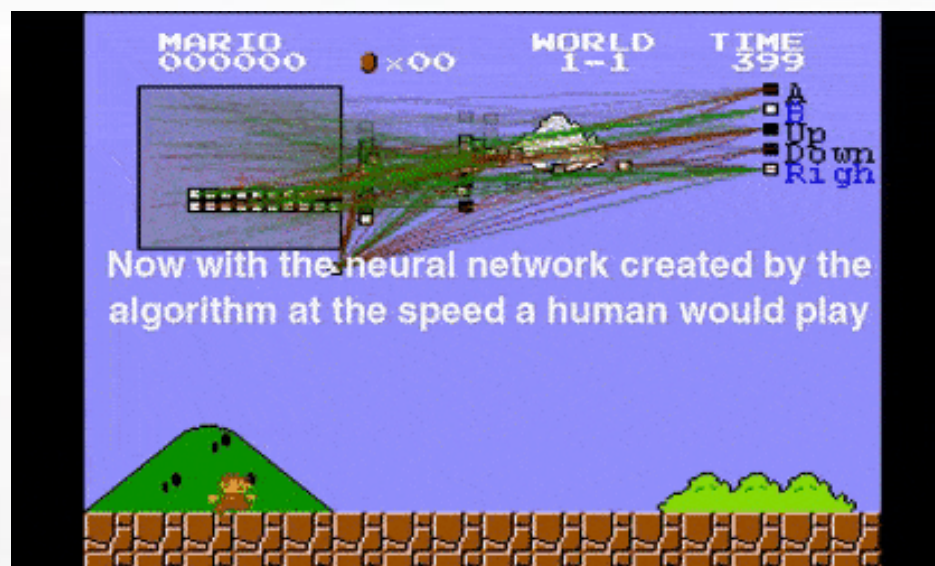
- Availability of outputs: unavailable, available, partially available



Data view: Learning Paradigm



Weight	Texture	Label
150g	Bumpy	Orange
170g	Bumpy	Orange
140g	Smooth	Apple
130g	Smooth	Apple




Action	↑	↓	→	←
Start	0	0	1	0
Idle	0	0	0	0
Correct Path	0	50	22	0
Wrong Path	15	0	18	0
End	0	0	1	0

Model view: **neural network**



Aspect	Neural Networks	Support Vector Machines (SVM)
Model Type	Non-linear, based on layers of neurons.	Linear or non-linear (with kernel tricks), based on maximizing margins.
Data Requirements	Requires large datasets for training.	Works well with small to medium-sized datasets.
Complexity	Can model highly complex, non-linear relationships.	Handles non-linearity through kernel functions.
Training Time	Computationally expensive, especially for deep learning.	Slower training for large datasets, quadratic complexity in training.
Scalability	Scales well with data and features but requires more computational power.	Struggles with very large datasets and high-dimensional data.
Interpretability	Black-box model; difficult to interpret.	Easier to interpret, especially with linear SVM.

Model view: Neural Network




Model view: neural network




- Large searching space and fitting capacity
- -> Scaling law
- -> Data-driven
- -> Fit everything
- -> Data is all you need and expert is gone

Computational Neuron and Perceptron




S. McCulloch – W. Pitts



- Adjustable Weights
- Weights are not Learned



F. Rosenblatt




- Learnable Weights and Threshold



B. Widrow – M. Hoff




M. Minsky – S. Papert



- XOR Problem




D. Rumelhart – G. Hinton – R. Williams



- Solution to nonlinearly separable problems
- Big computation, local optima and overfitting




V. Vapnik – C. Cortes



- Limitations of learning prior knowledge
- Kernel function: Human Intervention



G. Hinton – S. Ruslan






- Hierarchical feature Learning

Biological Neuron



Action potentials are the primary means of information transmission in biological neurons. They implement spatiotemporal information processing through an event-driven threshold-triggering mechanism (when the sum of synaptic inputs reaches a critical value), ensuring a balance between signal transmission reliability and energy efficiency.




Action potential illustration: A brief and specially shaped transmembrane potential pulse generated when the cell membrane at resting membrane potential is subjected to an appropriate stimulus.

Event-driven schematic: When an electrical signal (action potential) reaches the threshold, it triggers a pulse and transmits information directionally through synaptic release of chemical transmitters.


Biological Neuron




浙江大學
ZHEJIANG UNIVERSITY



Discrete Spiking representation :
An action potential is expressed in an **all-or-none** manner.



Asynchronous spiking information transmission: Whether the current neuron transmits information is **unrelated to whether other neurons transmit information**; it only depends on whether the threshold of the current neuron is triggered.






Equivalent circuit


Computational Neuron



浙江大學
ZHEJIANG UNIVERSITY



$$\varphi(\vec{X}) = \text{sgn}\left(\sum_{i=1..n} w_i \times x_i + b\right)$$



Example: Internet Traffic Prediction



- Suppose someone wants to make money through a video platform; they would care about whether the channel has traffic, so they would know their potential earnings.
- Assume that the backend can see a lot of relevant information, such as the number of people who like posts each day, the number of subscribers, and the number of views.
- Based on all the past information of a channel, it is possible to predict the number of views for tomorrow.
- Find a function whose input is the backend information and whose output is the total number of views the channel will have the next day.



Solution: Three Steps



■ Design a function with unknown parameters

□ model: f 、feature: x_1

□ parameter: b, w

□ weight: w 、bias: b

■ Define loss function: L

□ Mean Absolute Error, MAE: $e = |\hat{y} - y|$


□ Mean Squared Error, MSE: $e = (\hat{y} - y)^2$

□ cross entropy



■ Solve an optimization problem

□ global minimum and local minimum


$$y = b + wx_1$$




$$\hat{y} = 500 + 1x_1 \quad e_1 = |y - \hat{y}| = 400$$
$$e_2 = |y - \hat{y}| = 2100 \quad L = \frac{1}{N} \sum_n e_n$$




Advantage of Perception




$$y = b + wx_1$$



Multi-Layered Perceptron




S. McCulloch – W. Pitts



- Adjustable Weights
- Weights are not Learned



F. Rosenblatt




- Learnable Weights and Threshold



B. Widrow – M. Hoff




M. Minsky – S. Papert



- XOR Problem




D. Rumelhart – G. Hinton – R. Williams



- Solution to nonlinearly separable problems
- Big computation, local optima and overfitting




V. Vapnik – C. Cortes



- Limitations of learning prior knowledge
- Kernel function: Human Intervention



G. Hinton – S. Ruslan




- Hierarchical feature Learning

Multi-Layered Perceptron



浙江大學
ZHEJIANG UNIVERSITY





Multi-Layered Perceptron



- Understand MLP, understand everything
 - A few layers of neurons linked together.


Plain vanilla
(aka “multilayer perceptron”)





Multi-Layered Perceptron



- **Neuron:** a thing that holds a number
 - represent the inputs and outputs of our network (the images and digit predictions) in terms of these neuron values



0




$$28 \times 28 = 784$$

Multi-Layered Perceptron




what kind of digit does this network think it's looking at? How certain does it feel?




Multi-Layered Perceptron




■ More Compact Notation




Multi-Layered Perceptron




Multi-Layered Perceptron




Multi-Layered Perceptron



Multi-Layered Perceptron



Multi-Layered Perceptron



Multi-Layered Perceptron



- The result of the weighted sum like this can be any number, but for this network we want the activations to be values between 0 and 1.

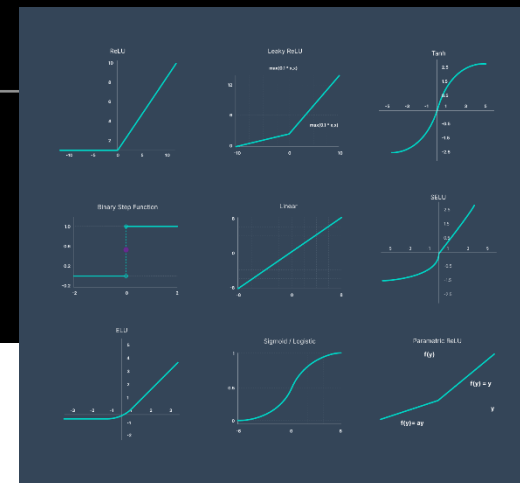

$$w_1a_1 + w_2a_2 + w_3a_3 + w_4a_4 + \cdots + w_na_n$$



Activations should be in this range

Sigmoid

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



Multi-Layered Perceptron



■ Formulated function

Superscript corresponds to the layer


$$a_0^{(1)} = \sigma(w_{0,0}a_0^{(0)} + w_{0,1}a_1^{(0)} + \dots + w_{0,n}a_n^{(0)} + b_0)$$

Subscript corresponds to a neuron in the layer


Neural network function



784 inputs

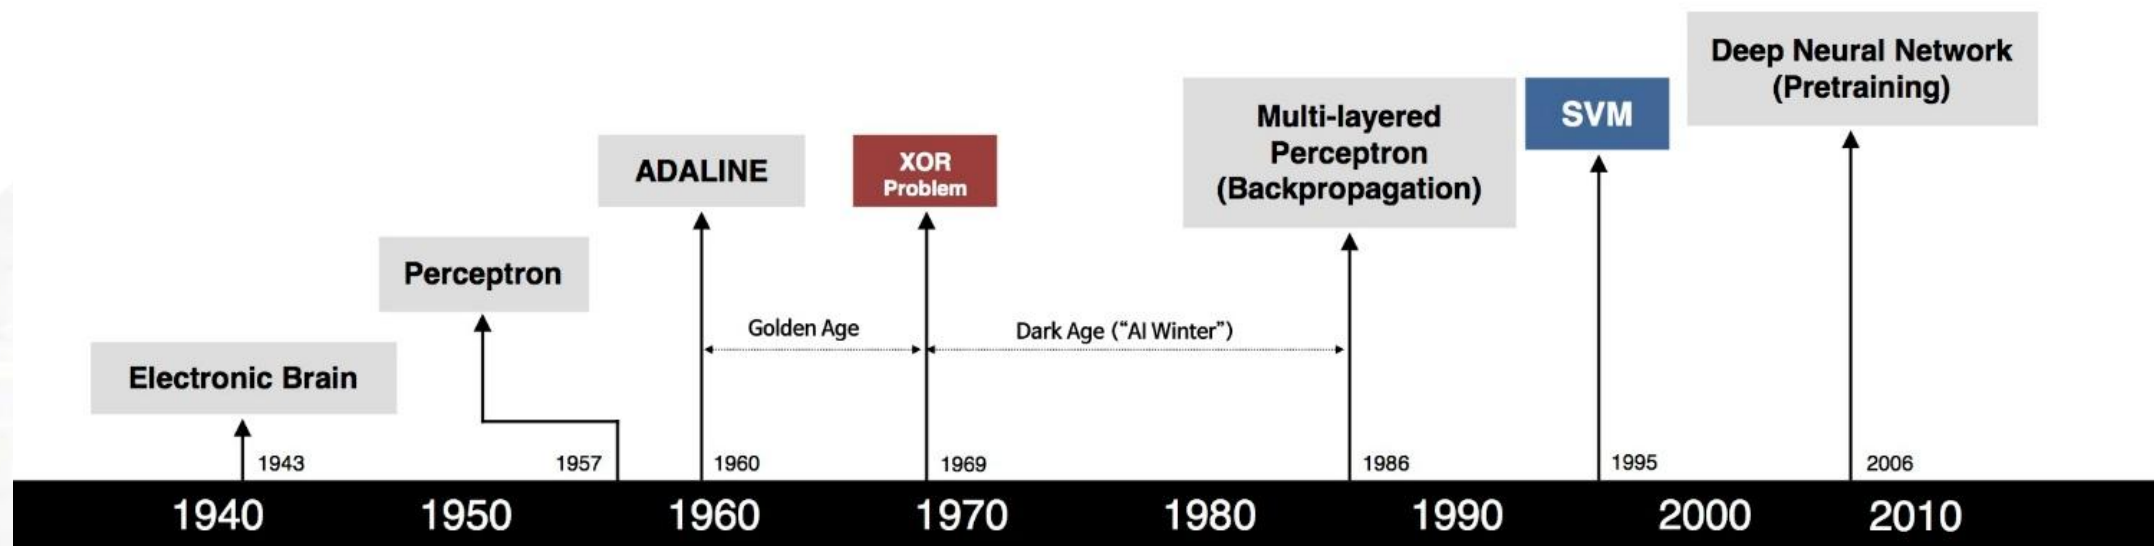


13,002 weights
and biases




10 outputs

Backpropagation




S. McCulloch – W. Pitts



- Adjustable Weights
- Weights are not Learned



F. Rosenblatt




- Learnable Weights and Threshold



B. Widrow – M. Hoff




M. Minsky – S. Papert



- XOR Problem




D. Rumelhart – G. Hinton – R. Williams



- Solution to nonlinearly separable problems
- Big computation, local optima and overfitting




V. Vapnik – C. Cortes



- Limitations of learning prior knowledge
- Kernel function: Human Intervention

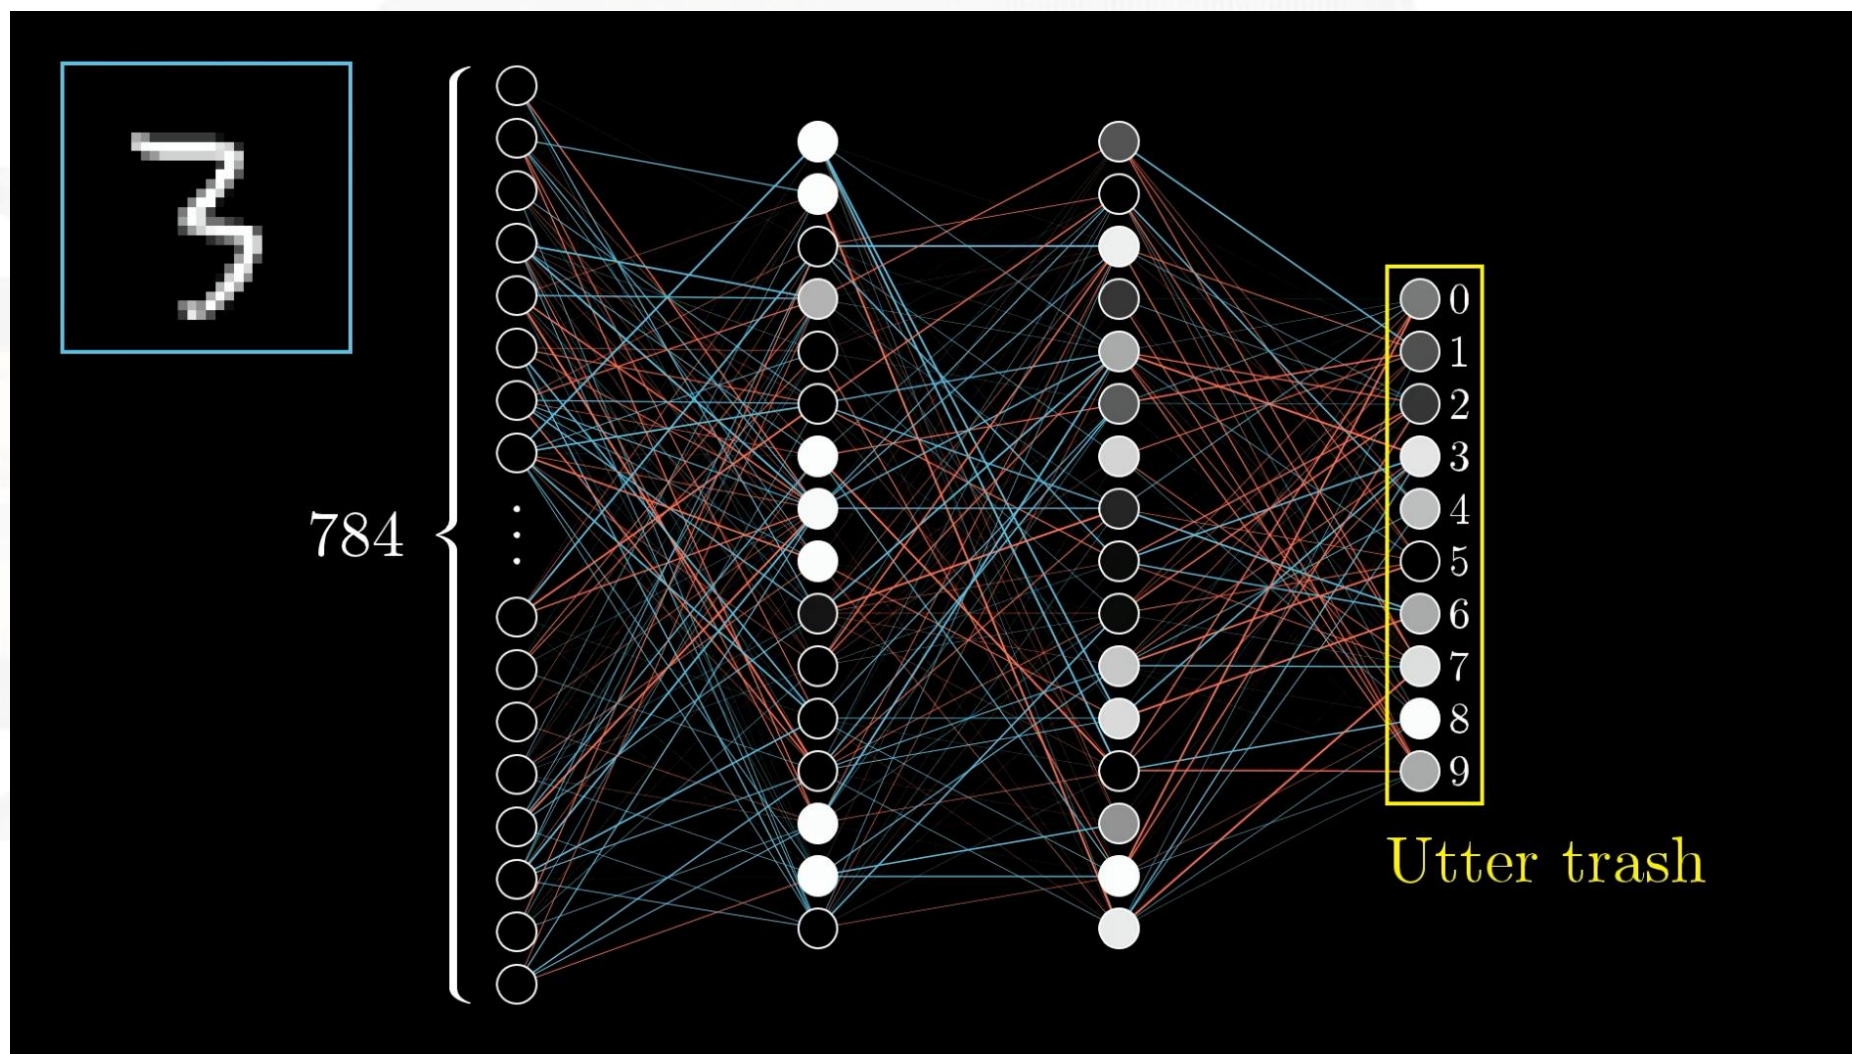


G. Hinton – S. Ruslan



- Hierarchical feature Learning



Optimization



Optimization



- The behavior of the network depends on its weights and biases.



■ Data Preparation

(0, 0) (6, 6) (3, 3) (6, 6) (7, 7) (8, 8) (0, 0) (9, 9)
(5, 5) (4, 4) (3, 3) (4, 6) (5, 5) (8, 8) (9, 9) (5, 5)
(4, 4) (4, 4) (7, 7) (2, 2) (0, 0) (3, 3) (2, 2) (8, 8)
(9, 9) (1, 1) (9, 9) (2, 2) (2, 2) (7, 7) (9, 9) (4, 4)
(8, 8) (7, 7) (4, 4) (1, 1) (3, 3) (1, 1) (5, 5) (3, 3)
(2, 2) (3, 3) (9, 9) (0, 0) (9, 9) (9, 9) (1, 1) (5, 5)
(8, 8) (4, 4) (1, 7) (7, 7) (4, 4) (4, 4) (4, 4) (2, 2)
(0, 0) (7, 7) (2, 2) (4, 4) (8, 8) (2, 2) (6, 6) (9, 9)
(9, 9) (2, 2) (8, 8) (7, 7) (6, 6) (1, 1) (1, 1) (2, 2)
(3, 3) (9, 9) (1, 1) (6, 6) (5, 5) (1, 1) (1, 1) (0, 0)

Optimization



Train on
these

5	→	5
0	→	0
4	→	4
1	→	1
9	→	9
2	→	2
1	→	1
3	→	3
1	→	1
4	→	4


Test on
these

3	→	3
5	→	5
3	→	3
6	→	6
1	→	1
7	→	7
2	→	2
8	→	8
6	→	6
9	→	9

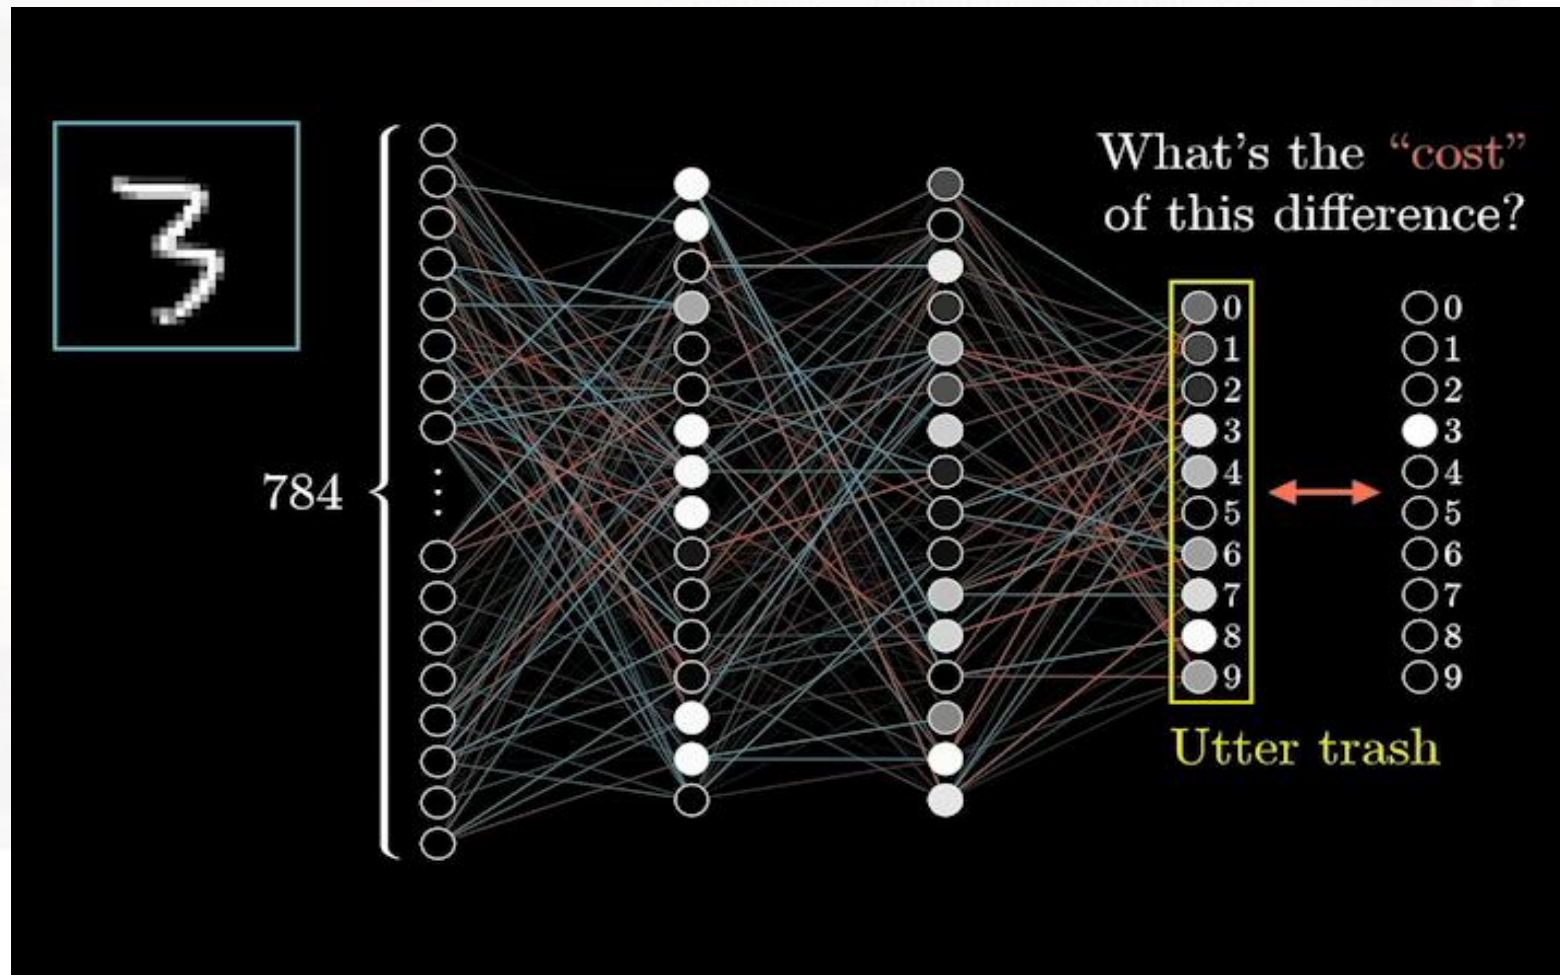
Optimization




- finding the minimum of a specific **function**



■ The Cost Function



Cost function



13,002 weights
and biases



(9, 9) (0, 0) (2, 2) (6, 6)
(0, 0) (4, 4) (6, 6) (7, 7)
(7, 7) (8, 8) (3, 3) (1, 1)
(1, 1) (1, 1) (6, 6) (3, 3)
(1, 1) (1, 1) (0, 0) (4, 4)

Lots of training data




3.37

One number

Optimization




浙江大學
ZHEJIANG UNIVERSITY



Optimization




- For a complicated cost function, computing the exact minimum directly isn't going to work.



Optimization




- By following the slope (moving in the downhill direction), we approach a local minimum.



Optimization




- As the slope gets shallower, take smaller steps to avoid overshooting the minimum.



Optimization




- Moving the input position according to the slope is a lot like a ball rolling down a hill.



Optimization



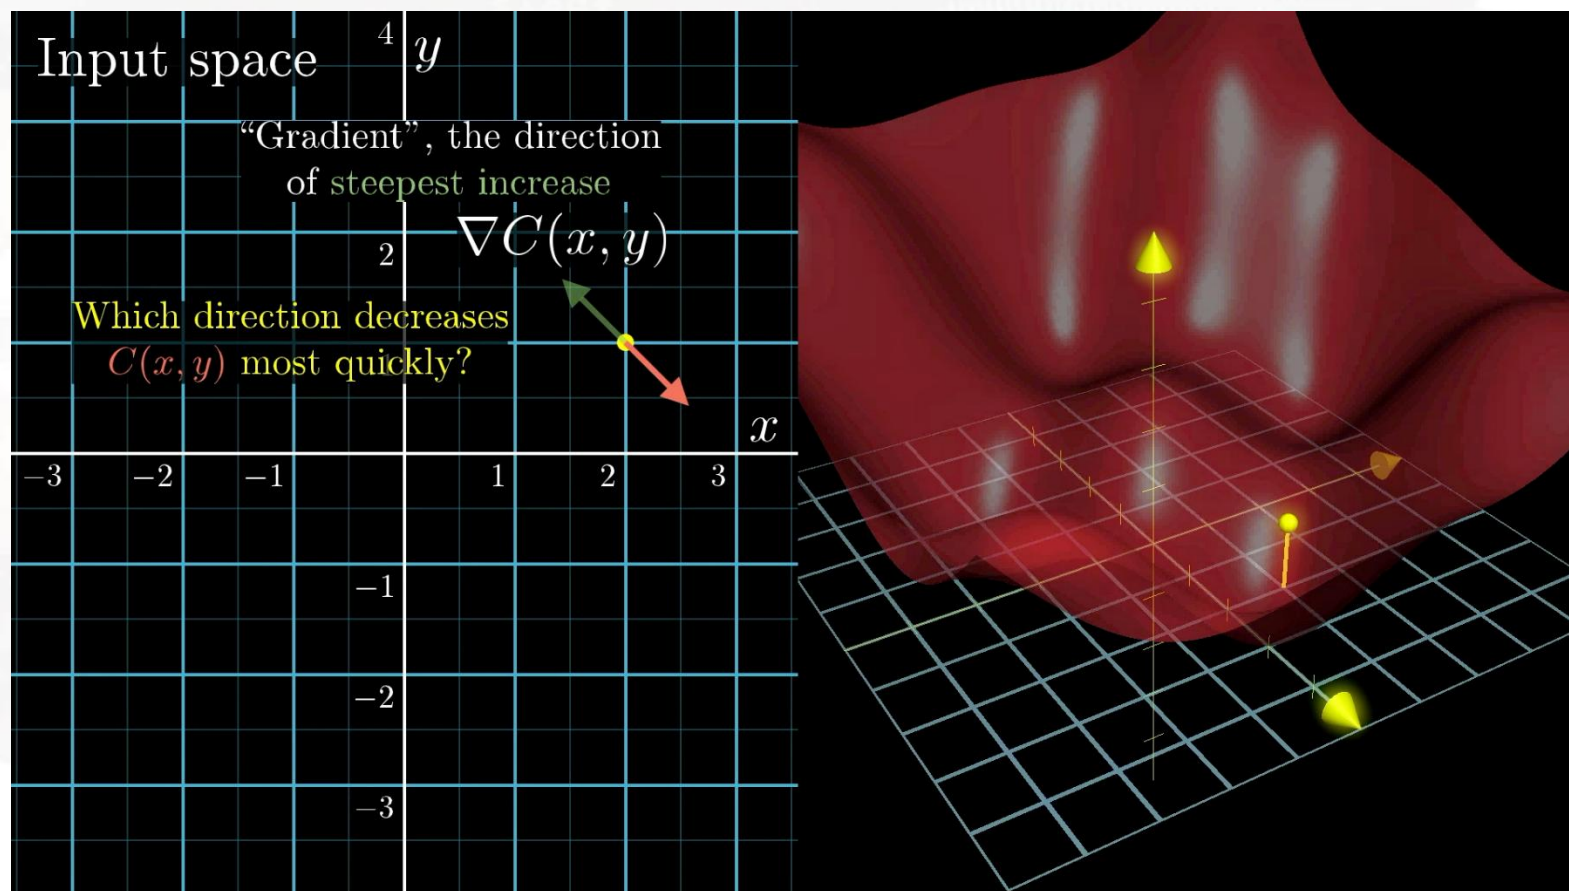
- We can imagine minimizing a function that takes two inputs
- Gradient descent just means walking in the downhill direction to minimize the cost function.



Optimization



- The gradient, ∇C , gives the uphill direction, so the negative of the gradient, $-\nabla C$, gives the downhill direction.



