

人工智慧投資分析



Tamkang
Universit
淡江大學

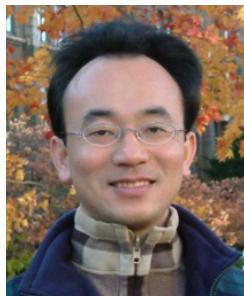
Artificial Intelligence for Investment Analysis

TensorFlow 深度學習財務時間序列預測 (Deep Learning for Financial Time Series Forecasting with TensorFlow)

1071AIIA09

EMBA, IMTKU (M2399) (8540)

Thu 12,13,14 (19:20-22:10) (D503)



Min-Yuh Day

戴敏育

Assistant Professor

專任助理教授

Dept. of Information Management, Tamkang University

淡江大學 資訊管理學系

<http://mail.tku.edu.tw/myday/>

2018-11-22; 2018-11-29; 2018-12-13



課程大綱 (Syllabus)

週次 (Week) 日期 (Date) 內容 (Subject/Topics)

- | | | |
|---|------------|---|
| 1 | 2018/09/13 | 人工智慧投資分析課程介紹
(Course Orientation on Artificial Intelligence for Investment Analysis) |
| 2 | 2018/09/20 | AI 金融科技：金融服務創新應用
(AI in FinTech: Financial Services Innovation and Application) |
| 3 | 2018/09/27 | 機器人理財顧問與AI交談機器人
(Robo-Advisors and AI Chatbots) |
| 4 | 2018/10/04 | 投資心理學與行為財務學
(Investing Psychology and Behavioral Finance) |
| 5 | 2018/10/11 | 財務金融事件研究法 (Event Studies in Finance) |
| 6 | 2018/10/18 | 人工智慧投資分析個案研究 I
(Case Study on Artificial Intelligence for Investment Analysis I) |

課程大綱 (Syllabus)

週次 (Week) 日期 (Date) 內容 (Subject/Topics)

- 7 2018/10/25 Python AI投資分析基礎
(Foundations of AI Investment Analysis in Python)
- 8 2018/11/01 Python Pandas量化投資分析
(Quantitative Investing with Pandas in Python)
- 9 2018/11/08 Python Scikit-Learn 機器學習
(Machine Learning with Scikit-Learn in Python)
- 10 2018/11/15 期中報告 (Midterm Project Report)
- 11 2018/11/22 TensorFlow 深度學習財務時間序列預測 I
(Deep Learning for Financial Time Series Forecasting with TensorFlow I)
- 12 2018/11/29 TensorFlow 深度學習財務時間序列預測 II
(Deep Learning for Financial Time Series Forecasting with TensorFlow II)

課程大綱 (Syllabus)

週次 (Week) 日期 (Date) 內容 (Subject/Topics)

13 2018/12/06 人工智慧投資分析個案研究 II

(Case Study on Artificial Intelligence for Investment Analysis II)

14 2018/12/13 TensorFlow 深度學習財務時間序列預測 III

(Deep Learning for Financial Time Series Forecasting
with TensorFlow III)

15 2018/12/20 投資組合最佳化與程式交易

(Portfolio Optimization and Algorithmic Trading)

16 2018/12/27 自然語言處理 (Natural Language Processing)

17 2019/01/03 期末報告 I (Final Project Presentation I)

18 2019/01/10 期末報告 II (Final Project Presentation II)

Outline

- Deep Learning for
Financial Time Series Forecasting
with TensorFlow
 - Deep Learning
 - Financial Time Series Forecasting
 - TensorFlow

AI, ML, DL

Artificial Intelligence (AI)

Machine Learning (ML)

Supervised
Learning

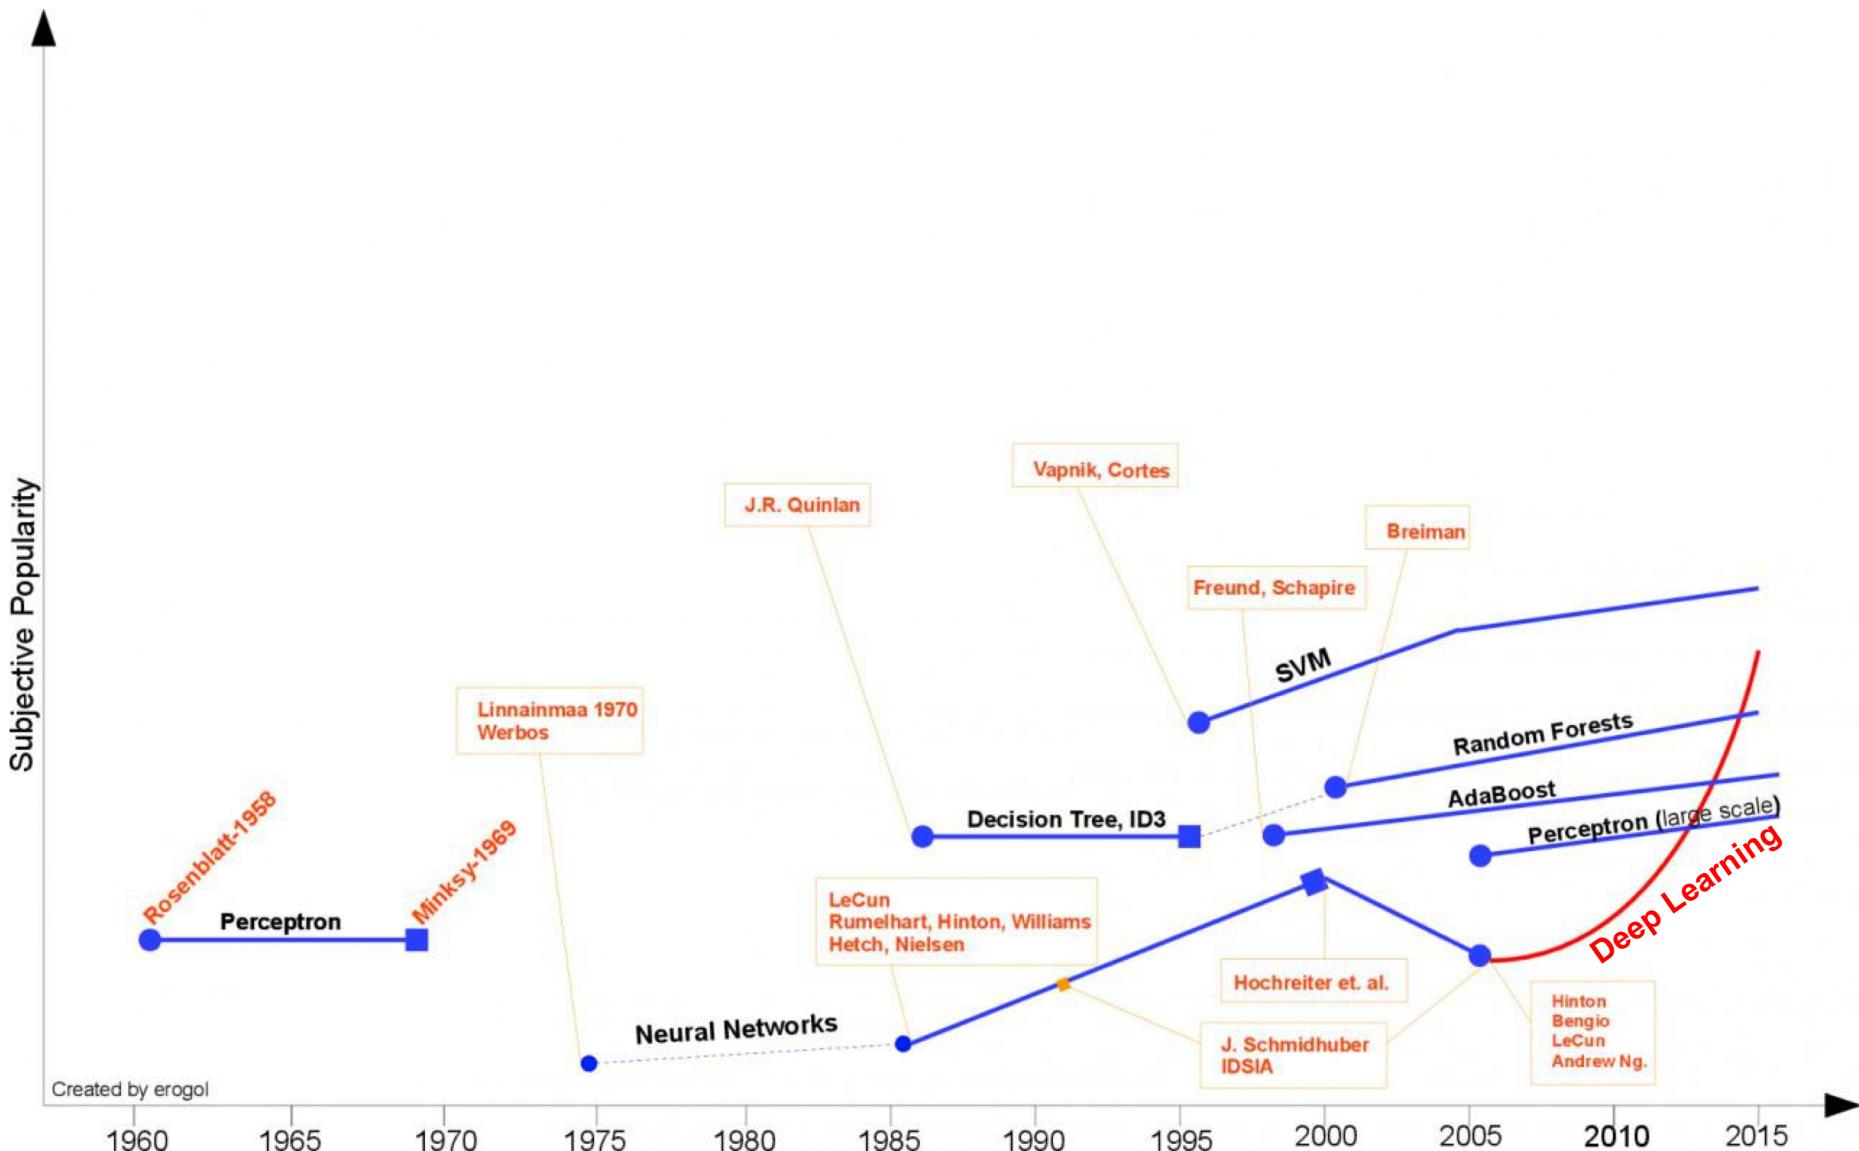
Unsupervised
Learning

Deep Learning (DL)
CNN
RNN LSTM GRU
GAN

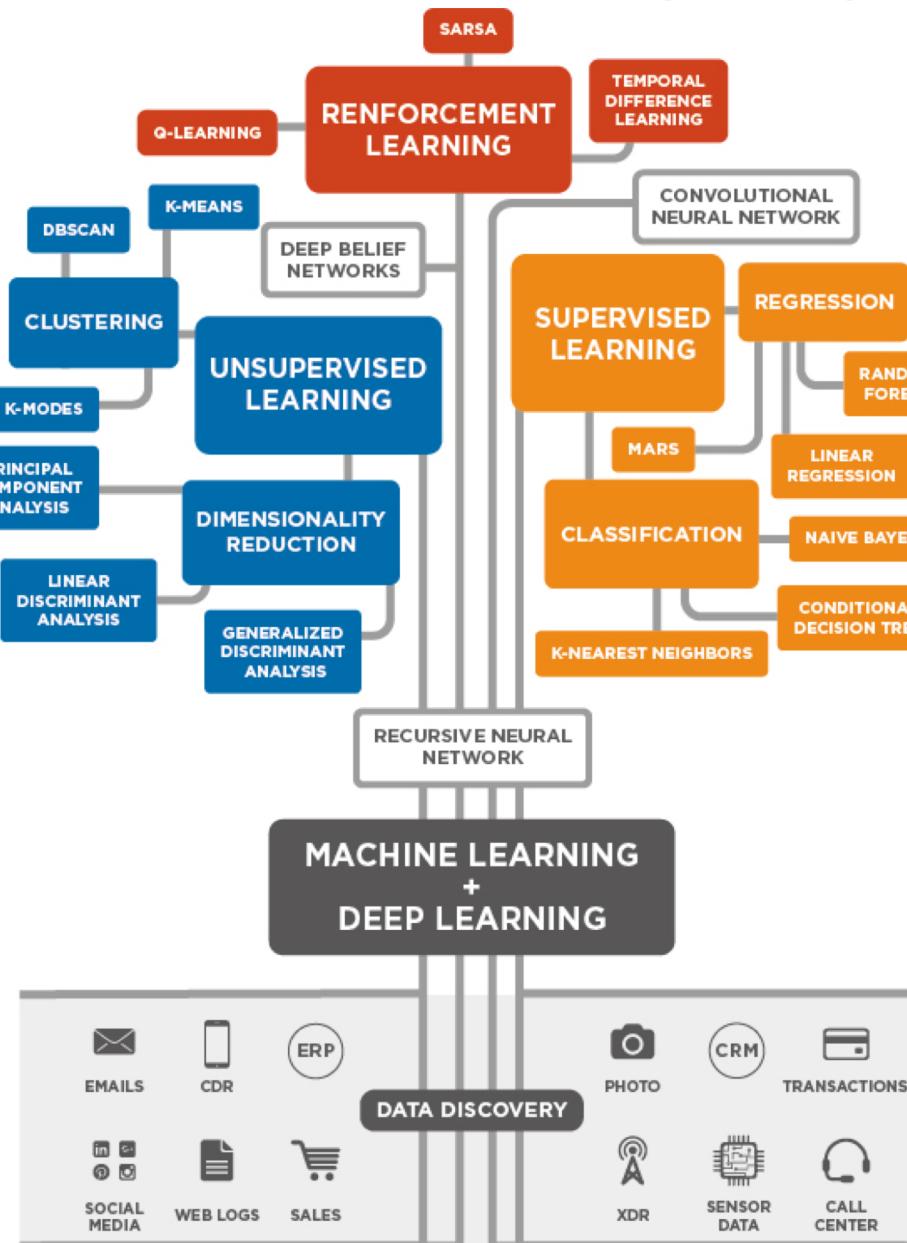
Semi-supervised
Learning

Reinforcement
Learning

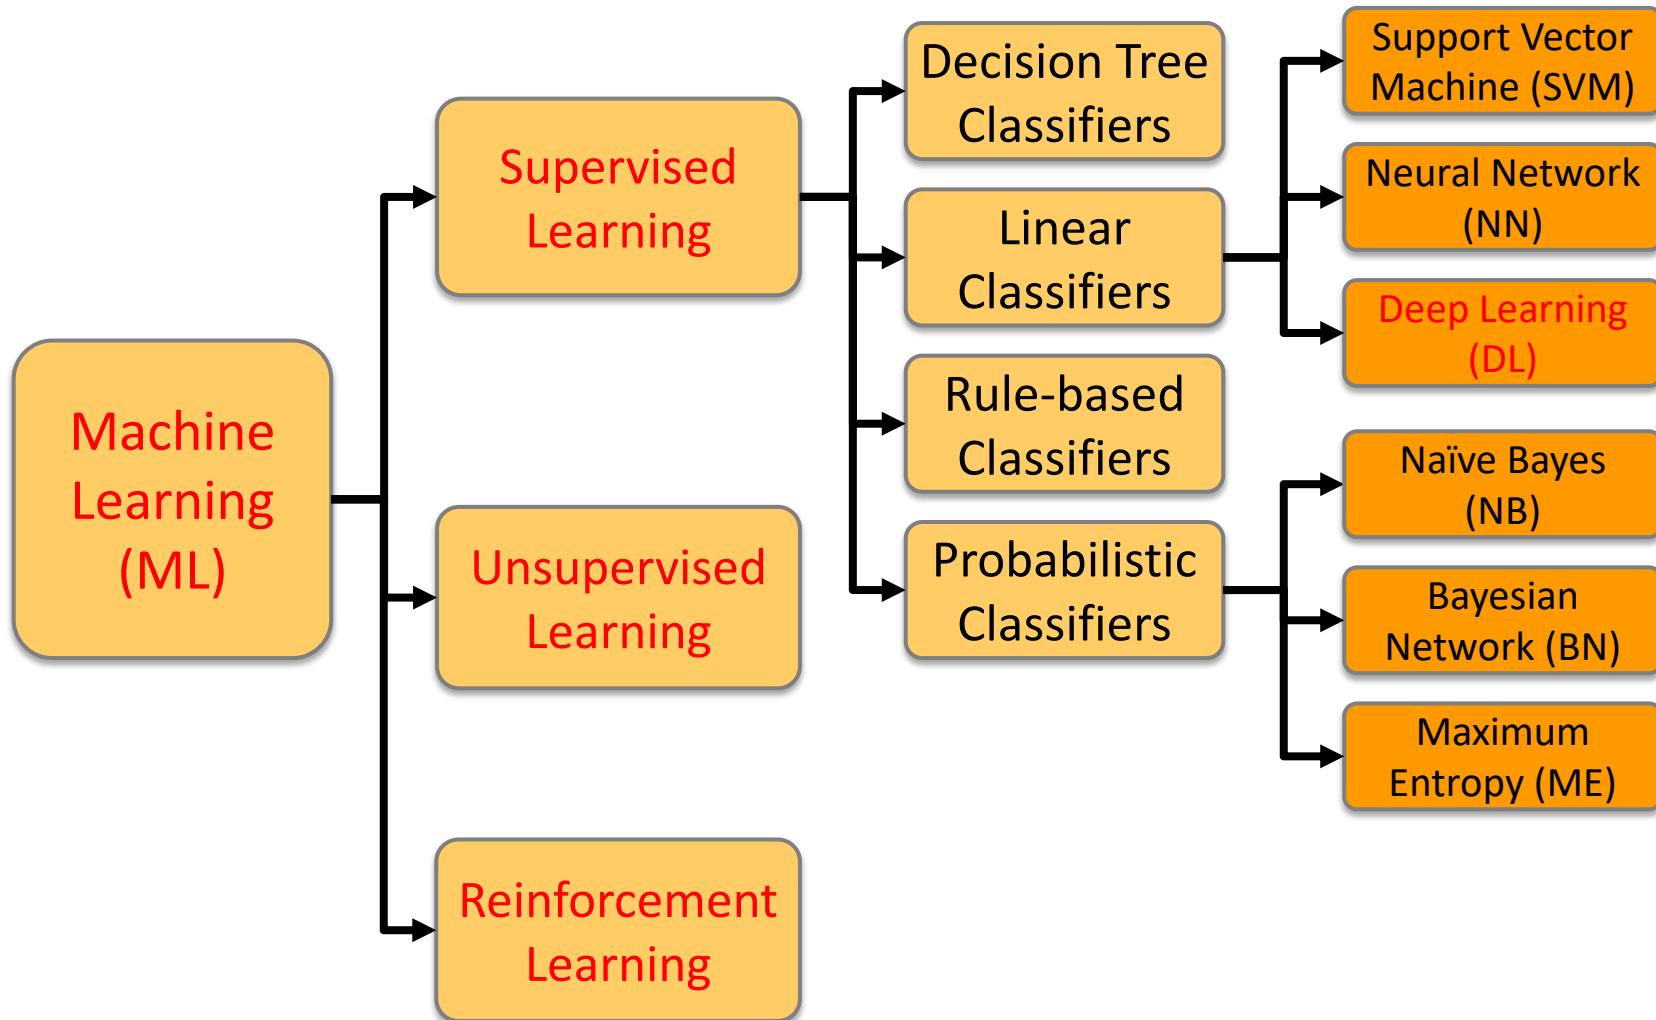
Deep Learning Evolution



3 Machine Learning Algorithms



Machine Learning (ML) / Deep Learning (DL)

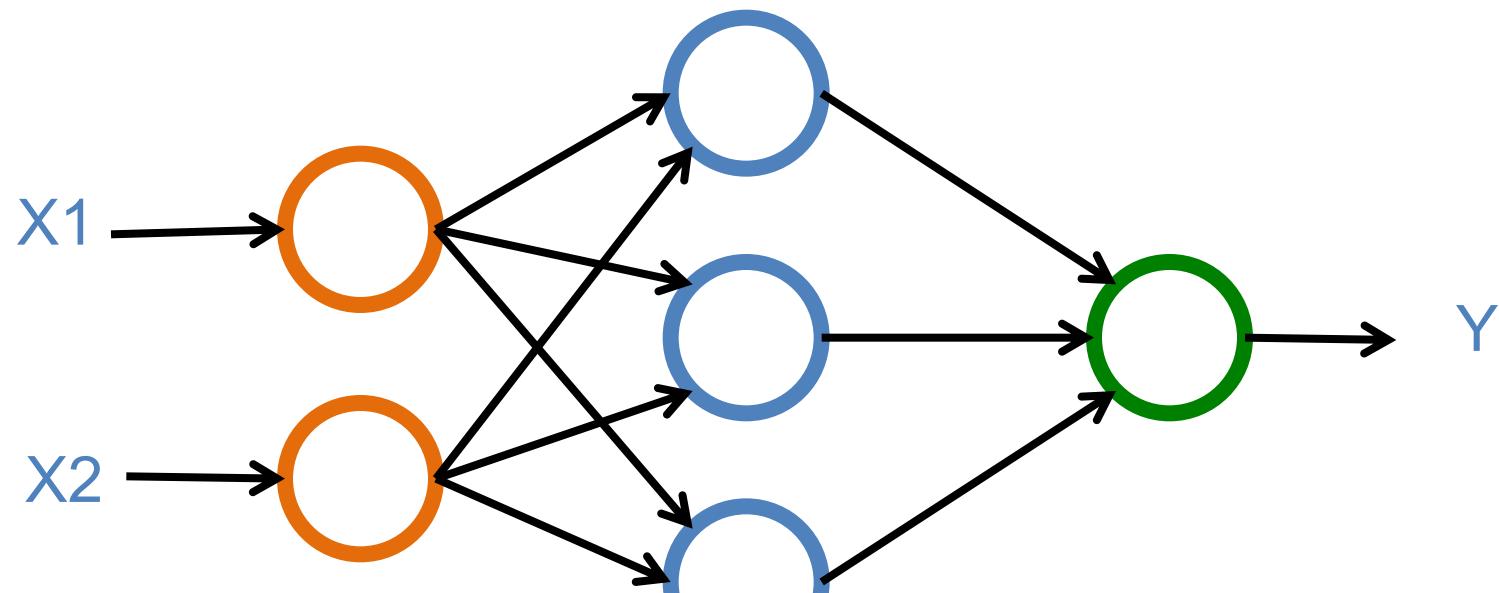


Deep Learning and Neural Networks

Deep Learning Foundations: Neural Networks

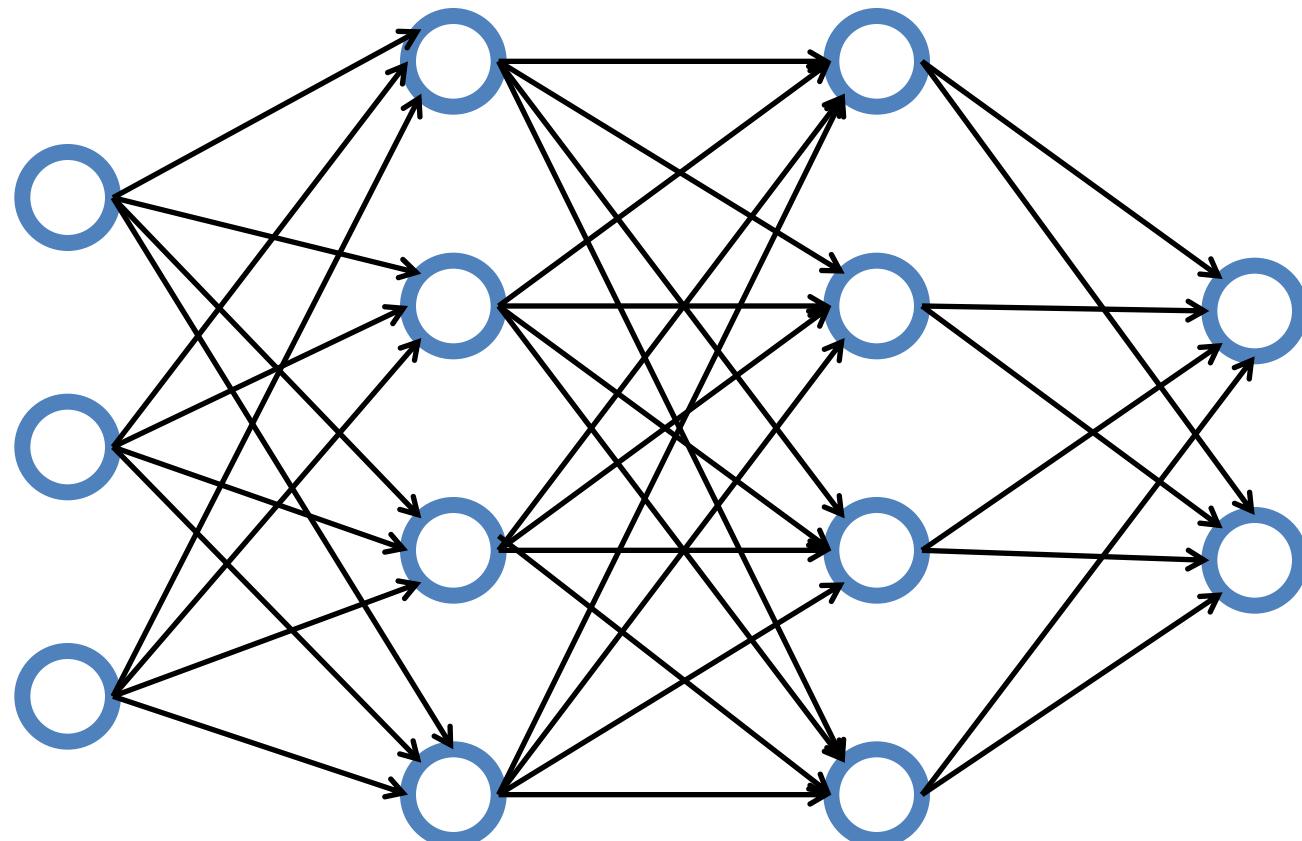
Deep Learning and Neural Networks

Input Layer Hidden Layer **Output Layer**
(X) (H) (Y)