■ Another Way to Think About The Gradient

13,002 weights and biases


$$\vec{W} = \begin{bmatrix} 2.25 \\ -1.57 \\ 1.98 \\ \vdots \\ -1.16 \\ 3.82 \\ 1.21 \end{bmatrix}$$

How to nudge all weights and biases

$$-\nabla C(\vec{W}) = \begin{bmatrix} 0.18 \\ 0.45 \\ -0.51 \\ \vdots \\ 0.40 \\ -0.32 \\ 0.82 \end{bmatrix}$$

$$-\nabla C(\underbrace{\dots}_{\text{All weights and biases}}) = \begin{bmatrix} 0.20 \\ 0.83 \\ -0.84 \\ 0.04 \\ \vdots \\ 1.57 \\ 1.59 \\ -1.50 \\ -1.17 \end{bmatrix}$$


Change by some small multiple of $-\nabla C(\dots)$



Optimization




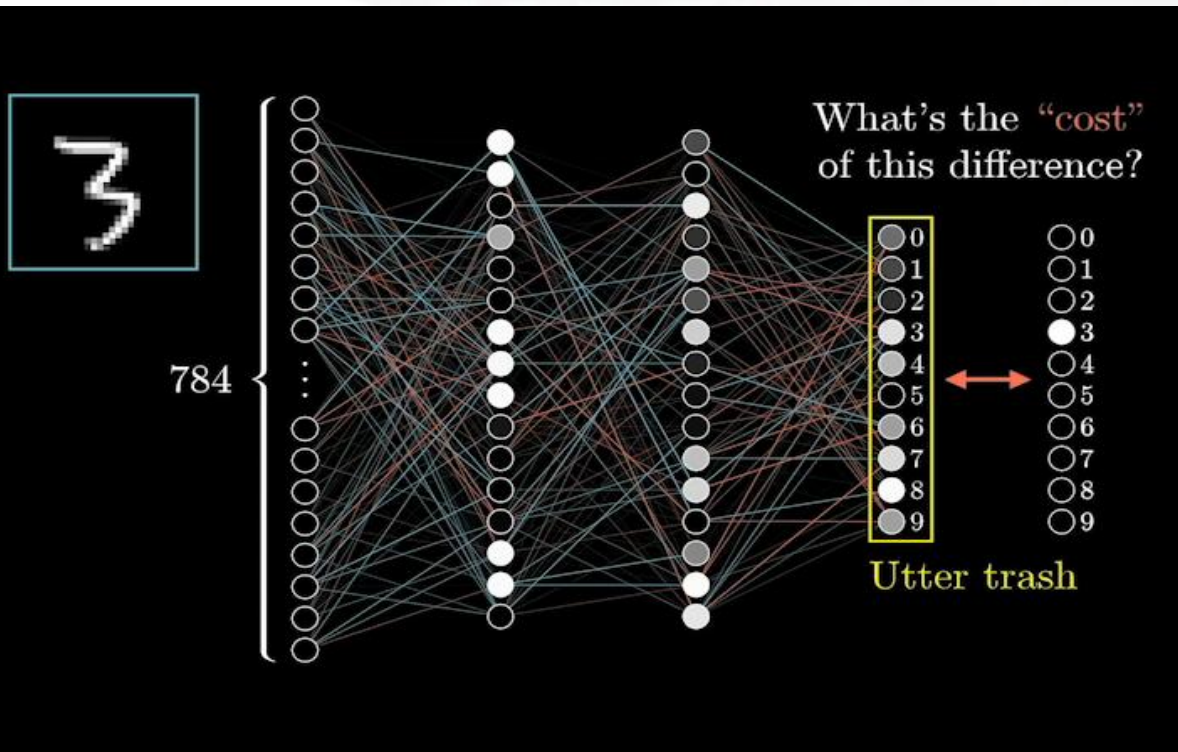
- Changing a weight that has a larger magnitude in the negative gradient vector has a bigger effect on the cost.



Forward and Backward




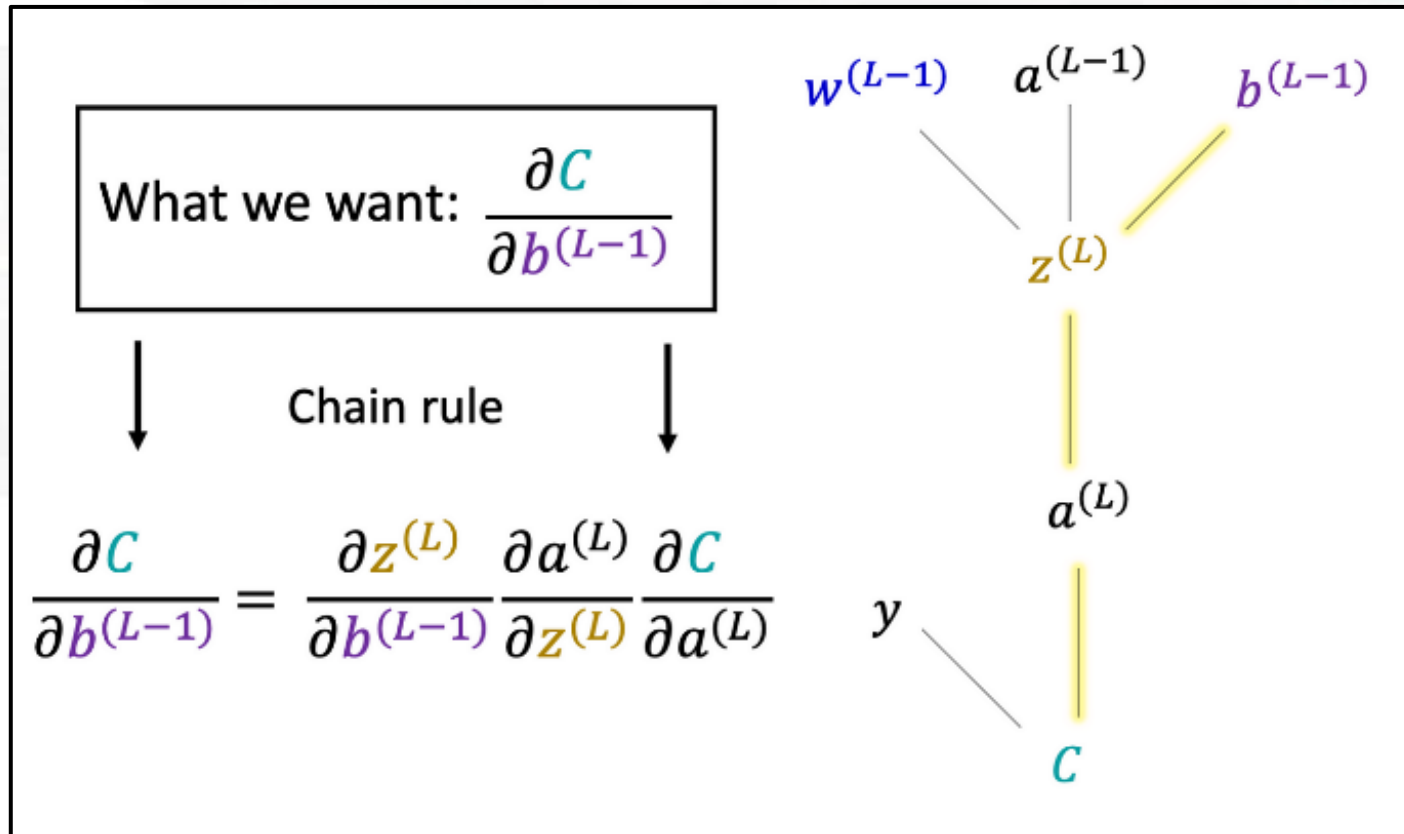
- Calculate the prediction error for training data
- Update model parameters based on the error




Backpropagation



- Approximate partial derivatives based on the chain rule
- Solving neural network parameters based on gradient descent



Chain Rule



Matrix operations form

sample-wise and parameter-wise differentiation

$$\frac{\partial C_0}{\partial w^{(L-1)}} = \frac{\partial z^{(L-1)}}{\partial w^{(L-1)}} \frac{\partial a^{(L-1)}}{\partial z^{(L-1)}} \frac{\partial z^{(L)}}{\partial a^{(L-1)}} \frac{\partial a^{(L)}}{\partial z^{(L)}} \frac{\partial C_0}{\partial a^{(L)}}$$

$$C_0 = (a^{(L)} - y)^2 \rightarrow \frac{\partial C_0}{\partial a^{(L)}} = 2(a^{(L)} - y)$$

$$a^{(L)} = \sigma(z^{(L)}) \rightarrow \frac{\partial a^{(L)}}{\partial z^{(L)}} = \sigma'(z^{(L)})$$

$$z^{(L)} = w^{(L)} a^{(L-1)} + b^{(L)} \rightarrow \frac{\partial z^{(L)}}{\partial a^{(L-1)}} = w^{(L)}$$

$$z^{(L)} = w^{(L)} a^{(L-1)} + b^{(L)} \rightarrow \frac{\partial z^{(L)}}{\partial w^{(L)}} = a^{(L-1)}$$

$$-\nabla C(\vec{W}) = \begin{bmatrix} 0.31 \\ 0.03 \\ -1.25 \\ \vdots \\ 0.78 \\ -0.37 \\ 0.16 \end{bmatrix}$$

w_0 should increase somewhat

w_1 should increase a little



w_2 should decrease a lot

$w_{13,000}$ should increase a lot

$w_{13,001}$ should decrease somewhat

$w_{13,002}$ should increase a little


Gradient Descent



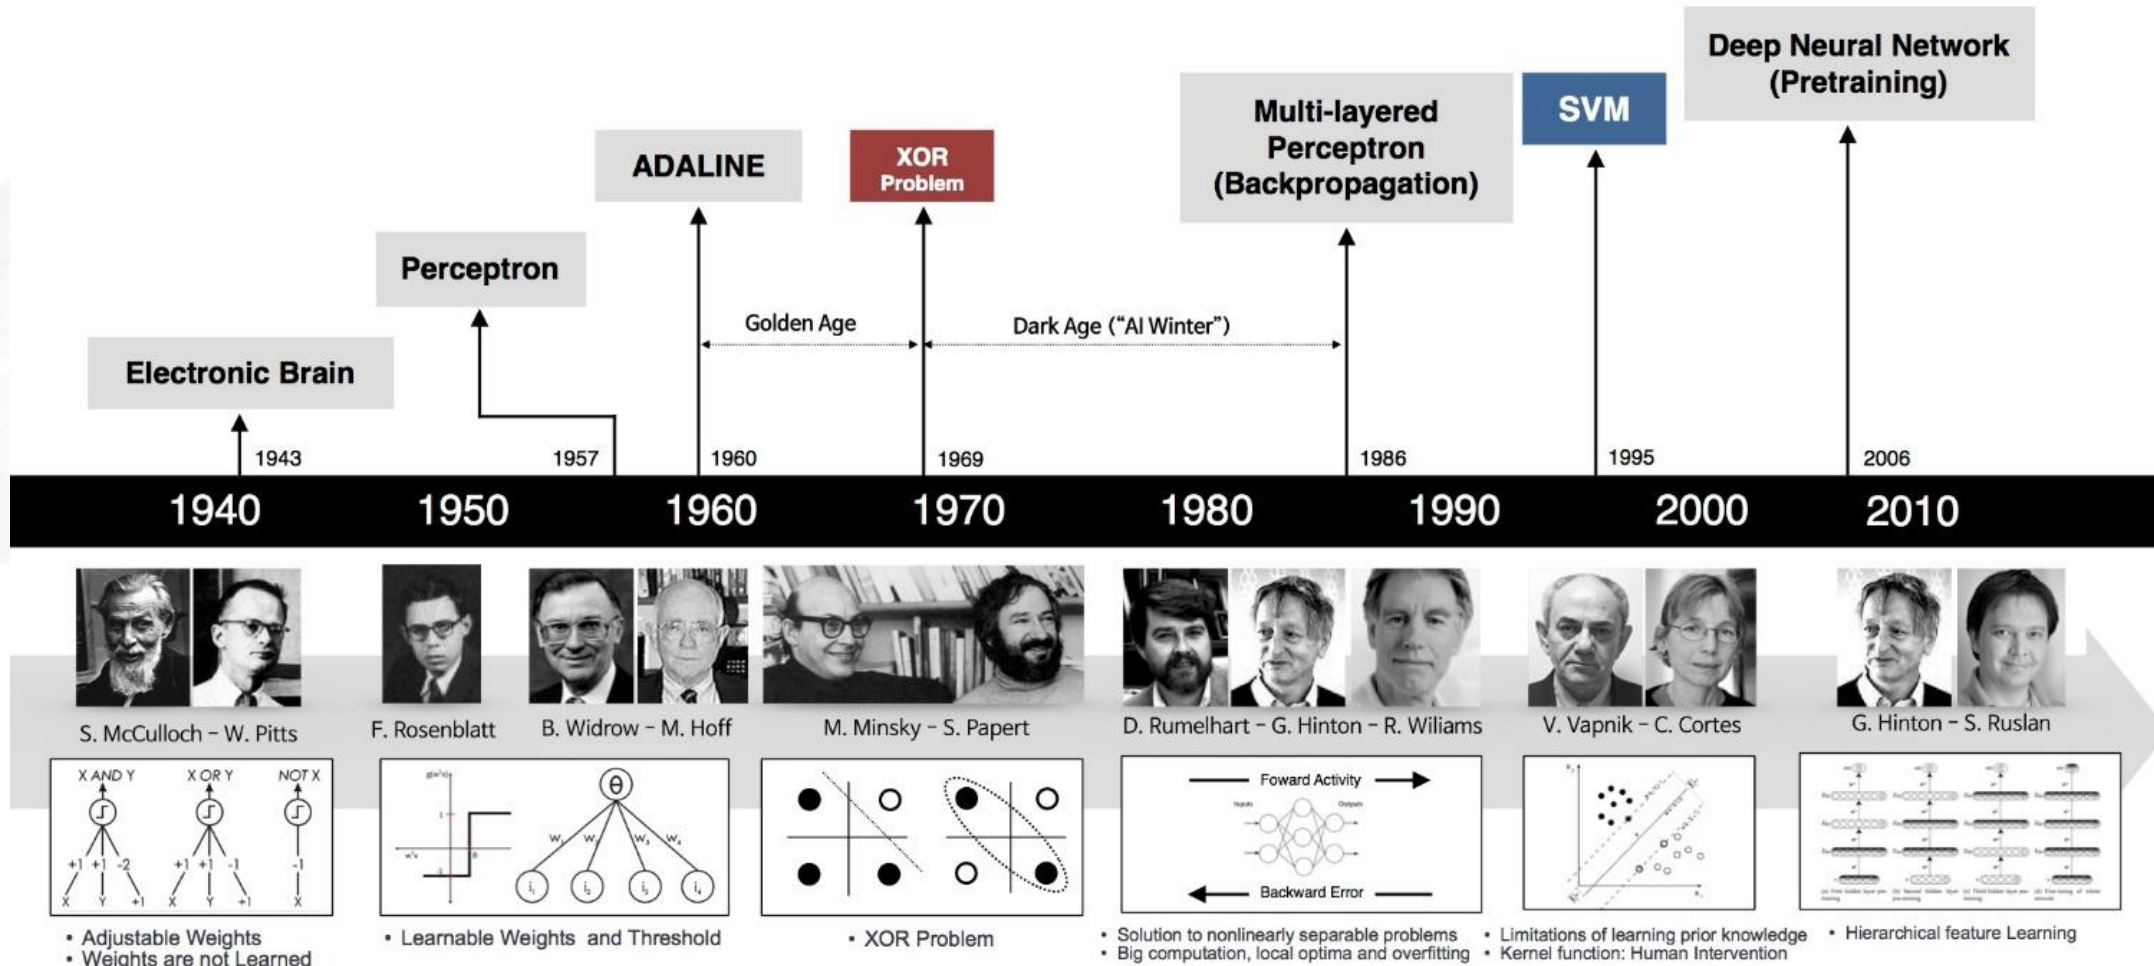
■ Algorithm

- Calculate ∇C based on chain rules
- Move all parameters toward $-\nabla C$ slightly
- Repeat all the steps above

- (randomly) select θ_0
- get gradient $g_0 = \nabla L_1(\theta_0)$
update $\theta_1 \leftarrow \theta_0 - \eta g_0$
- get gradient $g_1 = \nabla L_2(\theta_1)$
update $\theta_2 \leftarrow \theta_1 - \eta g_1$
- get gradient $g_2 = \nabla L_3(\theta_2)$
update $\theta_3 \leftarrow \theta_2 - \eta g_2$



The Problem of Backpropagation



The History of Backpropagation



■ The time when the relevant ideas were proposed

- ❑ 1986, the term was widely known
- ❑ 1974, the first time used to train neural network
- ❑ 1960, The basic knowledge of backpropagation has been accepted and widely used.
- ❑ 2006, Introduction of pre-training and fine-tuning mechanisms

■ Issues During Training

- ❑ Gradient vanishing
- ❑ Gradient exploding

Published: 09 October 1986

Learning representations by back-propagating errors

[David E. Rumelhart](#), [Geoffrey E. Hinton](#) & [Ronald J. Williams](#)

[Nature](#) 323, 533–536 (1986) | [Cite this article](#)

110k Accesses | 14712 Citations | 378 Altmetric | [Metrics](#)


Gradient Theory of Optimal Flight Paths

HENRY J. KELLEY¹

Grumman Aircraft Engineering Corp.
Bethpage, N. Y.

An analytical development of flight performance optimization according to the method of gradients or “method of steepest descent” is presented. Construction of a minimizing sequence of flight paths by a stepwise process of descent along the local gradient direction is described as a computational scheme. Numerical application of the technique is illustrated in a simple example of orbital transfer via solar sail propulsion. Successive approximations to minimum time planar flight paths from Earth’s orbit to the orbit of Mars are presented for cases corresponding to free and fixed boundary conditions on terminal velocity components.

Gradient Vanishing and Exploding



Gradient Vanishing

The weights are **almost 0**

The weights near the output layer **update quickly**, while the weights in the input layer **hardly update**.

Converges very slowly

Stop learning

Gradient Exploding

Weights with **NaN value**

Weights **increase explosively**

Instable performance

Performs poorly on the training set

Gradient Vanishing and Exploding



$$\frac{\partial C_0}{\partial w^{(L-1)}} = \frac{\partial z^{(L-1)}}{\partial w^{(L-1)}} \boxed{\frac{\partial a^{(L-1)}}{\partial z^{(L-1)}} \frac{\partial z^{(L)}}{\partial a^{(L-1)}}} \frac{\partial a^{(L)}}{\partial z^{(L)}} \frac{\partial C_0}{\partial a^{(L)}}$$

$$a^{(L)} = \sigma \left(z^{(L)} \right)$$

→


$$\boxed{\frac{\partial a^{(L)}}{\partial z^{(L)}} = \sigma' \left(z^{(L)} \right)}$$

$$z^{(L)} = w^{(L)} a^{(L-1)} + b^{(L)}$$


→

$$\boxed{\frac{\partial z^{(L)}}{\partial a^{(L-1)}} = w^{(L)}}$$

Sigmoid activation function



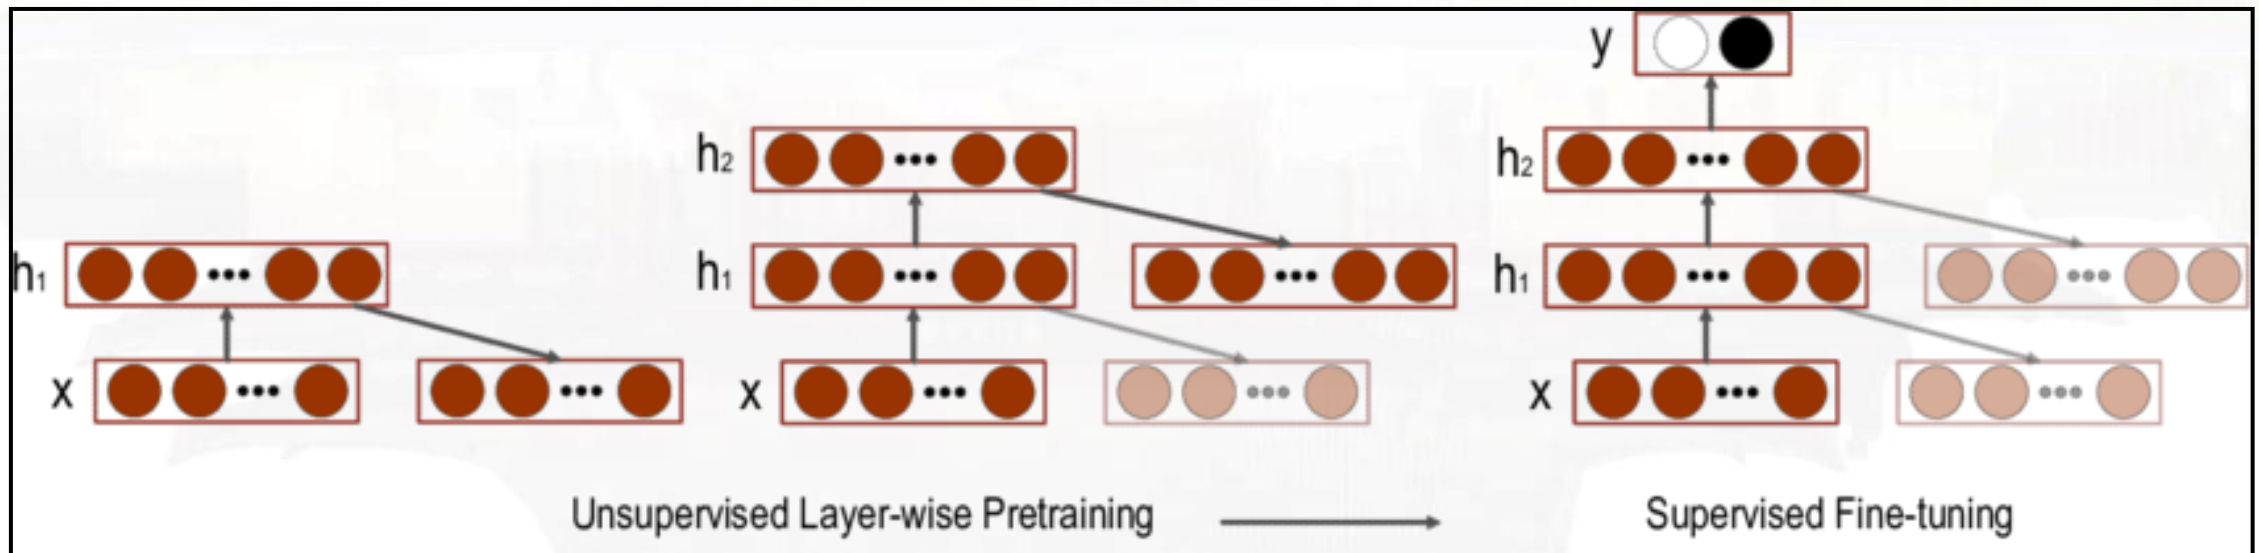
Derivative of the Sigmoid



Gradient Vanishing and Exploding




- Use **greedy unsupervised learning** to find a sensible set of weights one layer at a time. Then fine-tune with backpropagation.
 - The precious information in the labels is only used for the final fine-tuning.
 - We do not start backpropagation until we already have sensible weights that already do well at the task.



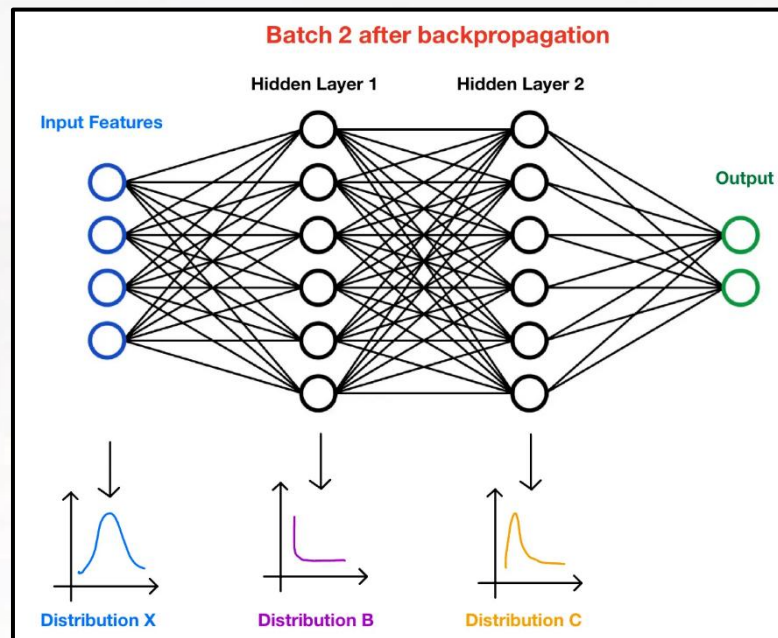
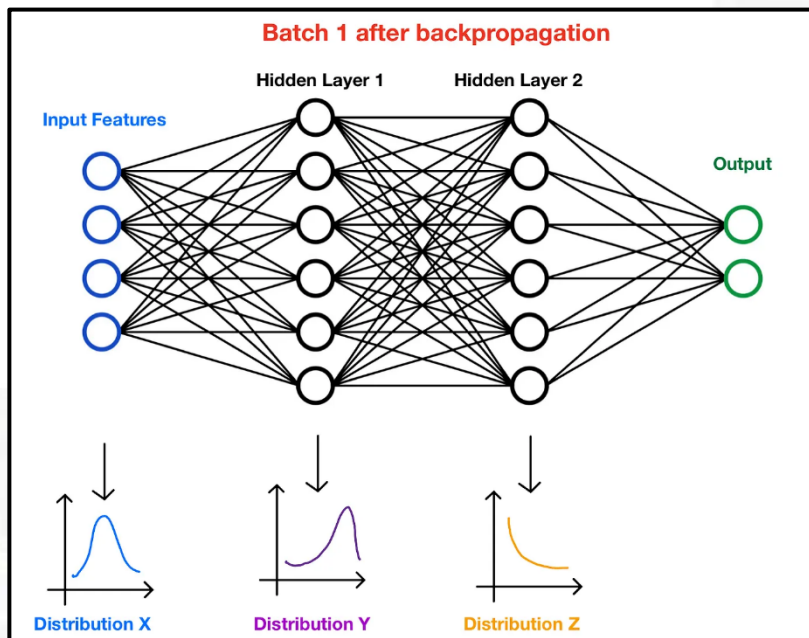
Gradient Vanishing and Exploding



1. *Zero Initialization*: Initialize all the weights and biases to zero. This is not generally used in deep learning as it leads to symmetry in the gradients, resulting in all the neurons learning the same feature.
2. *Random Initialization*: Initialize the weights and biases randomly from a uniform or normal distribution. This is the most common technique used in deep learning.
3. *Xavier Initialization*: Initialize the weights with a normal distribution with mean 0 and variance of $\sqrt{1/n}$, where n is the number of neurons in the previous layer. This is used for the sigmoid activation function.
4. *He Initialization*: Initialize the weights with a normal distribution with mean 0 and variance of $\sqrt{2/n}$, where n is the number of neurons in the previous layer. This is used for the ReLU activation function.
5. *Orthogonal Initialization*: Initialize the weights with an orthogonal matrix, which preserves the gradient norm during backpropagation.
6. *Uniform Initialization*: Initialize the weights with a uniform distribution. This is less commonly used than random initialization.
7. *Constant Initialization*: Initialize the weights and biases with a constant value. This is rarely used in deep learning.



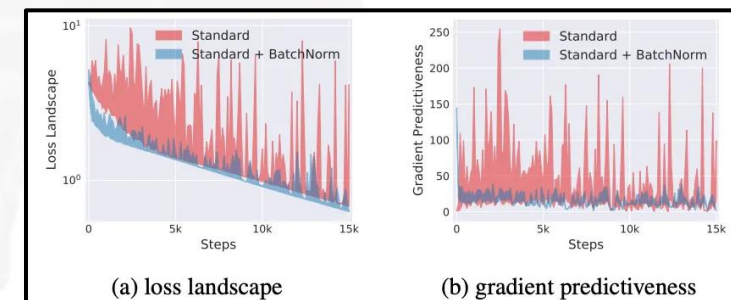

Gradient Vanishing and Exploding



$$\mu_{\mathcal{B}} \leftarrow \frac{1}{m} \sum_{i=1}^m x_i \quad // \text{ mini-batch mean}$$
$$\sigma_{\mathcal{B}}^2 \leftarrow \frac{1}{m} \sum_{i=1}^m (x_i - \mu_{\mathcal{B}})^2 \quad // \text{ mini-batch variance}$$
$$\hat{x}_i \leftarrow \frac{x_i - \mu_{\mathcal{B}}}{\sqrt{\sigma_{\mathcal{B}}^2 + \epsilon}} \quad // \text{ normalize}$$
$$y_i \leftarrow \gamma \hat{x}_i + \beta \equiv \text{BN}_{\gamma, \beta}(x_i) \quad // \text{ scale and shift}$$

Internal Covariate Shift

- Slow Convergence
- Gradient Vanishing




Gradient Vanishing and Exploding



ResNets @ ILSVRC & COCO 2015 Competitions

• 1st places in all five main tracks

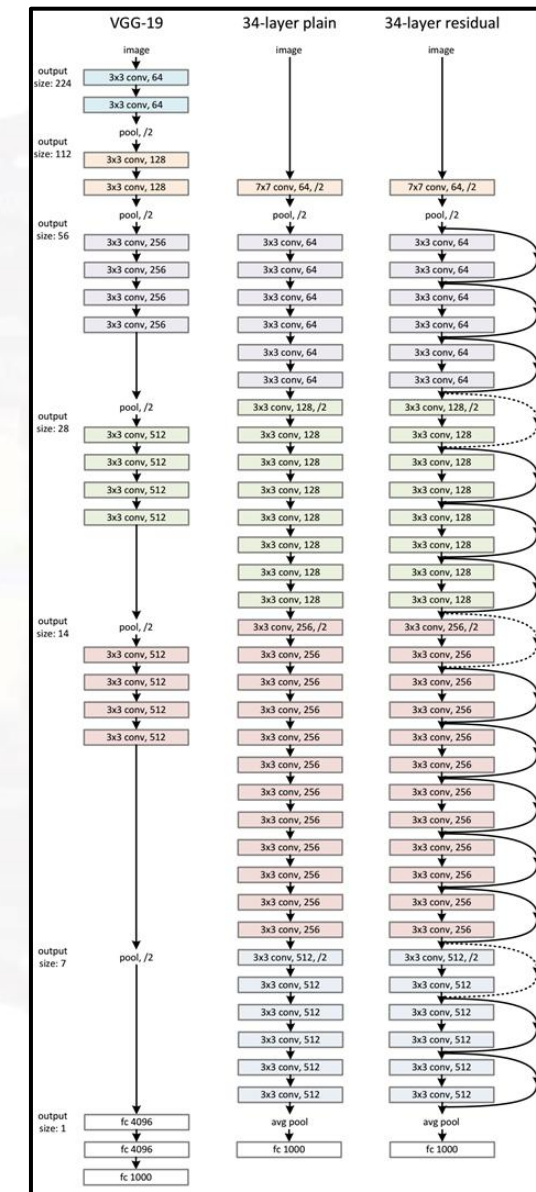


- ImageNet Classification: "Ultra-deep" 152-layer nets
- ImageNet Detection: 16% better than 2nd
- ImageNet Localization: 27% better than 2nd
- COCO Detection: 11% better than 2nd
- COCO Segmentation: 12% better than 2nd



$$a^{(L)} = a^{(L-2)} + H(x)$$

$$\frac{\partial a^{(L)}}{\partial a^{(L-2)}} = \boxed{1} + \frac{\partial H(x)}{\partial a^{(L-2)}}$$

$$\frac{\partial C_0}{\partial w^{(L-2)}} = \frac{\partial z^{(L-2)}}{\partial w^{(L-2)}} \cdot \boxed{\frac{\partial a^{(L-2)}}{\partial z^{(L-2)}} \cdot \frac{\partial z^{(L-1)}}{\partial a^{(L-2)}} \cdot \frac{\partial a^{(L-1)}}{\partial z^{(L-1)}} \cdot \frac{\partial z^{(L)}}{\partial a^{(L-1)}} \cdot \frac{\partial a^{(L)}}{\partial z^{(L)}}} \cdot \frac{\partial C_0}{\partial a^{(L)}}$$



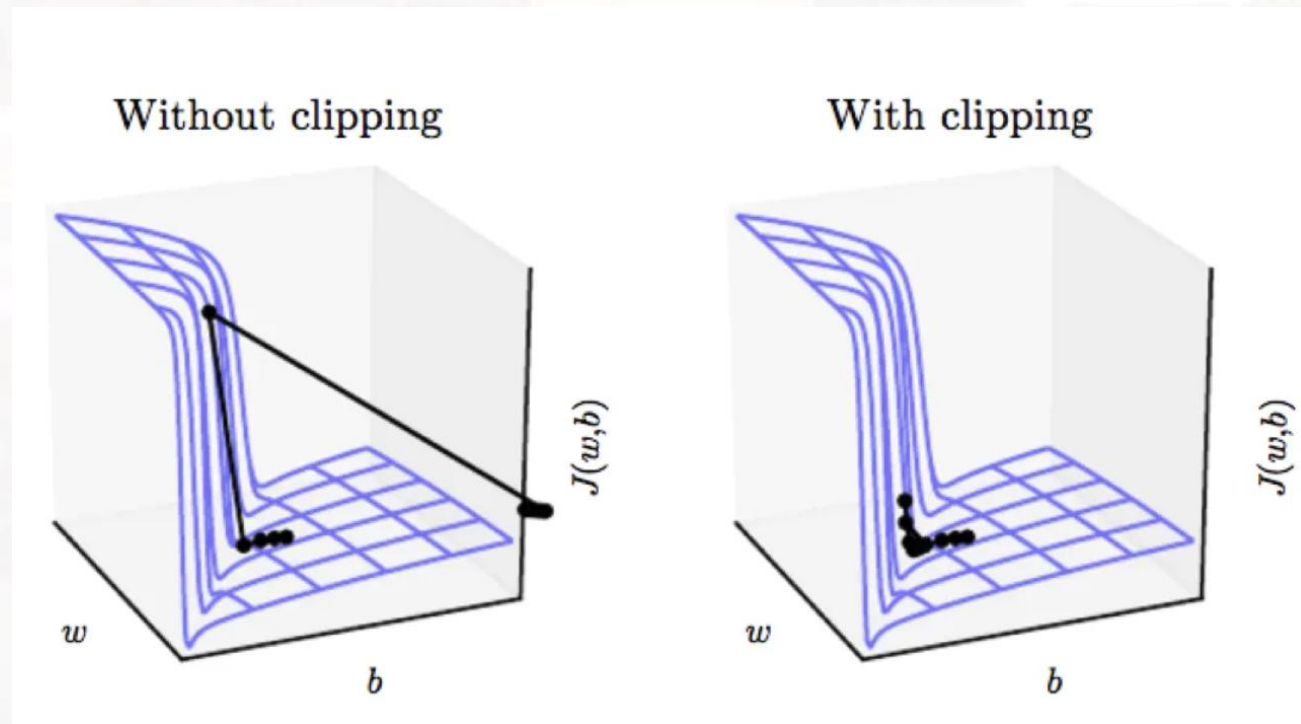

Gradient Vanishing and Exploding



■ ReLU

■ Gradient clipping (gradient exploding) $\nabla C = \eta \frac{\nabla C}{|\nabla C|_2}$

$$\frac{\partial C_0}{\partial w^{(L-1)}} = \frac{\partial z^{(L-1)}}{\partial w^{(L-1)}} \boxed{\frac{\partial a^{(L-1)}}{\partial z^{(L-1)}}} \frac{\partial z^{(L)}}{\partial a^{(L-1)}} \frac{\partial a^{(L)}}{\partial z^{(L)}} \frac{\partial C_0}{\partial a^{(L)}}$$



Gradient Vanishing and Exploding



■ Batch gradient descent



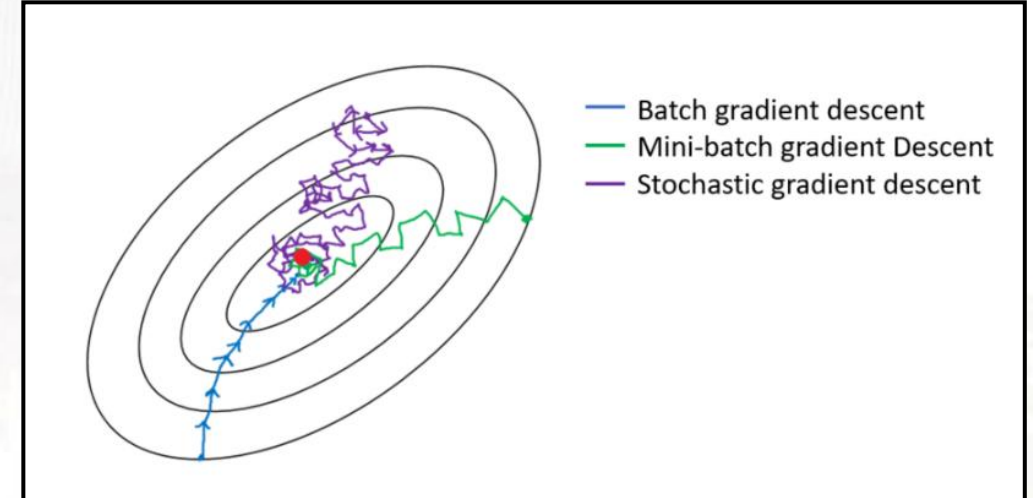
$$W = W - \eta \cdot \nabla_W C(W)$$

■ Stochastic gradient descent

$$W = W - \eta \cdot \nabla_W C(W; x^{(i)}; y^{(i)})$$

■ Mini-batch gradient descent

$$W = W - \eta \cdot \nabla_W C(W; x^{(i:i+n)}; y^{(i:i+n)})$$



■ Momentum

- SGD has the problem of slow convergence in 'valley'-shaped spaces
- SGDM speeds up SGD by modifying the direction.

- Add an update direction of a past moment

$$W \leftarrow W - \eta \frac{\partial L}{\partial W}$$



$$V_t \leftarrow \beta V_{t-1} - \eta \frac{\partial L}{\partial W}$$

$$W \leftarrow W + V_t$$



Image 2: SGD without momentum




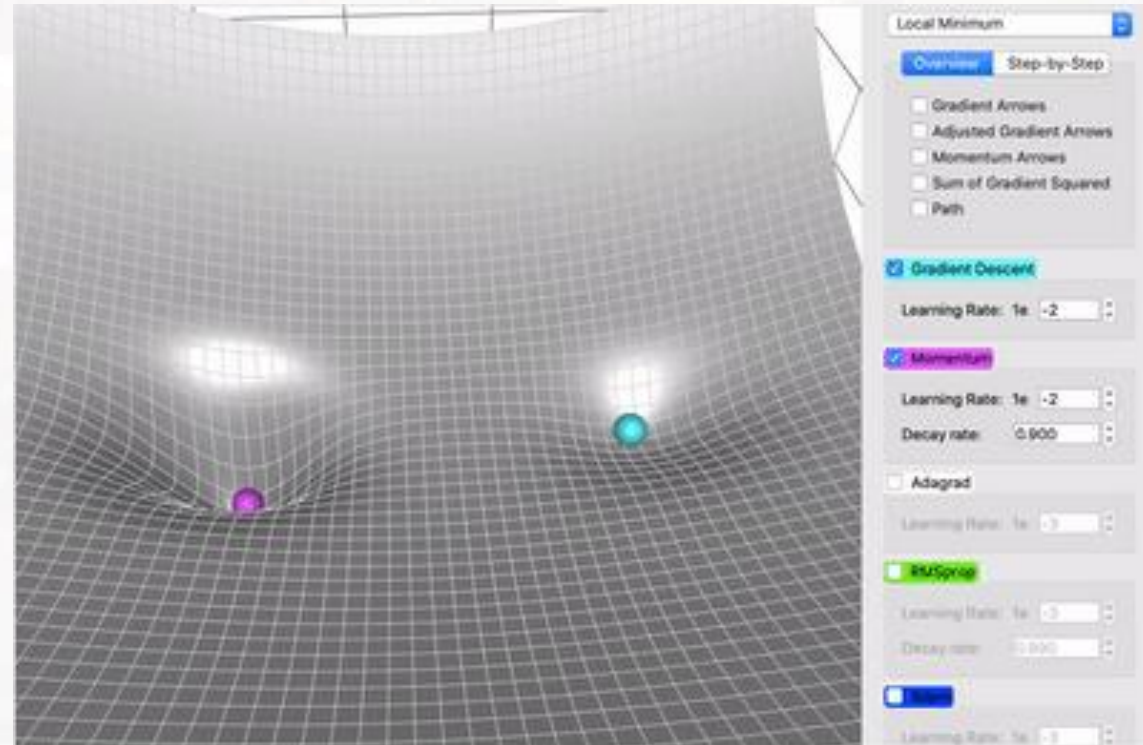

Image 3: SGD with momentum

■ Advantage

- faster convergence
- chance of escaping the local minimum

$$\theta = \theta - \eta \cdot \nabla_{\theta} J(\theta; x^{(i:i+n)}; y^{(i:i+n)}).$$


$$v_t = \gamma v_{t-1} + \eta \nabla_{\theta} J(\theta)$$
$$\theta = \theta - v_t$$




AdaGrad



$$\theta_{t+1,i} = \theta_{t,i} - \frac{\eta}{\sqrt{G_{t,i} + \epsilon}} \cdot g_{t,i}$$

$$G_t^{(i,i)} = \sum_{\tau}^t (g_{\tau}^{(i)})^2$$

$$g_{t,i} = \nabla_{\theta} J(\theta_{t,i})$$



RMSProp




$$\theta_{t+1,i} = \theta_{t,i} - \frac{\eta}{\sqrt{G_{t,i} + \epsilon}} \cdot g_{t,i}$$

$$G_t^{(i,i)} = \sum_{\tau}^t (g_{\tau}^{(i)})^2$$

$$g_{t,i} = \nabla_{\theta} J(\theta_{t,i})$$



$$E[g^2]_t = 0.9E[g^2]_{t-1} + 0.1g_t^2$$
$$\theta_{t+1} = \theta_t - \frac{\eta}{\sqrt{E[g^2]_t + \epsilon}} g_t$$



Adam(Adaptive Moment Estimation)



■ Combine SGDM and RMSProp

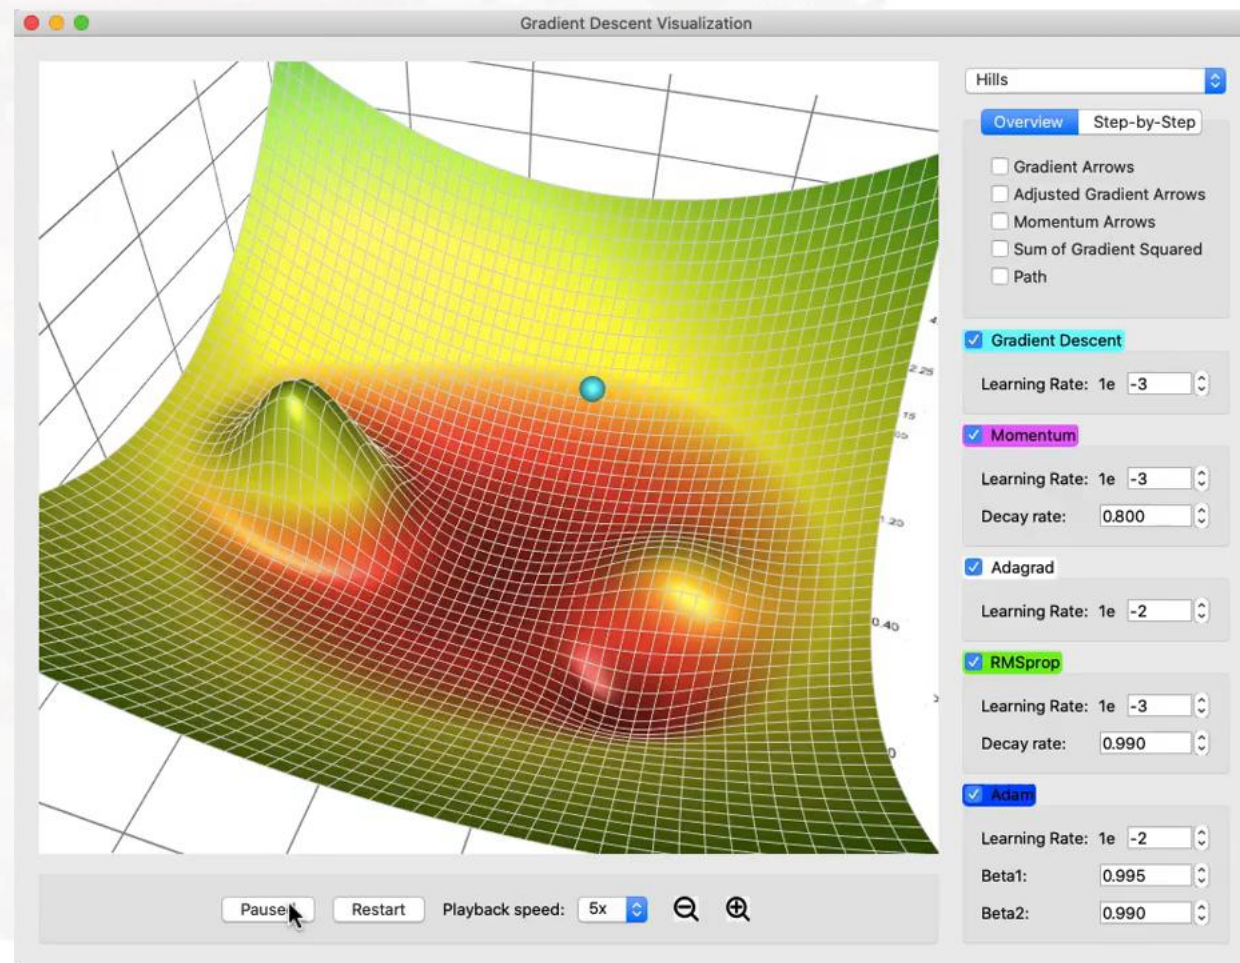
$$\theta_{t+1} = \theta_t - \frac{\eta}{\sqrt{\hat{v}_t} + \epsilon} \hat{m}_t.$$

$$m_t = \beta_1 m_{t-1} + (1 - \beta_1) g_t$$

$$v_t = \beta_2 v_{t-1} + (1 - \beta_2) g_t^2$$

$$v_t = \gamma v_{t-1} + \eta \nabla_{\theta} J(\theta)$$

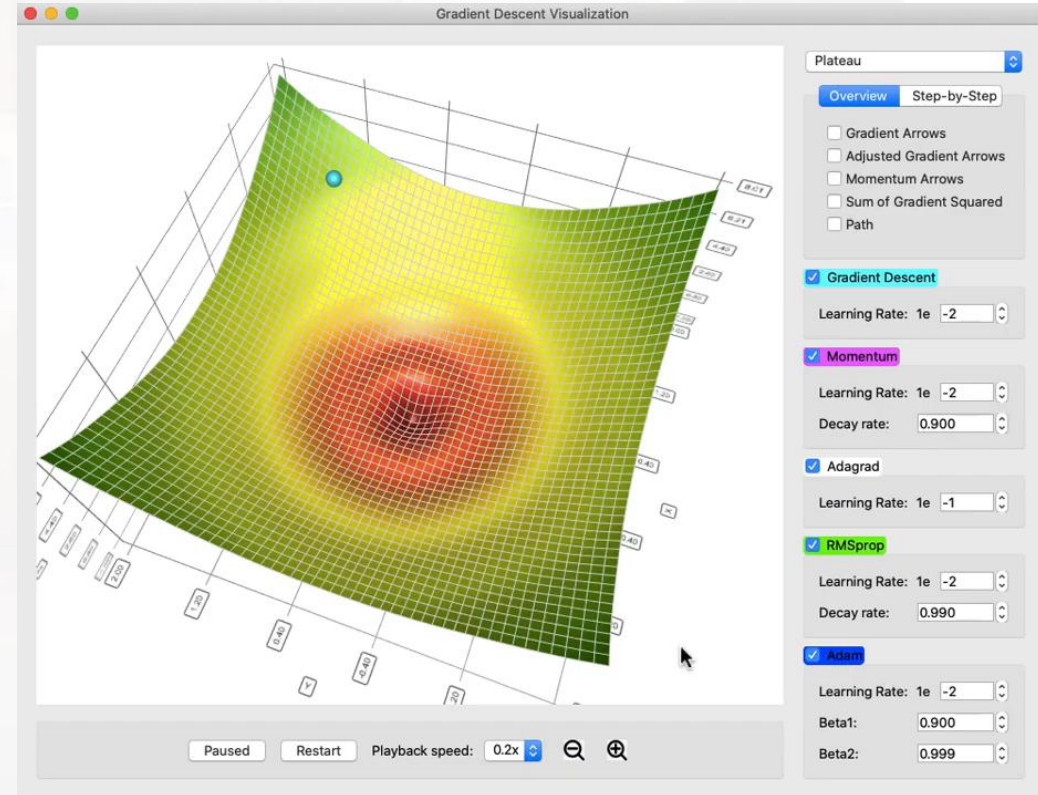

$$E[g^2]_t = 0.9 E[g^2]_{t-1} + 0.1 g_t^2$$



The Advantage of Adam



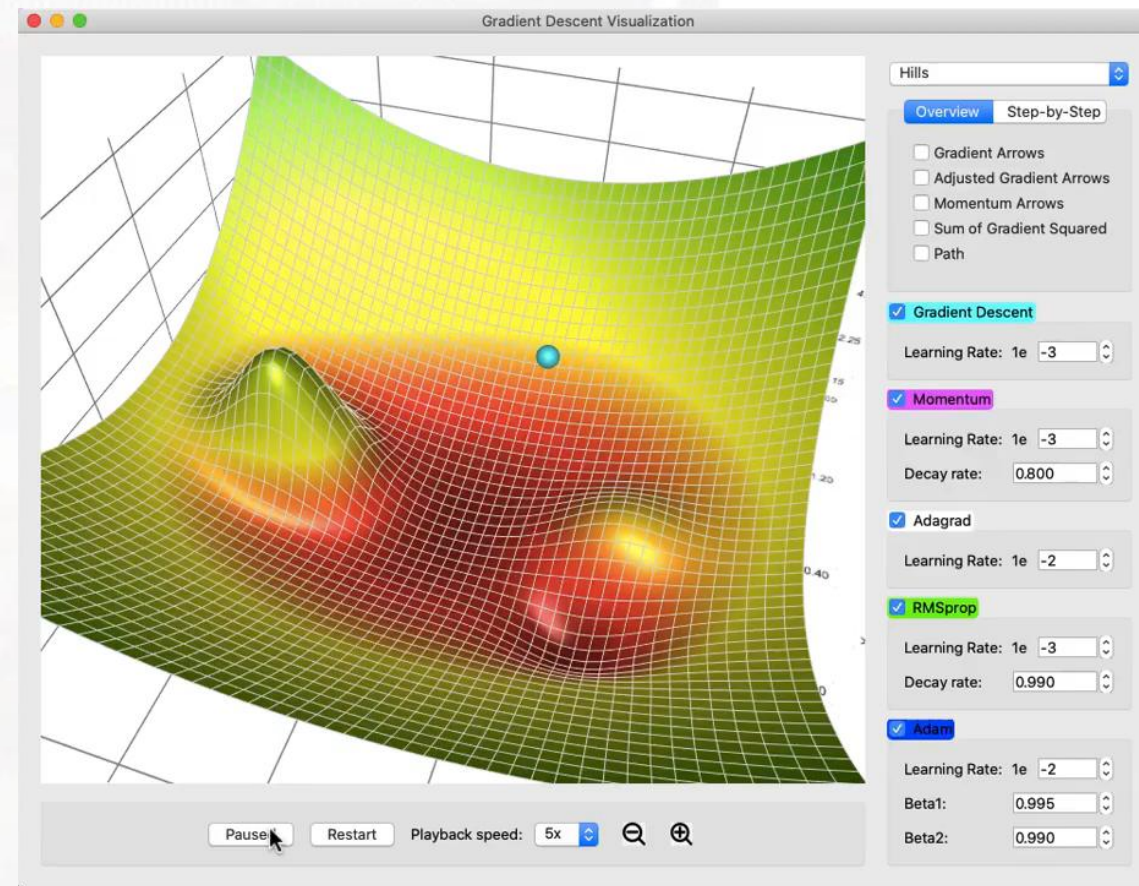


- The advantage of Adaptive
- The advantage of Momentum



The Problem of Adam



Overfitting



Summary





■ Adaptive vs Non Adaptive

Adaptive Method	Non-Adaptive Method
Adam, AdaGrad, RMSProp	SGDM, SGD
Difficult data, complex networks, hard to converge	Good initialization and learning rate scheduling scheme

	SGDM	Adam
Training Speed	Slow	Fast
Convergence	Good	Poor
Stability	Good	Poor
Generalization	Good	Poor

If you are interested in visualizing these or other optimization algorithms, refer to [this useful tutorial](#).






浙江大學
ZHEJIANG UNIVERSITY







Fit Everything?

- We can have good approximation with sufficient pieces.



The universal approximation brought by nonlinearity

Neural Network



$$y = b + wx_1$$

$$y = b + \sum_i c_i \text{sigmoid}(b_i + w_i x_1)$$

$$y = b + \sum_j w_j x_j$$

$$y = b + \sum_i c_i \text{sigmoid}\left(b_i + \sum_j w_{ij} x_j\right)$$

How to represent
this function?


Hard Sigmoid

Sigmoid Function


$$y = c \frac{1}{1 + e^{-(b+wx_1)}}$$

$$= c \text{sigmoid}(b + wx_1)$$


How to
represent this
function?




Rectified
Linear Unit
(ReLU)



piecewise
linear

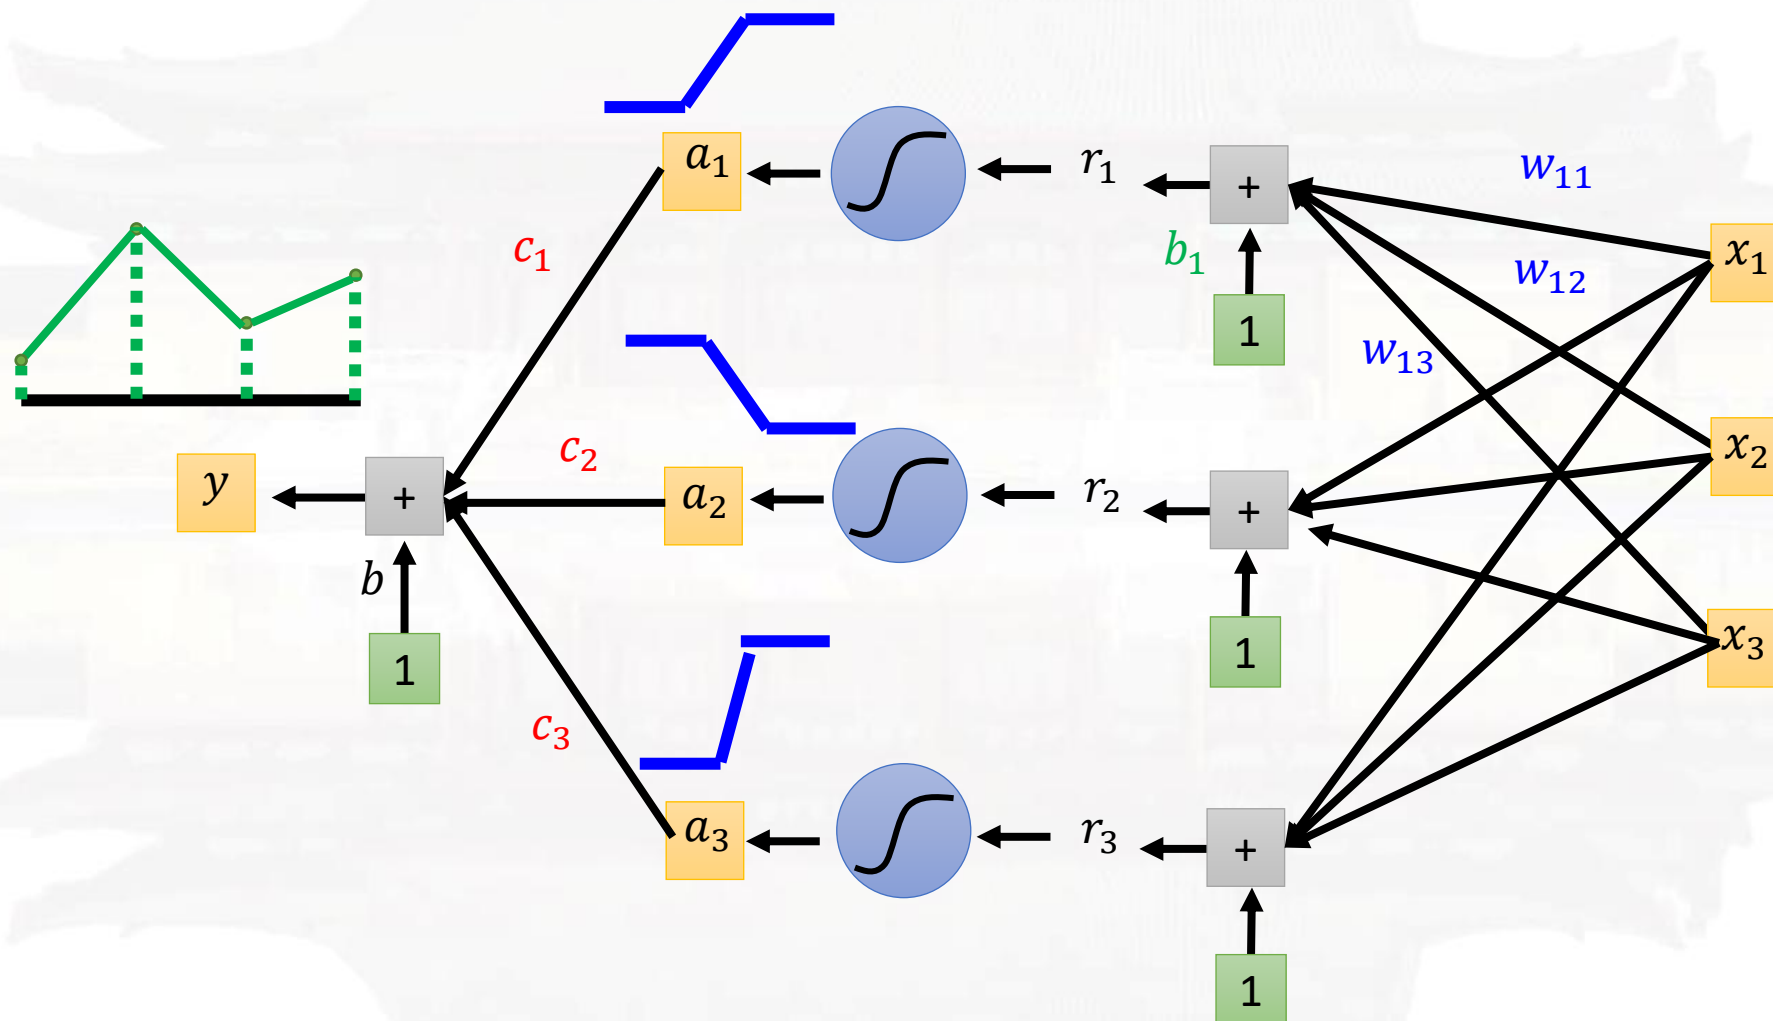


= constant +
sum of a set of




The universal approximation brought by nonlinearity


Neural Network



■ Why Use Layers?





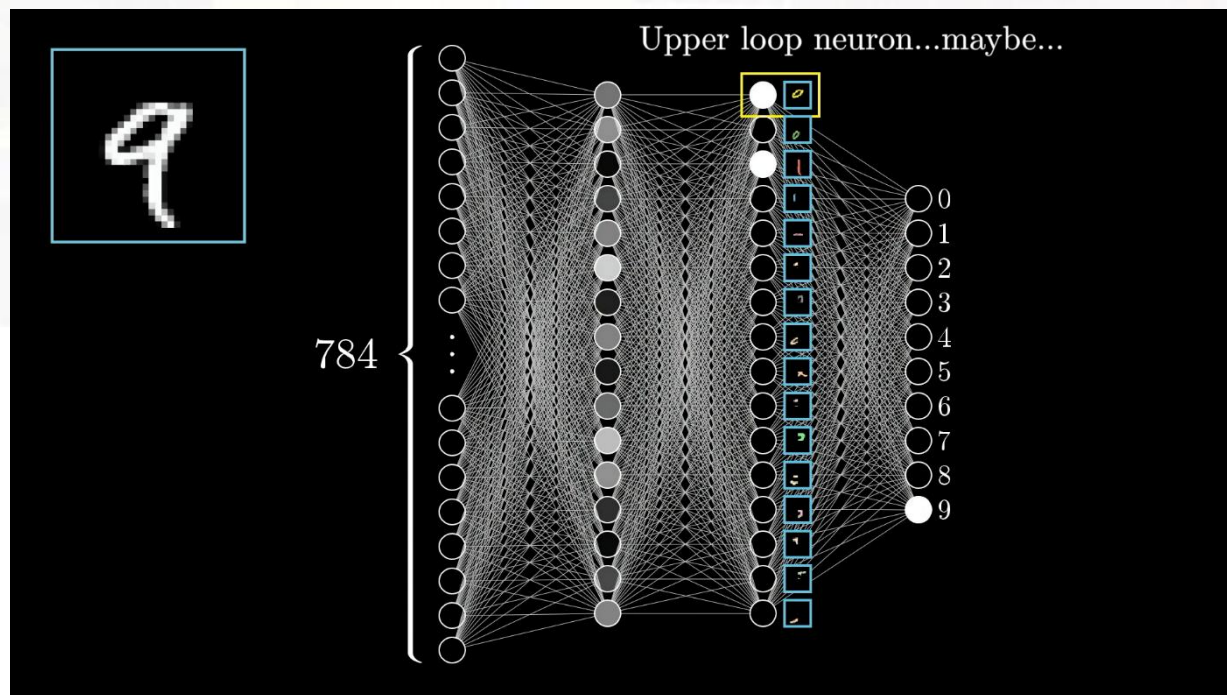
- we piece together various components like loops and lines



Neural Network



- In a perfect world, we might hope that each neuron in the second-to-last layer corresponds to one of these subcomponents.



■ Layers Break Problems Into Bite-Sized Pieces

- Edge detection is a useful step for all kinds of image-recognition problems.
- beyond image recognition





→ r e c o g n i t i o n → re·cog·ni·tion → recognition

Raw audio

Neural Network






- Rank the four images (A, B, C, and D) based on how much they would activate that neuron:



■ We hope ... but

Does the network
actually do this?