Deep Learning and Neural Networks

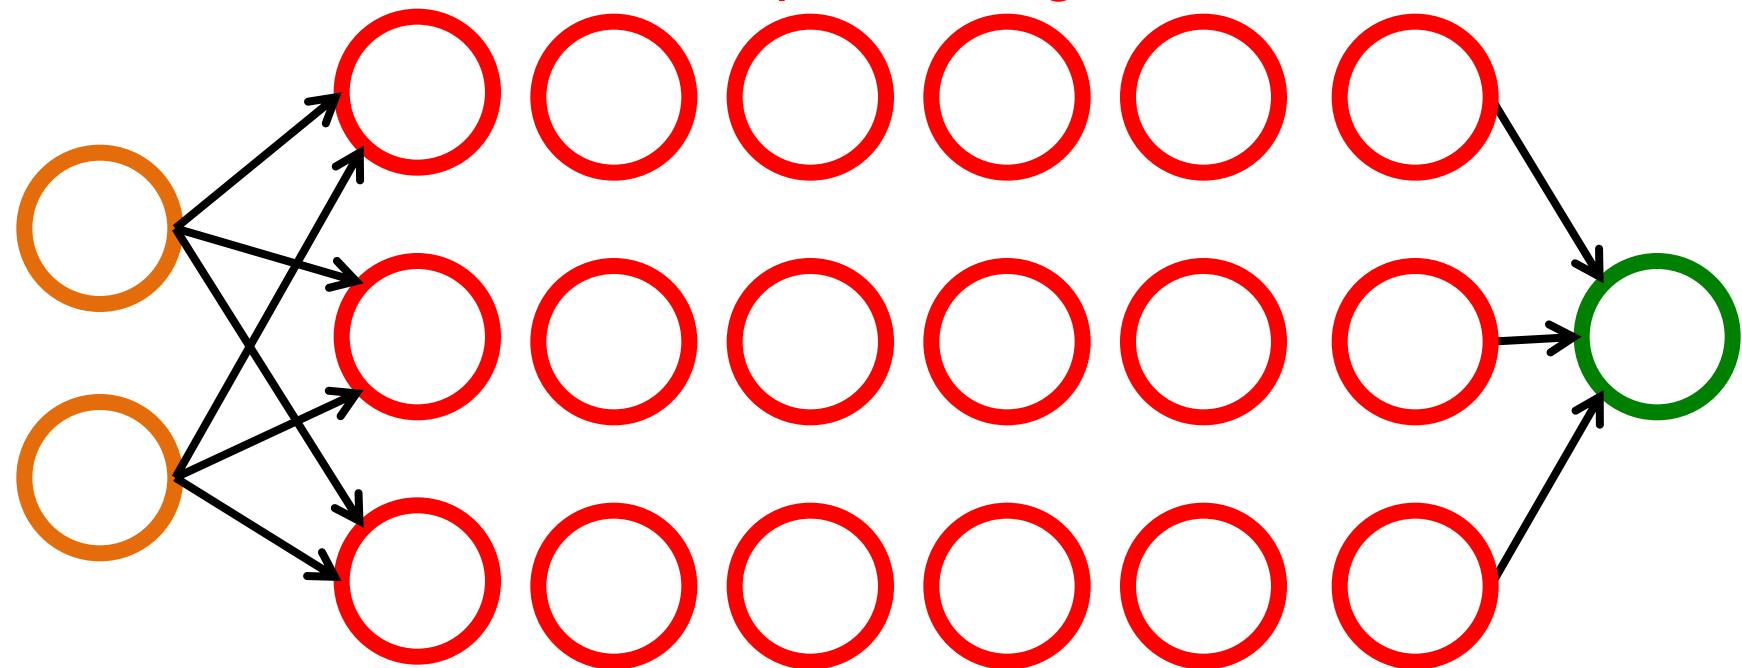
Input Layer (X) Hidden Layer (H) Output Layer (Y)



Deep Learning and Neural Networks

Input Layer (X) Hidden Layers (H) Output Layer (Y)

Deep Neural Networks
Deep Learning



Deep Learning and Deep Neural Networks

LeCun, Yann,
Yoshua Bengio,
and Geoffrey Hinton.

"Deep learning."
Nature 521, no. 7553 (2015): 436-
444.

Deep learning

Yann LeCun^{1,2}, Yoshua Bengio³ & Geoffrey Hinton^{4,5}

Deep learning allows computational models that are composed of multiple processing layers to learn representations of data with multiple levels of abstraction. These methods have dramatically improved the state-of-the-art in speech recognition, visual object recognition, object detection and many other domains such as drug discovery and genomics. Deep learning discovers intricate structure in large data sets by using the backpropagation algorithm to indicate how a machine should change its internal parameters that are used to compute the representation in each layer from the representation in the previous layer. Deep convolutional nets have brought about breakthroughs in processing images, video, speech and audio, whereas recurrent nets have shone light on sequential data such as text and speech.

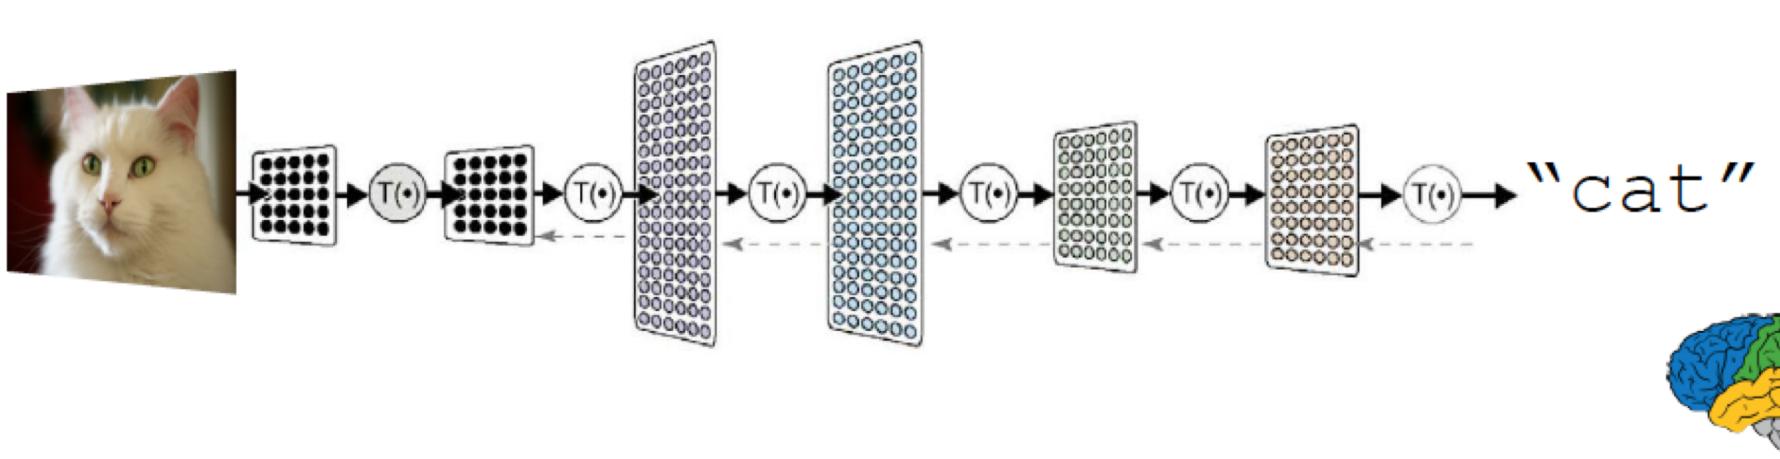
Machine-learning technology powers many aspects of modern society: from web searches to content filtering on social networks to recommendations on e-commerce websites, and it is increasingly present in consumer products such as cameras and smartphones. Machine-learning systems are used to identify objects in images, transcribe speech into text, match news items, posts or products with users' interests, and select relevant results of search. Increasingly, these applications make use of a class of techniques called deep learning.

Conventional machine-learning techniques were limited in their ability to process natural data in their raw form. For decades, con-

intricate structures in high-dimensional data and is therefore applicable to many domains of science, business and government. In addition to beating records in image recognition^{1–4} and speech recognition^{5–7}, it has beaten other machine-learning techniques at predicting the activity of potential drug molecules⁸, analysing particle accelerator data^{9,10}, reconstructing brain circuits¹¹, and predicting the effects of mutations in non-coding DNA on gene expression and disease^{12,13}. Perhaps more surprisingly, deep learning has produced extremely promising results for various tasks in natural language understanding¹⁴, particularly topic classification, sentiment analysis, question answering¹⁵ and language translation^{16,17}.

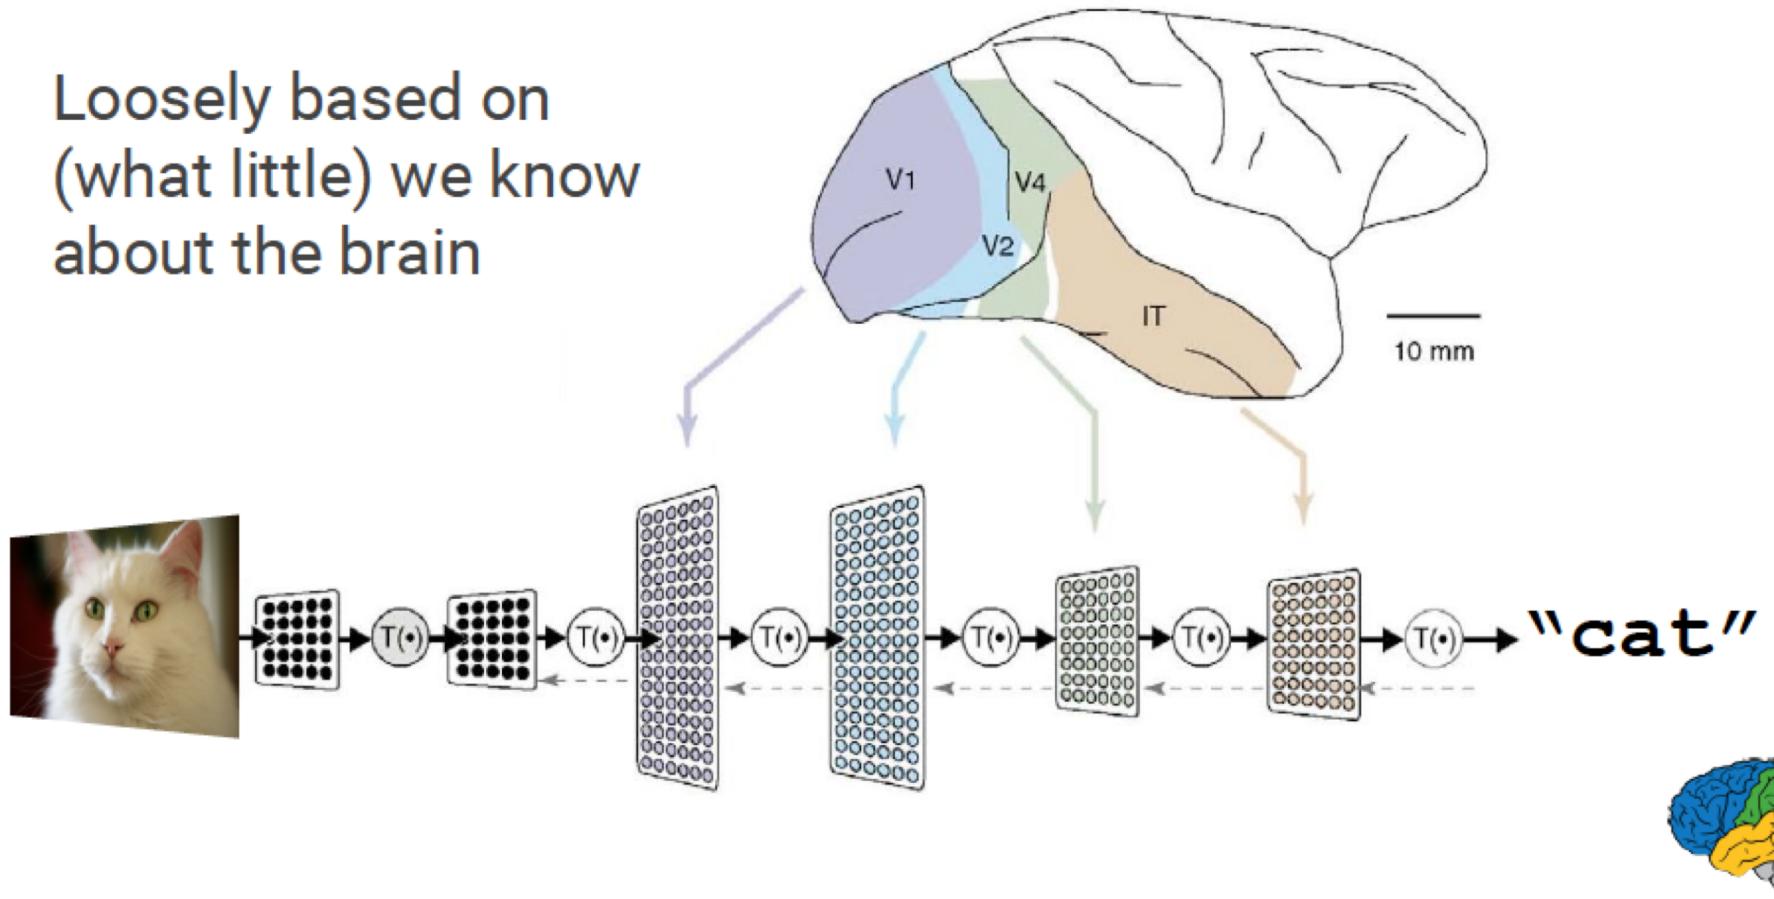
Deep Learning

- A powerful class of machine learning model
- Modern reincarnation of artificial neural networks
- Collection of simple, trainable mathematical functions
- Compatible with many variants of machine learning

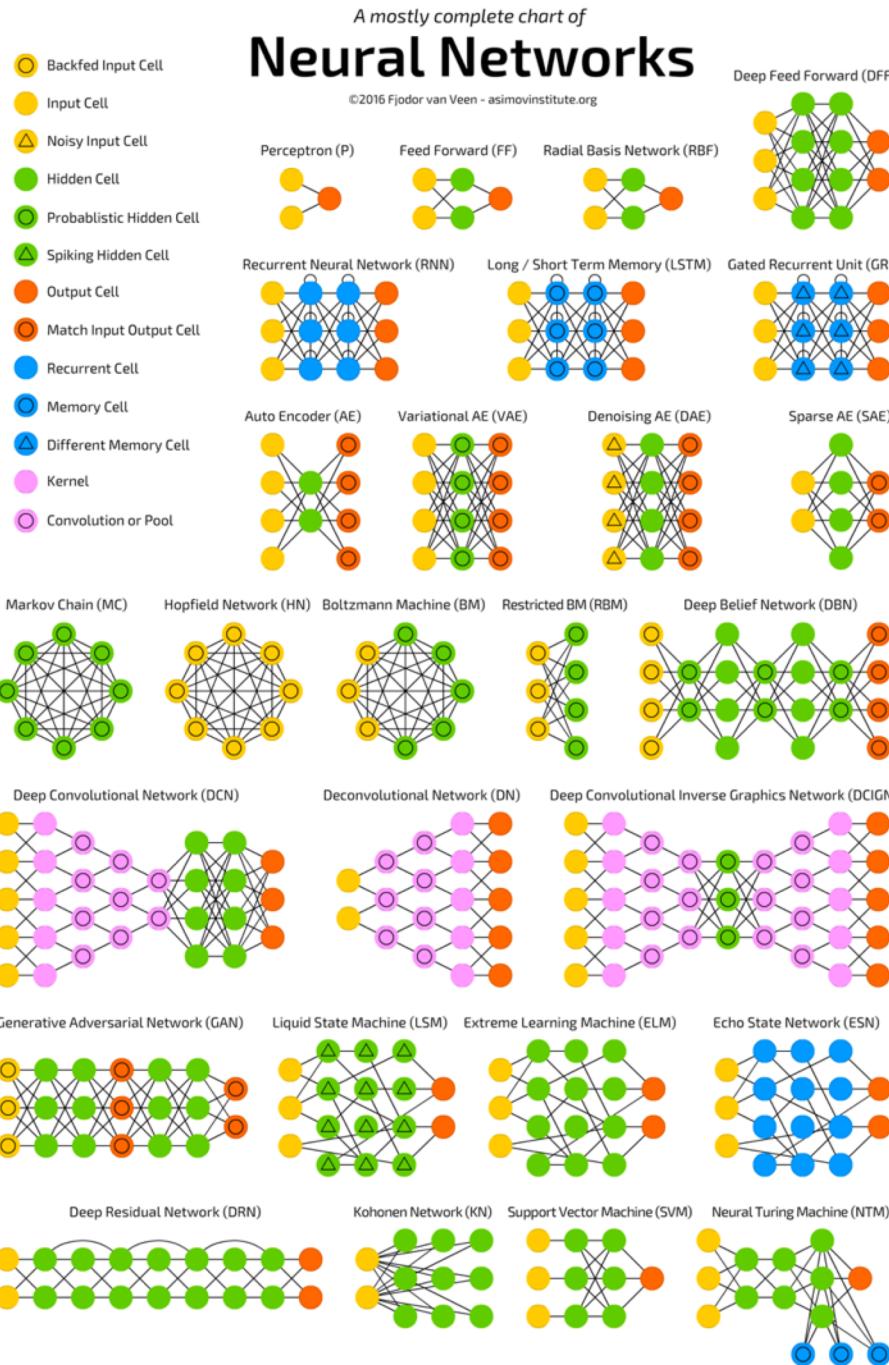


What is Deep Learning?

- Loosely based on (what little) we know about the brain



Neural Networks (NN)

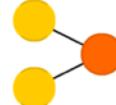


A mostly complete chart of
Neural Networks

©2016 Fjodor van Veen - asimovinstitute.org

- (○) Backfed Input Cell
- (○) Input Cell
- (△) Noisy Input Cell
- (●) Hidden Cell
- (○) Probabilistic Hidden Cell
- (△) Spiking Hidden Cell
- (○) Output Cell
- (○) Match Input Output Cell
- (●) Recurrent Cell
- (○) Memory Cell
- (△) Different Memory Cell
- (●) Kernel
- (○) Convolution or Pool

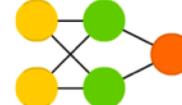
Perceptron (P)



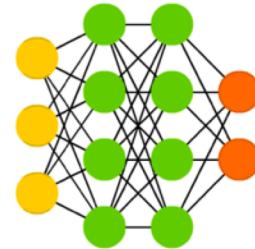
Feed Forward (FF)



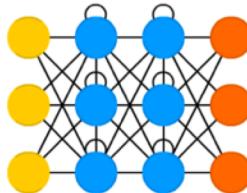
Radial Basis Network (RBF)



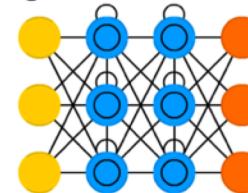
Deep Feed Forward (DFF)



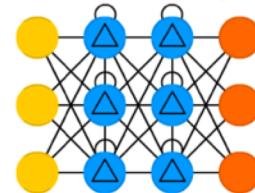
Recurrent Neural Network (RNN)



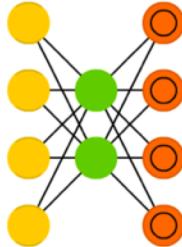
Long / Short Term Memory (LSTM)



Gated Recurrent Unit (GRU)



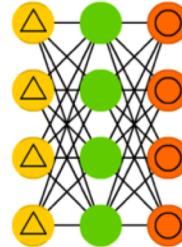
Auto Encoder (AE)



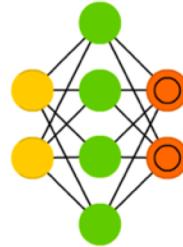
Variational AE (VAE)



Denoising AE (DAE)



Sparse AE (SAE)



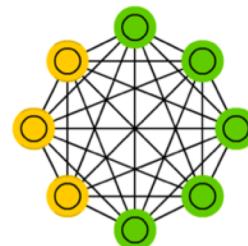
Markov Chain (MC)



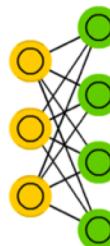
Hopfield Network (HN)



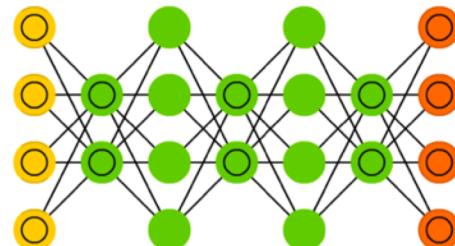
Boltzmann Machine (BM)



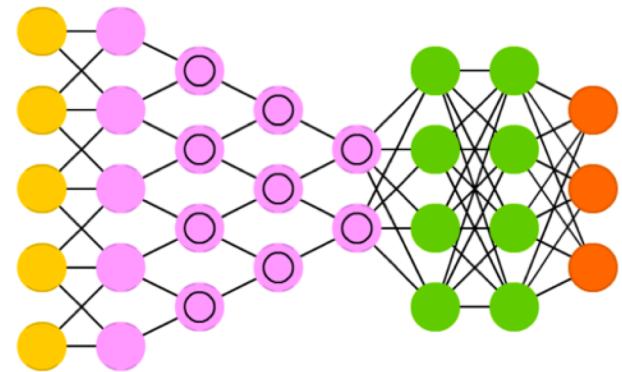
Restricted BM (RBM)



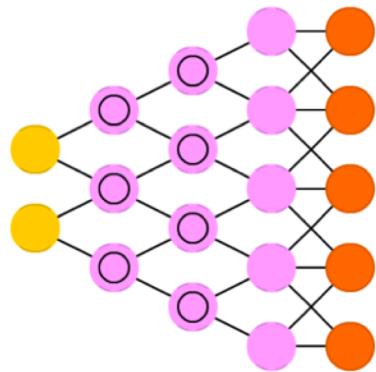
Deep Belief Network (DBN)



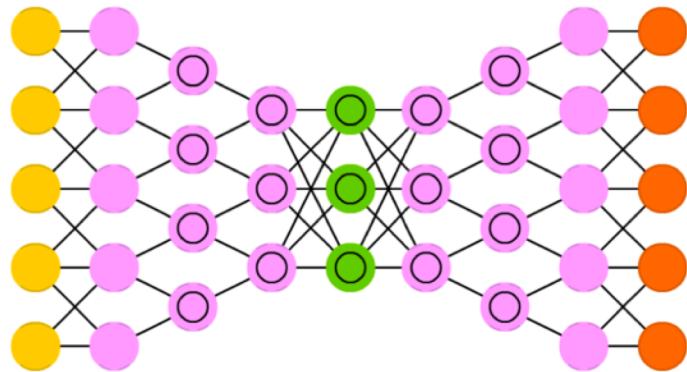
Deep Convolutional Network (DCN)



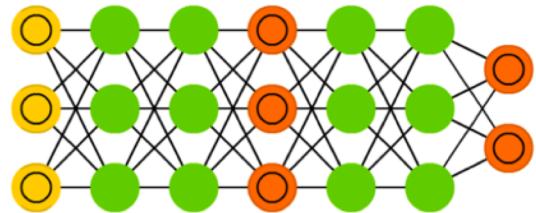
Deconvolutional Network (DN)



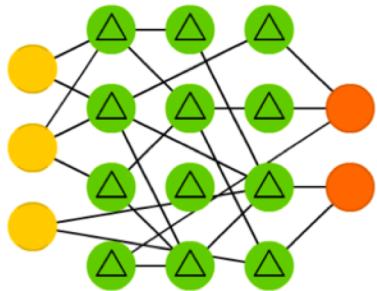
Deep Convolutional Inverse Graphics Network (DCIGN)