Pixels



Edges



Shapes




Digits

Overfitting and Underfitting



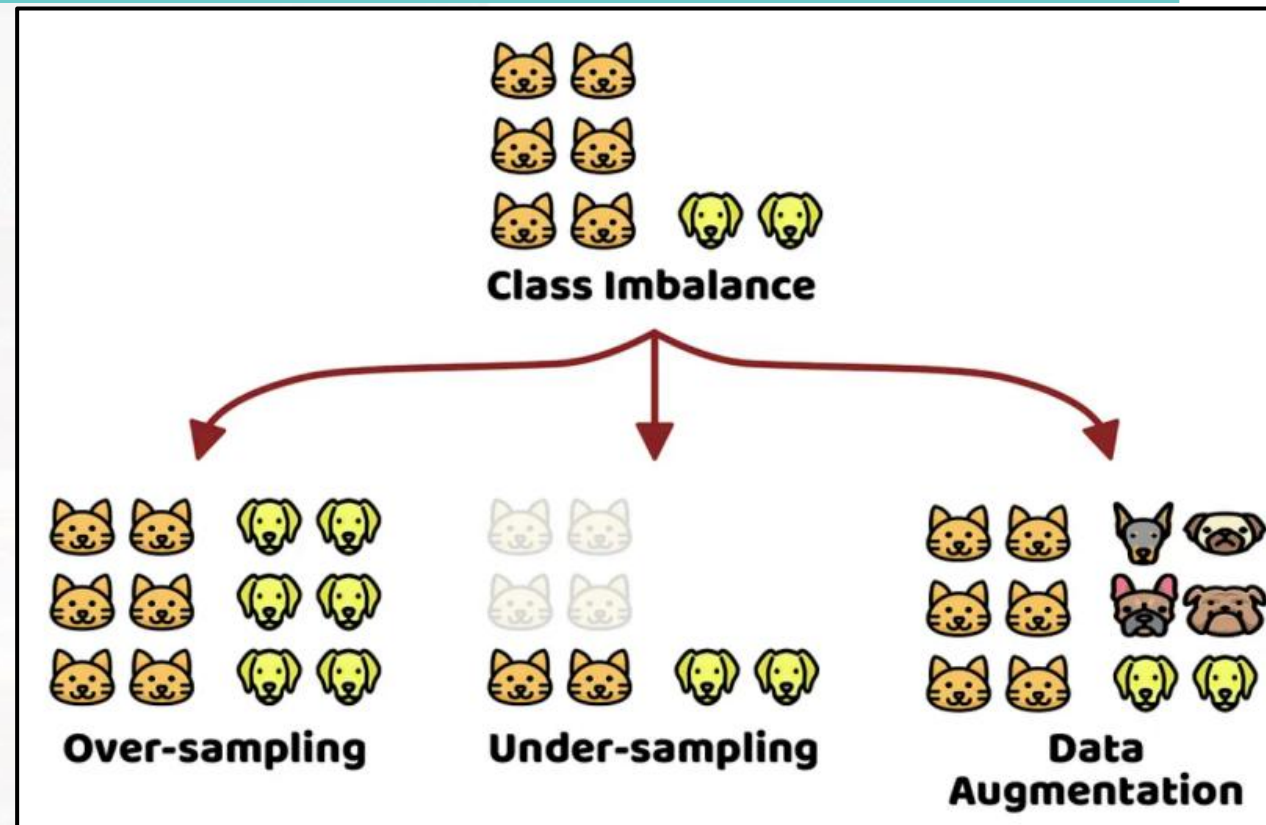
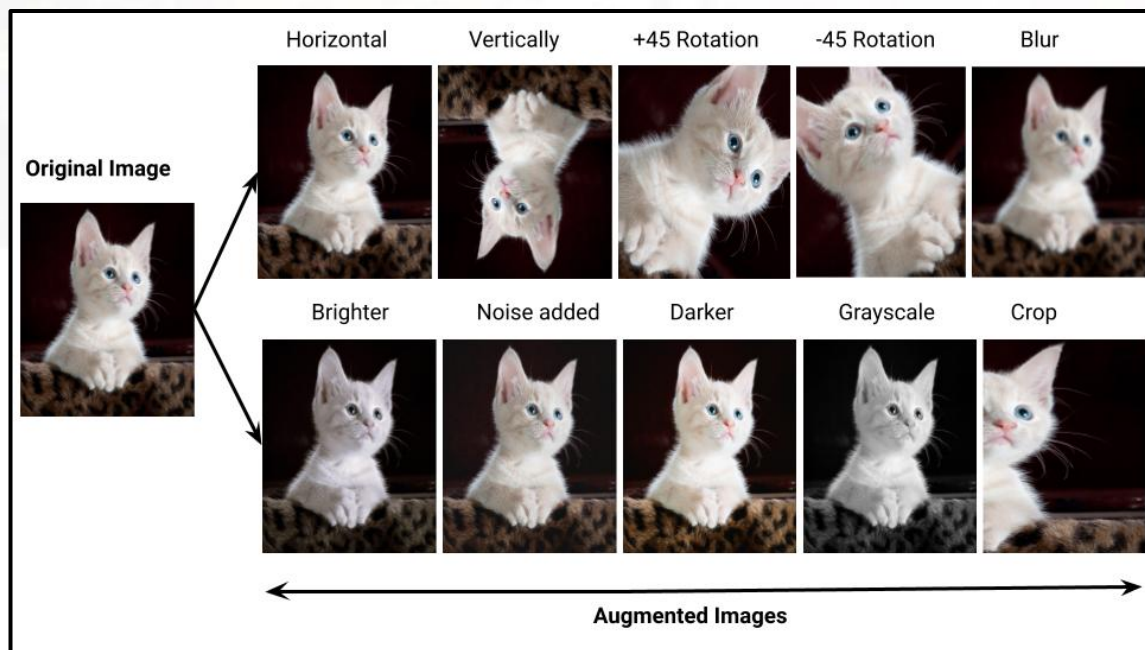
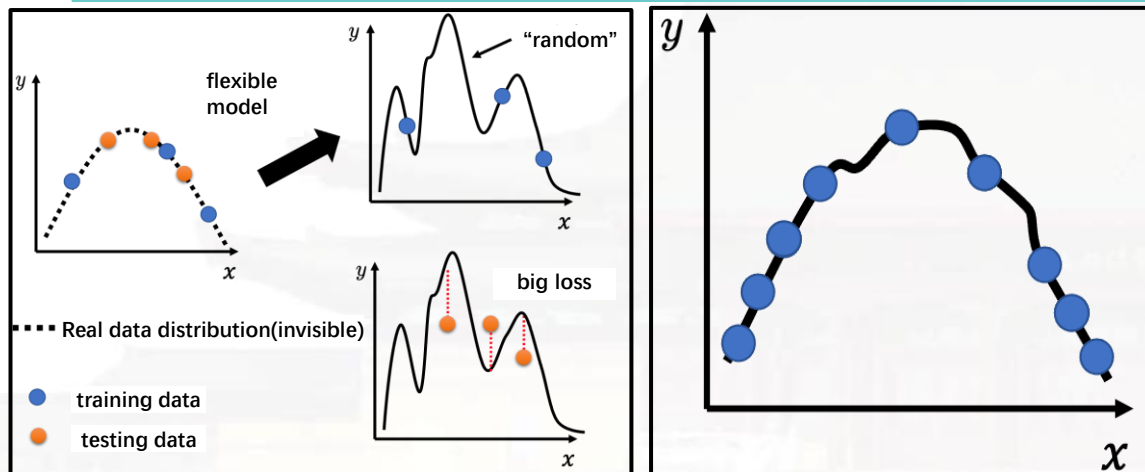
- The balance between data (knowns or constraints) and parameters (unknowns)



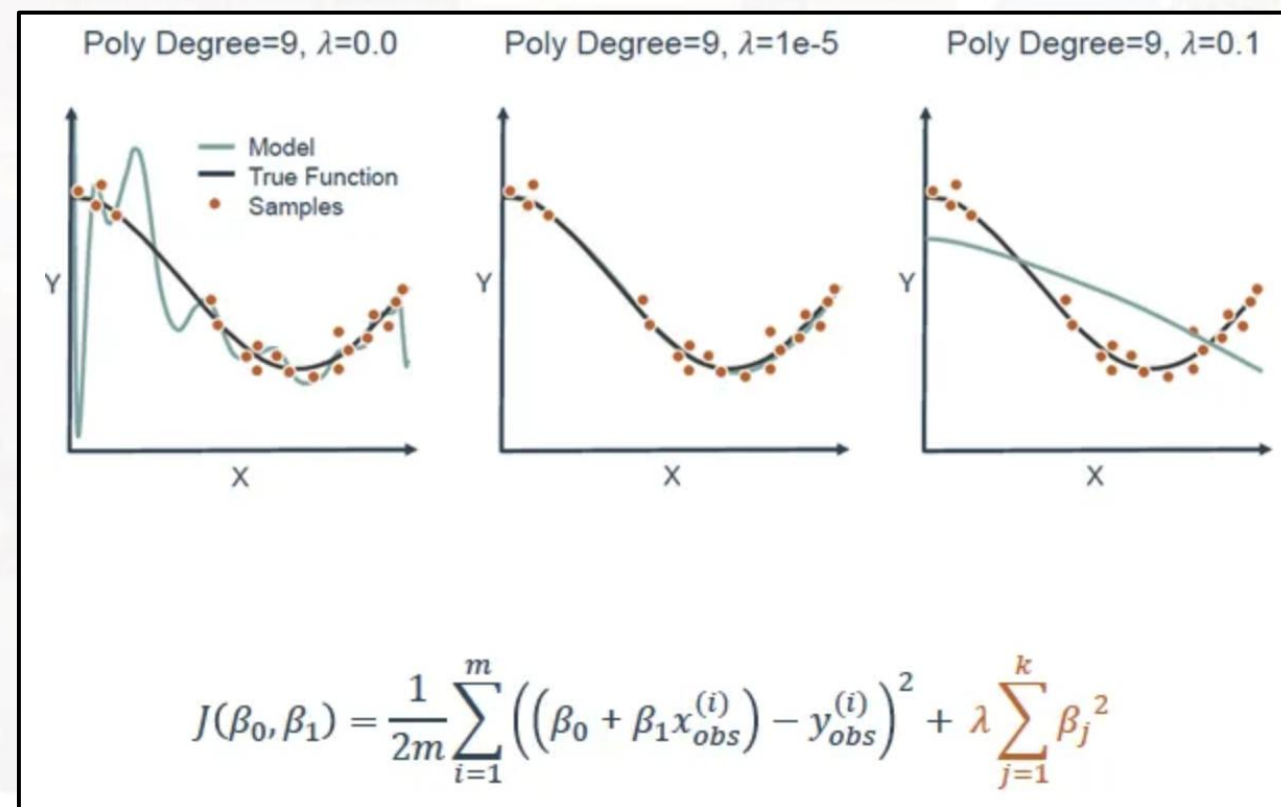






Overfitting: Data Augmentation



浙江大學
ZHEJIANG UNIVERSITY




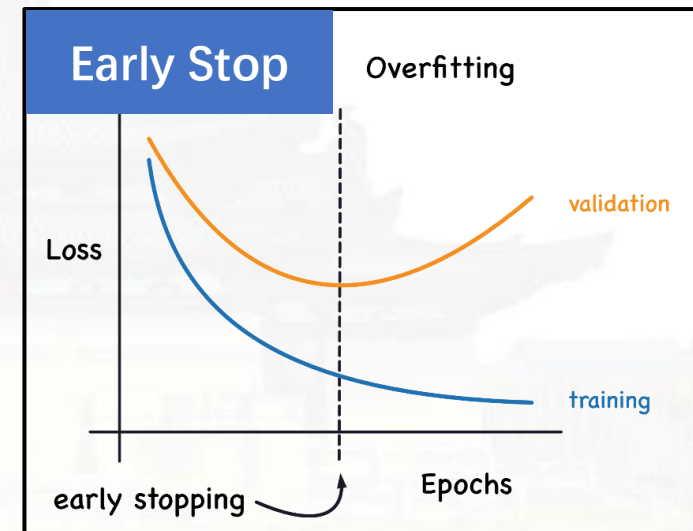

Overfitting: Regularization



Overfitting



浙江大學
ZHEJIANG UNIVERSITY



■ Solutions to overfitting

- Data augmentation
- Model constraints: regularization, architecture, dropout
- Early stop



浙江大學
ZHEJIANG UNIVERSITY



Architectures

The Problem of MLP



■ Numerous Parameters

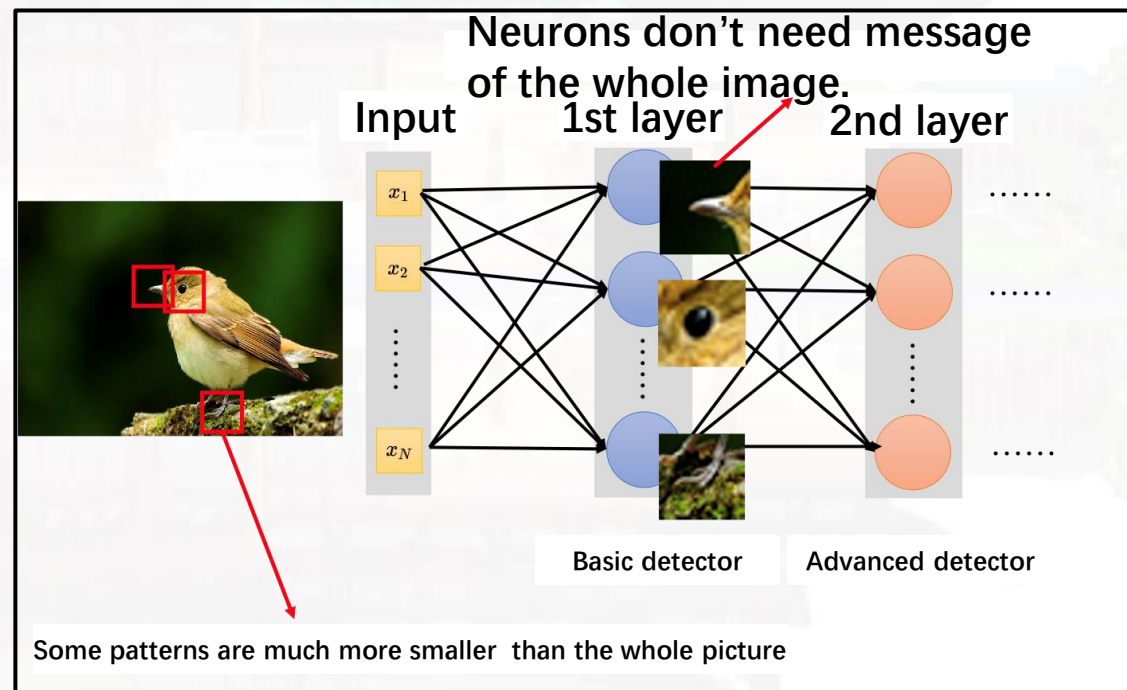

- Poor Explanation due to Large Scale Minimum
- Large Scale Minimum due to Large Scale Parameters



The Problem of MLP



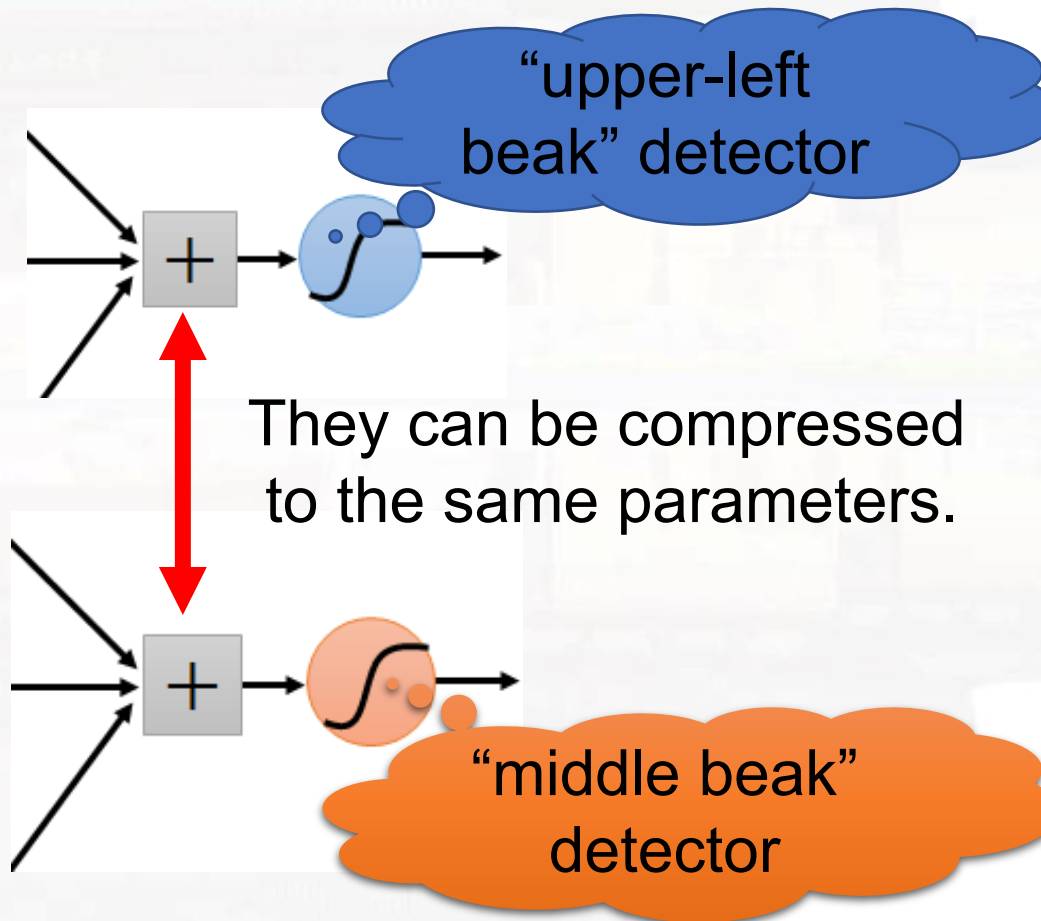
■ Poor Flexibility



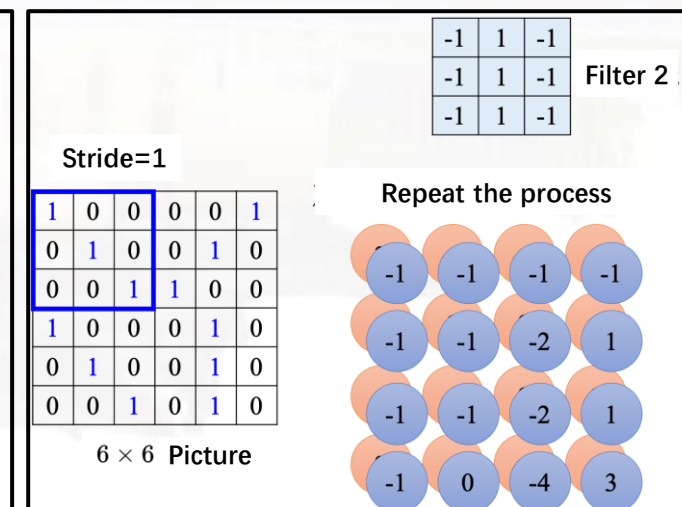


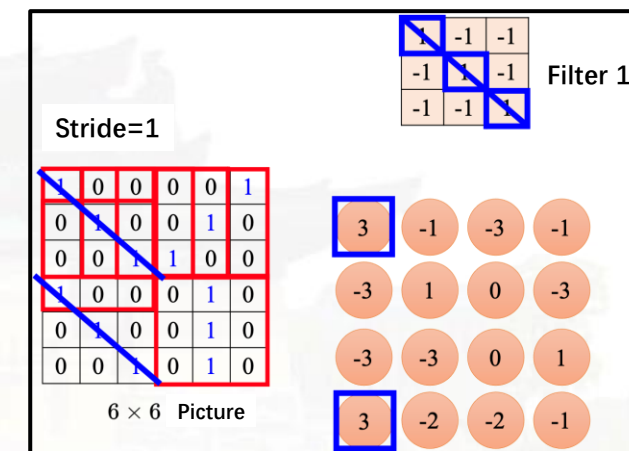


The Problem of MLP





■ Poor Flexibility



From MLP to CNN



Pooling and Stride



stride=1

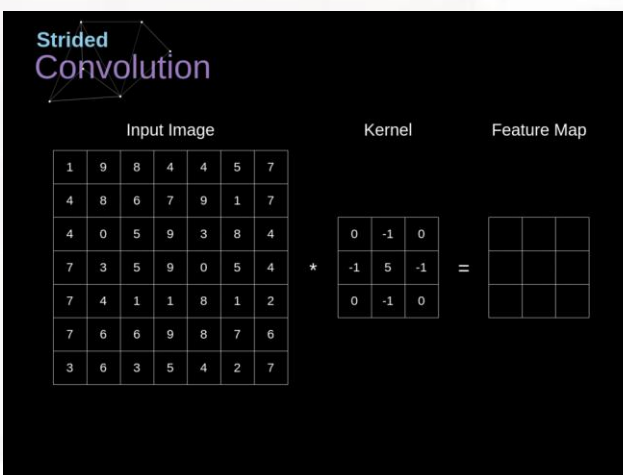


1	0	0	0	0	1
0	1	0	0	1	0
0	0	1	1	0	0
1	0	0	0	1	0
0	1	0	0	1	0
0	0	1	0	1	0

Dot product



If stride=2

1	0	0	0	0	1
0	1	0	0	1	0
0	0	1	1	0	0
1	0	0	0	1	0
0	1	0	0	1	0
0	0	1	0	1	0





6 x 6 image

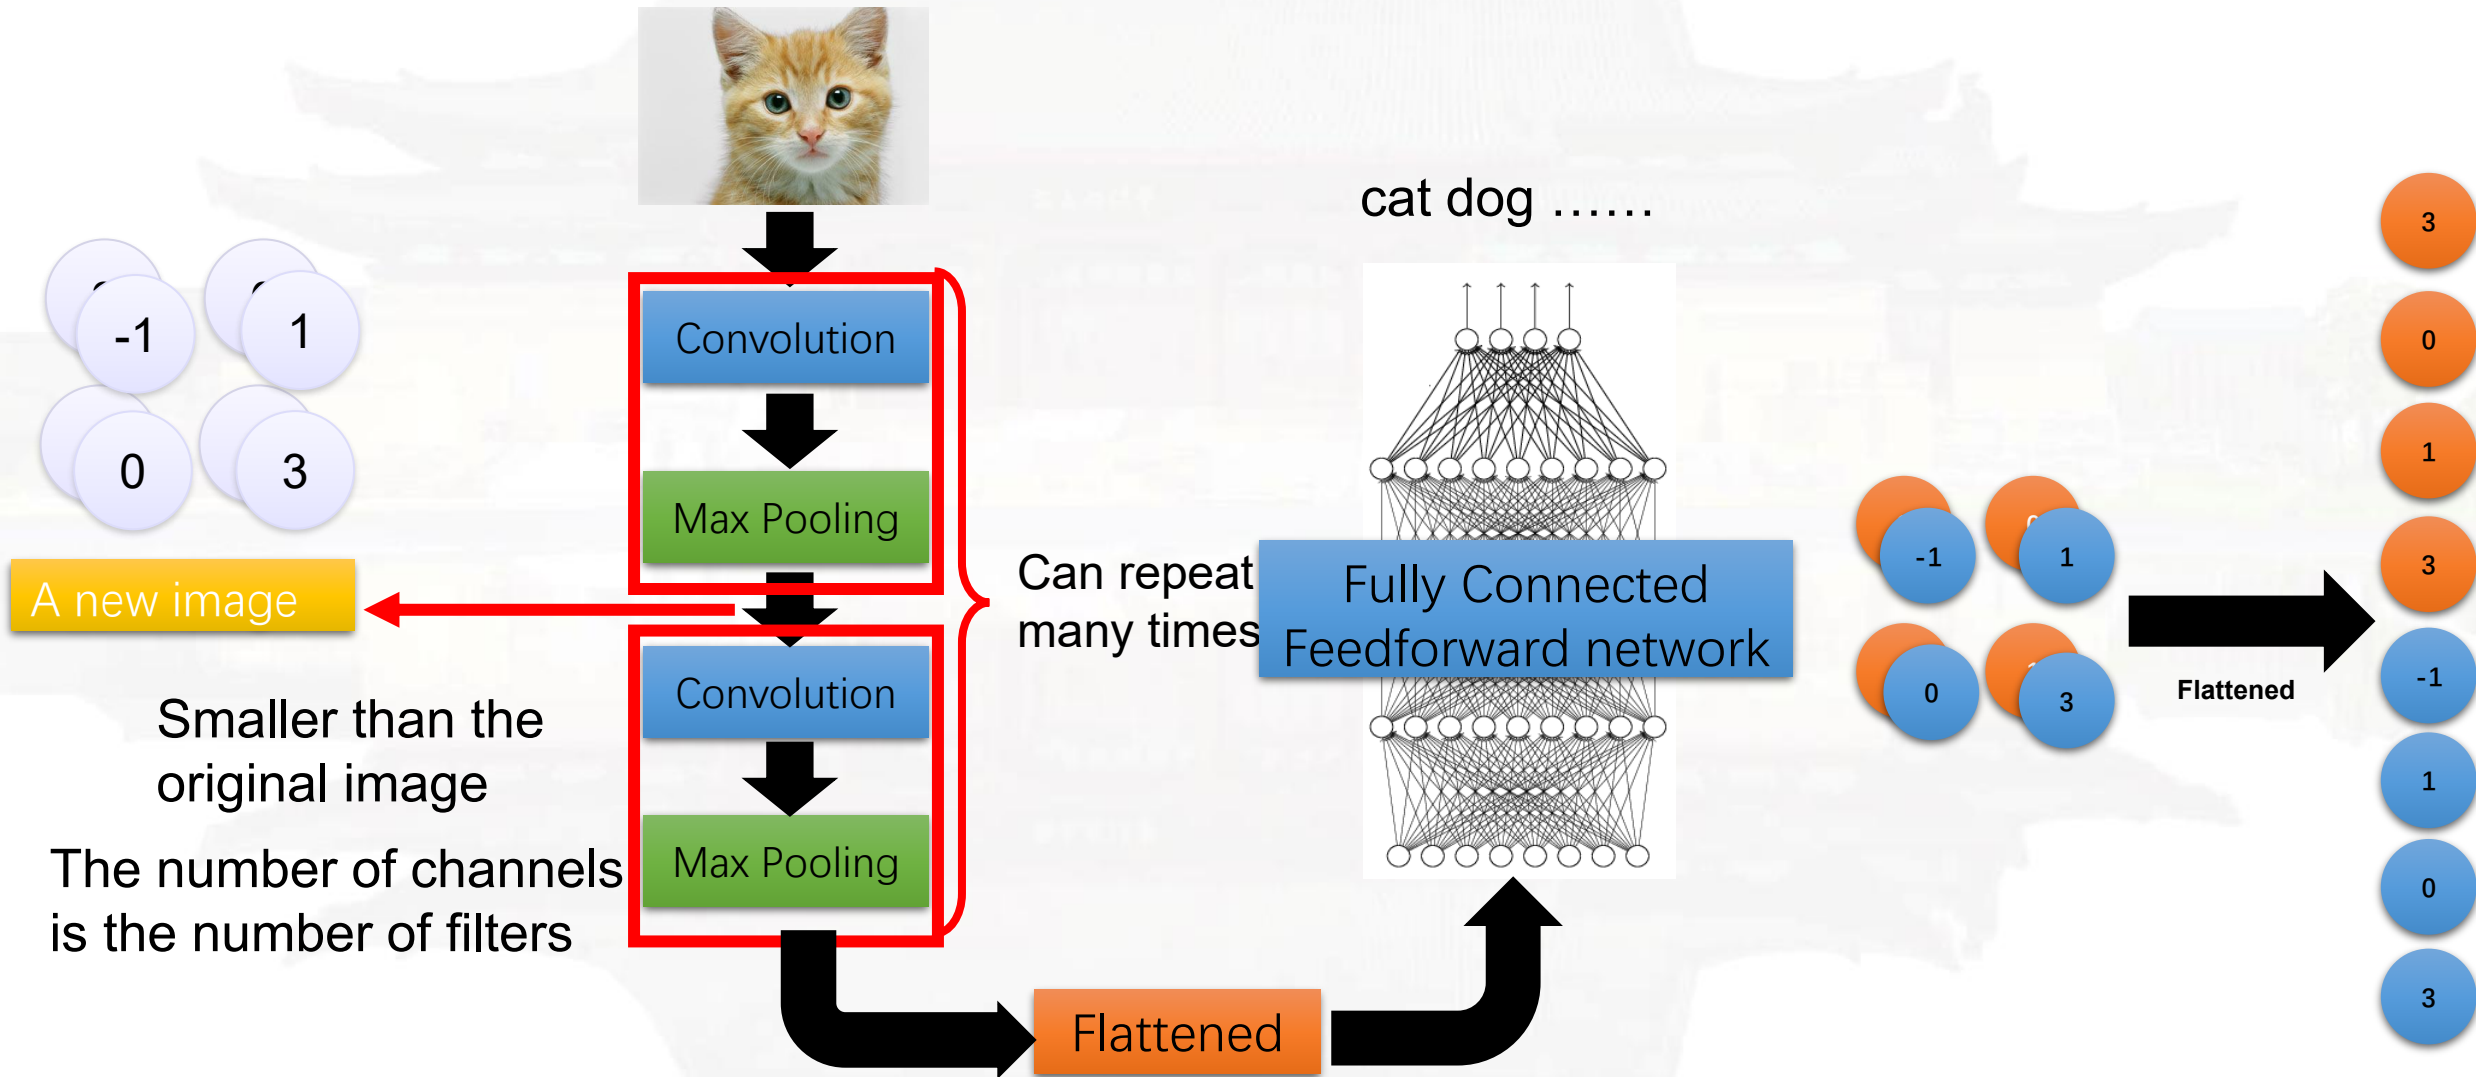
Increase Receptive Field



Receptive Field in Convolutional Networks




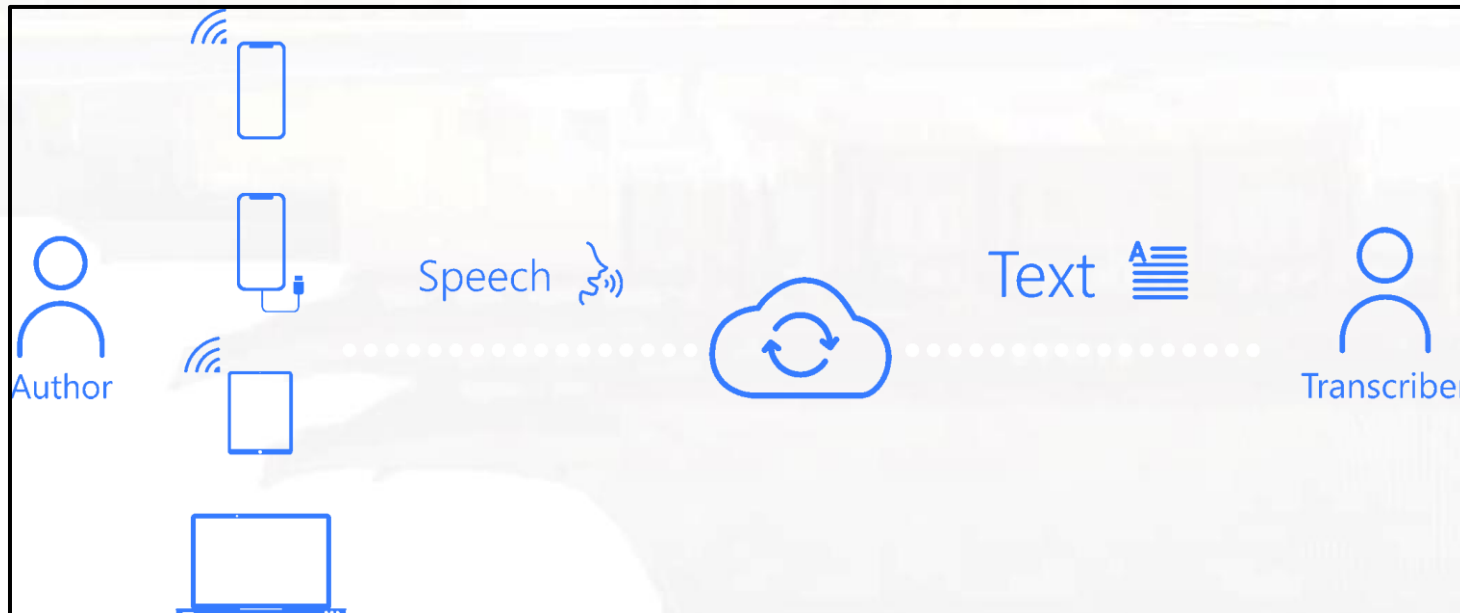
Convolutional Neural Networks



Feedforward Neural Networks



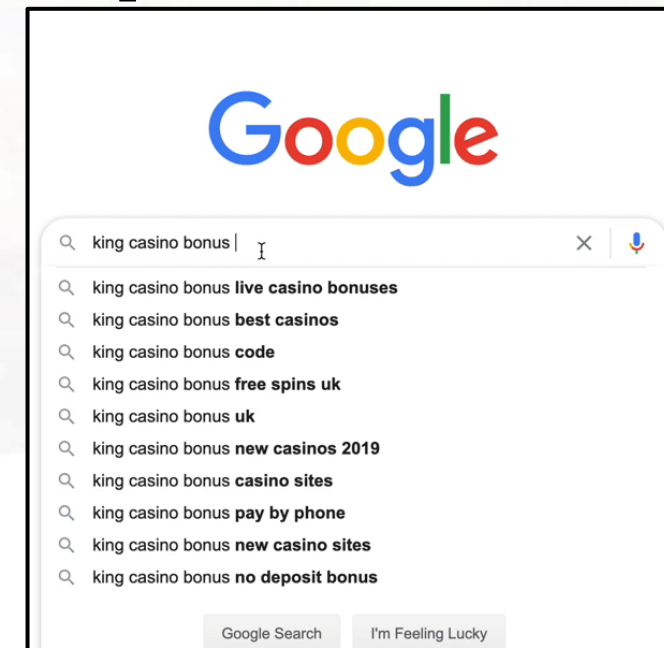

- A feedforward neural network (FNN) is one of the two broad types of artificial neural network, characterized by direction of the flow of information between its layers
- Human brain deals with information streams. Most data is obtained, processed, and generated sequentially.