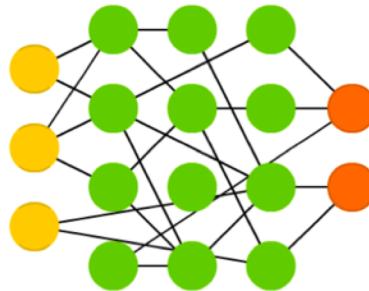
Generative Adversarial Network (GAN)



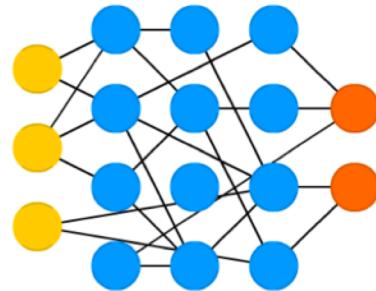
Liquid State Machine (LSM)



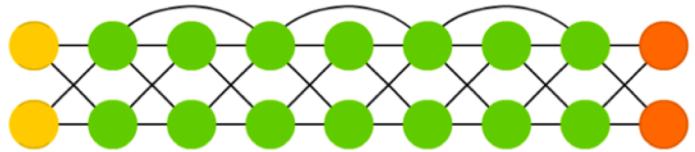
Extreme Learning Machine (ELM)



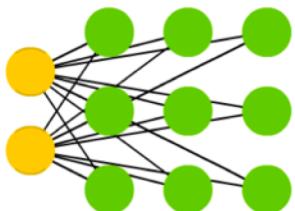
Echo State Network (ESN)



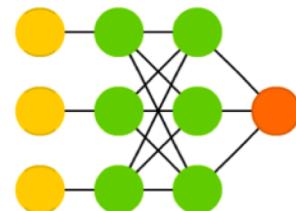
Deep Residual Network (DRN)



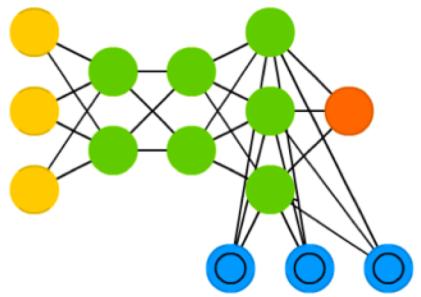
Kohonen Network (KN)



Support Vector Machine (SVM)

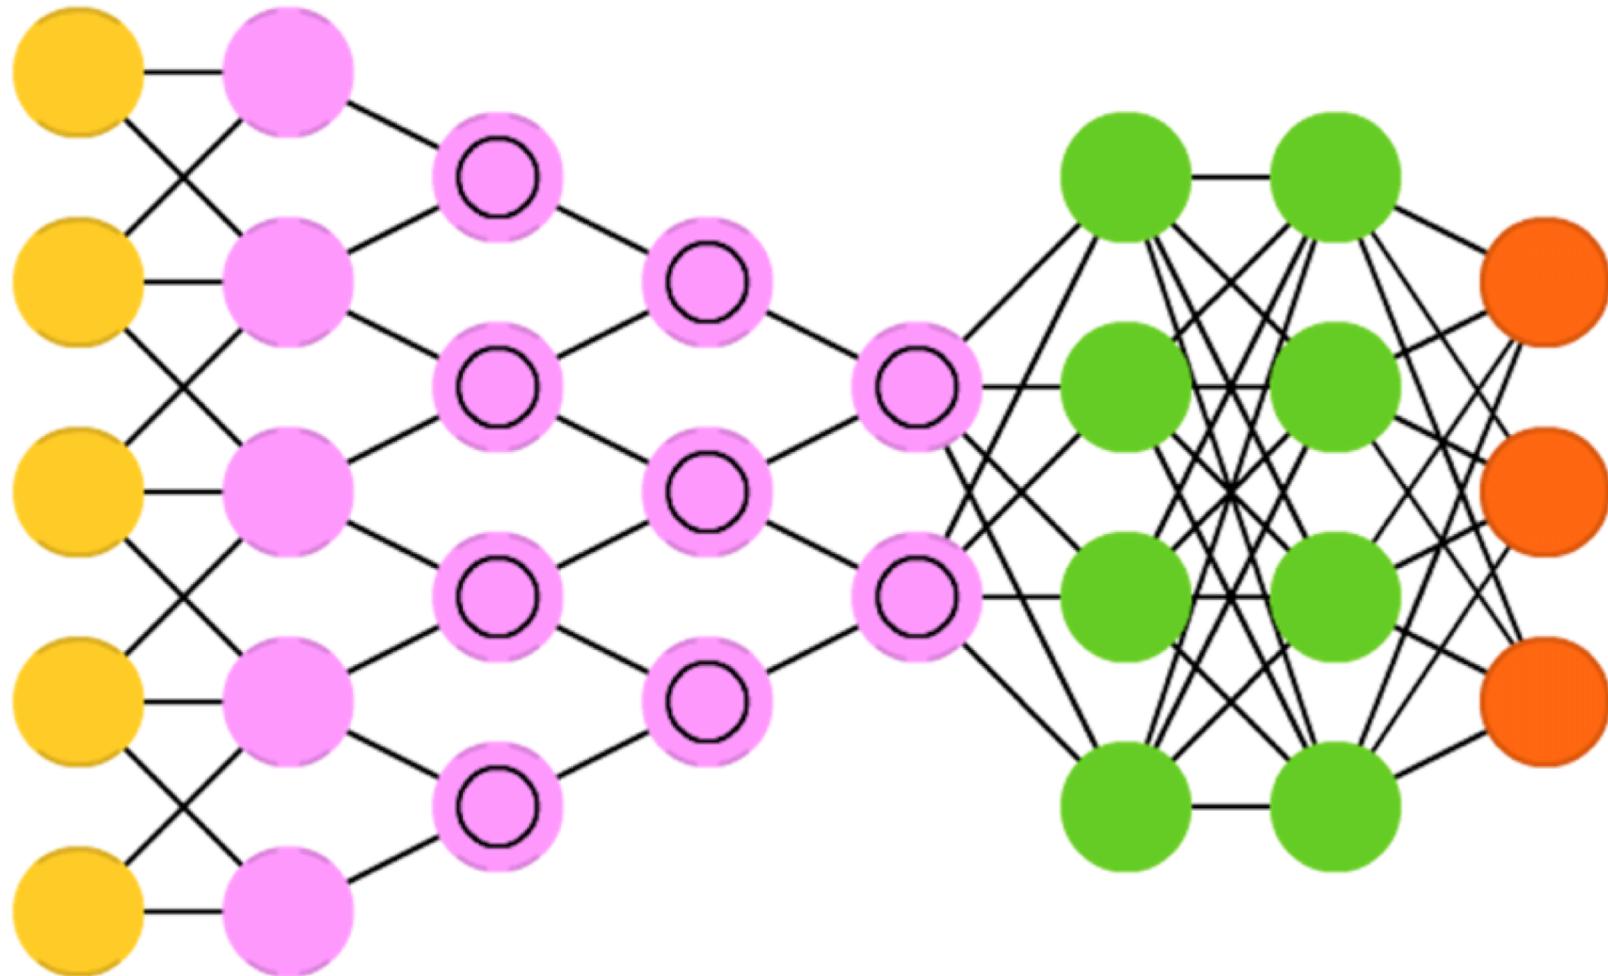


Neural Turing Machine (NTM)

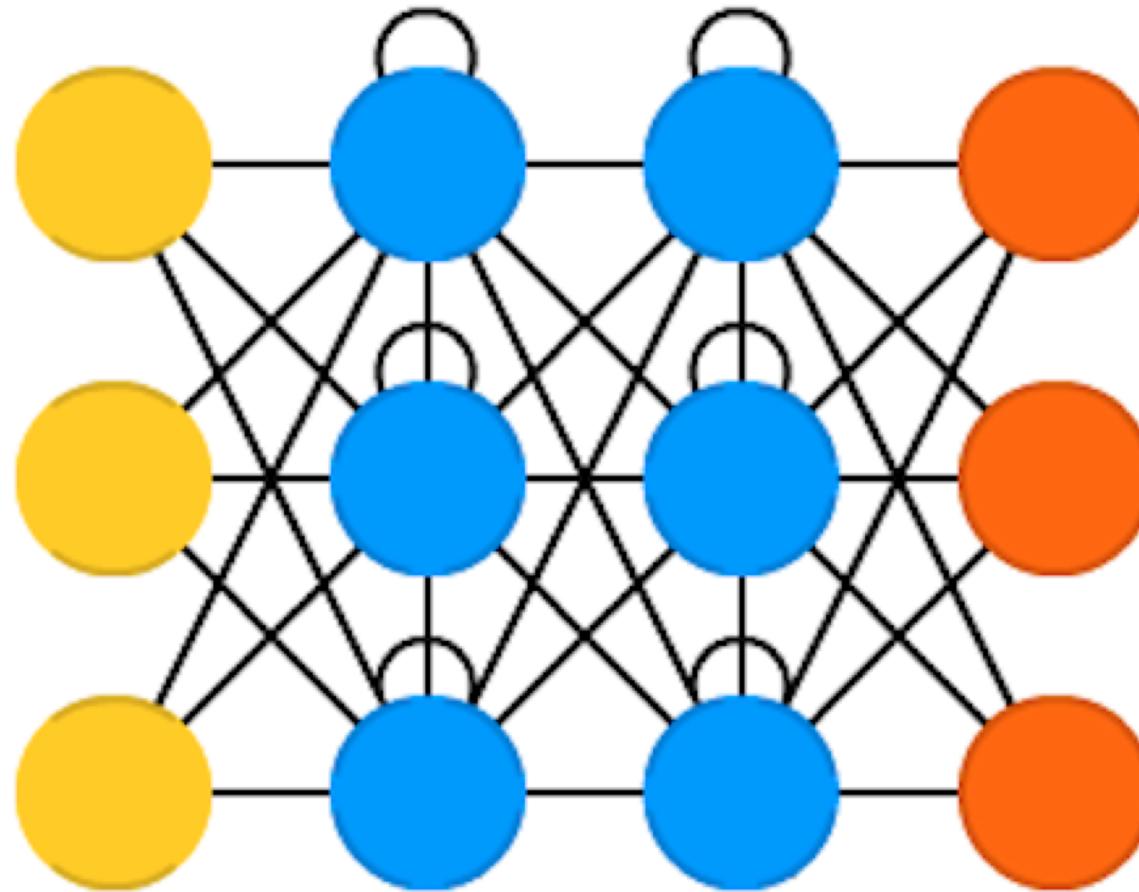


Convolutional Neural Networks

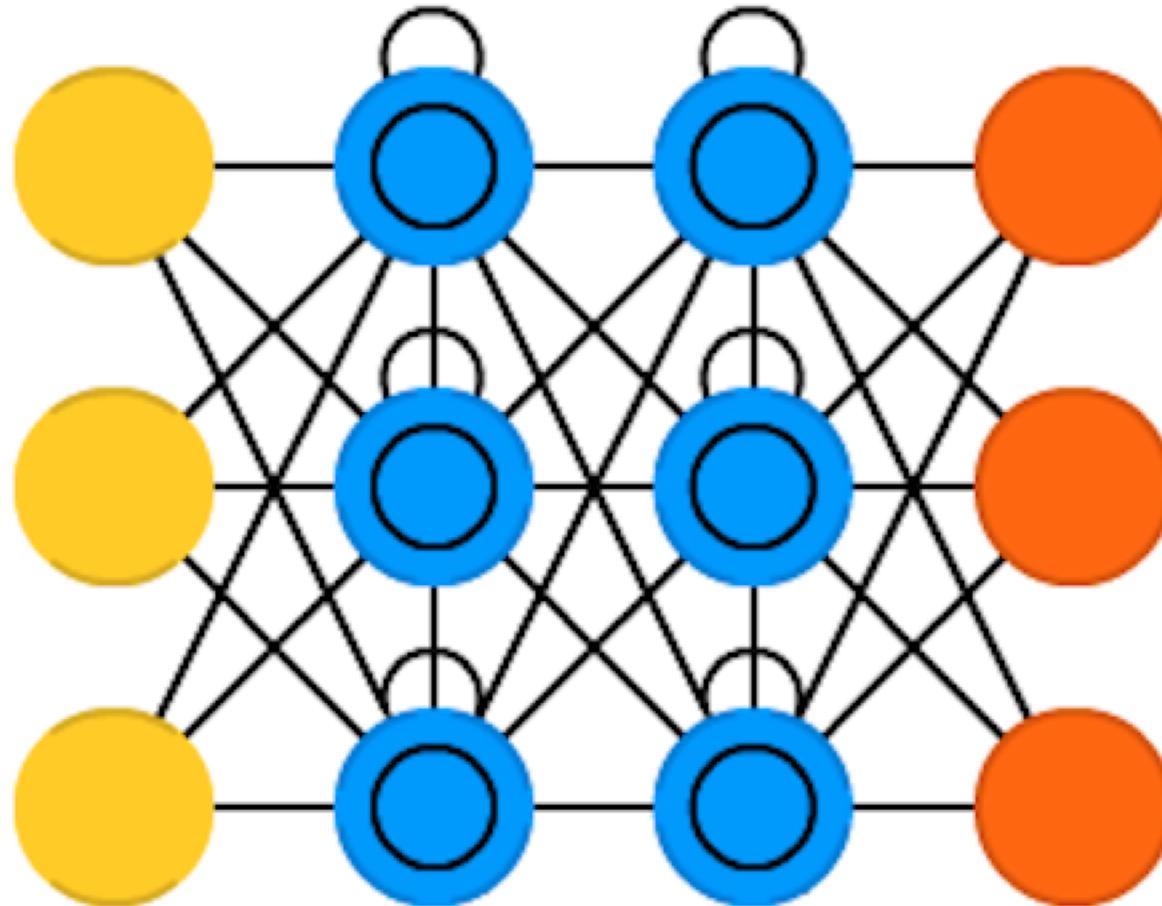
(CNN or Deep Convolutional Neural Networks, DCNN)



Recurrent Neural Networks (RNN)



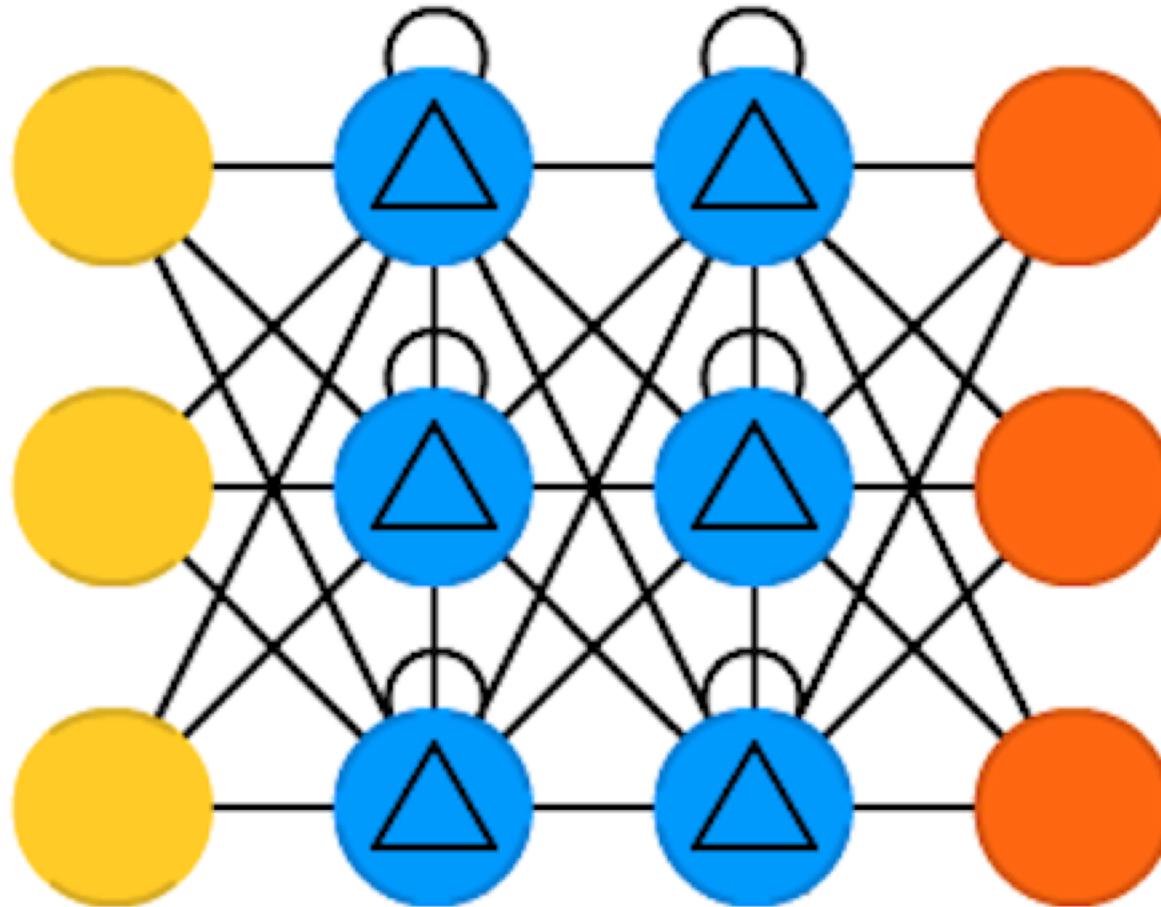
Long / Short Term Memory (LSTM)



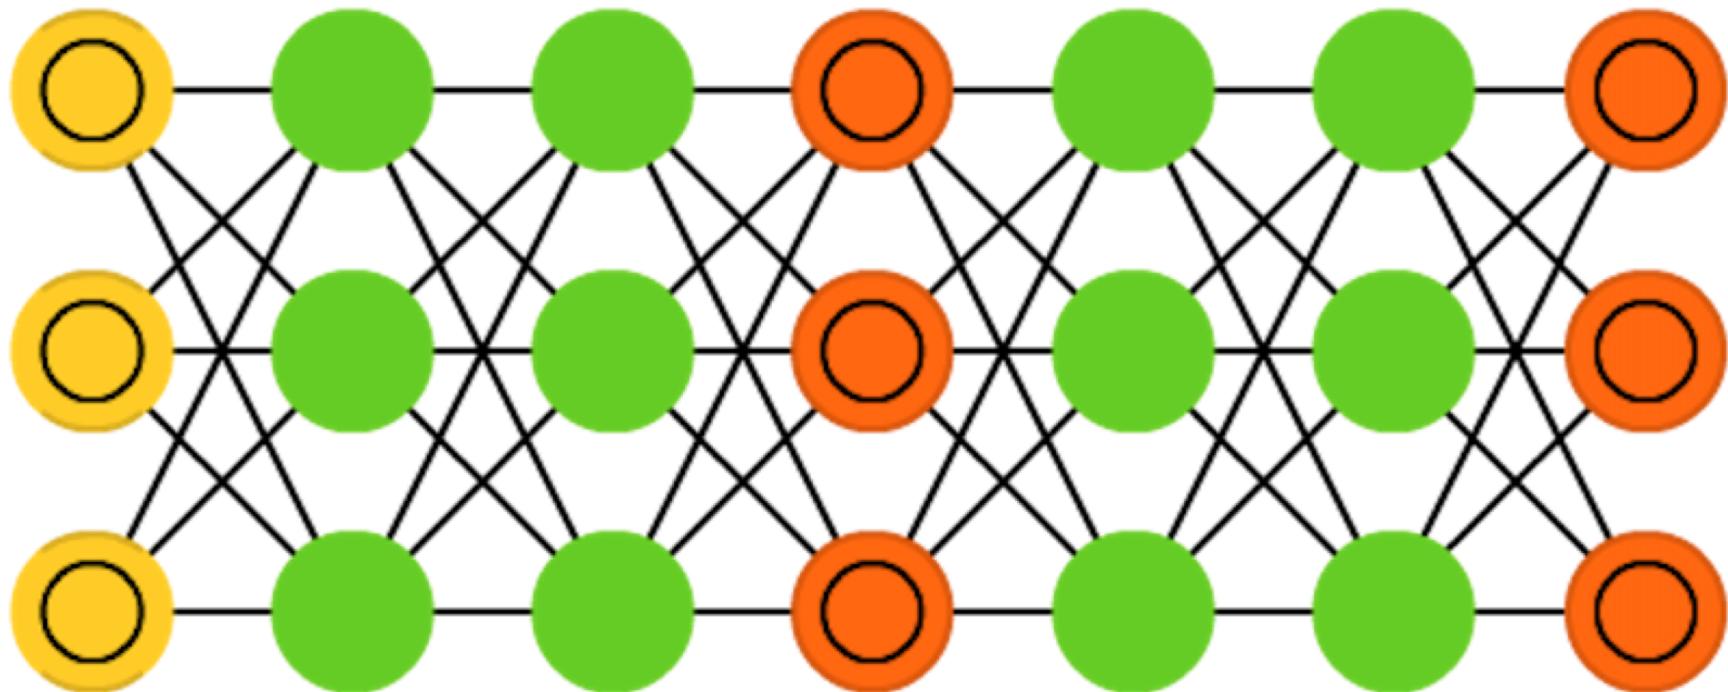
Hochreiter, Sepp, and Jürgen Schmidhuber. "Long short-term memory." Neural computation 9.8 (1997): 1735-1780.

Source: <http://www.asimovinstitute.org/neural-network-zoo/>

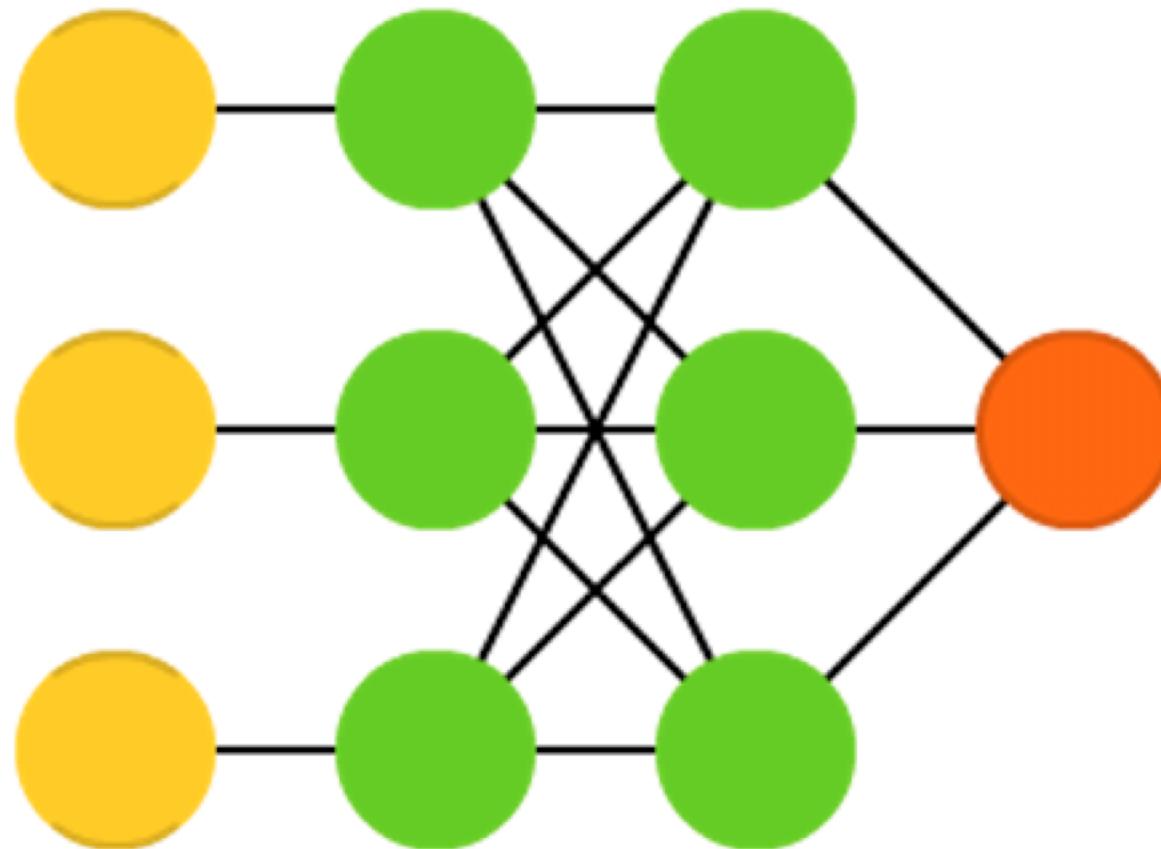
Gated Recurrent Units (GRU)



Generative Adversarial Networks (GAN)



Support Vector Machines (SVM)

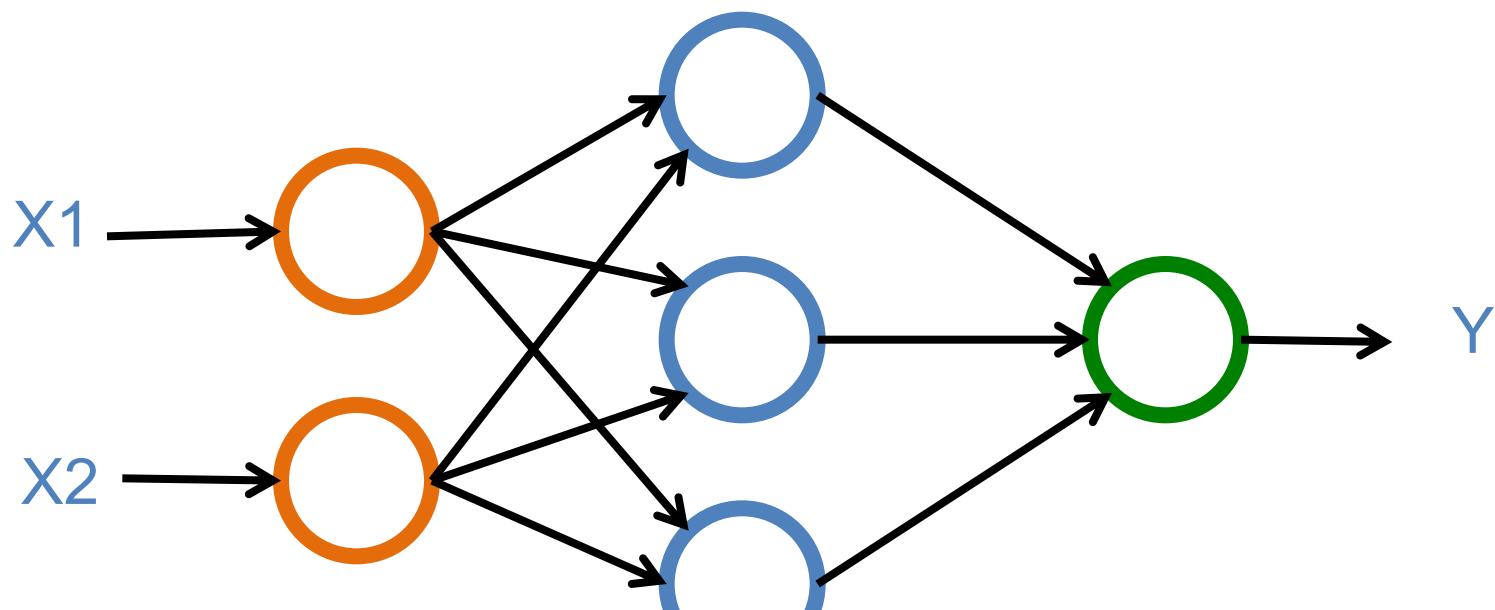


Cortes, Corinna, and Vladimir Vapnik. "Support-vector networks." *Machine learning* 20.3 (1995): 273-297.