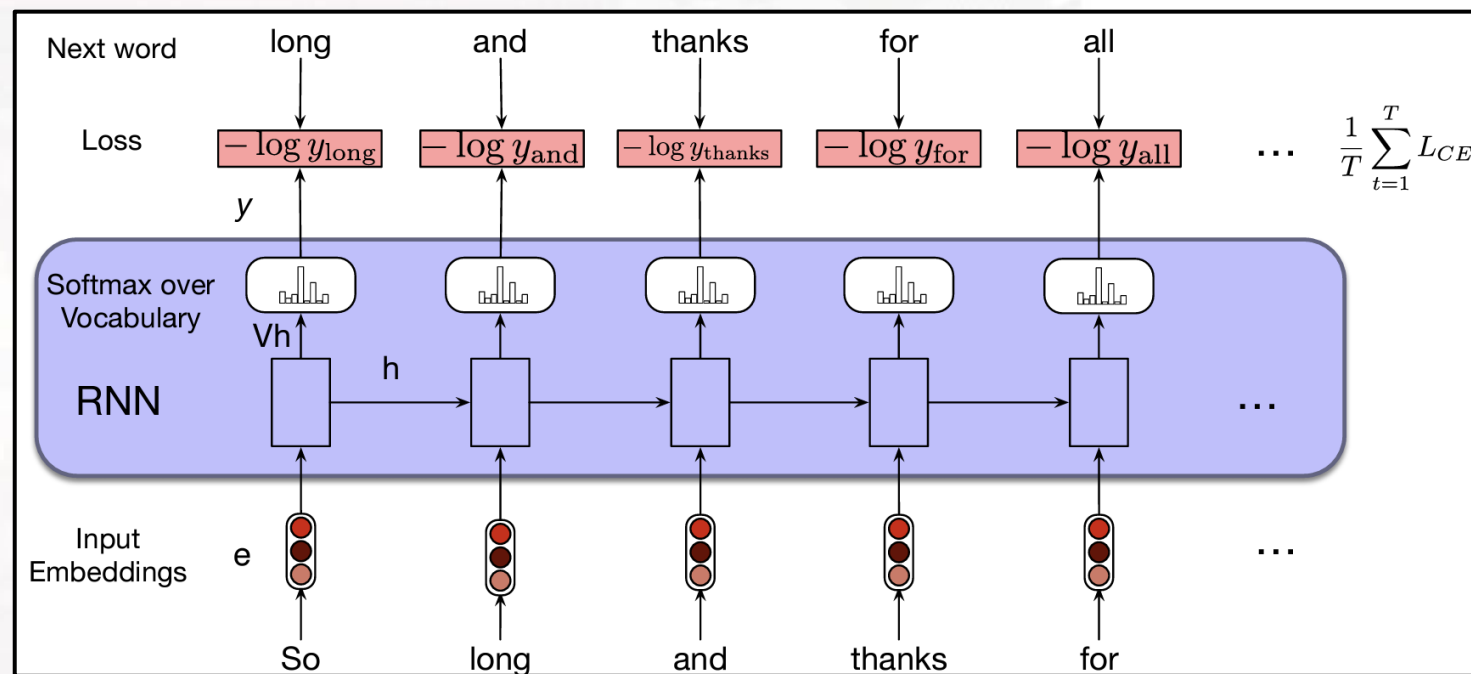

The Problem of FNNs



- Human thoughts have persistence; humans don't start their thinking from scratch every second.
- The applications of standard FNNs are limited due to:
 - They only accepted a fixed-size vector as input (e.g., an image) and produce a fixed-size vector as output
 - These models use a fixed amount of computational steps



Recurrent Neural Networks


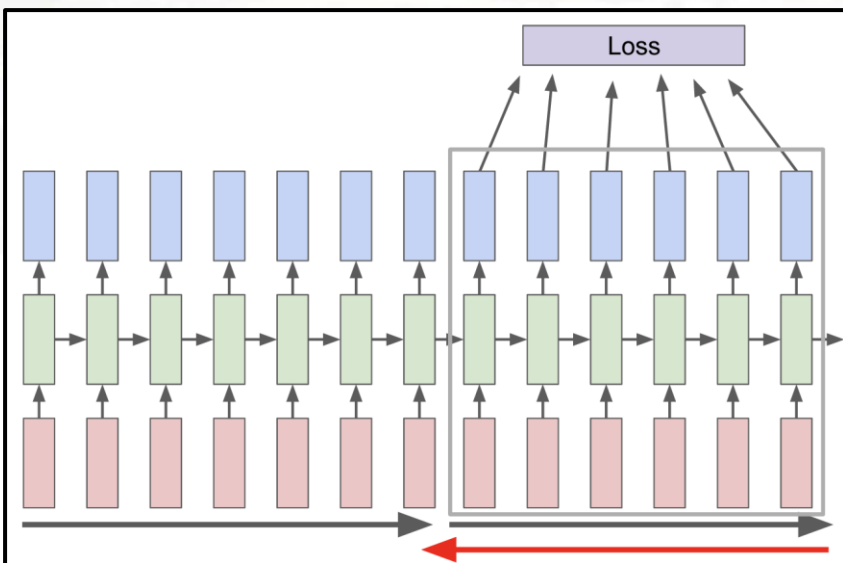
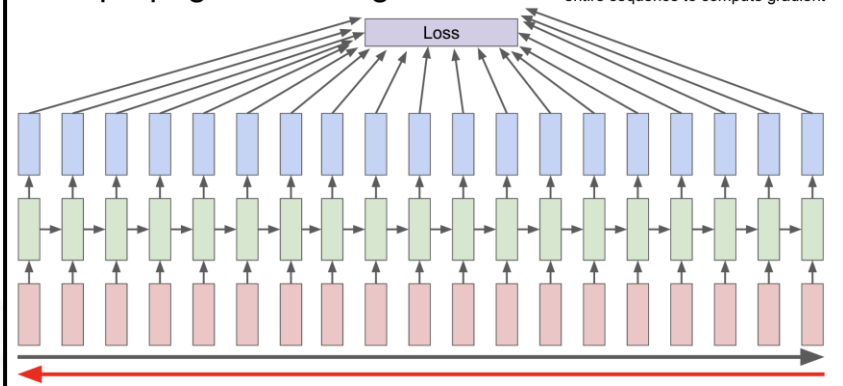


Training RNN




Backpropagation through time

Forward through entire sequence to compute loss, then backward through entire sequence to compute gradient



Gradient Vanishing and Exploding



$$w = 1 \Rightarrow y^{1000} = 1$$

$$w = 1.01 \Rightarrow y^{1000} \approx 20000$$

$$w = 0.99 \Rightarrow y^{1000} \approx 0$$



$$w = 0.01 \Rightarrow y^{1000} \approx 0$$

Large
 $\partial L / \partial w$

Small Learning
rate?

small
 $\partial L / \partial w$

Large Learning
rate?




Gradient Vanishing




Vanilla RNN Gradient Flow

Bengio et al, "Learning long-term dependencies with gradient descent is difficult", IEEE Transactions on Neural Networks, 1994
Pascanu et al, "On the difficulty of training recurrent neural networks", ICML 2013



$$\begin{aligned}h_t &= \tanh(W_{hh}h_{t-1} + W_{xh}x_t) \\&= \tanh\left((W_{hh} \quad W_{hx}) \begin{pmatrix} h_{t-1} \\ x_t \end{pmatrix}\right) \\&= \tanh\left(W \begin{pmatrix} h_{t-1} \\ x_t \end{pmatrix}\right)\end{aligned}$$

$$\frac{\partial h_t}{\partial h_{t-1}} = \tanh'(W_{hh}h_{t-1} + W_{xh}x_t)W_{hh}$$



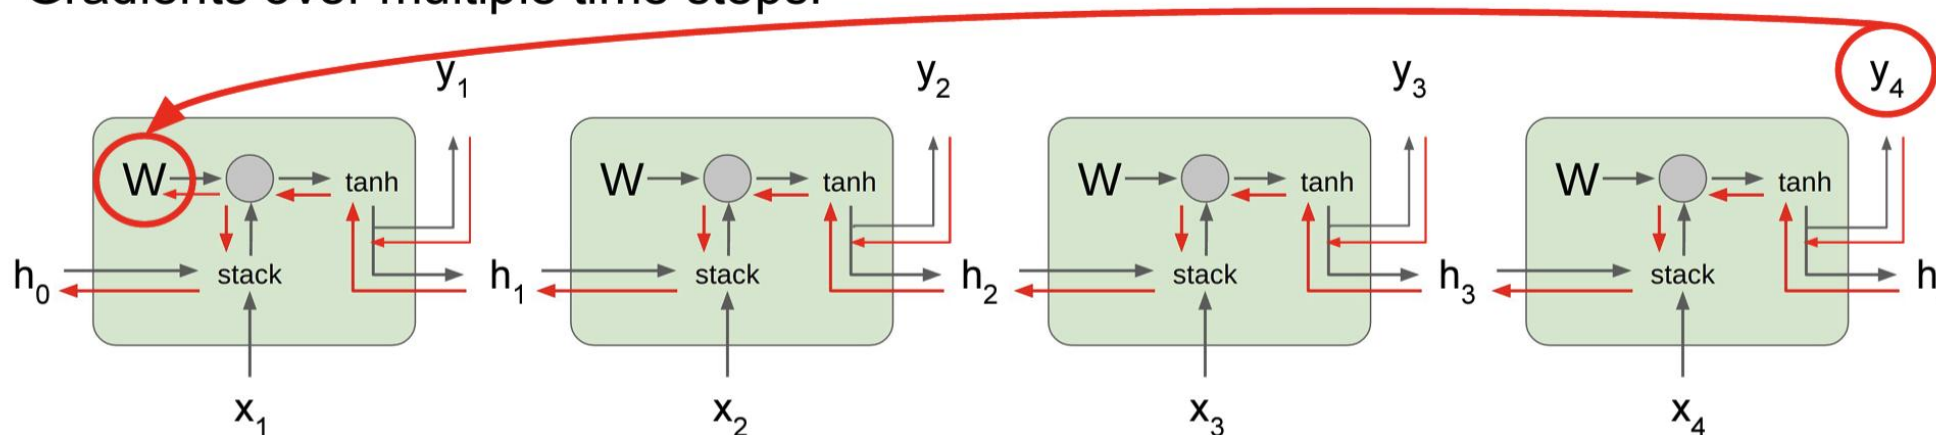
Gradient Vanishing



Vanilla RNN Gradient Flow

Gradients over multiple time steps:

Bengio et al, "Learning long-term dependencies with gradient descent is difficult", IEEE Transactions on Neural Networks, 1994
Pascanu et al, "On the difficulty of training recurrent neural networks", ICML 2013



$$\frac{\partial L}{\partial W} = \sum_{t=1}^T \frac{\partial L_t}{\partial W}$$

$$\frac{\partial h_t}{\partial h_{t-1}} = \tanh'(W_{hh} h_{t-1} + W_{xh} x_t) W_{hh}$$

$$\frac{\partial L_T}{\partial W} = \frac{\partial L_T}{\partial h_T} \frac{\partial h_t}{\partial h_{t-1}} \cdots \frac{\partial h_1}{\partial W} = \frac{\partial L_T}{\partial h_T} \left(\prod_{t=2}^T \frac{\partial h_t}{\partial h_{t-1}} \right) \frac{\partial h_1}{\partial W}$$

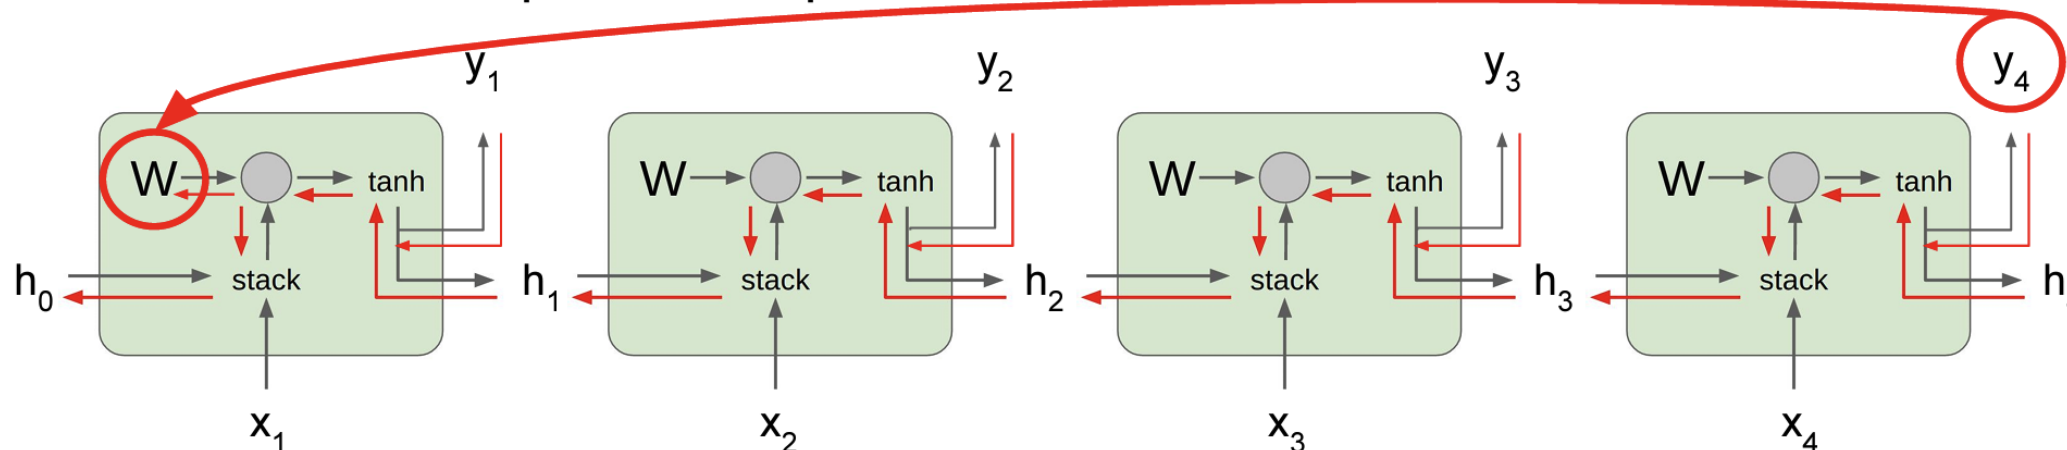
Gradient Vanishing



Vanilla RNN Gradient Flow

Gradients over multiple time steps:

Bengio et al, "Learning long-term dependencies with gradient descent is difficult", IEEE Transactions on Neural Networks, 1994
Pascanu et al, "On the difficulty of training recurrent neural networks", ICML 2013



What if we assumed no non-linearity?

$$\frac{\partial L}{\partial W} = \sum_{t=1}^T \frac{\partial L_t}{\partial W}$$

$$\frac{\partial L_T}{\partial W} = \frac{\partial L_T}{\partial h_T} \boxed{W^{T-1}_{hh}} \frac{\partial h_1}{\partial W}$$



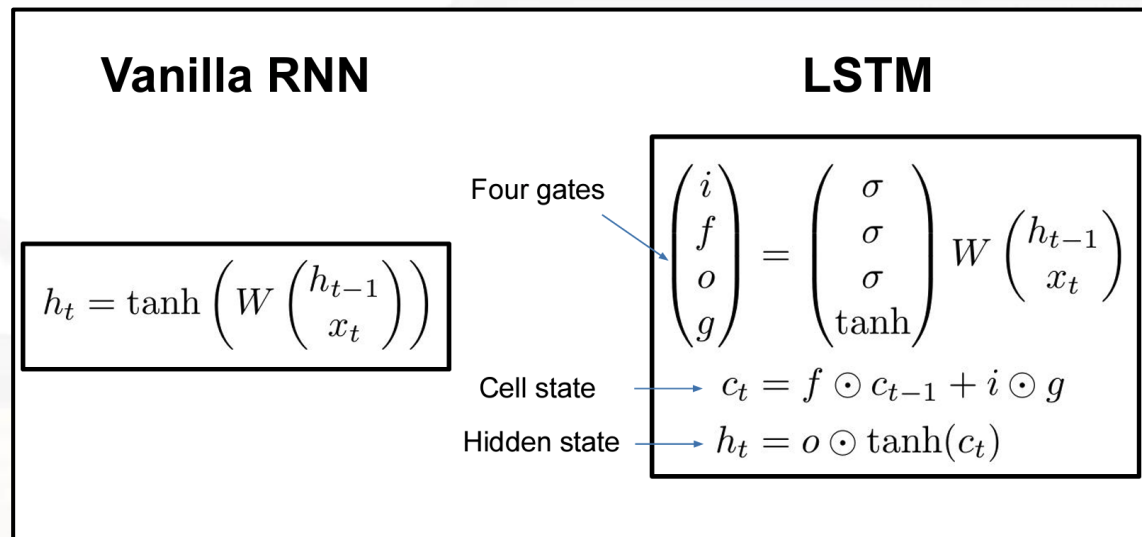
Largest singular value > 1:
Exploding gradients

Largest singular value < 1:
Vanishing gradients

Gradient clipping:
Scale gradient if its norm is too big

```
grad_norm = np.sum(grad * grad)
if grad_norm > threshold:
    grad *= (threshold / grad_norm)
```

RNN and LSTM




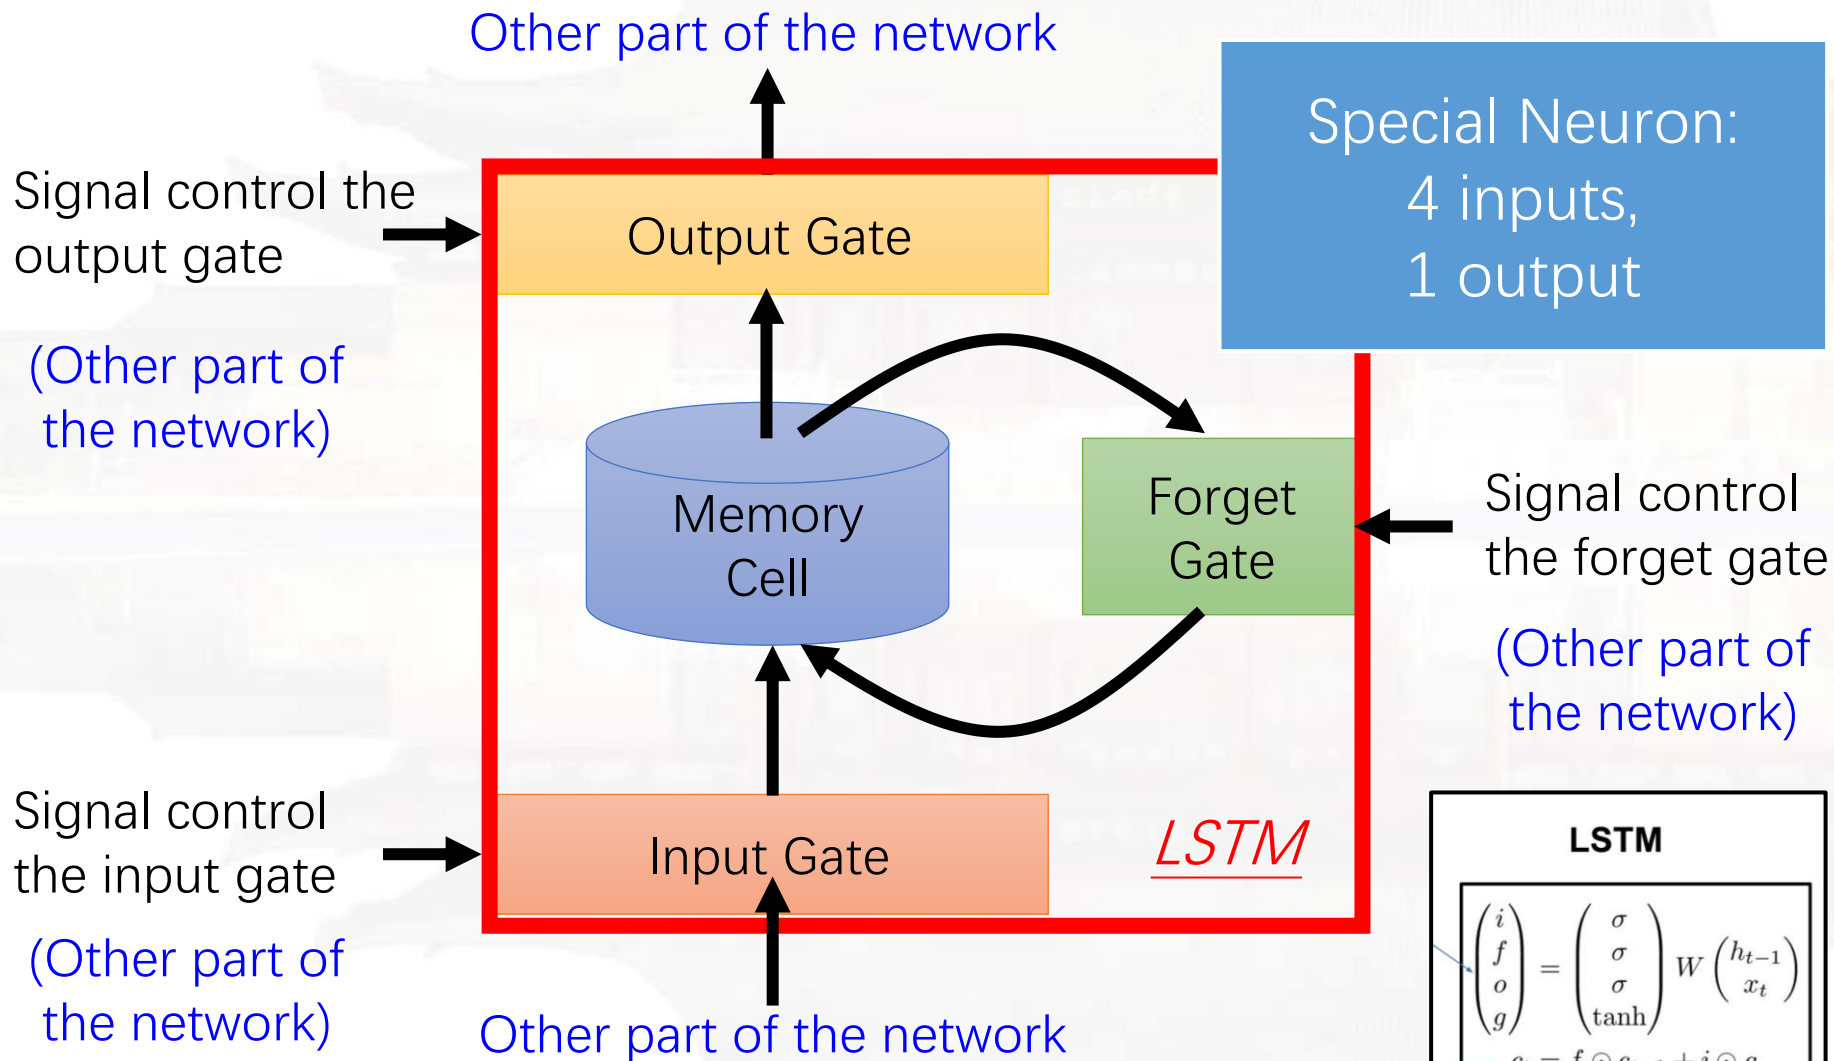
i: Input gate, whether to write to cell

f: Forget gate, Whether to erase cell

o: Output gate, How much to reveal cell

g: Gate gate (?), How much to write to cell

RNN and LSTM



LSTM

$$\begin{pmatrix} i \\ f \\ o \\ g \end{pmatrix} = \begin{pmatrix} \sigma \\ \sigma \\ \sigma \\ \tanh \end{pmatrix} W \begin{pmatrix} h_{t-1} \\ x_t \end{pmatrix}$$

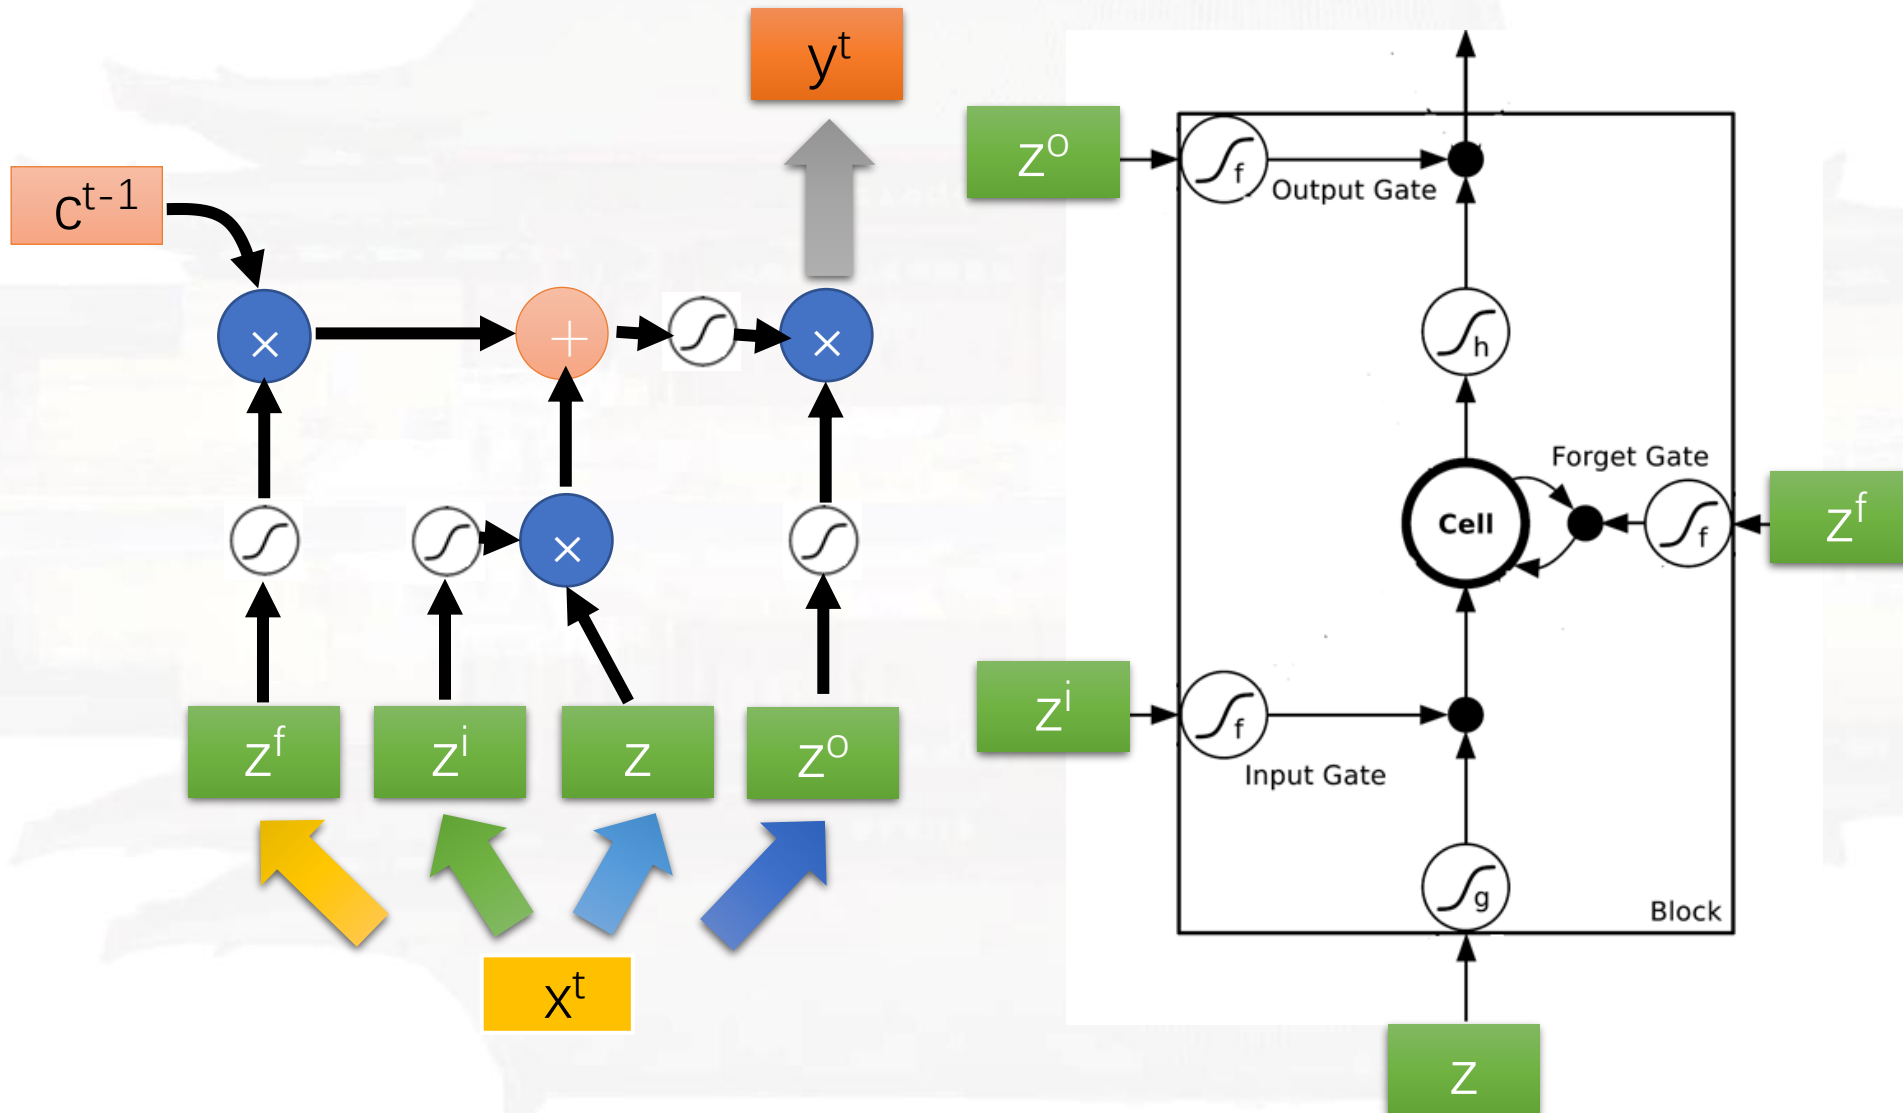
$\rightarrow c_t = f \odot c_{t-1} + i \odot g$