**Neural networks
(NN)
1960**

Neural Networks

Input Layer **Hidden Layer** **Output Layer**
(X) **(H)** **(Y)**



Multilayer Perceptrons (MLP)

1985

Support Vector Machine (SVM)

1995



Hinton presents the **Deep Belief Network (DBN)**

**New interests in deep learning
and RBM**

State of the art MNIST

2005

Deep Recurrent Neural Network (RNN) 2009

Convolutional DBN

2010

Max-Pooling CDBN

2011

Deep Learning

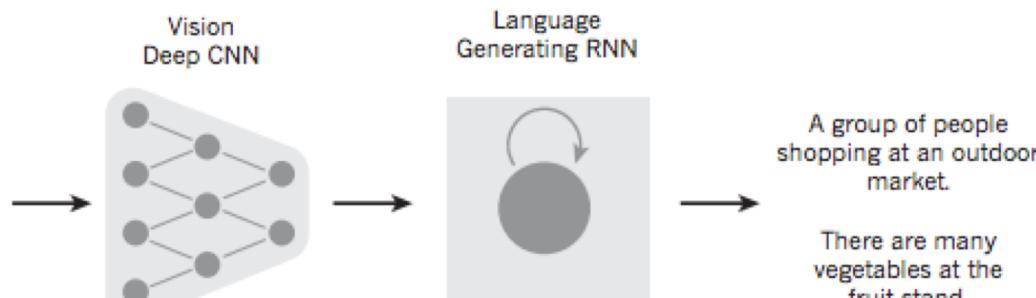
Geoffrey Hinton

Yann LeCun

Yoshua Bengio

Andrew Y. Ng

From image to text



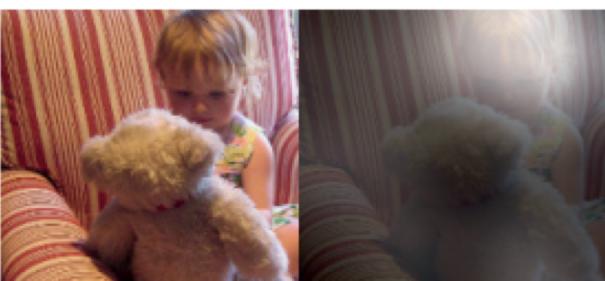
A woman is throwing a **frisbee** in a park.



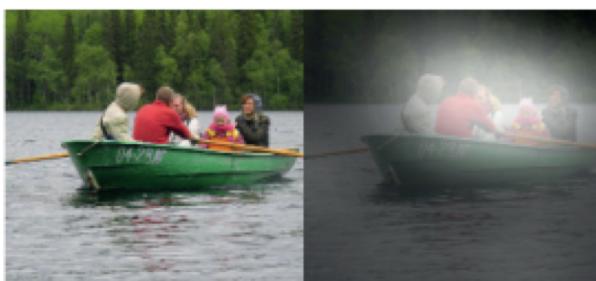
A **dog** is standing on a hardwood floor.



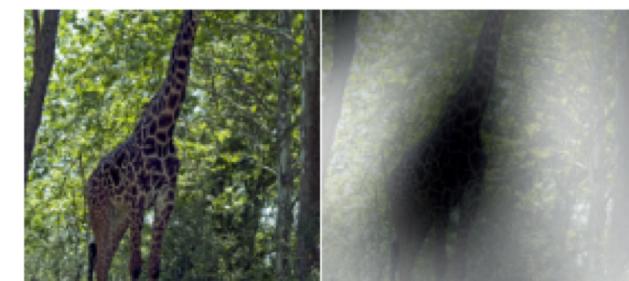
A **stop** sign is on a road with a mountain in the background



A **little girl** sitting on a bed with a **teddy bear**.



A group of **people** sitting on a boat in the water.



A **giraffe** standing in a forest with **trees** in the background.

From image to text

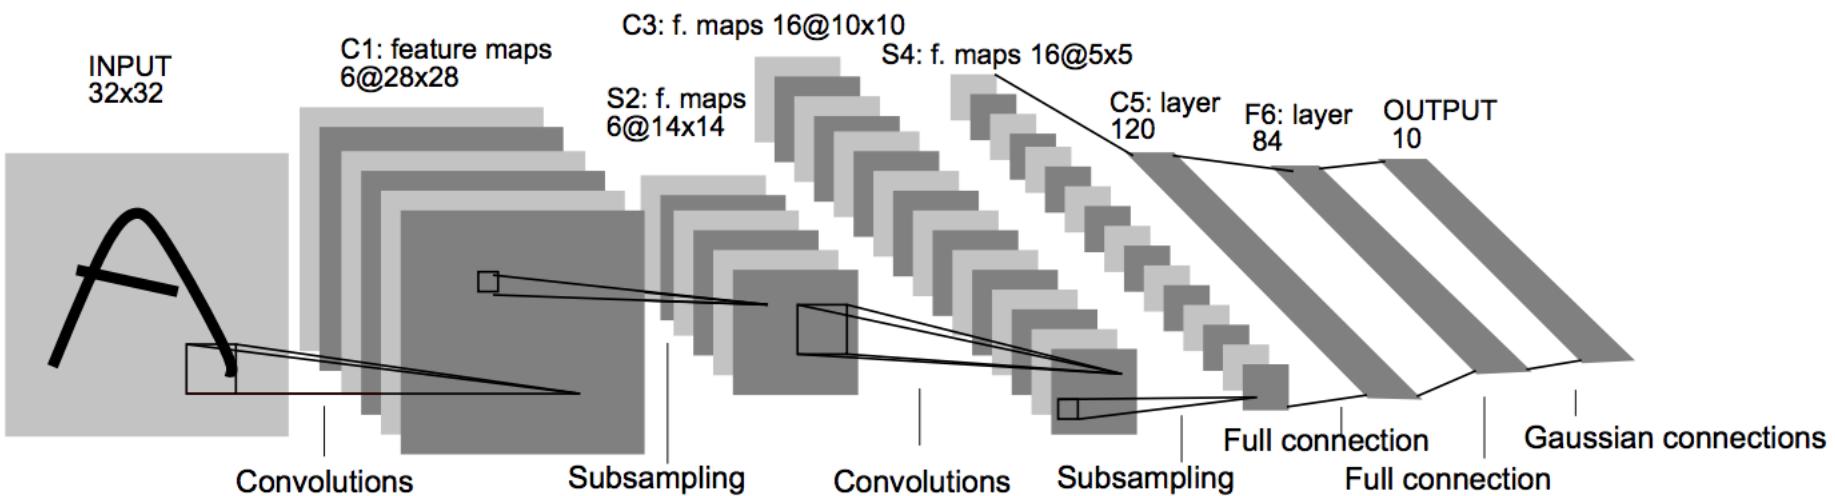
Image: deep convolution neural network (CNN)
Text: recurrent neural network (RNN)



A group of **people** sitting on a boat in the water.

Convolutional Neural Networks (CNN)

Convolutional Neural Networks (CNN)



Architecture of LeNet-5 (7 Layers) (LeCun et al., 1998)

Source: <http://yann.lecun.com/exdb/publis/pdf/lecun-01a.pdf>

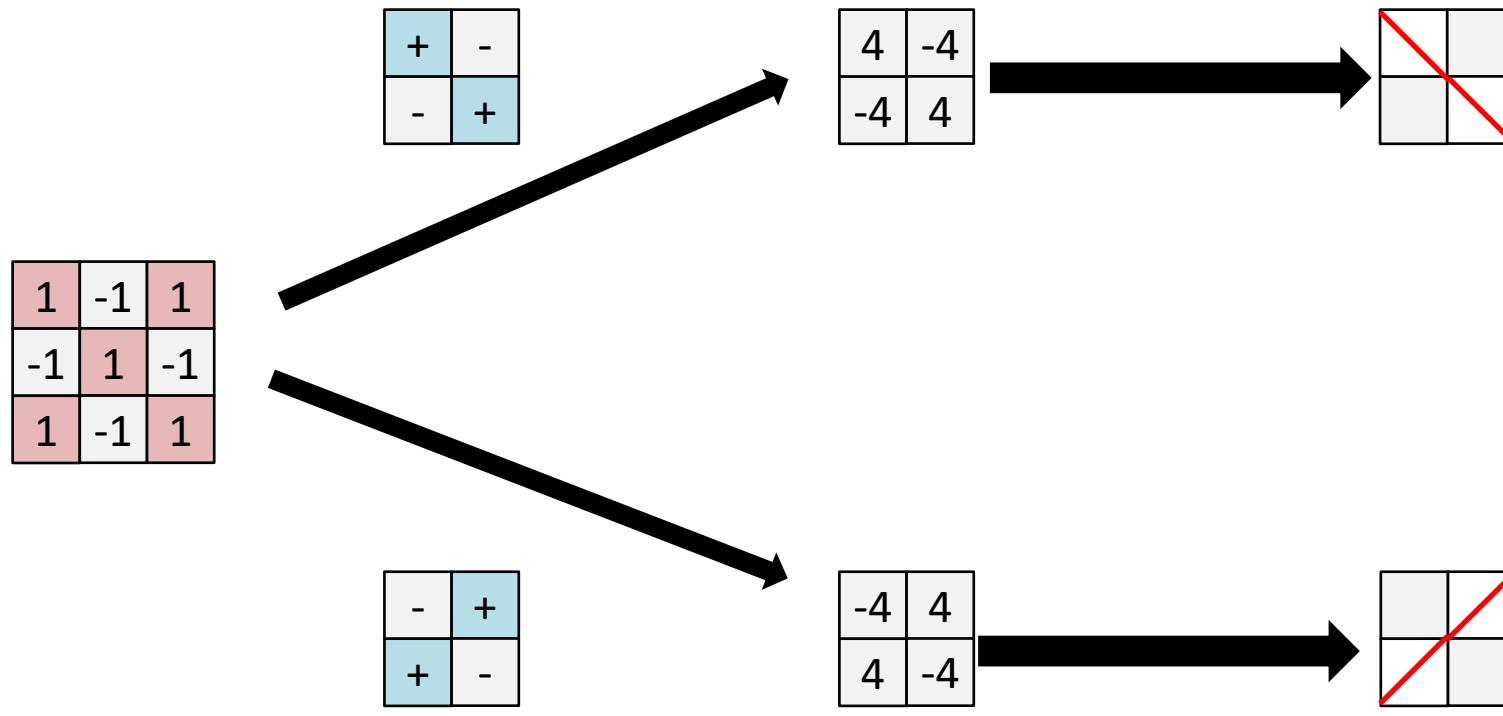
Source: LeCun, Yann, Léon Bottou, Yoshua Bengio, and Patrick Haffner.

"Gradient-based learning applied to document recognition." *Proceedings of the IEEE* 86, no. 11 (1998): 2278-2324.

Convolutional Neural Networks (CNN)

- Convolution
- Pooling
- Fully Connection (FC) (Flattening)

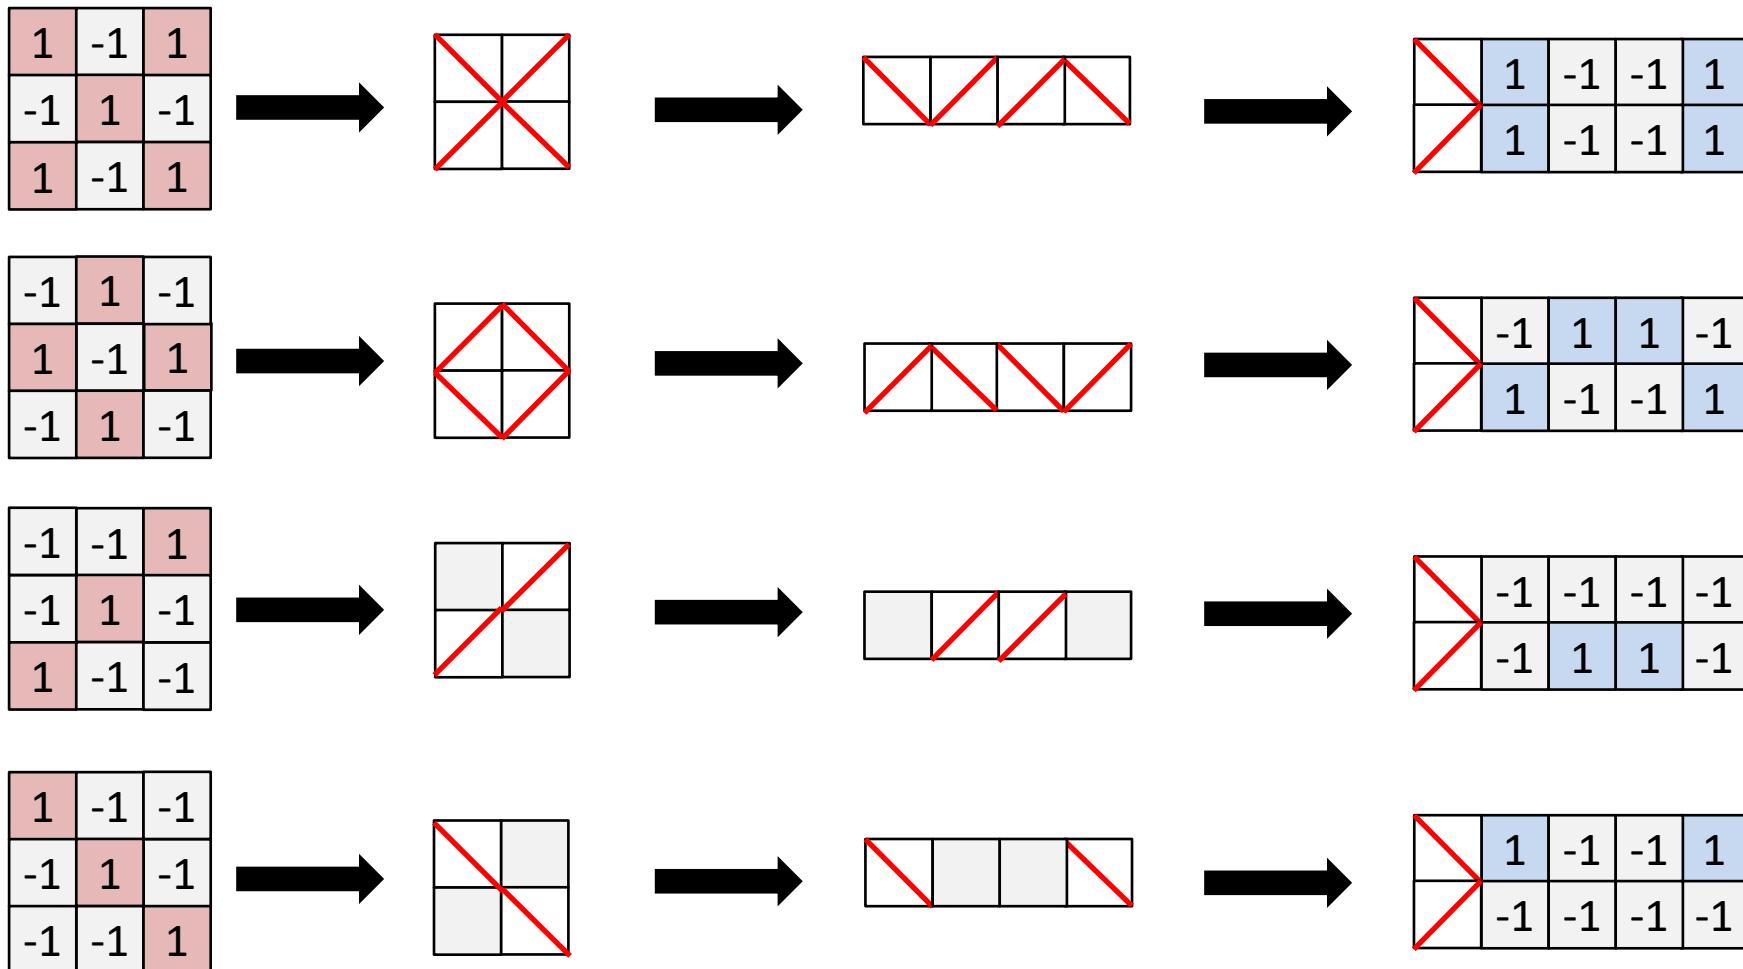
A friendly introduction to Convolutional Neural Networks and Image Recognition



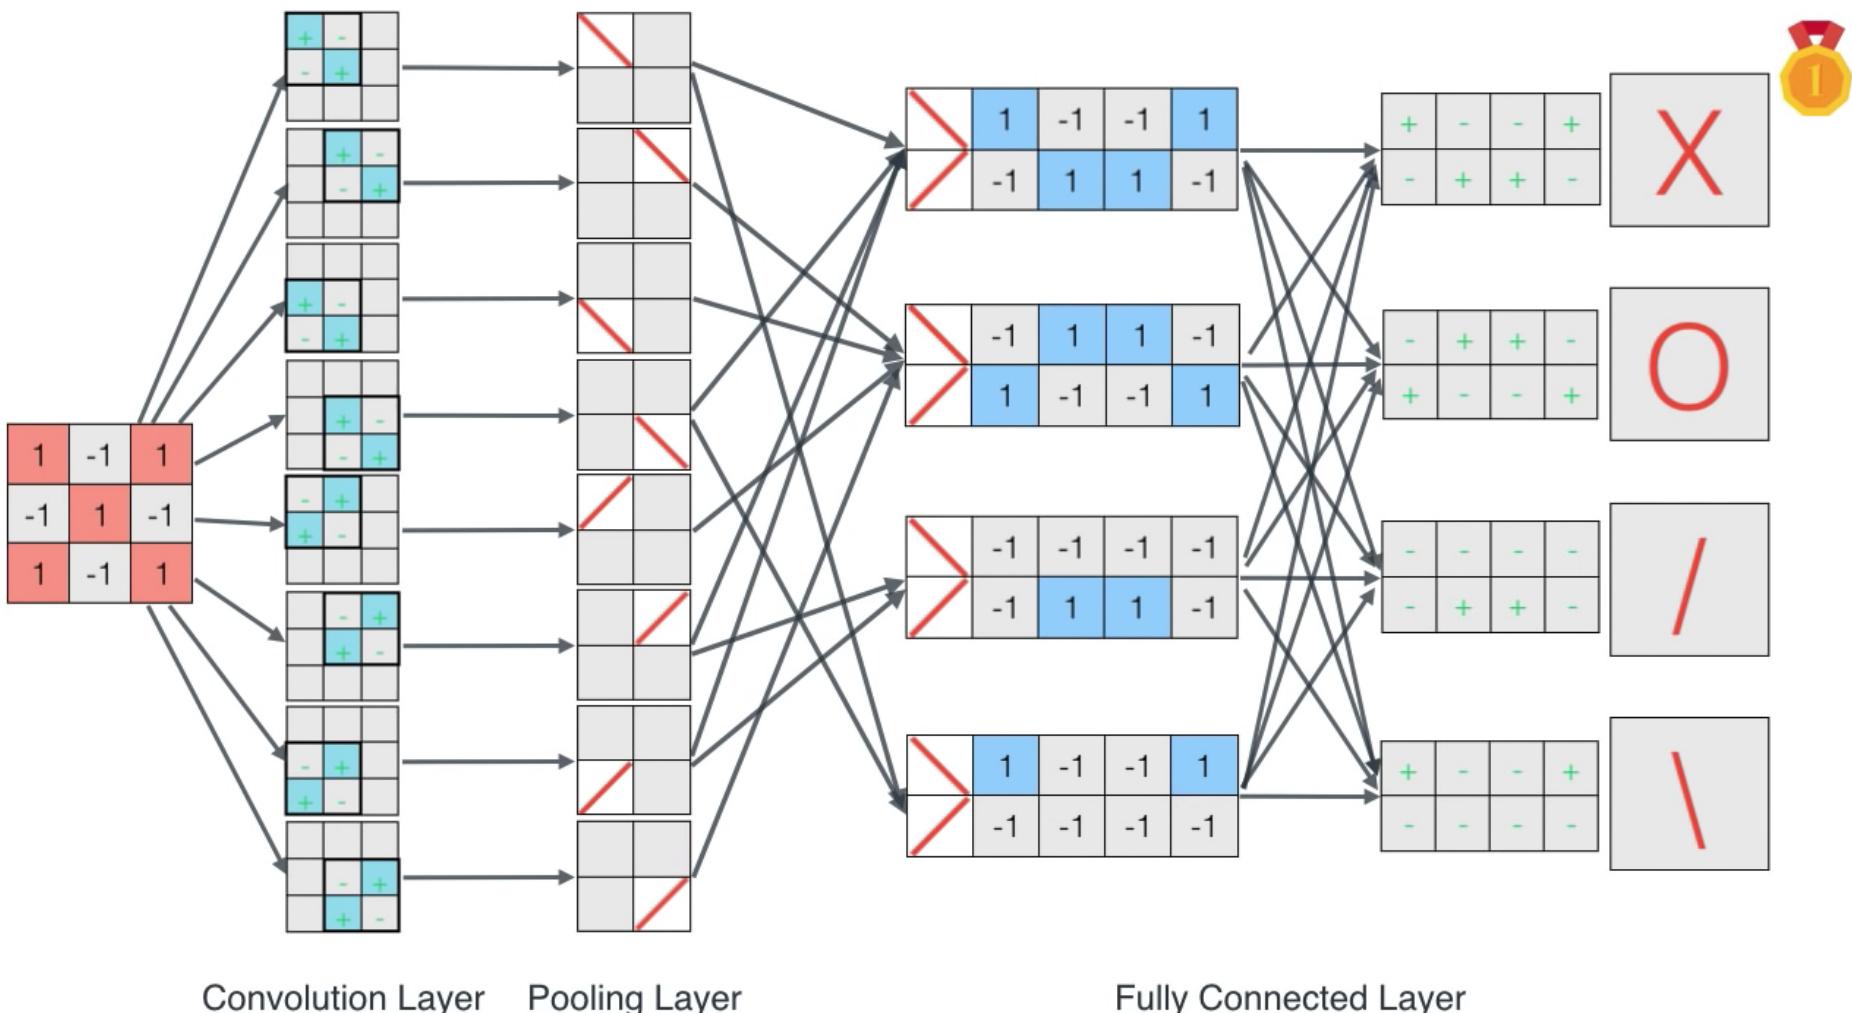
Convolution Layer

Pooling Layer

A friendly introduction to Convolutional Neural Networks and Image Recognition



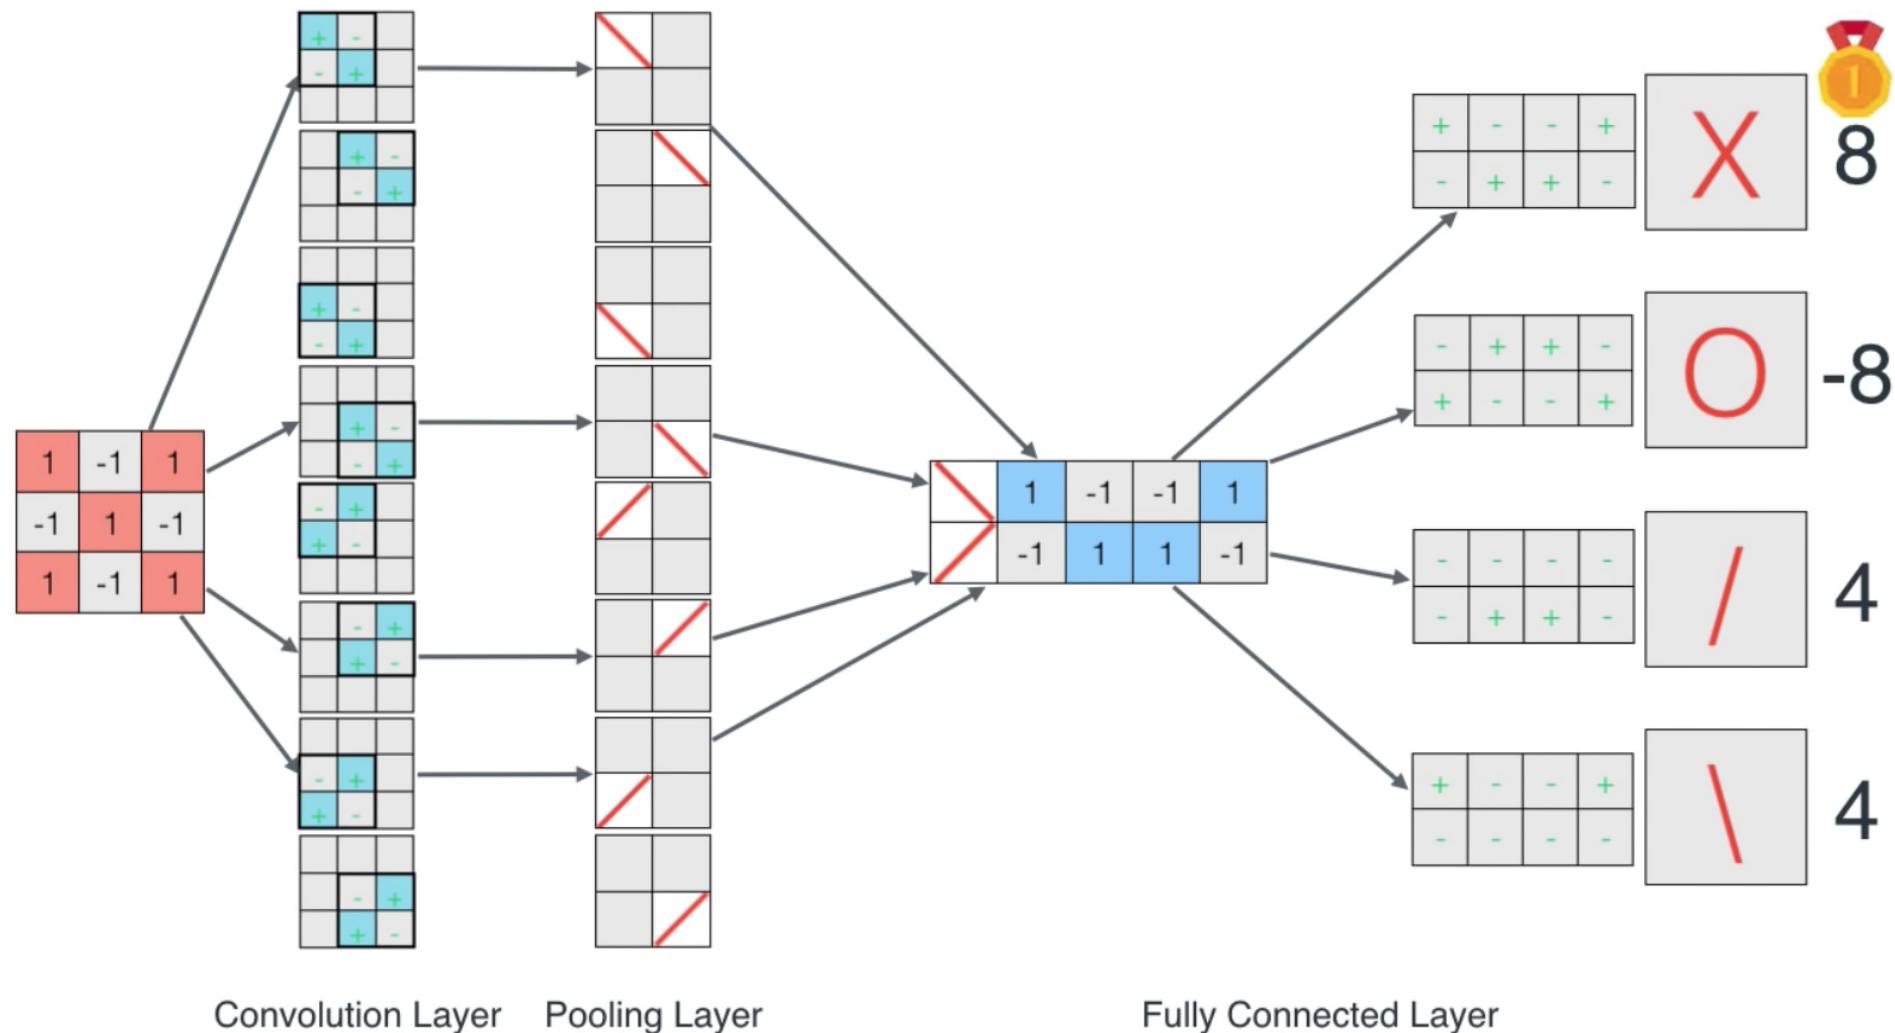
A friendly introduction to Convolutional Neural Networks and Image Recognition



Convolution Layer Pooling Layer

Fully Connected Layer

A friendly introduction to Convolutional Neural Networks and Image Recognition

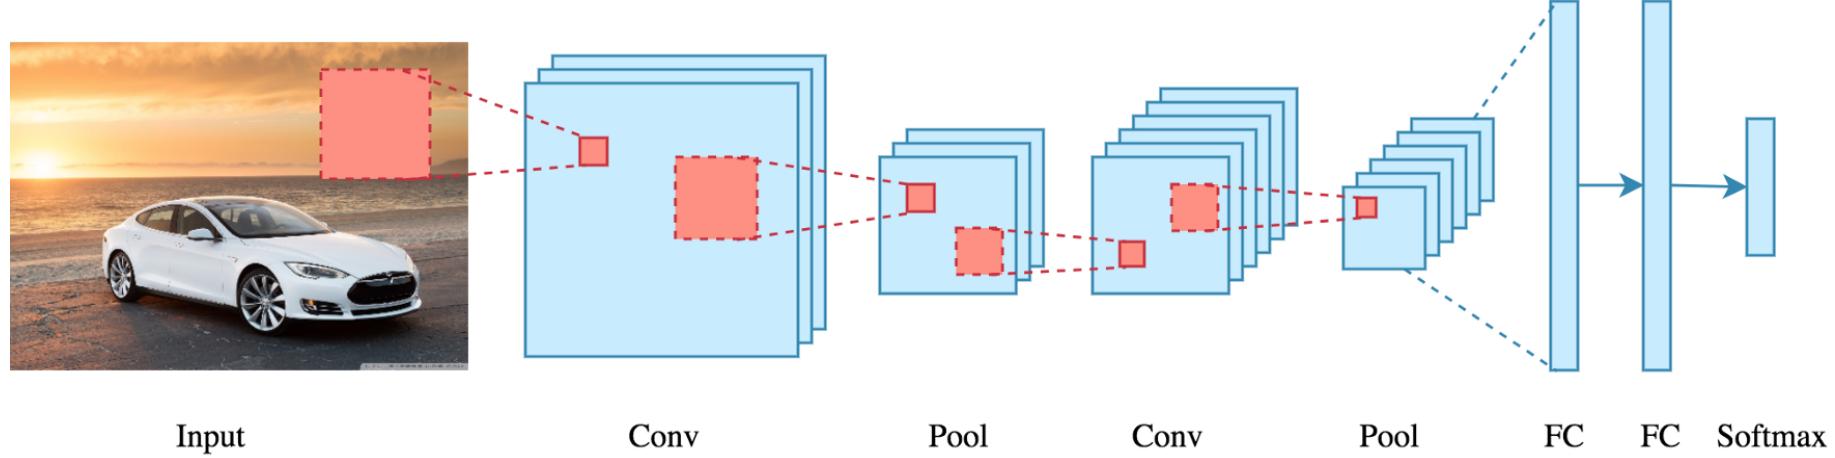


Convolution Layer

Pooling Layer

Fully Connected Layer

CNN Architecture



Source: Arden Dertat (2017), Applied Deep Learning - Part 4: Convolutional Neural Networks,
<https://towardsdatascience.com/applied-deep-learning-part-4-convolutional-neural-networks-584bc134c1e2>

CNN Convolution Layer

Convolution is a mathematical operation to merge two sets of information

3x3 convolution

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

Input

1	0	1
0	1	0
1	0	1

Filter / Kernel

CNN Convolution Layer

Input x Filter --> Feature Map

receptive field: 3x3

1x1	1x0	1x1	0	0
0x0	1x1	1x0	1	0
0x1	0x0	1x1	1	1
0	0	1	1	0
0	1	1	0	0

Input x Filter

4		

Feature Map

CNN Convolution Layer

Input x Filter --> Feature Map

receptive field: 3x3

1	1x1	1x0	0x1	0
0	1x0	1x1	1x0	0
0	0x1	1x0	1x1	1
0	0	1	1	0
0	1	1	0	0

Input x Filter

4	3	

Feature Map

CNN Convolution Layer

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

Input

1	0	1
0	1	0
1	0	1

Filter / Kernel

1x1	1x0	1x1	0	0
0x0	1x1	1x0	1	0
0x1	0x0	1x1	1	1
0	0	1	1	0
0	1	1	0	0

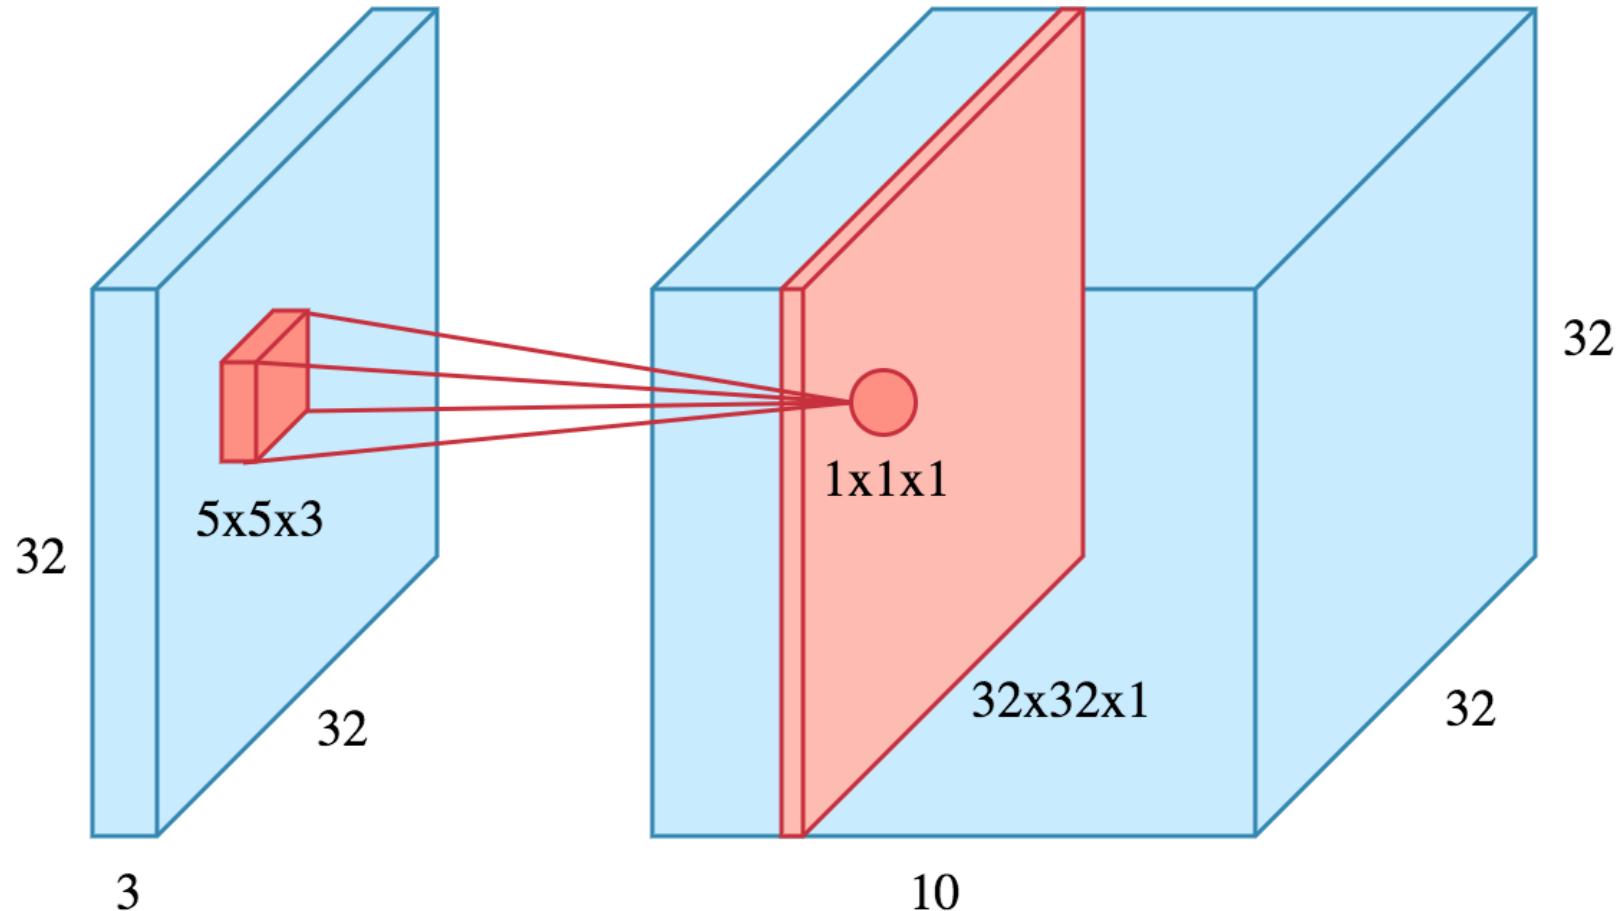
4		

Example convolution operation shown in 2D using a 3x3 filter

Source: Arden Dertat (2017), Applied Deep Learning - Part 4: Convolutional Neural Networks,
<https://towardsdatascience.com/applied-deep-learning-part-4-convolutional-neural-networks-584bc134c1e2>

CNN Convolution Layer

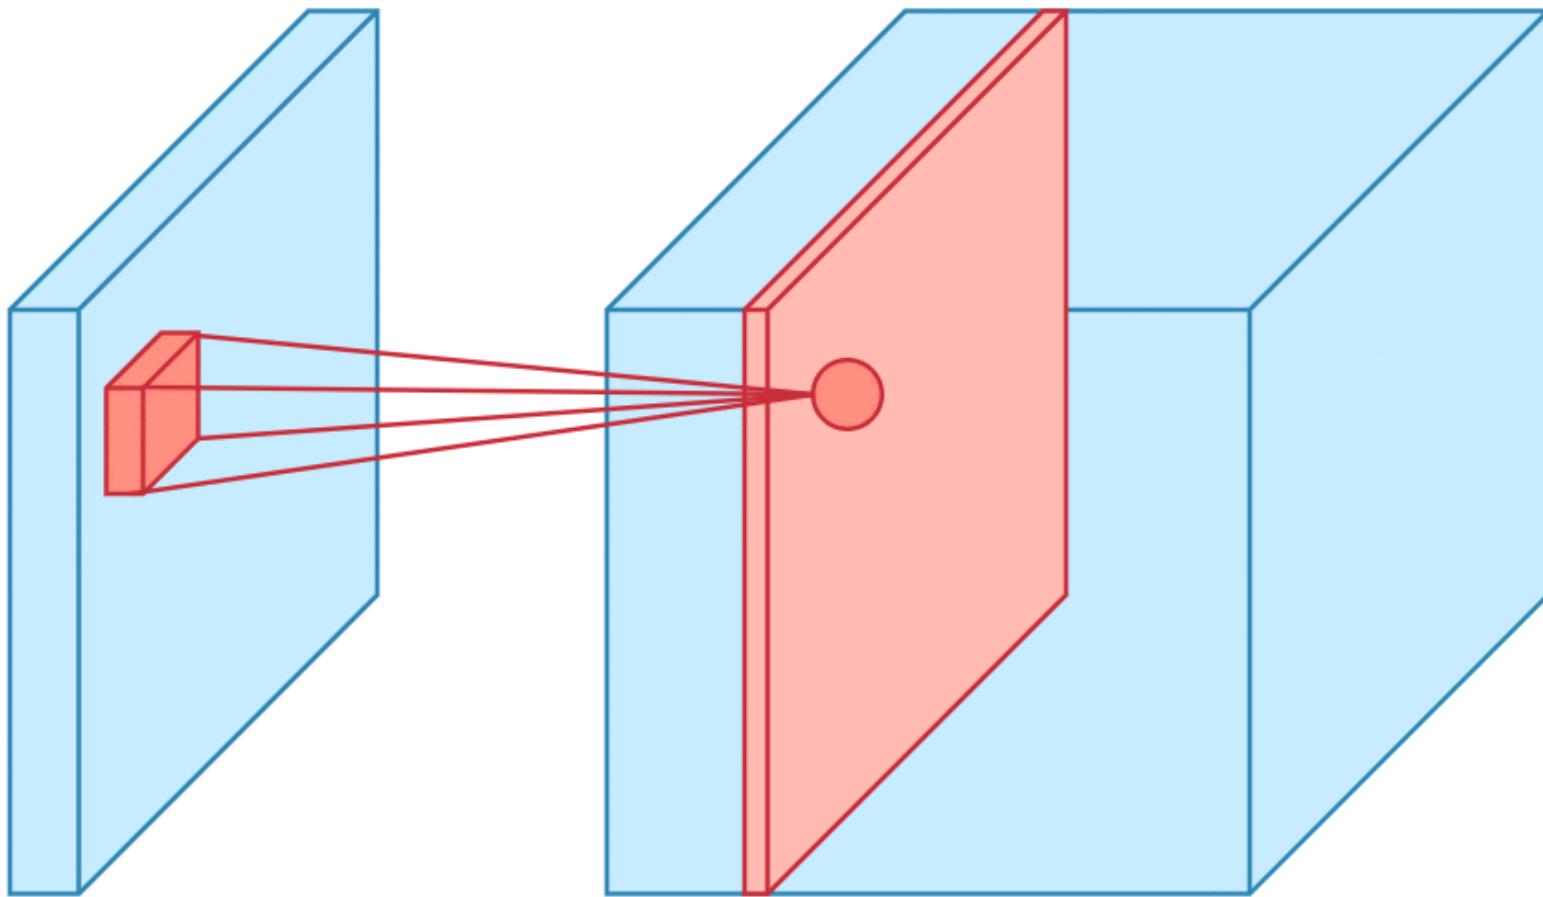
10 different filters 10 feature maps of size $32 \times 32 \times 1$