$\rightarrow h_t = o \odot \tanh(c_t)$

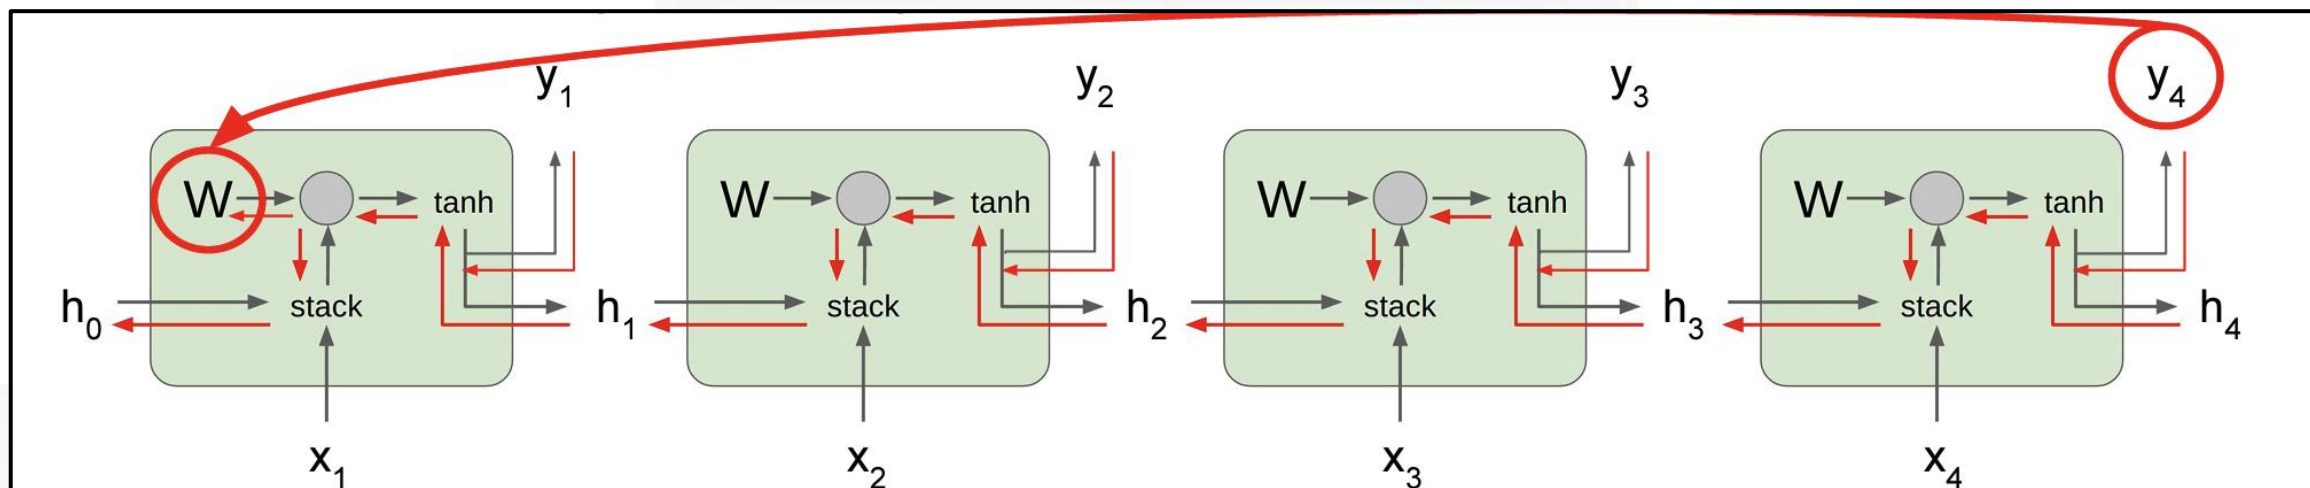
RNN and LSTM



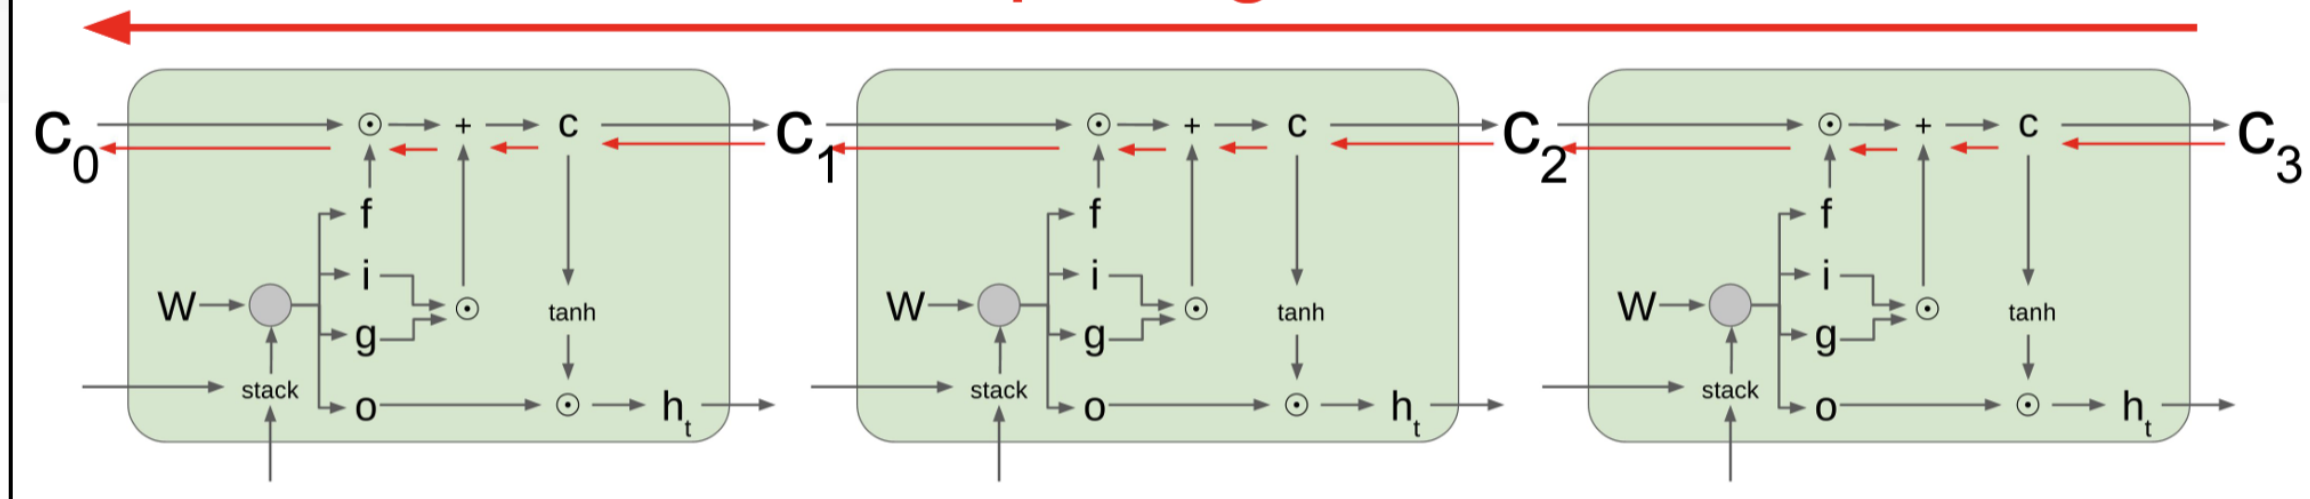
浙江大學
ZHEJIANG UNIVERSITY




RNN and LSTM



Uninterrupted gradient flow!



From RNN to Attention



<https://arxiv.org/abs/2006.16036>


Transformers are RNNs: Fast Autoregressive Transformers with Linear Attention

From RNN to Attention



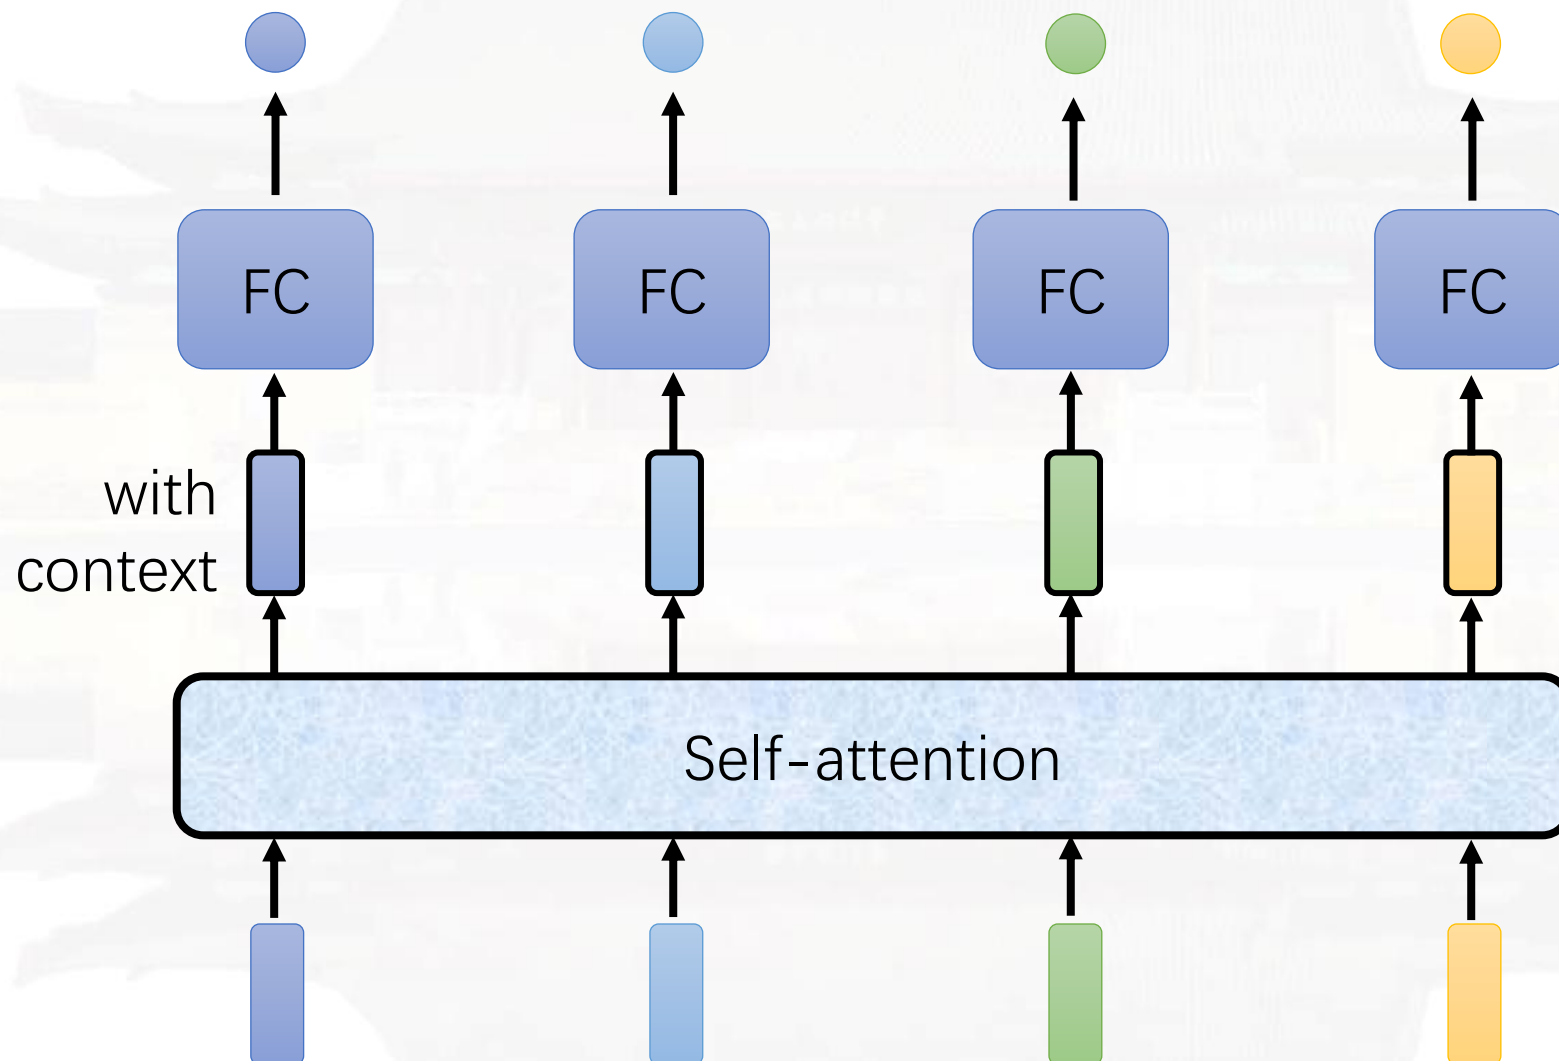
b^i is obtained based on the whole input sequence.

b^1, b^2, b^3, b^4 can be parallelly computed.




You can try to replace any thing that has been done by RNN with self-attention.

Self-Attention




Self-Attention



Can be either **input** or a **hidden layer**

Self-Attention




Find the relevant vectors in a sequence


Self-Attention




Dot-product




Additive




Self-Attention




Self-Attention




Self-Attention



Self-Attention



Self-Attention







Self-Attention



$$q^i = W^q a^i \quad \begin{matrix} q^1 & q^2 & q^3 & q^4 \\ Q \end{matrix} = \begin{matrix} W^q \\ I \end{matrix} \begin{matrix} a^1 & a^2 & a^3 & a^4 \\ I \end{matrix}$$

$$k^i = W^k a^i \quad \begin{matrix} k^1 & k^2 & k^3 & k^4 \\ K \end{matrix} = \begin{matrix} W^k \\ I \end{matrix} \begin{matrix} a^1 & a^2 & a^3 & a^4 \\ I \end{matrix}$$

$$v^i = W^v a^i \quad \begin{matrix} v^1 & v^2 & v^3 & v^4 \\ V \end{matrix} = \begin{matrix} W^v \\ I \end{matrix} \begin{matrix} a^1 & a^2 & a^3 & a^4 \\ I \end{matrix}$$




Self-Attention




$$\begin{aligned}\alpha_{1,1} &= k^1 q^1 & \alpha_{1,2} &= k^2 q^1 \\ \alpha_{1,3} &= k^3 q^1 & \alpha_{1,4} &= k^4 q^1\end{aligned}$$

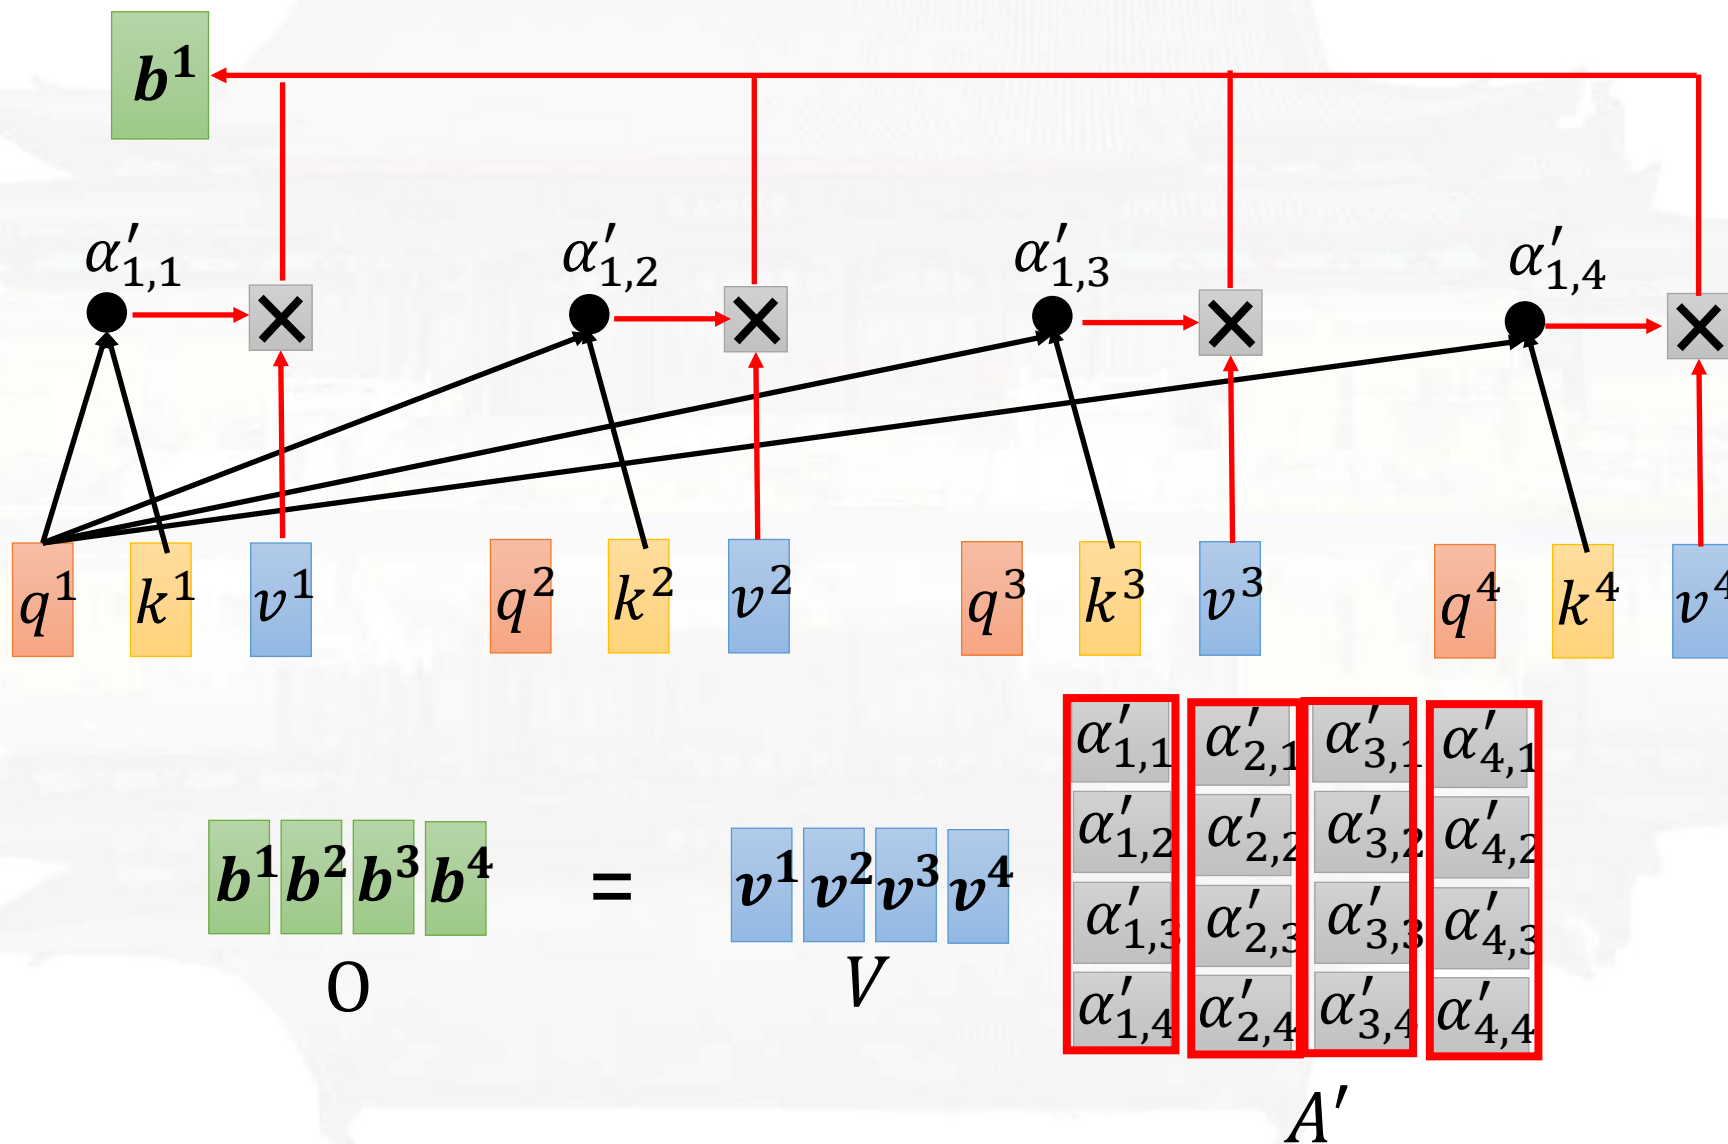
$$\begin{aligned}\alpha_{1,1} \\ \alpha_{1,2} \\ \alpha_{1,3} \\ \alpha_{1,4}\end{aligned} = \begin{aligned}k^1 \\ k^2 \\ k^3 \\ k^4\end{aligned} q^1$$