final output of the convolution layer:
a volume of size $32 \times 32 \times 10$

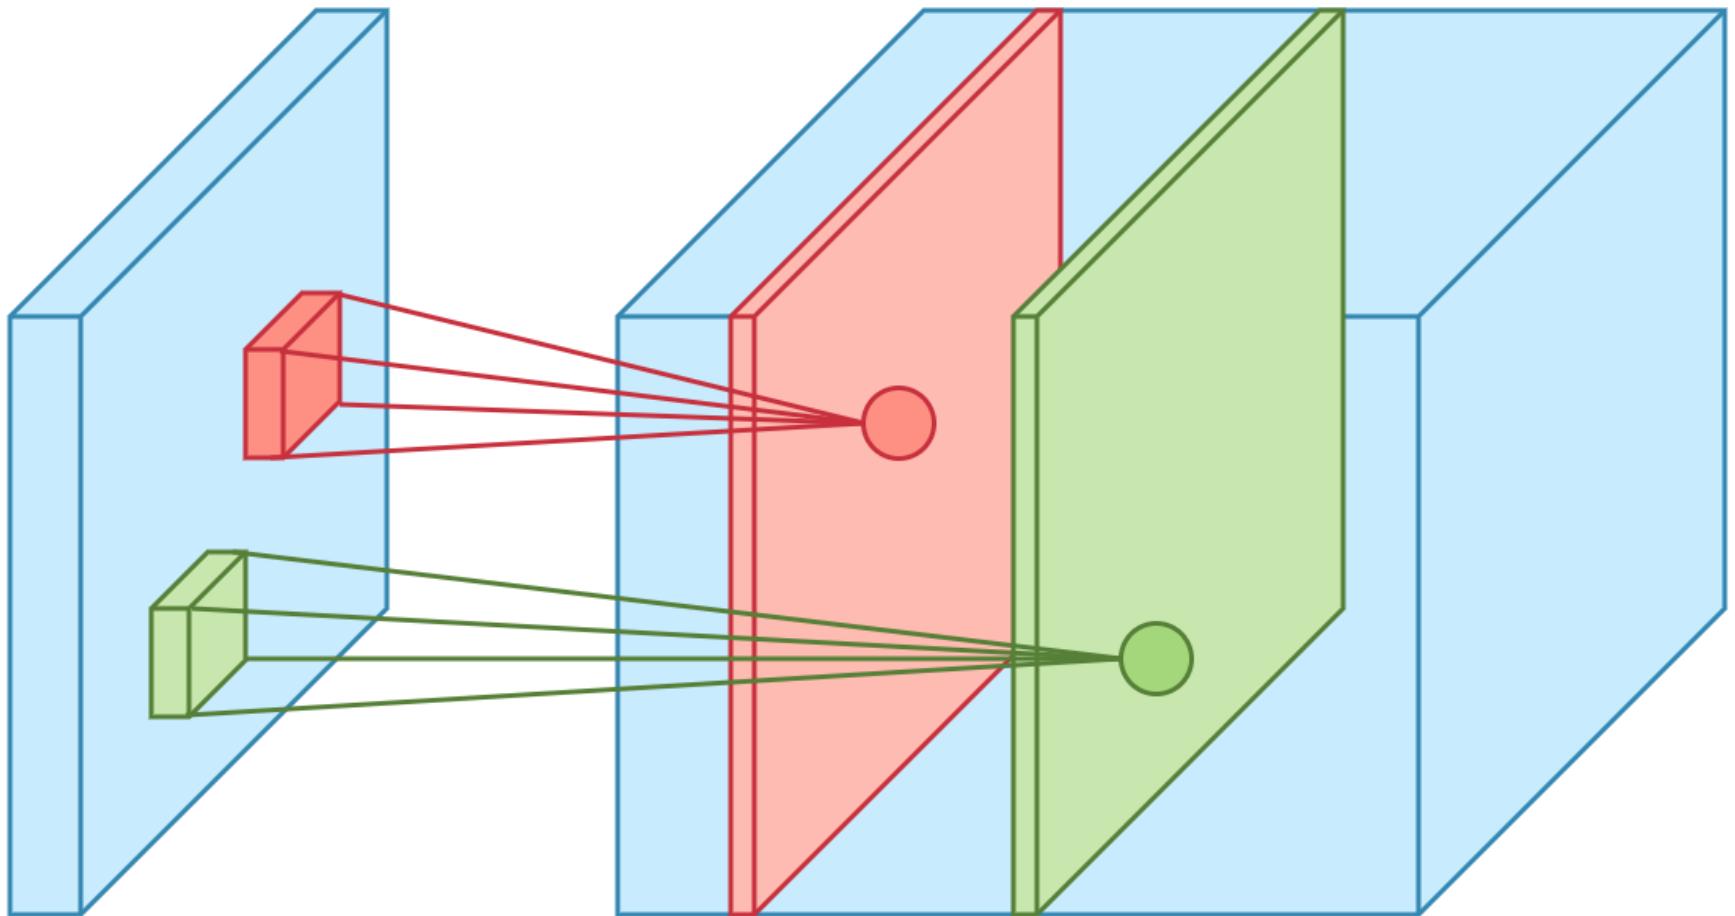
CNN Convolution Layer

Sliding operation at 4 locations



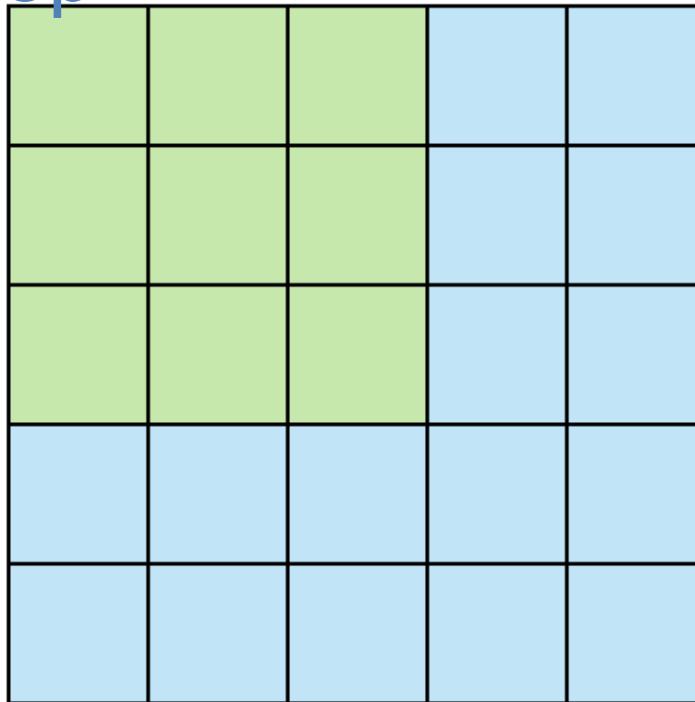
CNN Convolution Layer

two feature maps

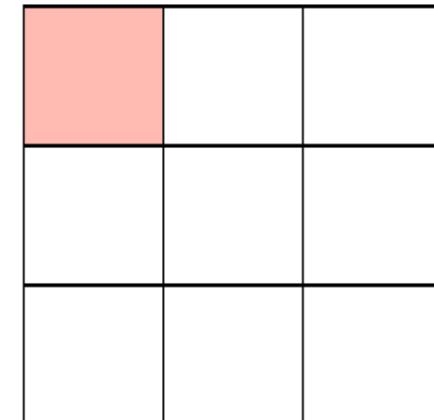


CNN Convolution Layer

Stride specifies how much we move the convolution filter at each step



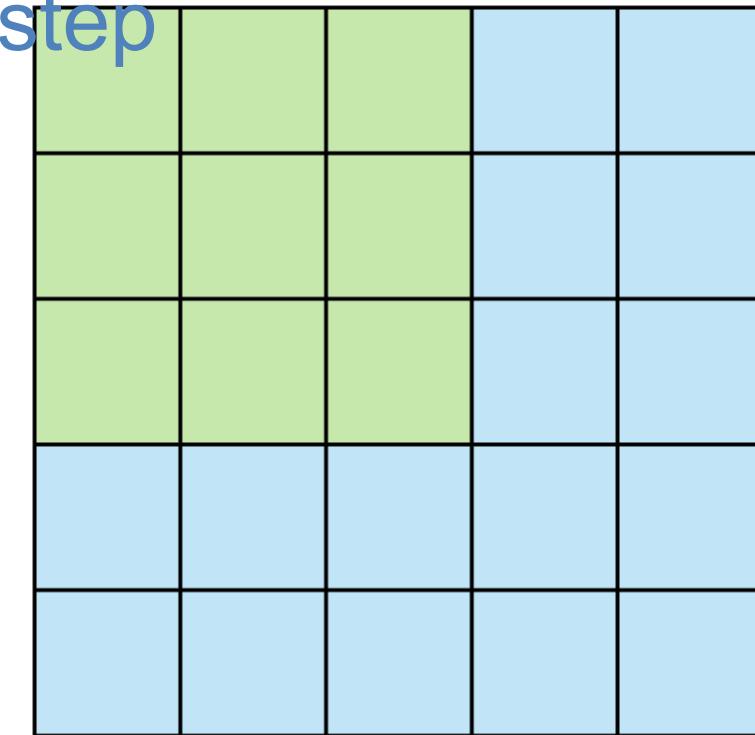
Stride 1



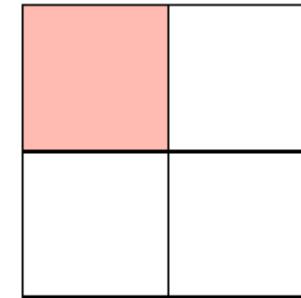
Feature Map

CNN Convolution Layer

Stride specifies how much we move the convolution filter at each step



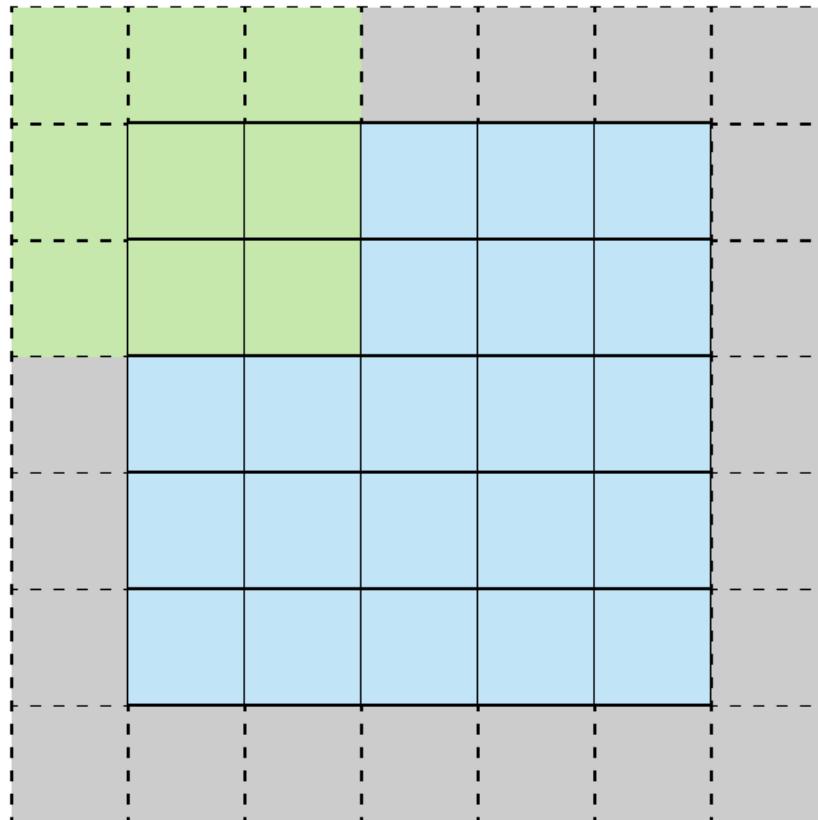
Stride 2



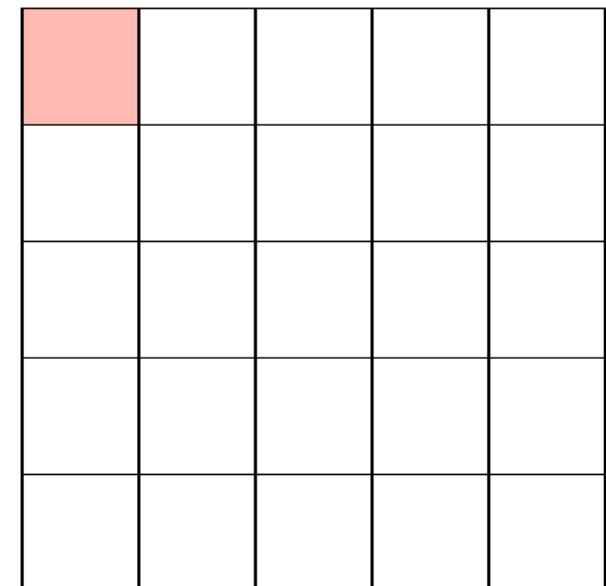
Feature Map

CNN Convolution Layer

Stride 1 with Padding



Stride 1 with Padding



Feature Map

CNN Pooling Layer

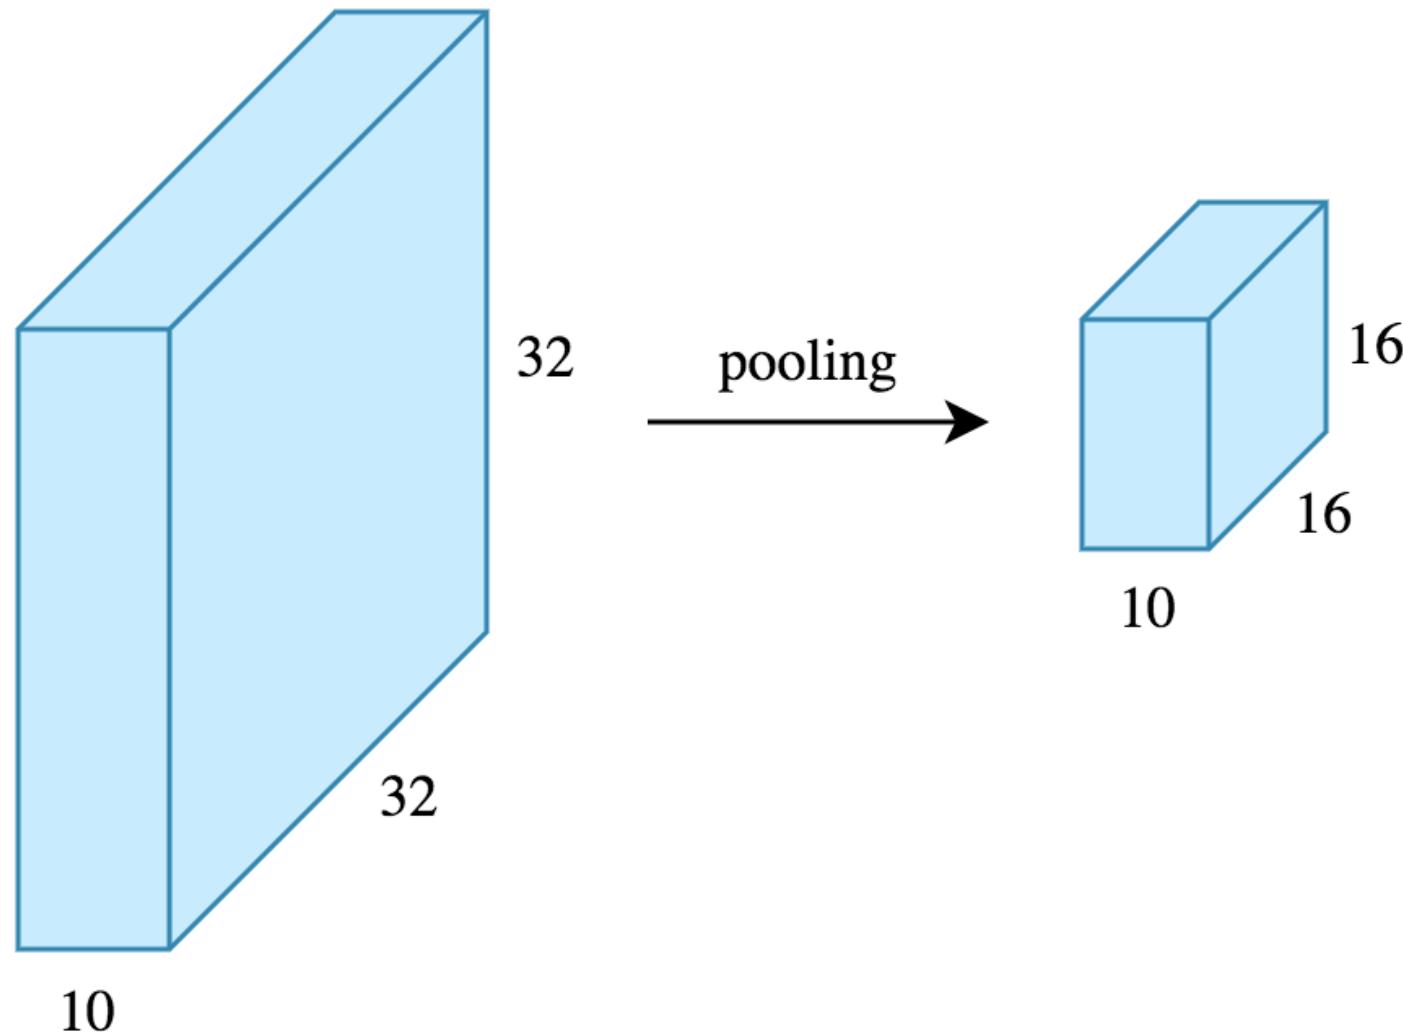
Max Pooling

1	1	2	4
5	6	7	8
3	2	1	0
1	2	3	4

max pool with 2x2
window and stride 2

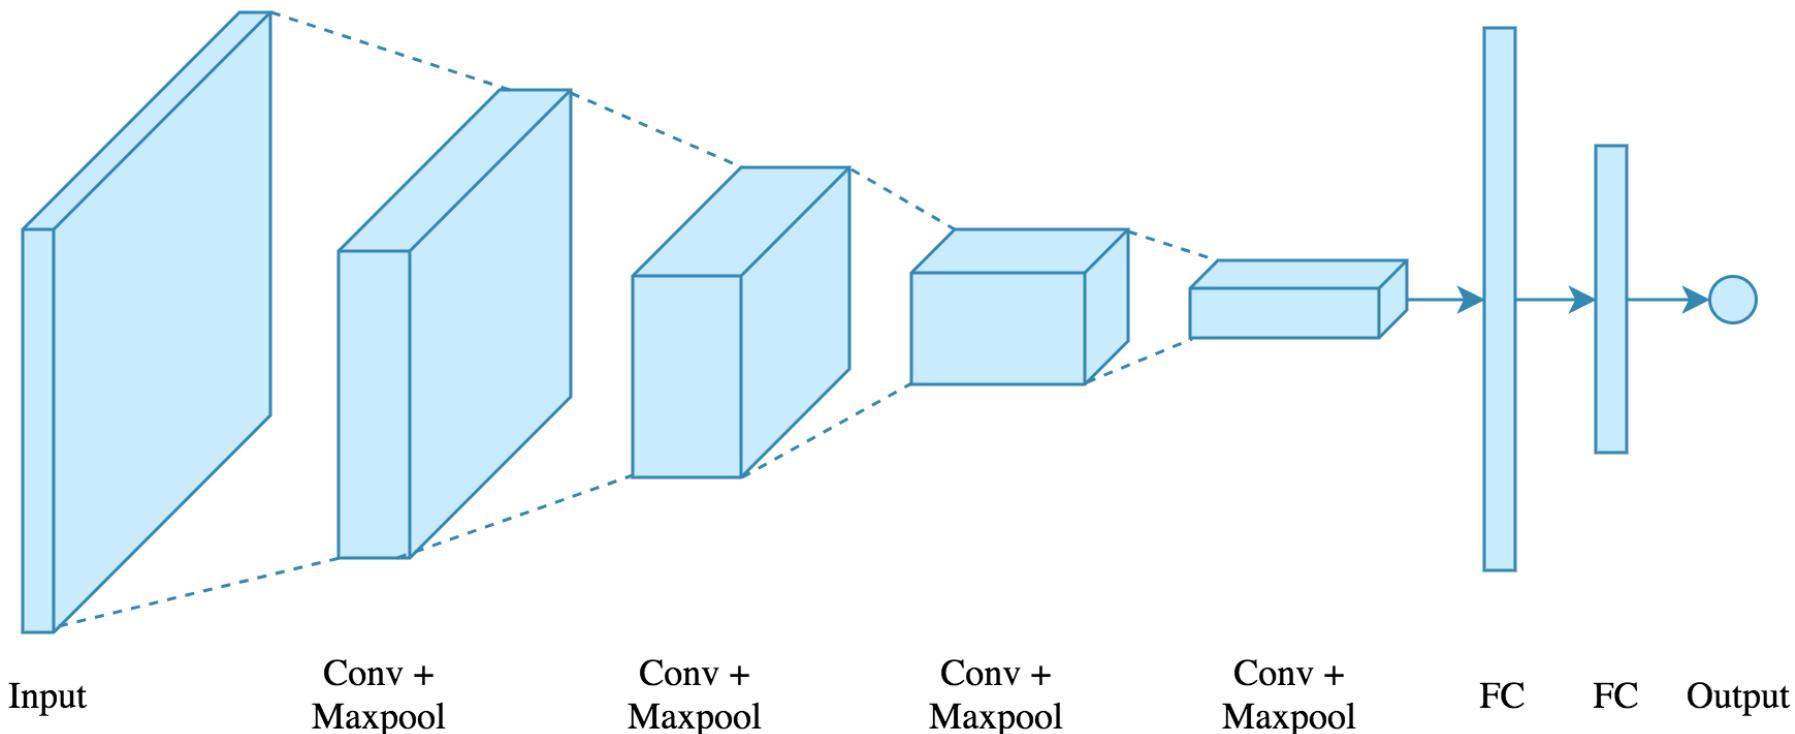
6	8
3	4

CNN Pooling Layer



CNN Architecture

4 convolution + pooling layers, followed by 2 fully connected layers



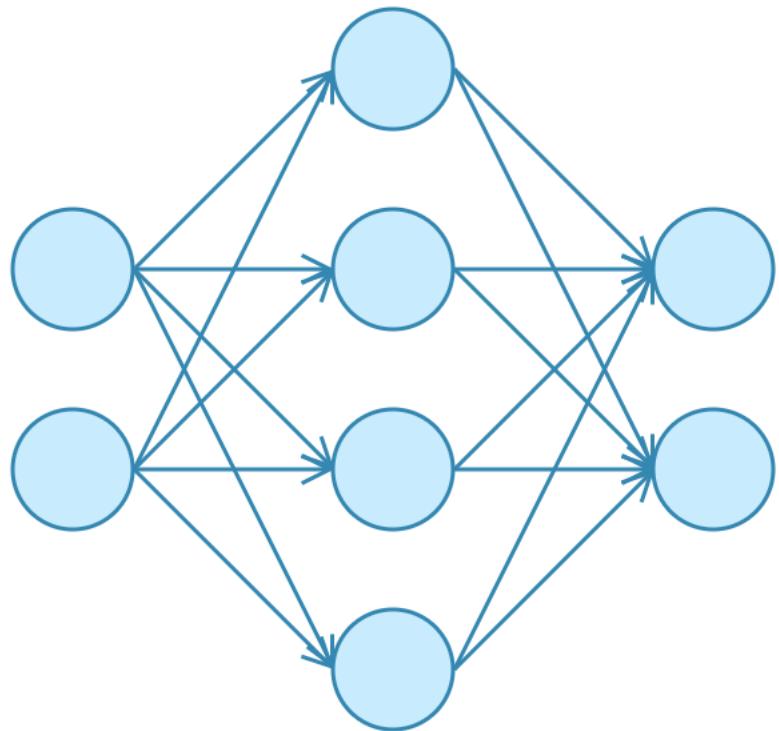
CNN Architecture

4 convolution + pooling layers, followed by 2 fully connected layers

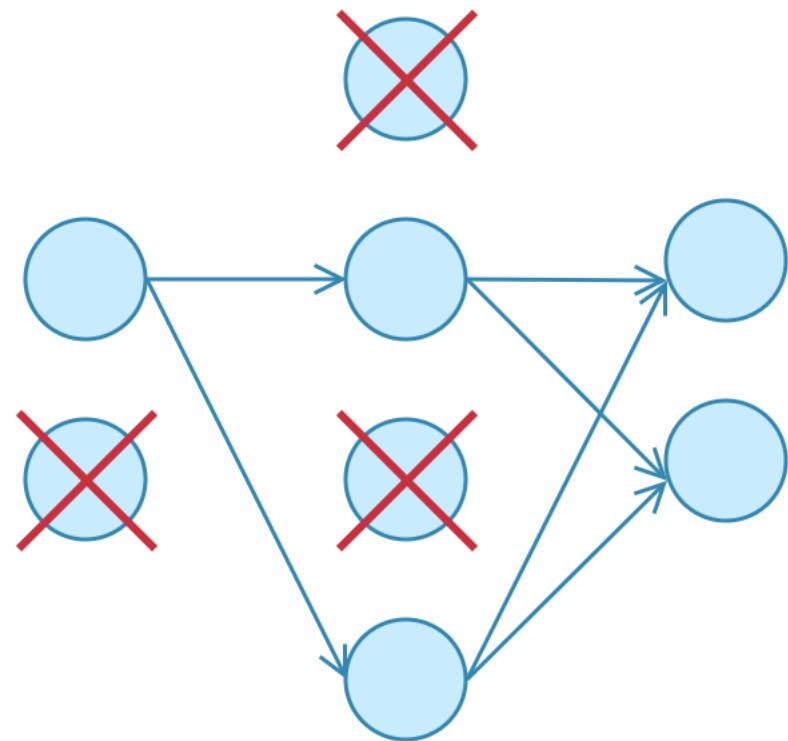
<https://gist.github.com/ardendertat/0fc5515057c47e7386fe04e9334504e3>

```
model = Sequential()
model.add(Conv2D(32, (3, 3), activation='relu', padding='same', name='conv_1',
                input_shape=(150, 150, 3)))
model.add(MaxPooling2D((2, 2), name='maxpool_1'))
model.add(Conv2D(64, (3, 3), activation='relu', padding='same', name='conv_2'))
model.add(MaxPooling2D((2, 2), name='maxpool_2'))
model.add(Conv2D(128, (3, 3), activation='relu', padding='same', name='conv_3'))
model.add(MaxPooling2D((2, 2), name='maxpool_3'))
model.add(Conv2D(128, (3, 3), activation='relu', padding='same', name='conv_4'))
model.add(MaxPooling2D((2, 2), name='maxpool_4'))
model.add(Flatten())
model.add(Dropout(0.5))
model.add(Dense(512, activation='relu', name='dense_1'))
model.add(Dense(128, activation='relu', name='dense_2'))
model.add(Dense(1, activation='sigmoid', name='output'))
```

Dropout

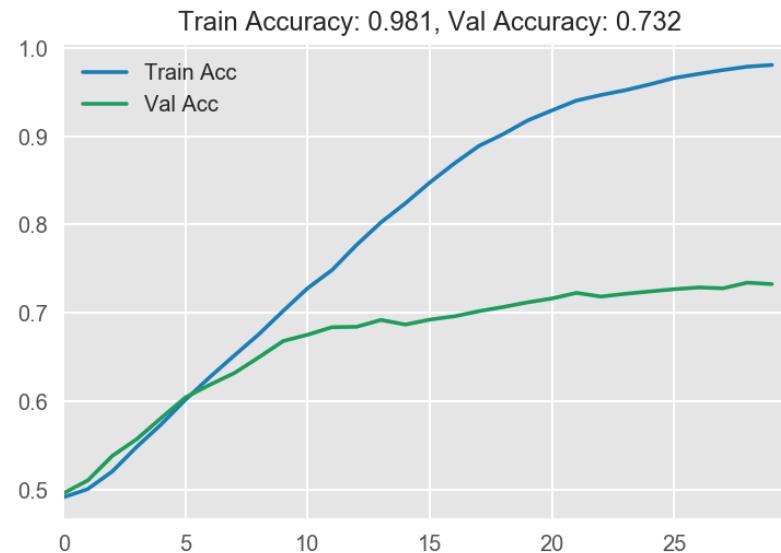
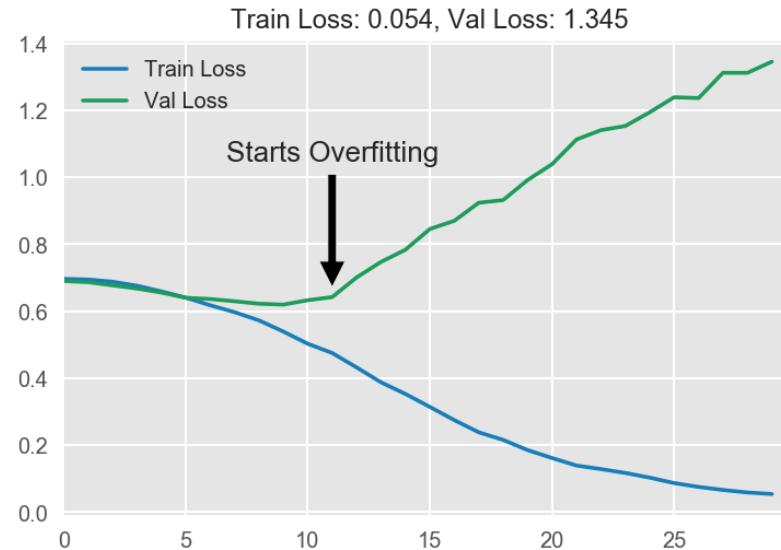


No Dropout



With Dropout

Model Performance



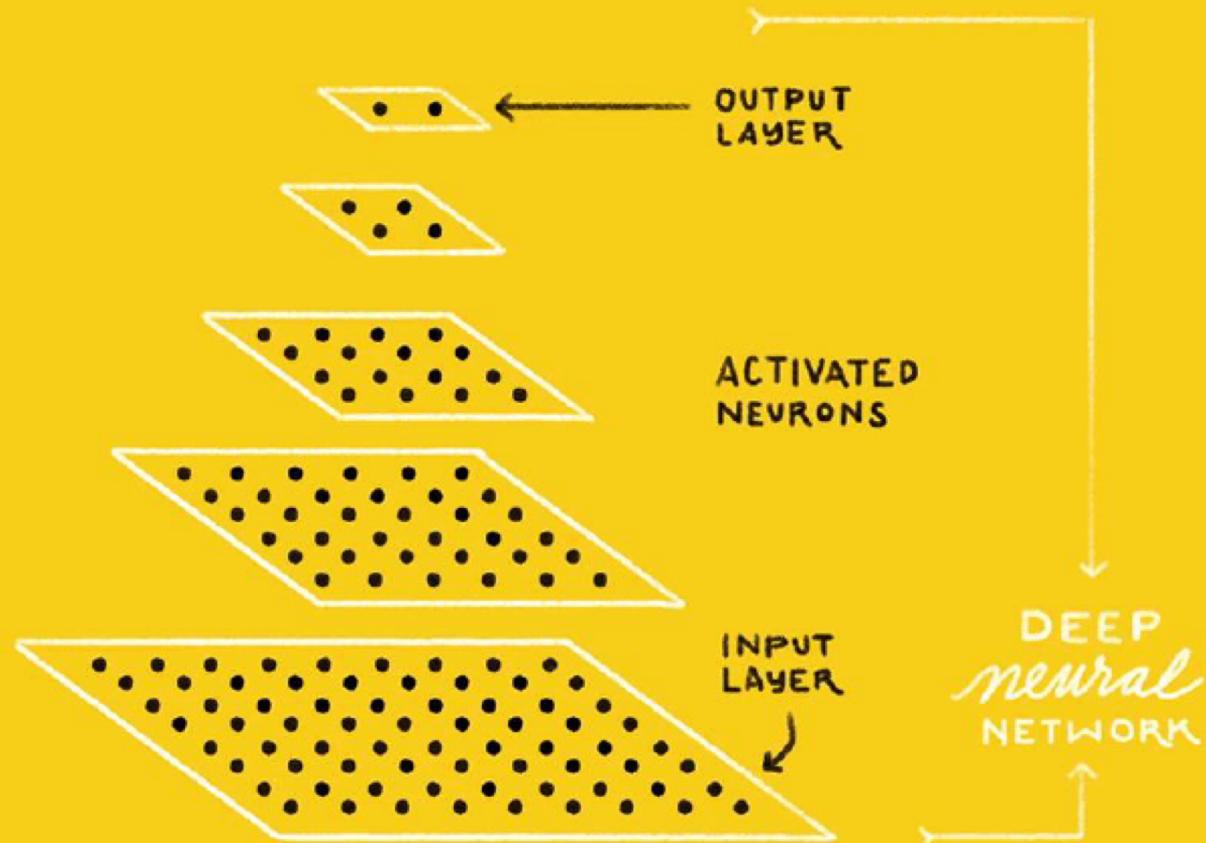
Visual Recognition

Image Classification

IS THIS A
CAT or DOG?