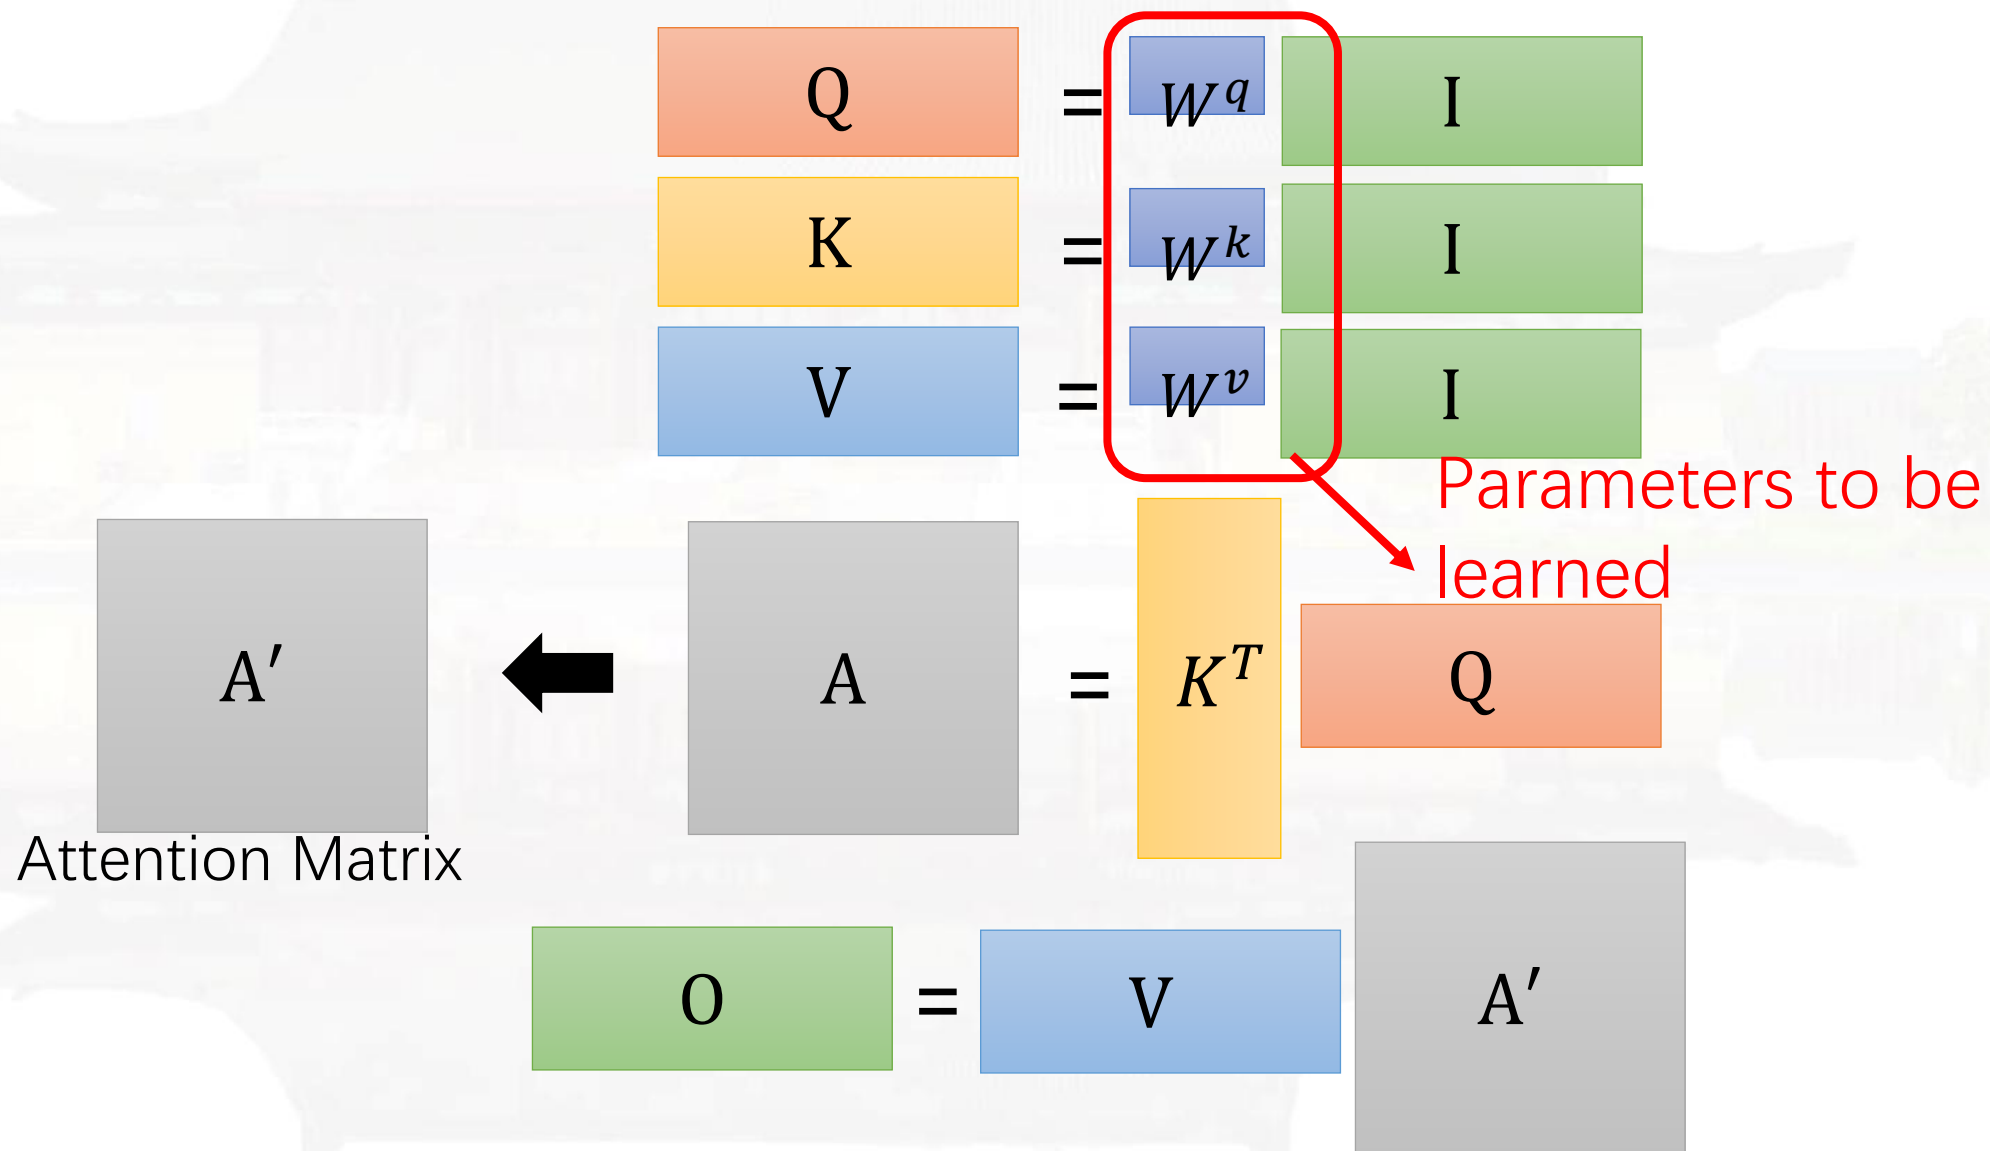
Self-Attention




Self-Attention



Self-Attention



Multi-head Self-attention

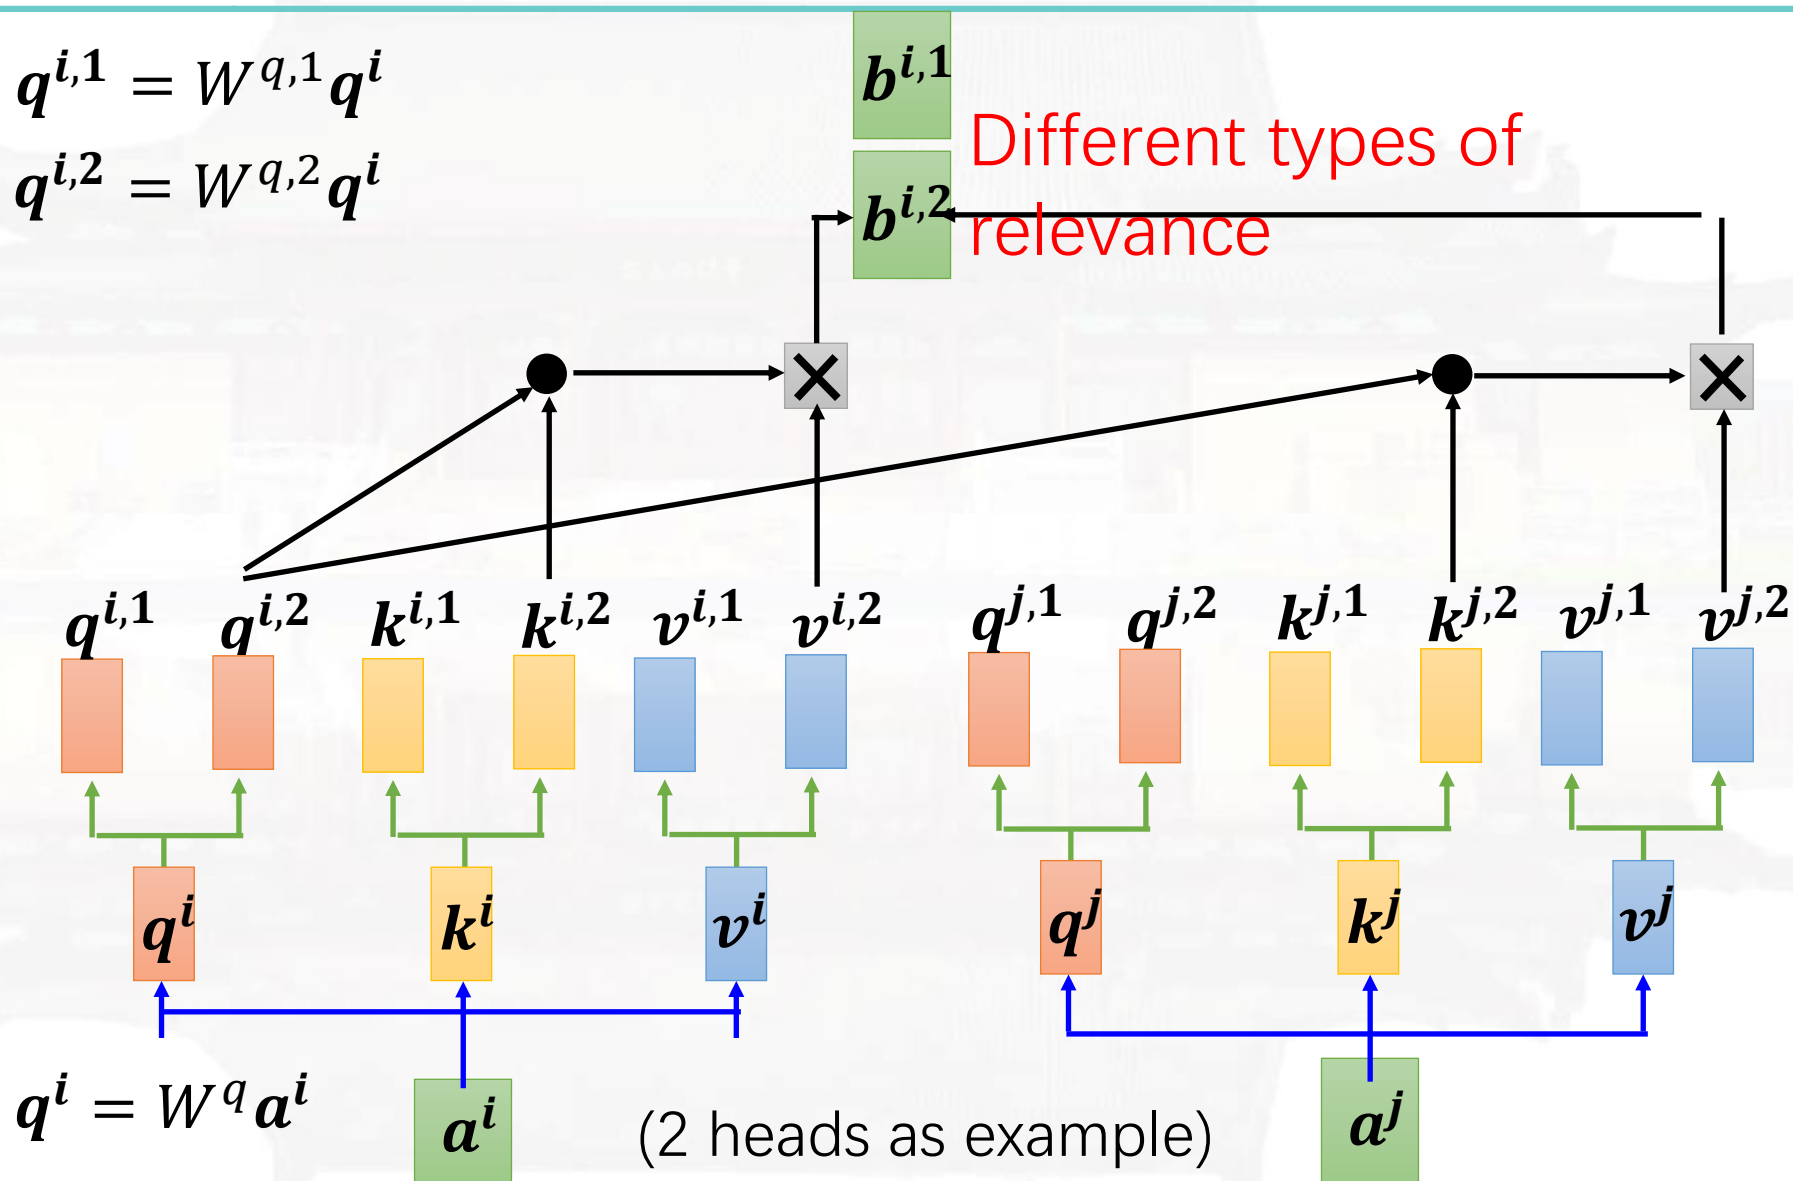


Multi-head Self-attention



$$q^{i,1} = W^{q,1} q^i$$

$$q^{i,2} = W^{q,2} q^i$$




Multi-head Self-attention



$$b^i = W^o \begin{bmatrix} b^{i,1} \\ b^{i,2} \end{bmatrix}$$


Different types of relevance




Positional Encoding



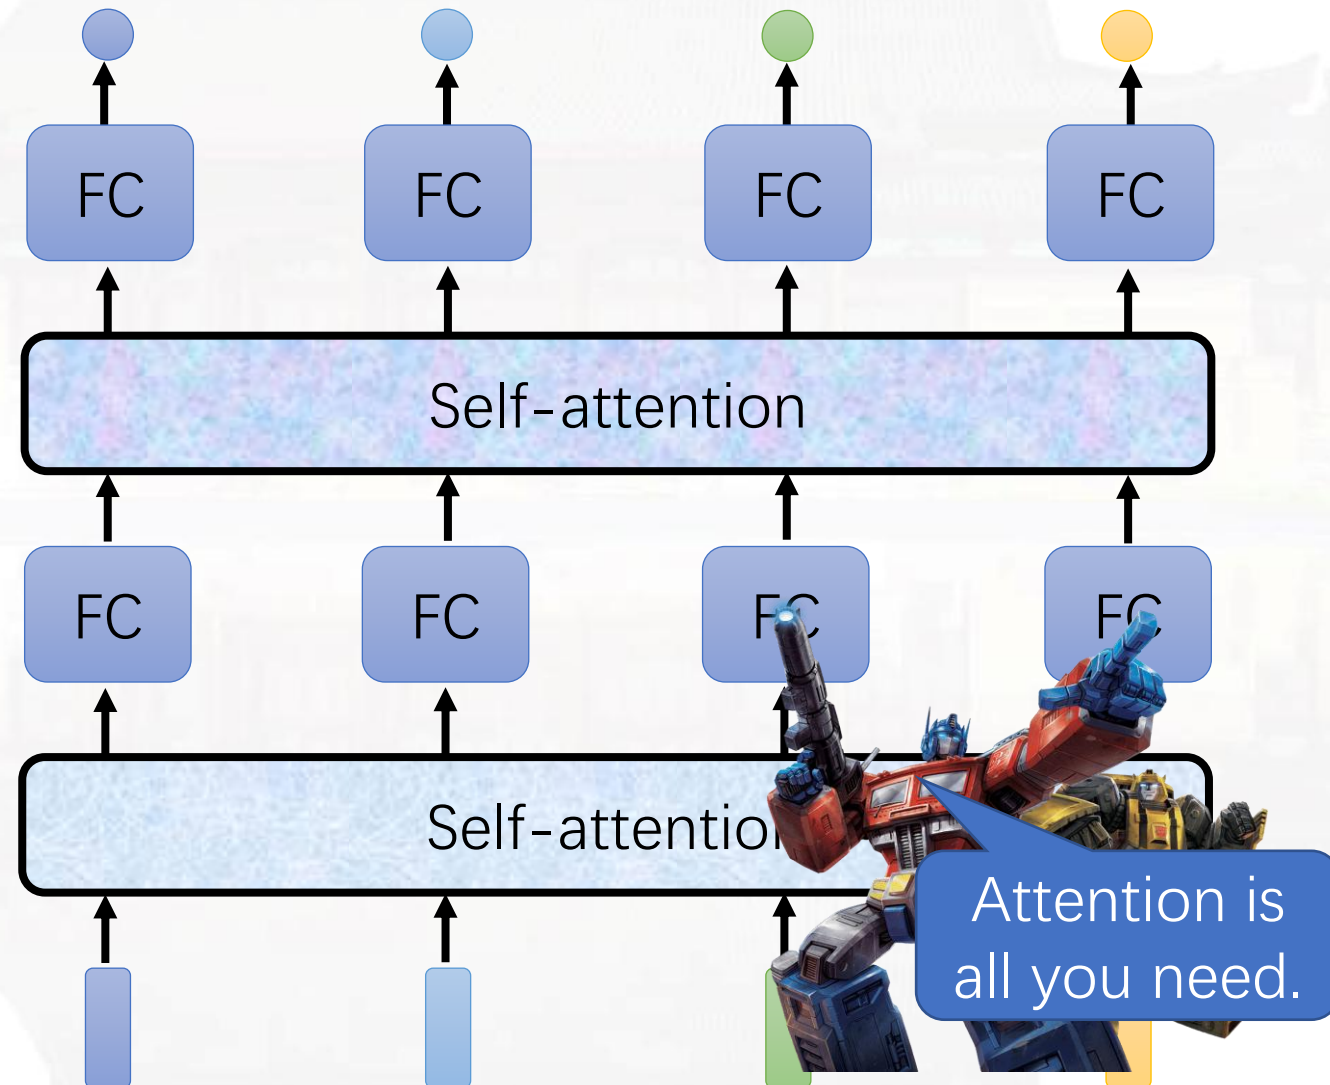
- No position information in self-attention.
- Each position has a unique positional vector e^i
- hand-crafted
- learned from data



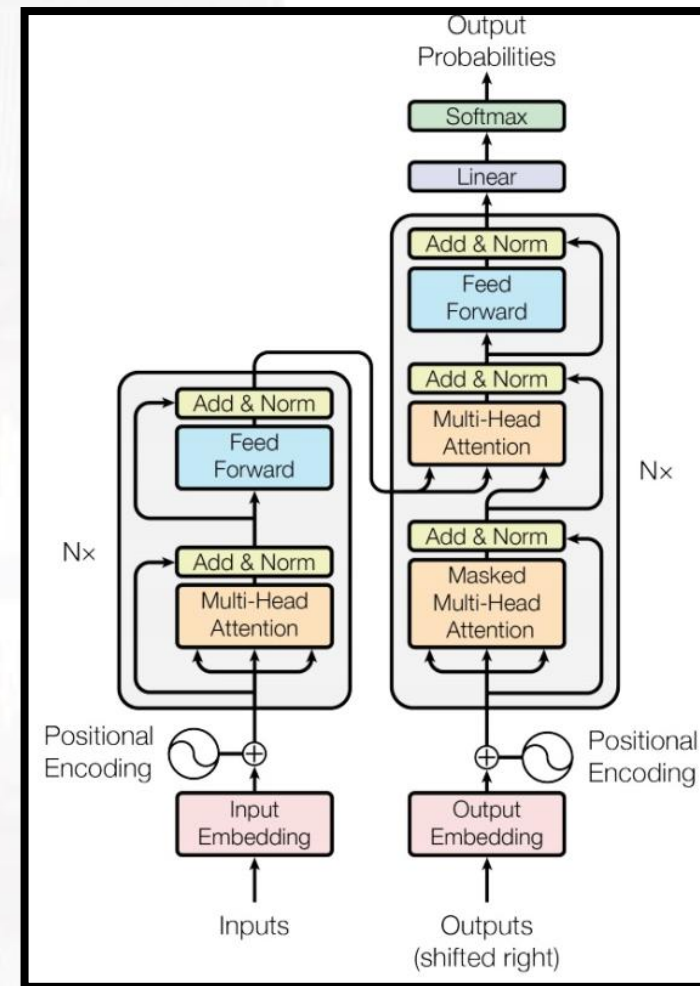

Each column represents a positional vector e^i



Transformer



Transformer

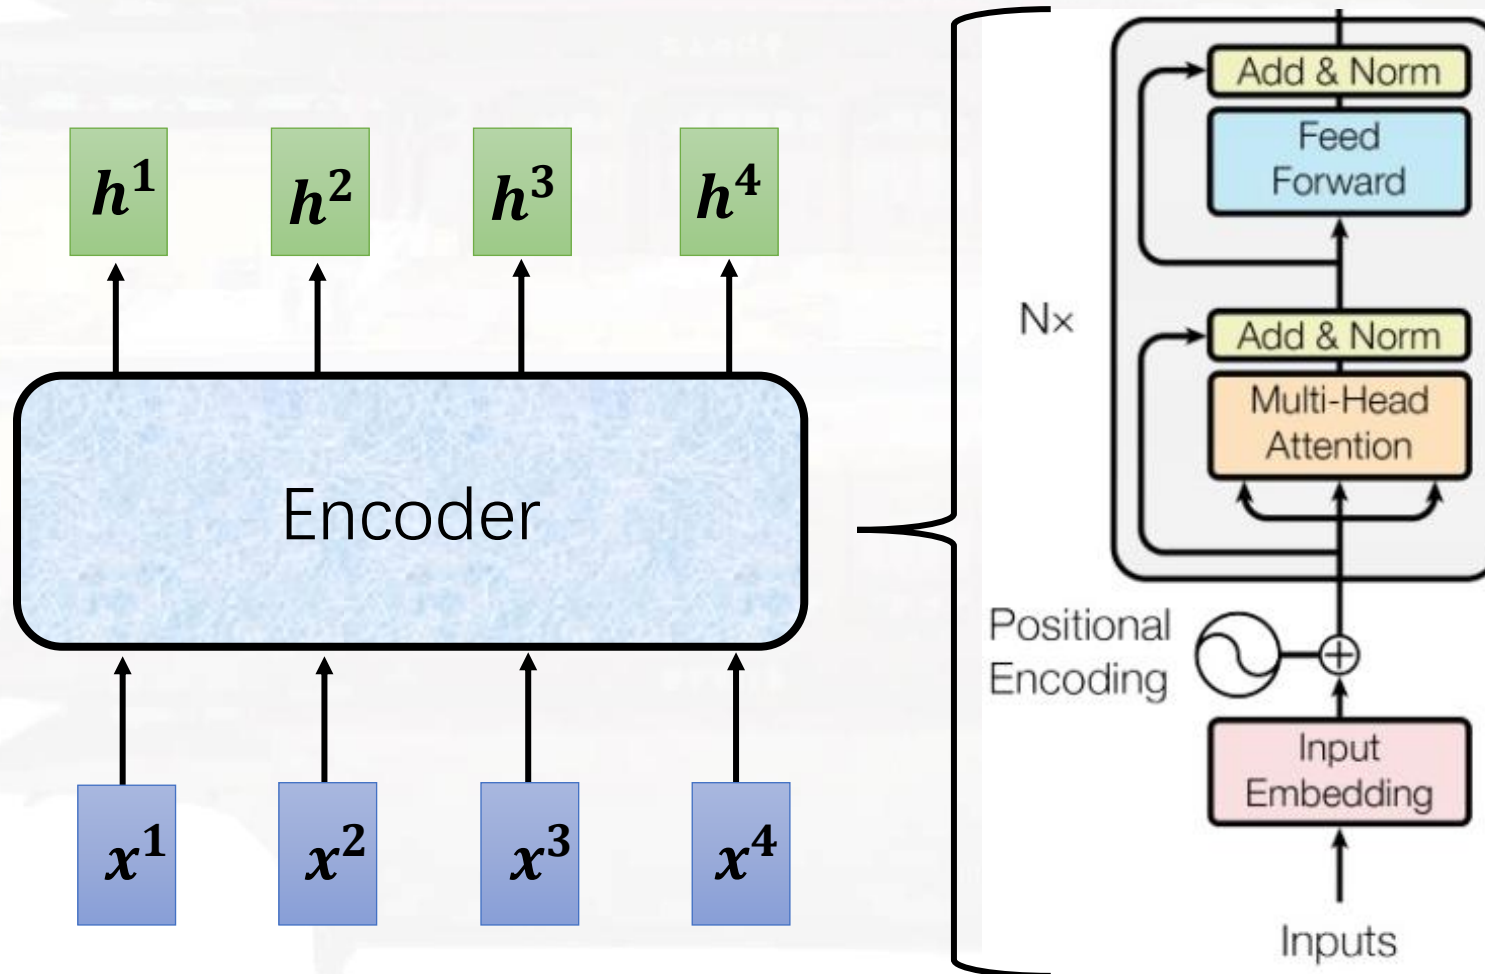


Transformer

Transformer




- Transformer's Encoder



Transformer



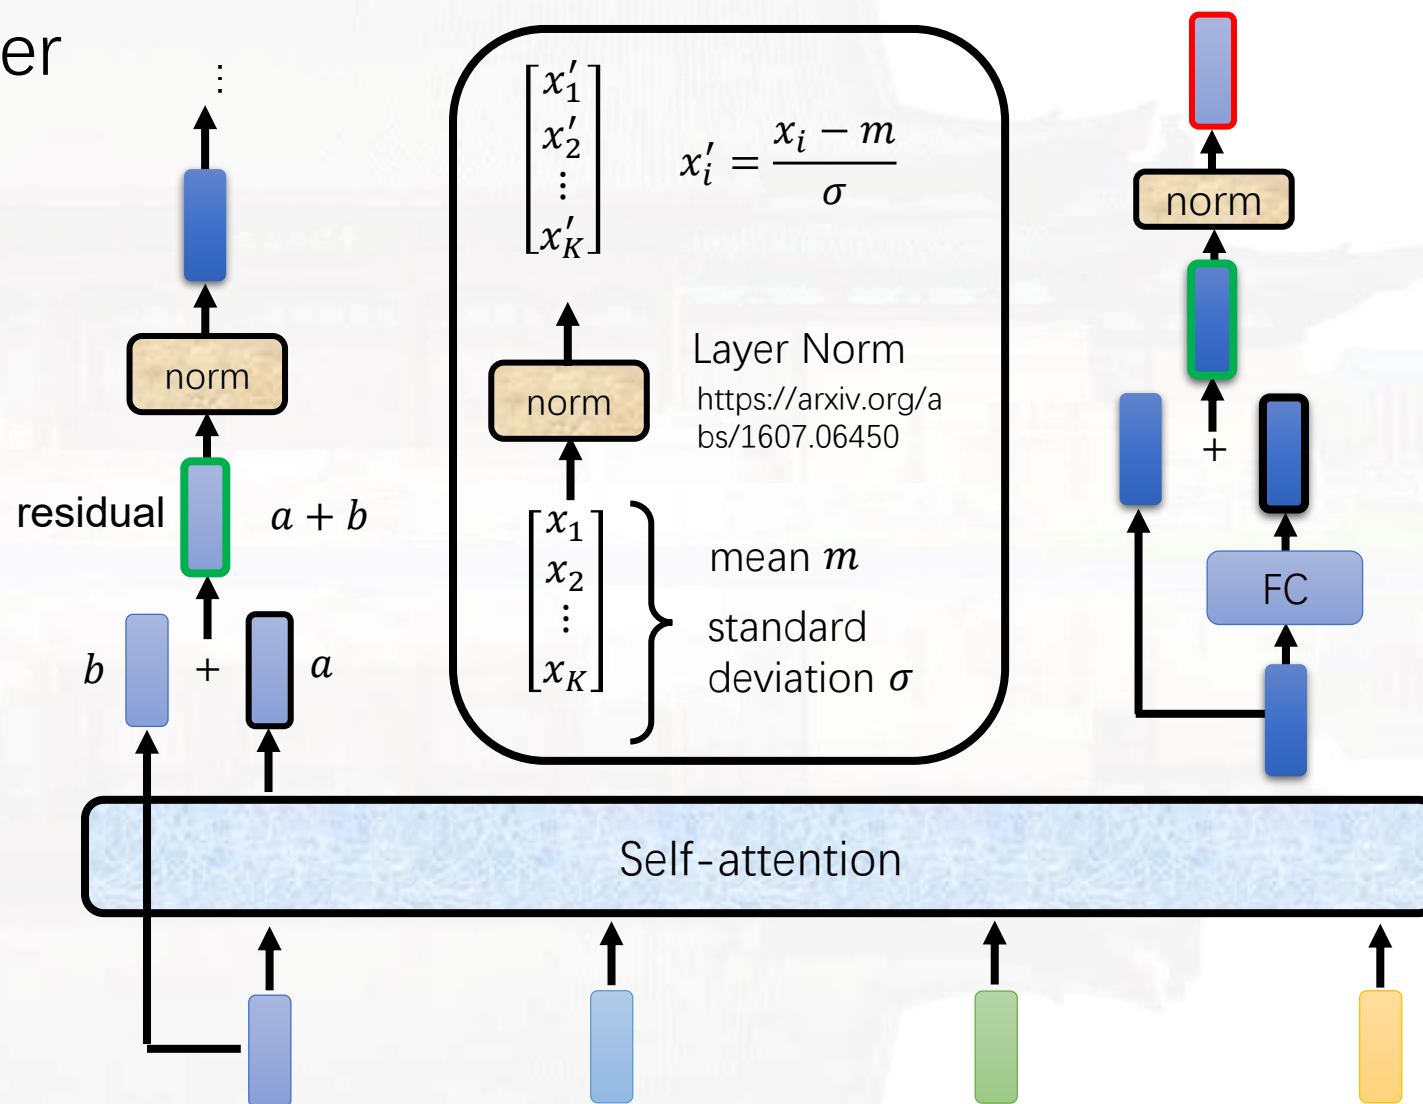
- Transformer's Encoder



Transformer



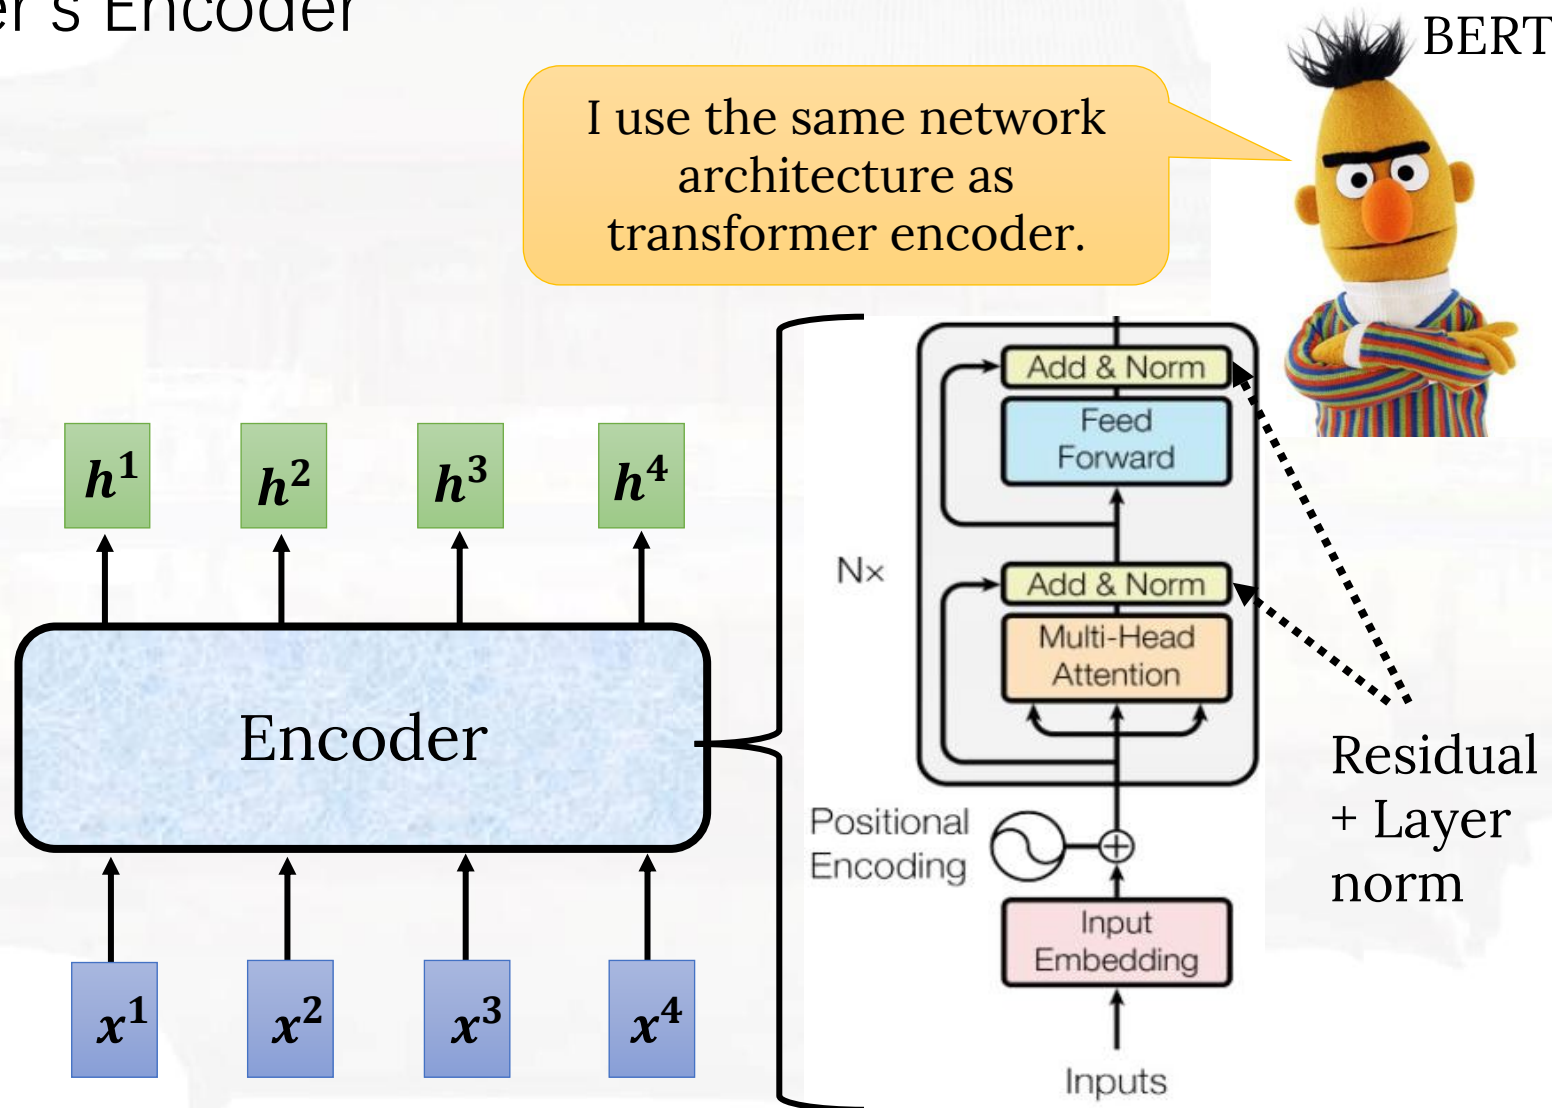
- Transformer's Encoder



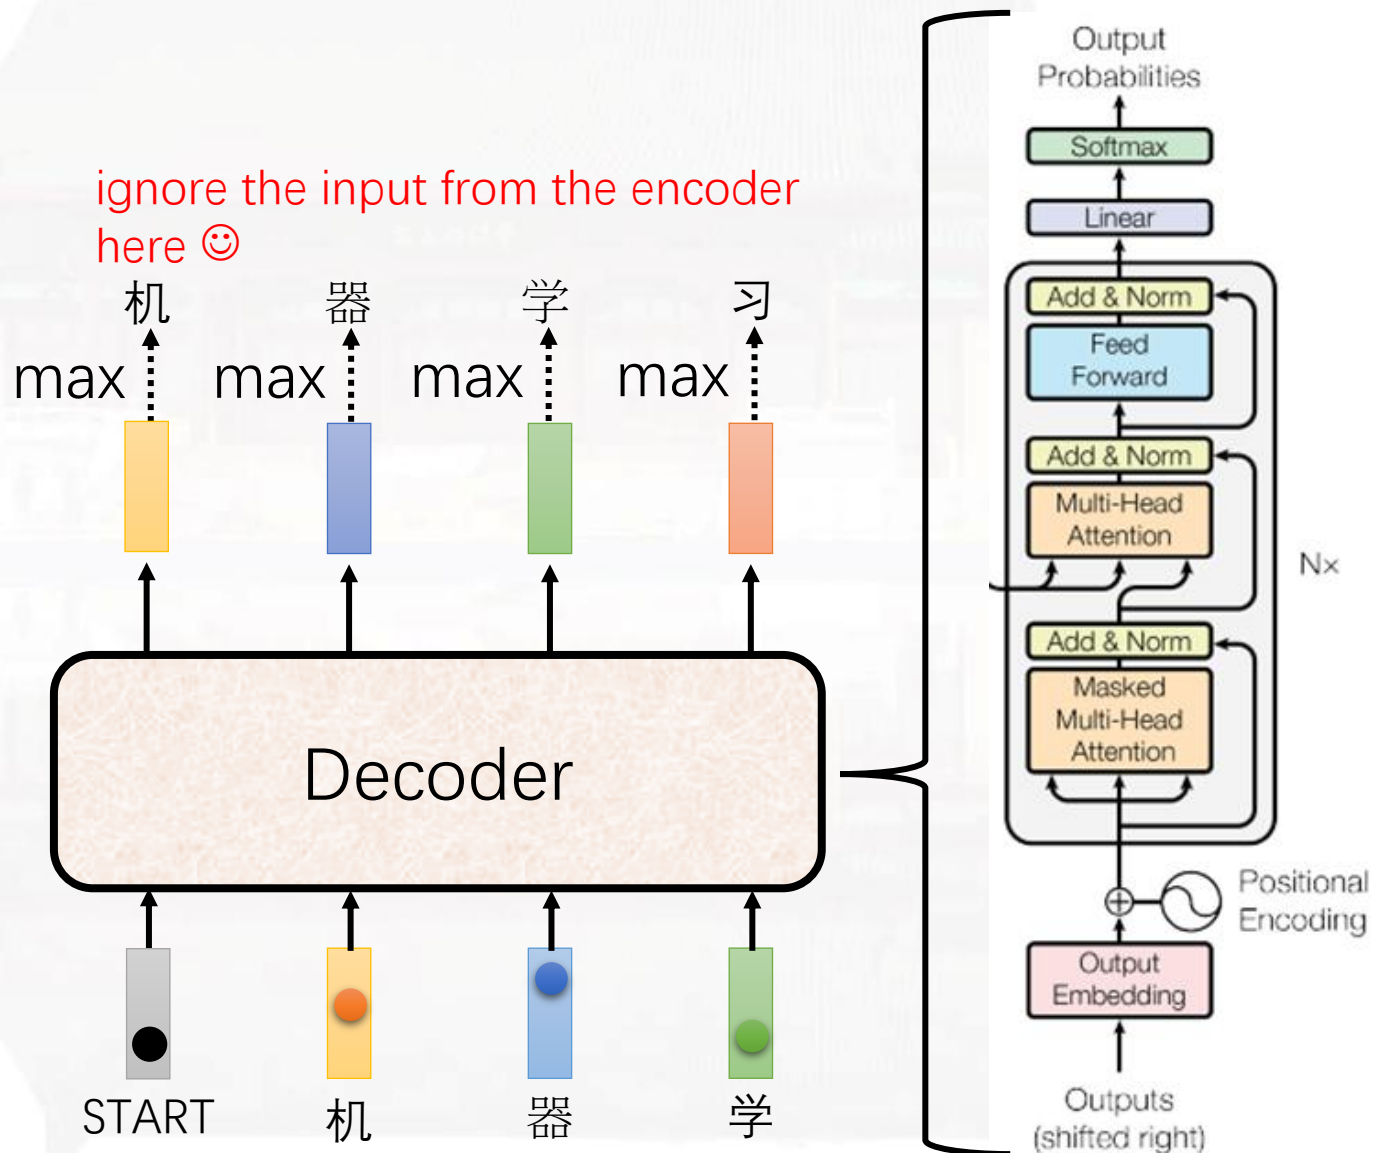
Transformer



- Transformer's Encoder



Transformer

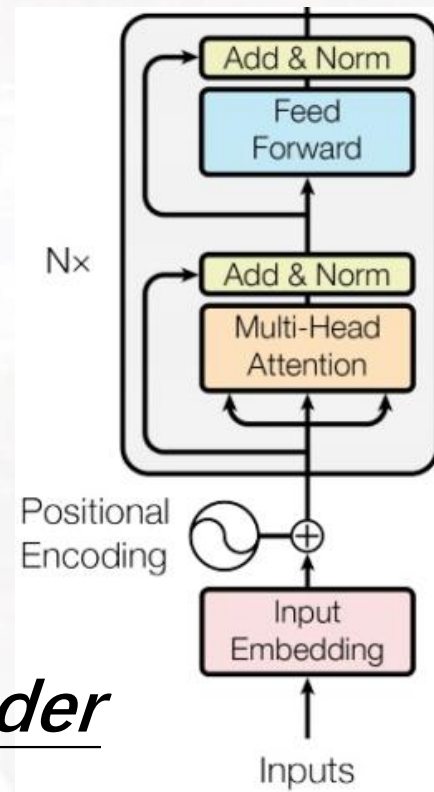


Transformer

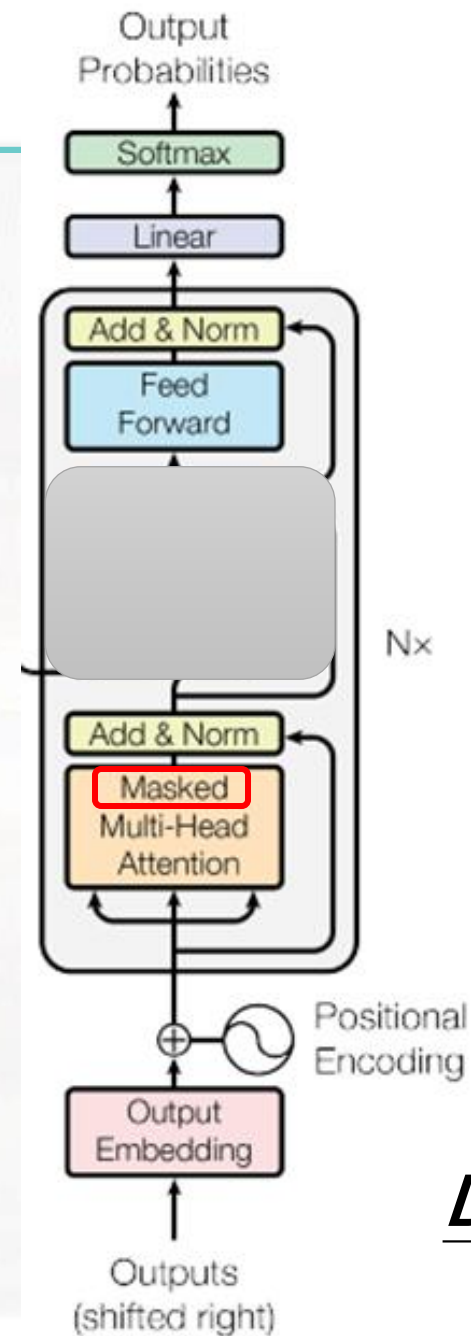


浙江大學
ZHEJIANG UNIVERSITY

Encoder




Decoder



Transformer



- Masked Self-attention




Can be either **input** or a **hidden layer**

Transformer

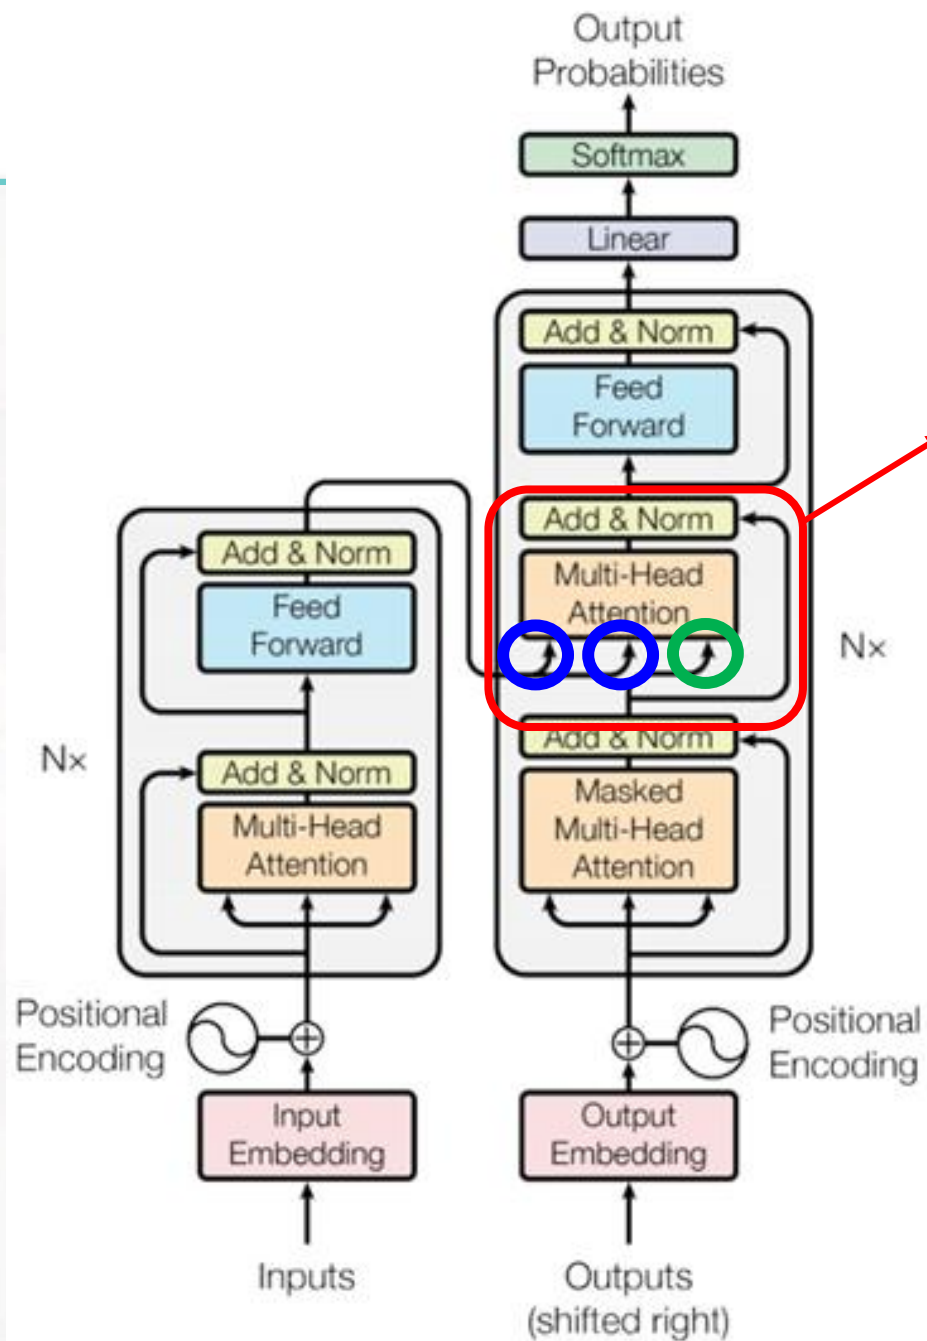


- Masked Self-attention

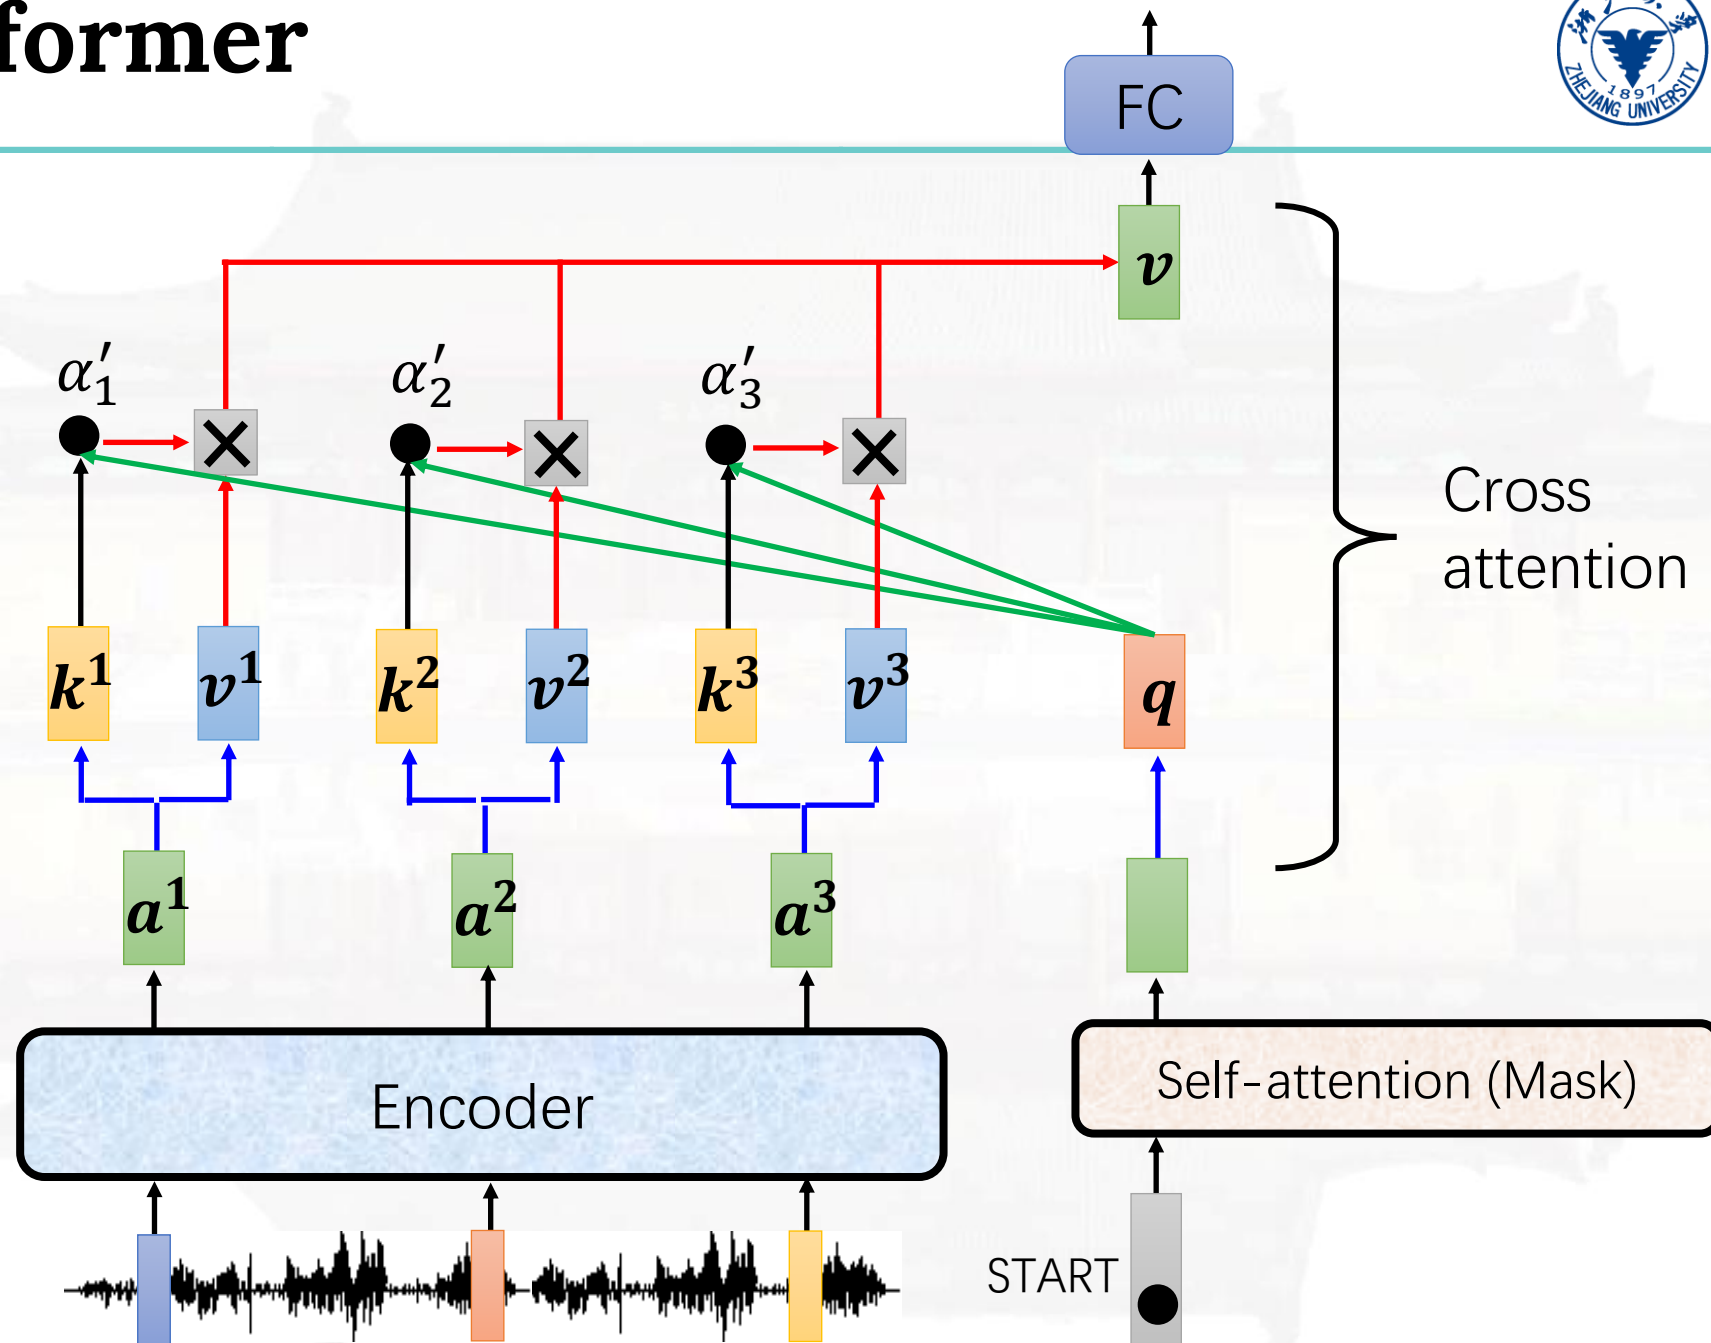


Why masked? Consider how does decoder work

Transformer



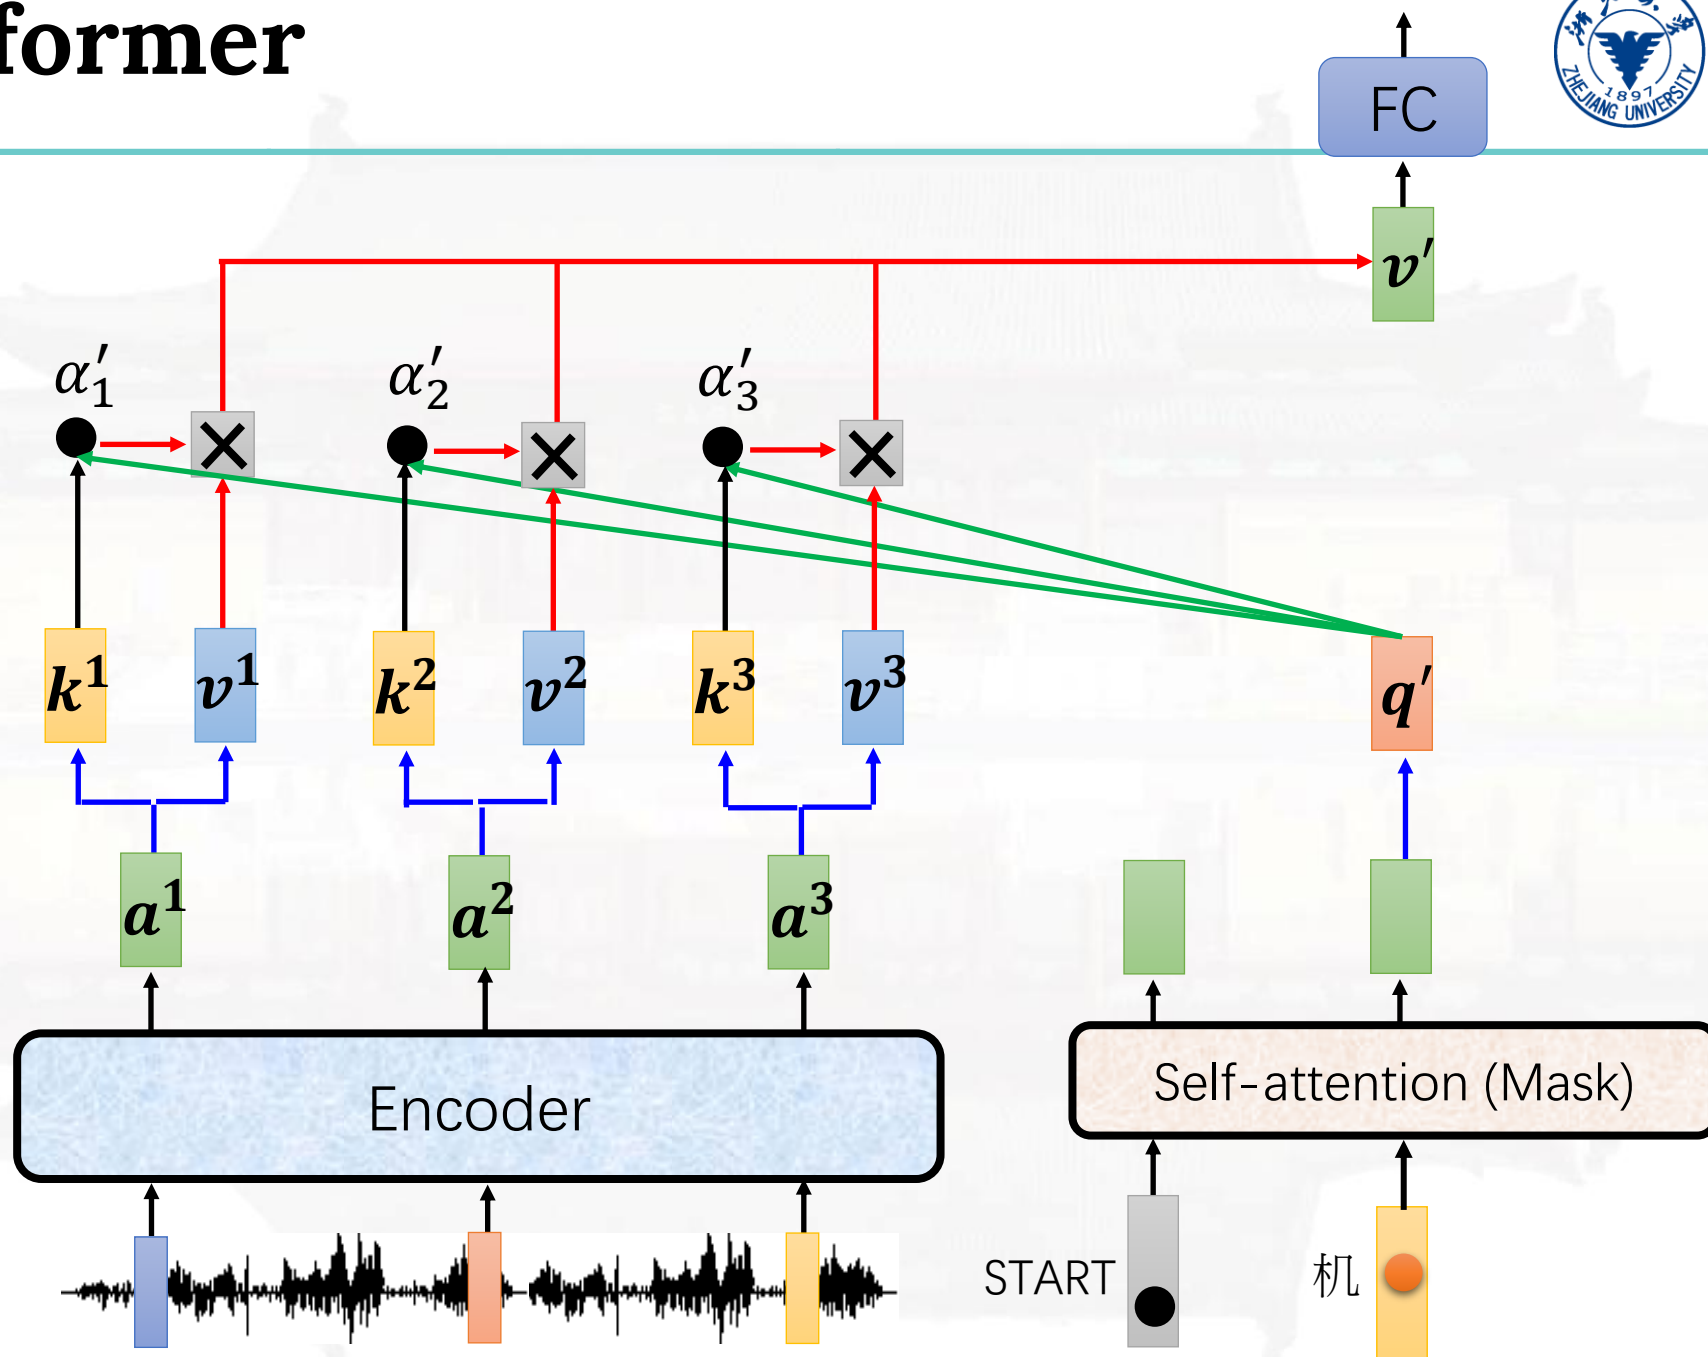
Transformer



Transformer



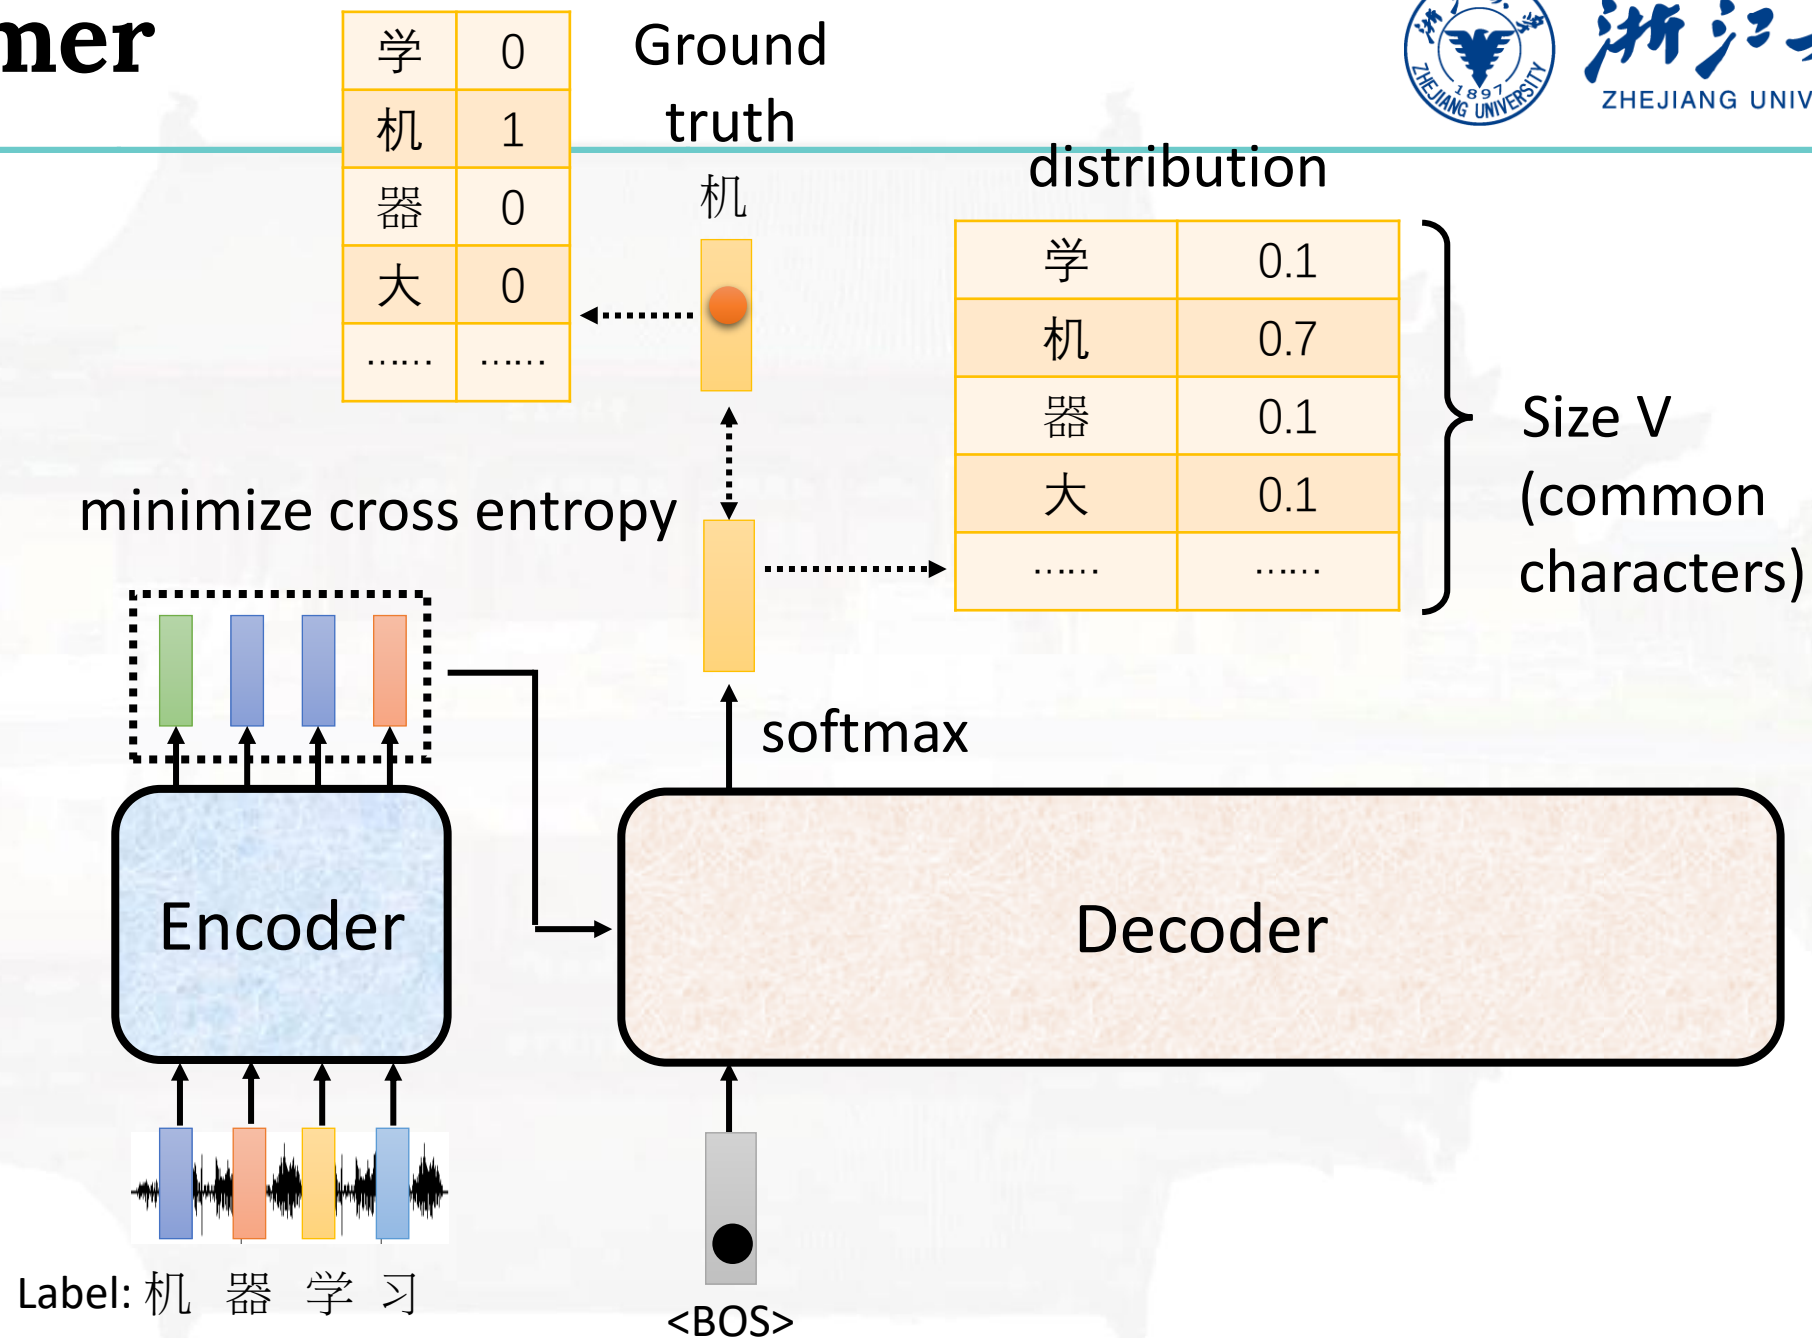
浙江大学
ZHEJIANG UNIVERSITY



Transformer



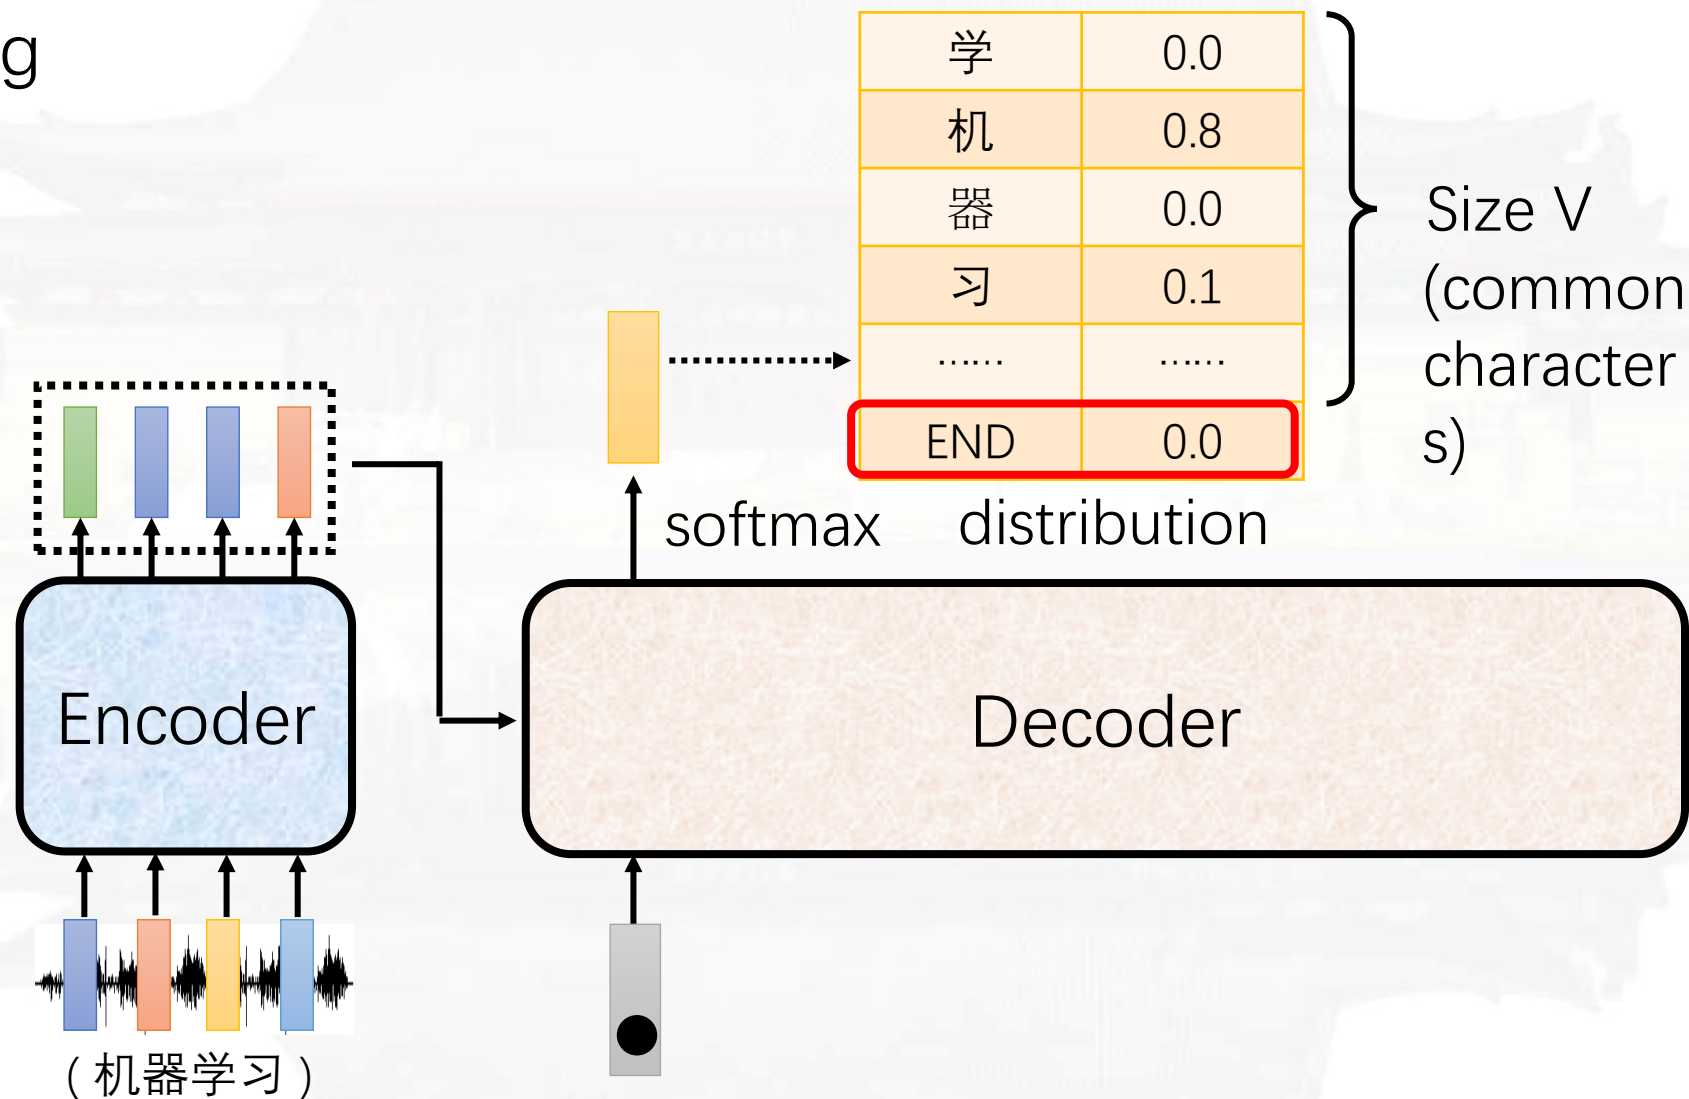
- Training



Transformer



- Training



Transformer



- Training