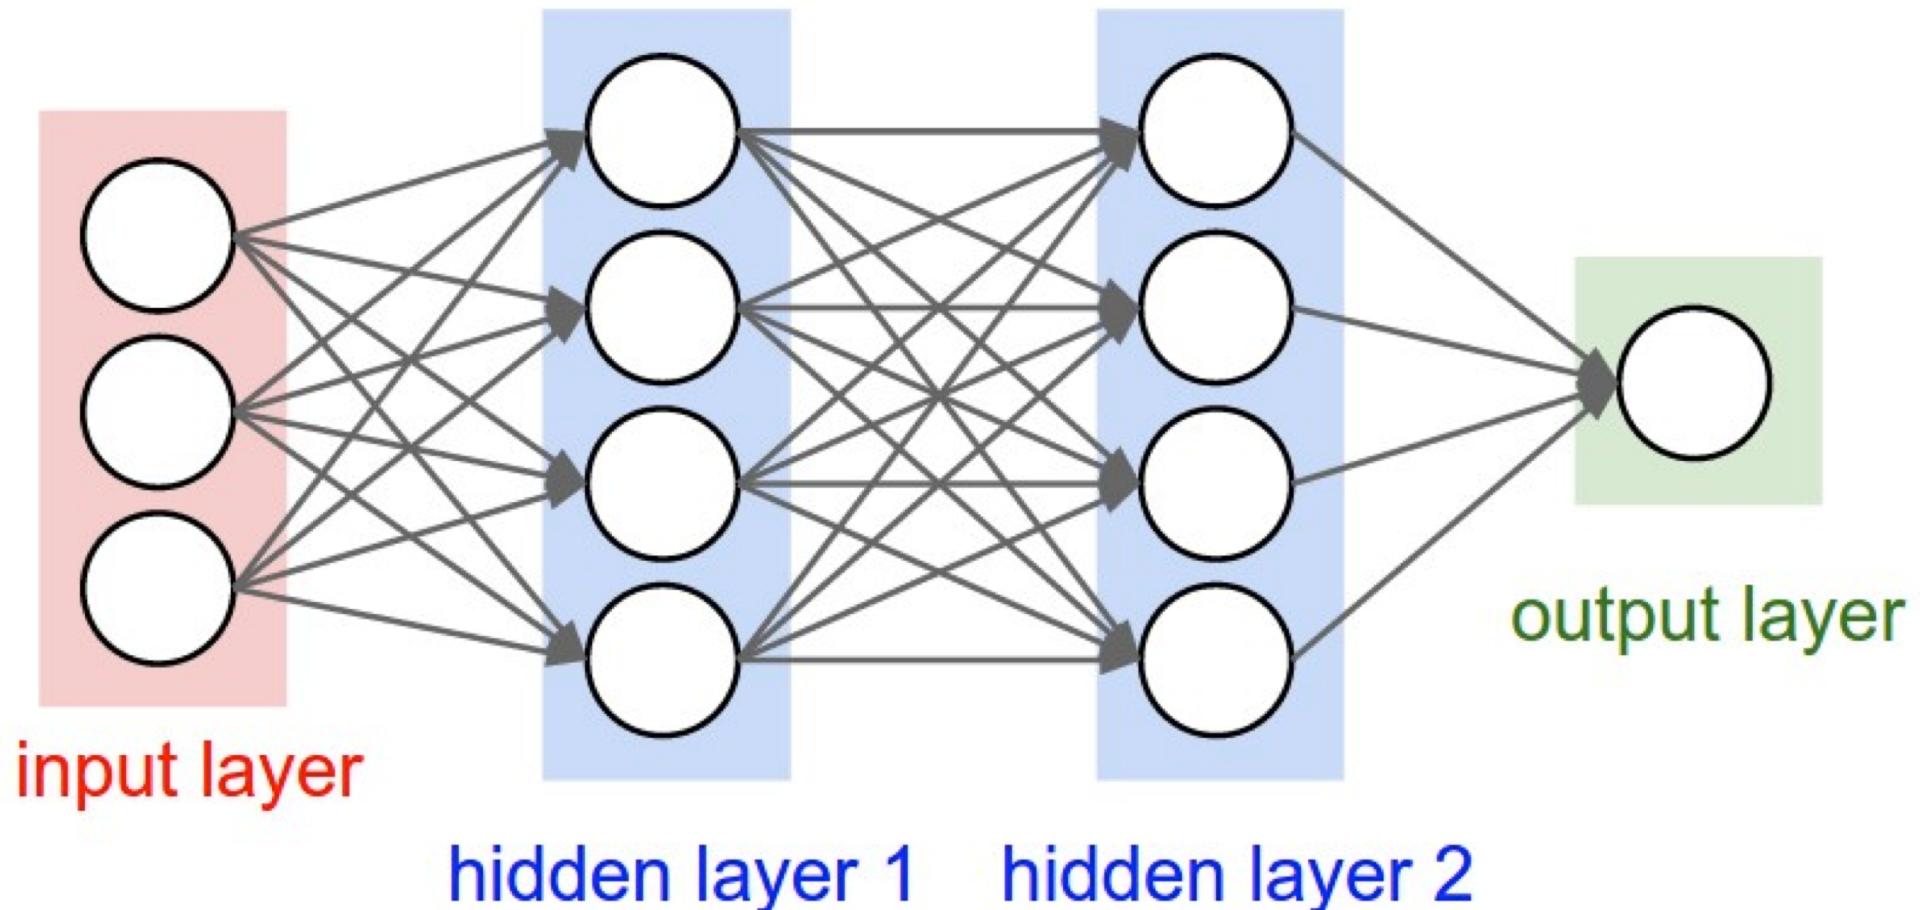
CAT DOG



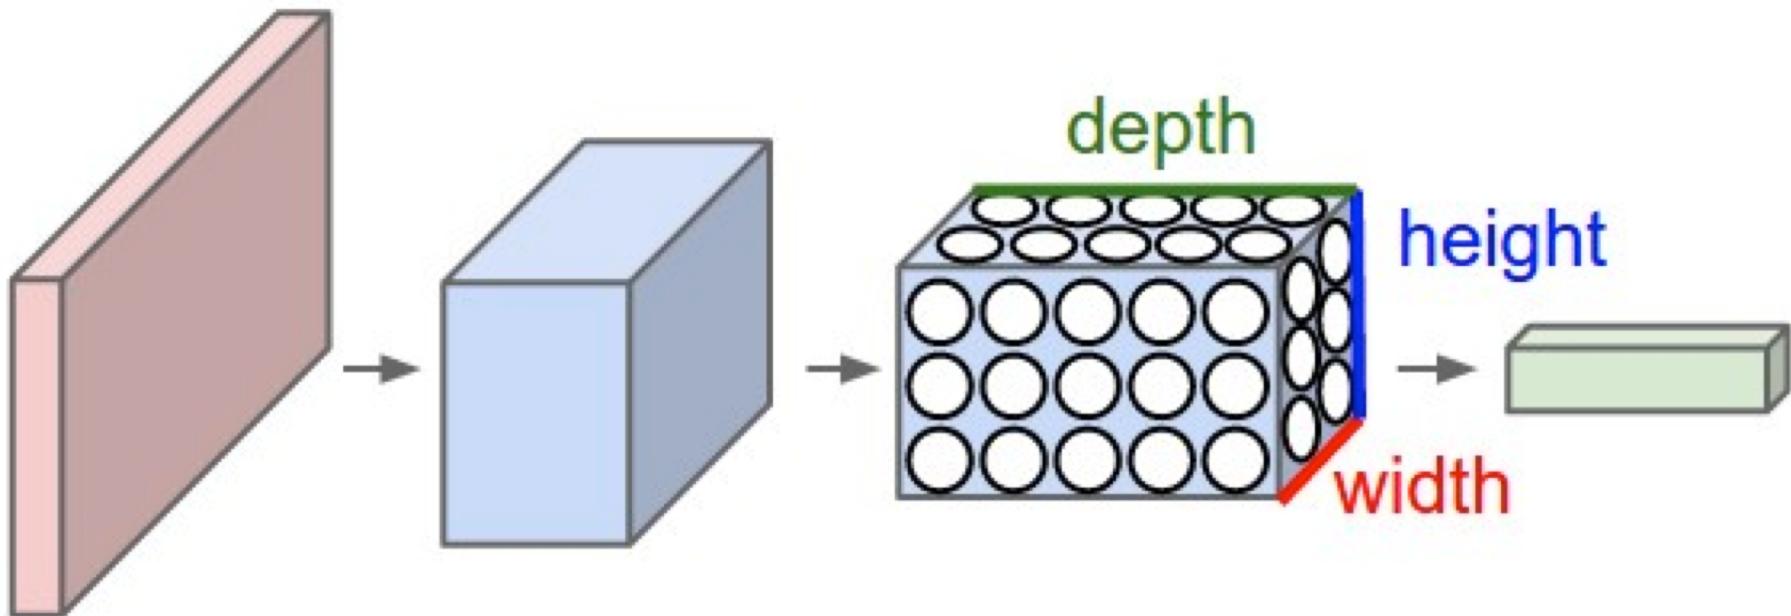
Convolutional Neural Networks (CNNs / ConvNets)

<http://cs231n.github.io/convolutional-networks/>

A regular 3-layer Neural Network



A ConvNet arranges its neurons in three dimensions (width, height, depth)

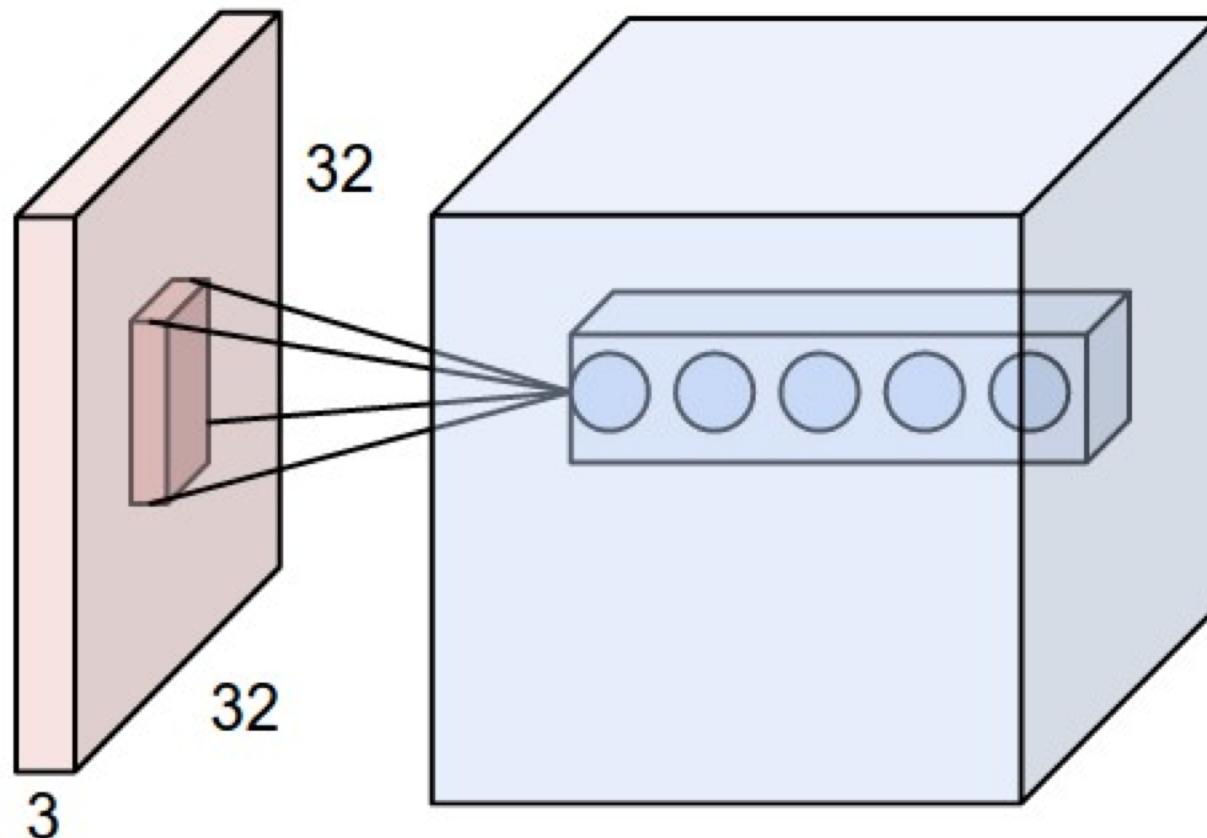


The activations of an example ConvNet architecture.



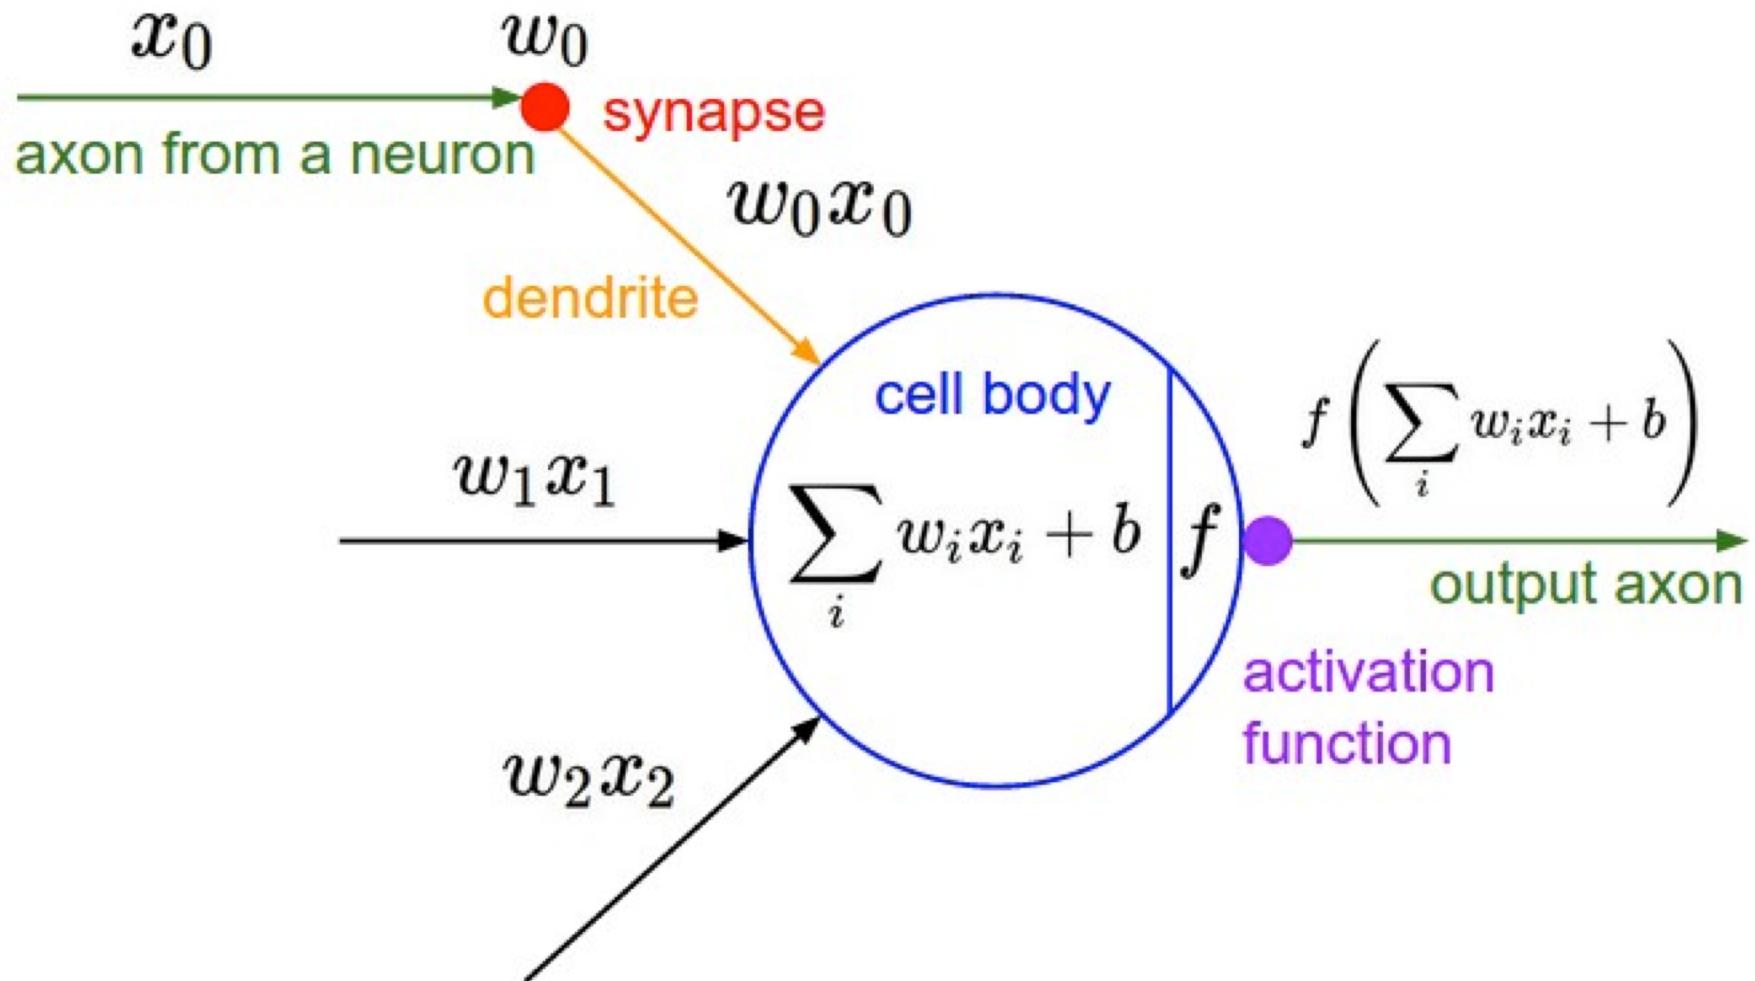
ConvNets

32x32x3 CIFAR-10 image



first Convolutional layer

ConvNets



Convolution Demo

Input Volume (+pad 1) (7x7x3)

$x[:, :, 0]$

0	0	0	0	0	0	0	0
0	1	2	0	2	1	0	
0	2	2	2	1	1	0	
0	2	2	2	0	1	0	
0	2	2	1	2	1	0	
0	2	1	2	0	1	0	
0	0	0	0	0	0	0	

Filter W0 (3x3x3)

$w0[:, :, 0]$

-1	-1	0
1	1	1
-1	0	1
0	0	1
0	1	0
0	1	0
0	0	1

Filter W1 (3x3x3)

$w1[:, :, 0]$

1	-1	0
0	1	1
0	-1	1
-1	1	0
-1	-1	1
0	0	0

Output Volume (3x3x2)

$o[:, :, 0]$

6	3	6
7	-1	-2
2	3	-2
-1	1	-3
4	3	2
-1	0	-1

$o[:, :, 1]$

7	-1	-3
4	3	2
-1	0	-1

$w1[:, :, 2]$

1	0	-1
0	0	-1
1	0	1

Bias b0 (1x1x1)

$b0[:, :, 0]$

1

0

toggle movement

$x[:, :, 1]$

0	0	0	0	0	0	0	0
0	0	2	2	1	2	0	
0	1	2	0	0	2	0	
0	0	1	2	1	0	0	
0	2	2	2	2	0	0	
0	2	2	2	0	2	0	
0	0	0	0	0	0	0	

$w0[:, :, 2]$

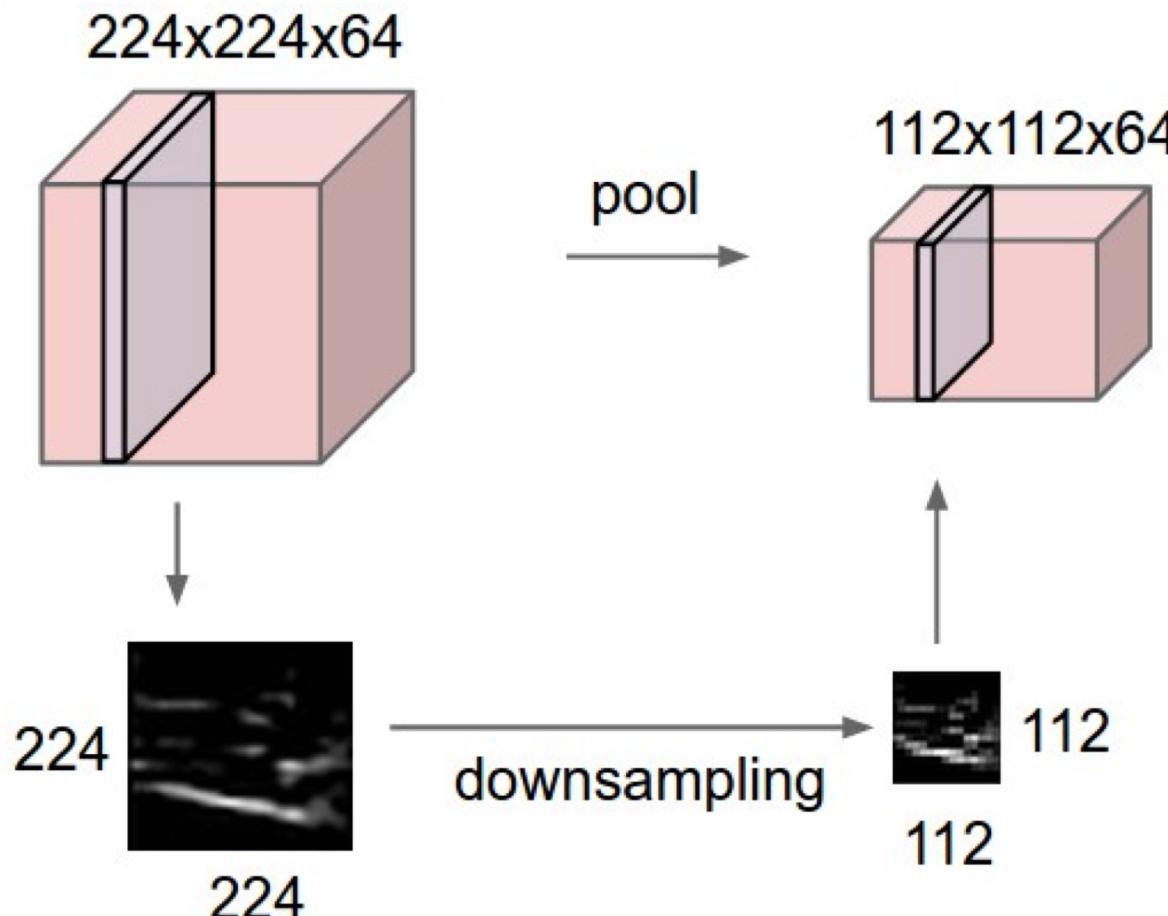
-1	-1	0
1	0	-1
-1	0	-1
0	0	0

$x[:, :, 2]$

0	0	0	0	0	0	0	0
0	1	0	0	1	0	0	
0	0	2	0	0	0	0	
0	0	0	1	1	1	0	
0	2	2	2	1	2	0	
0	1	2	0	0	2	0	
0	0	0	0	0	0	0	

ConvNets

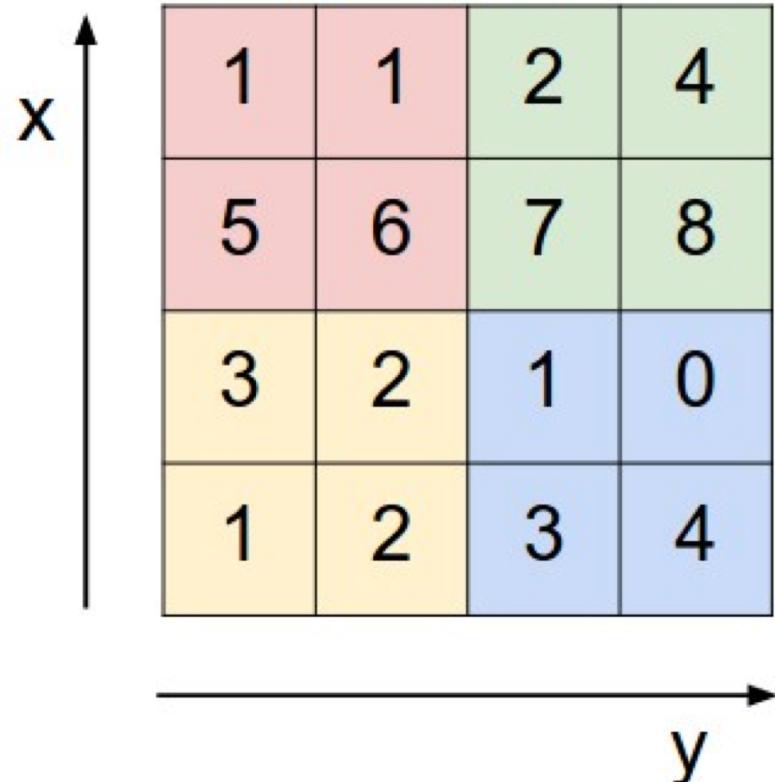
input volume of size [224x224x64]
is pooled with **filter size 2, stride 2**
into output volume of size [112x112x64]



ConvNets

max pooling

Single depth slice

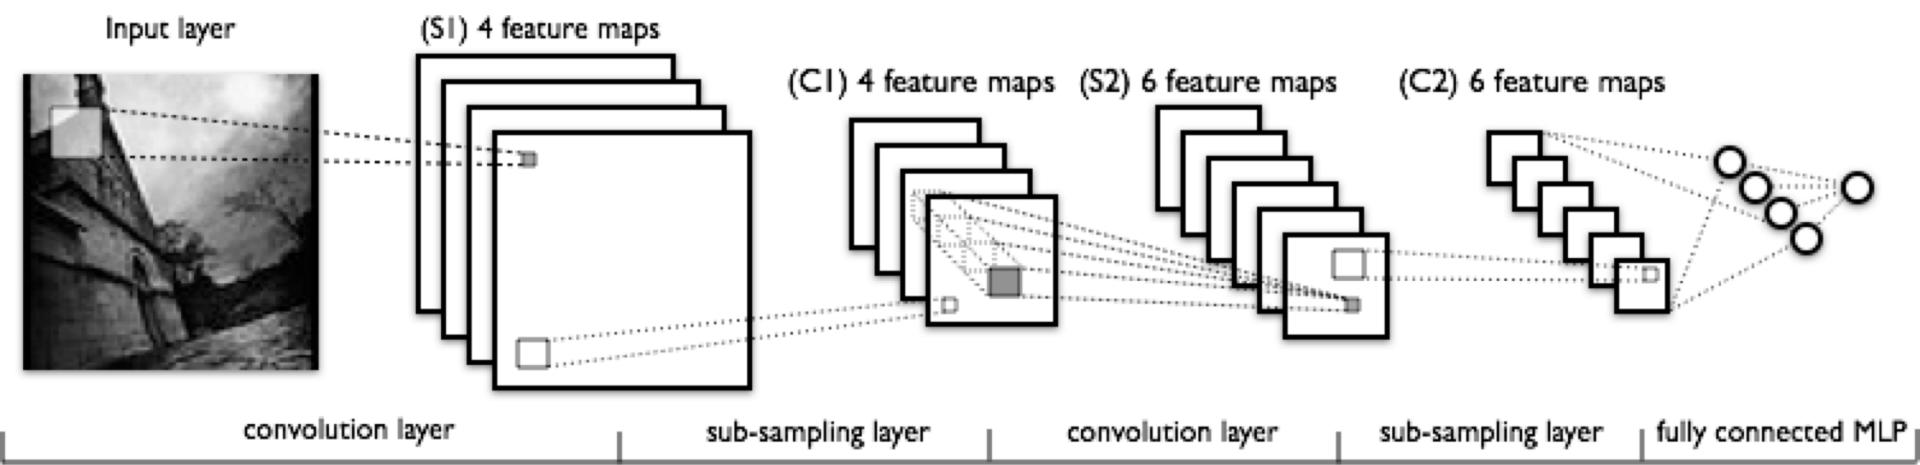


max pool with 2x2 filters
and stride 2



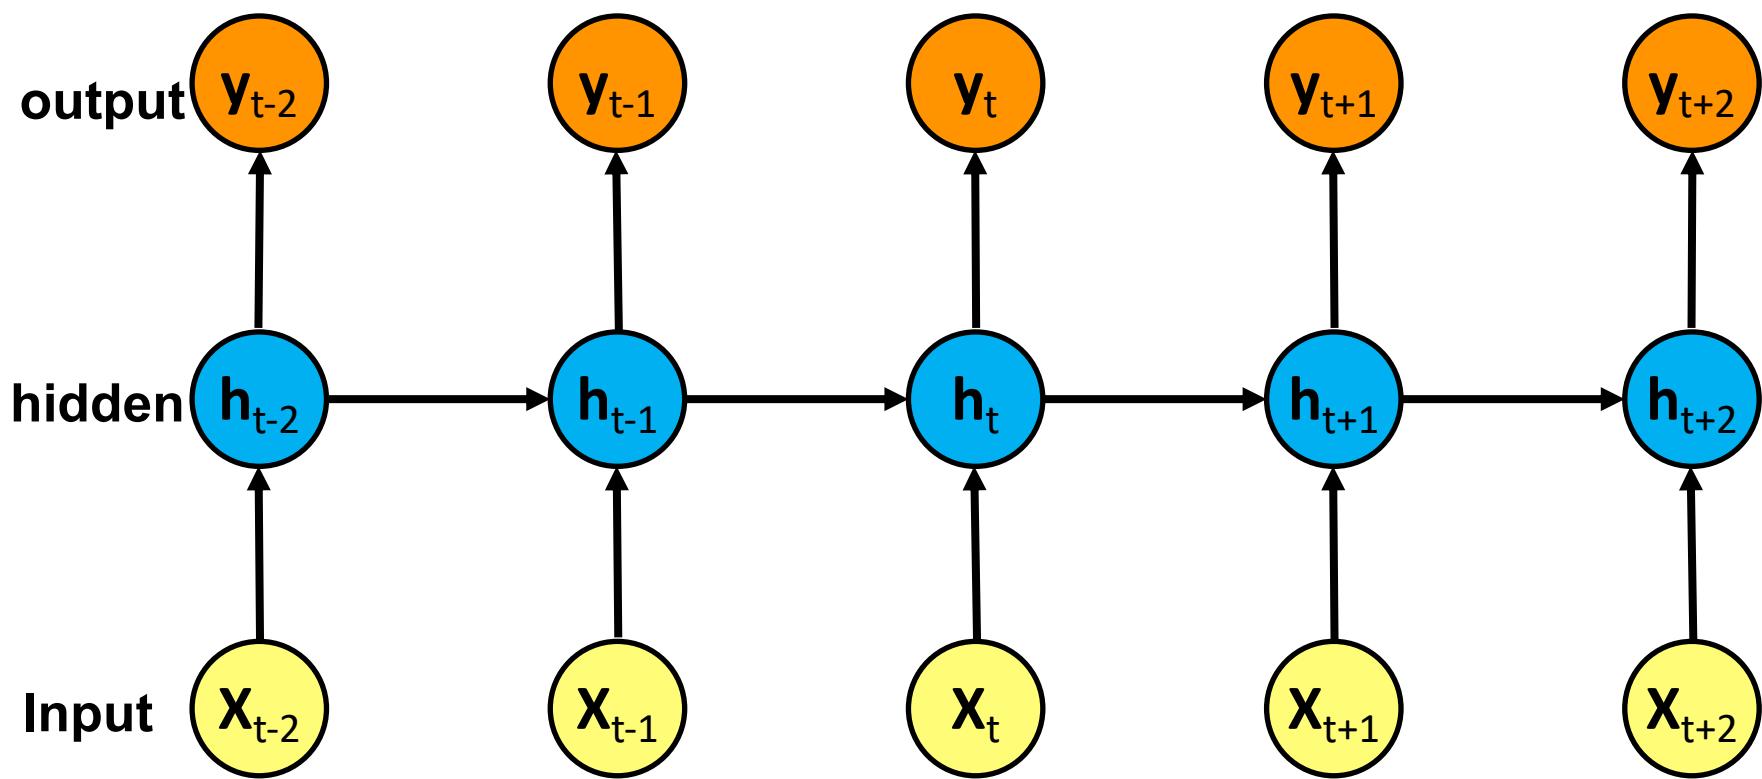
6	8
3	4

Convolutional Neural Networks (CNN) (LeNet)



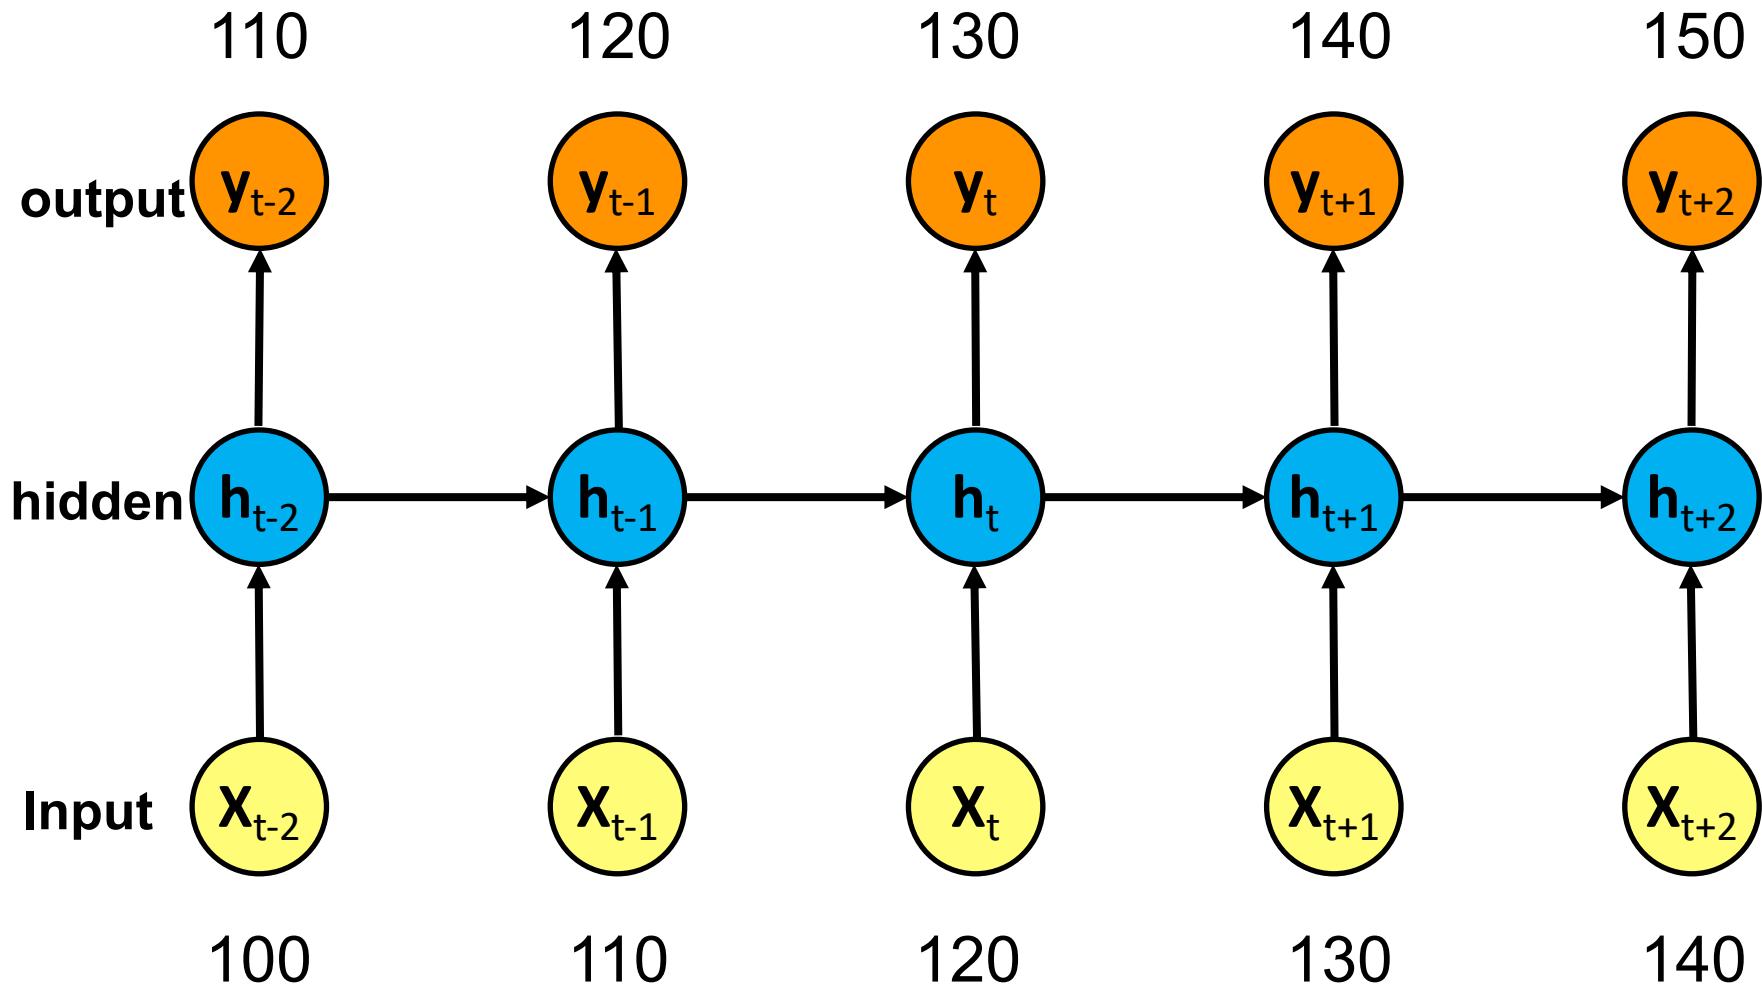
Recurrent Neural Networks (RNN)

Recurrent Neural Networks (RNN)



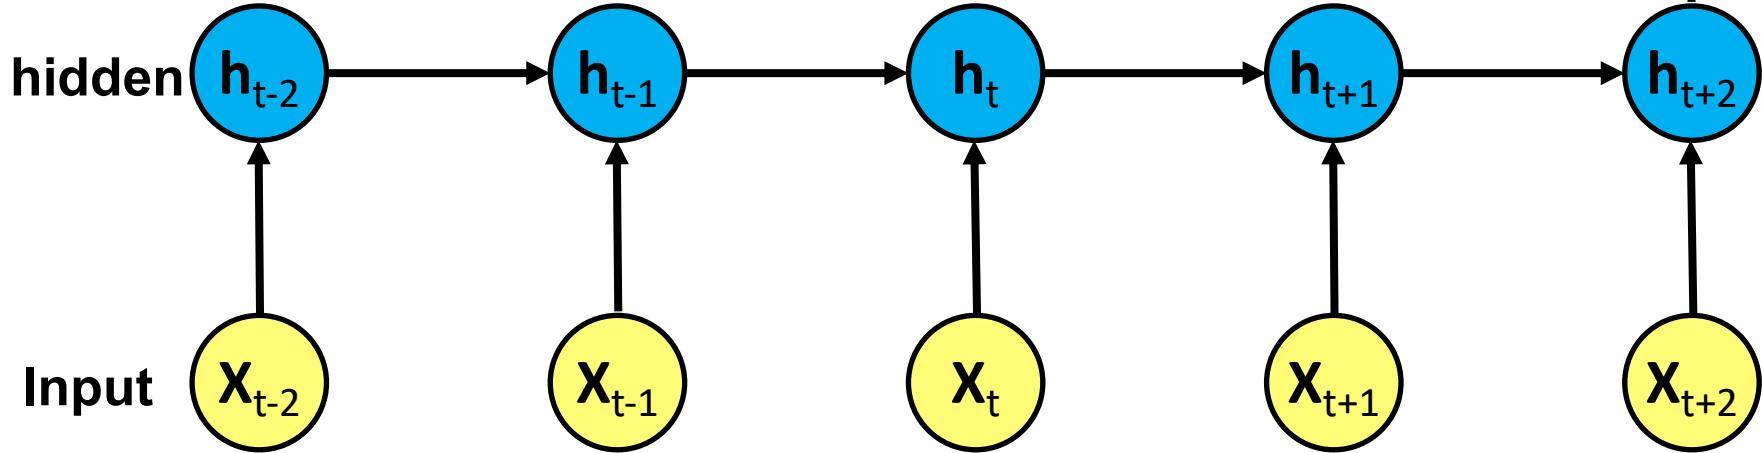
Recurrent Neural Networks (RNN)

Time Series Forecasting



Recurrent Neural Networks (RNN)

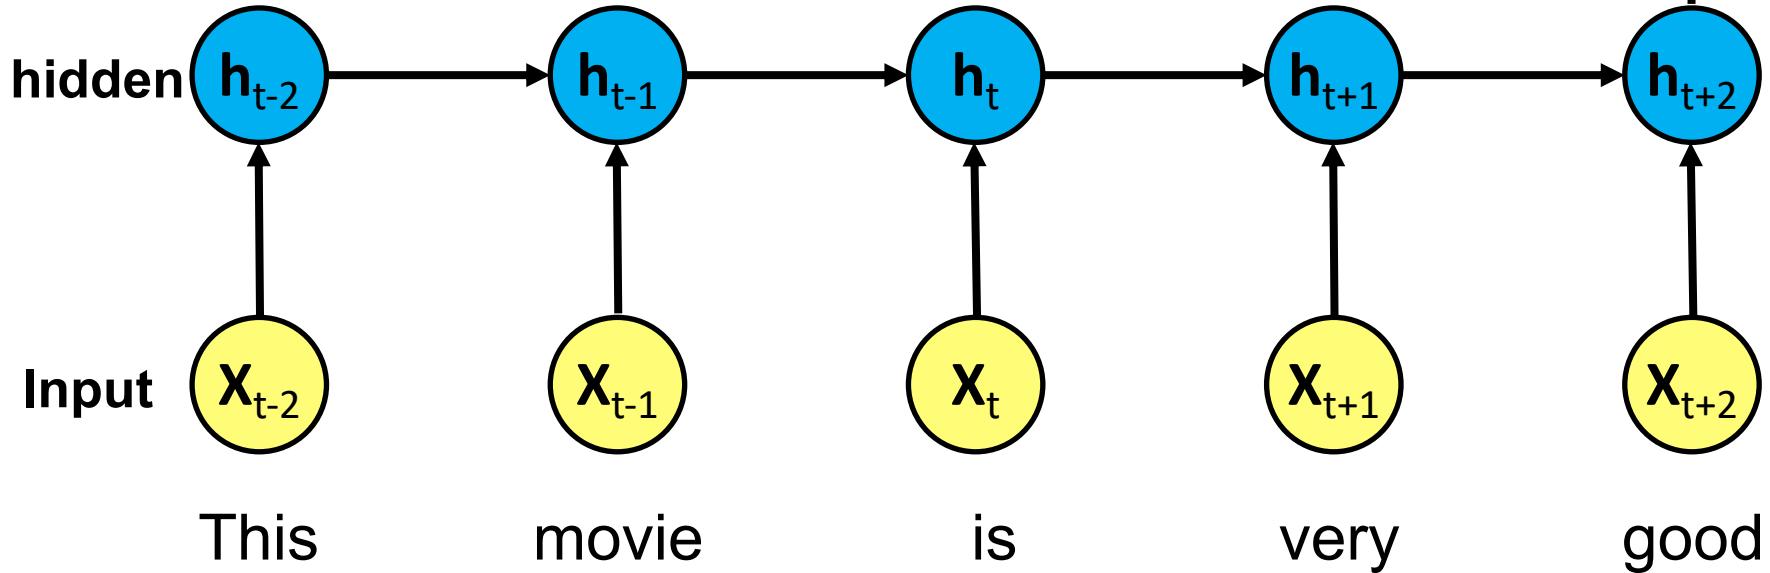
output



Recurrent Neural Networks (RNN)

Sentiment Analysis

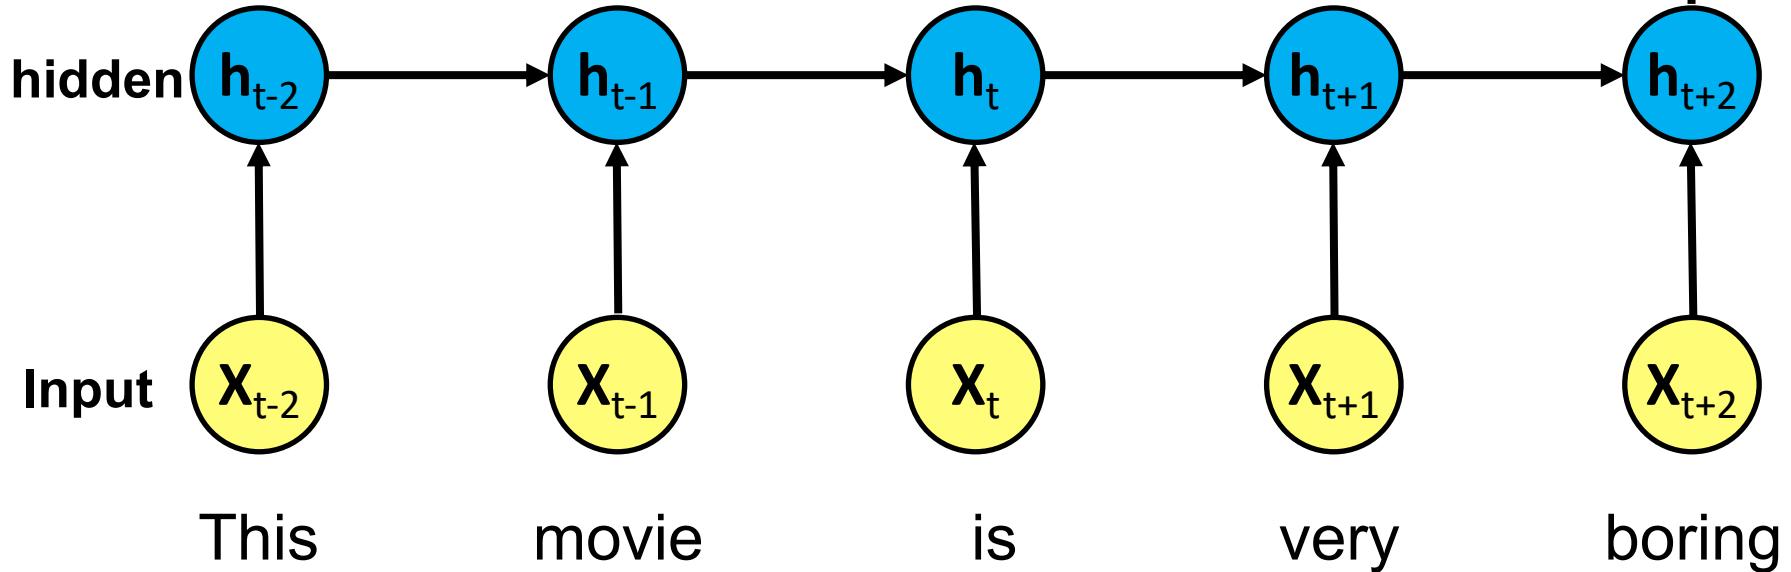
output



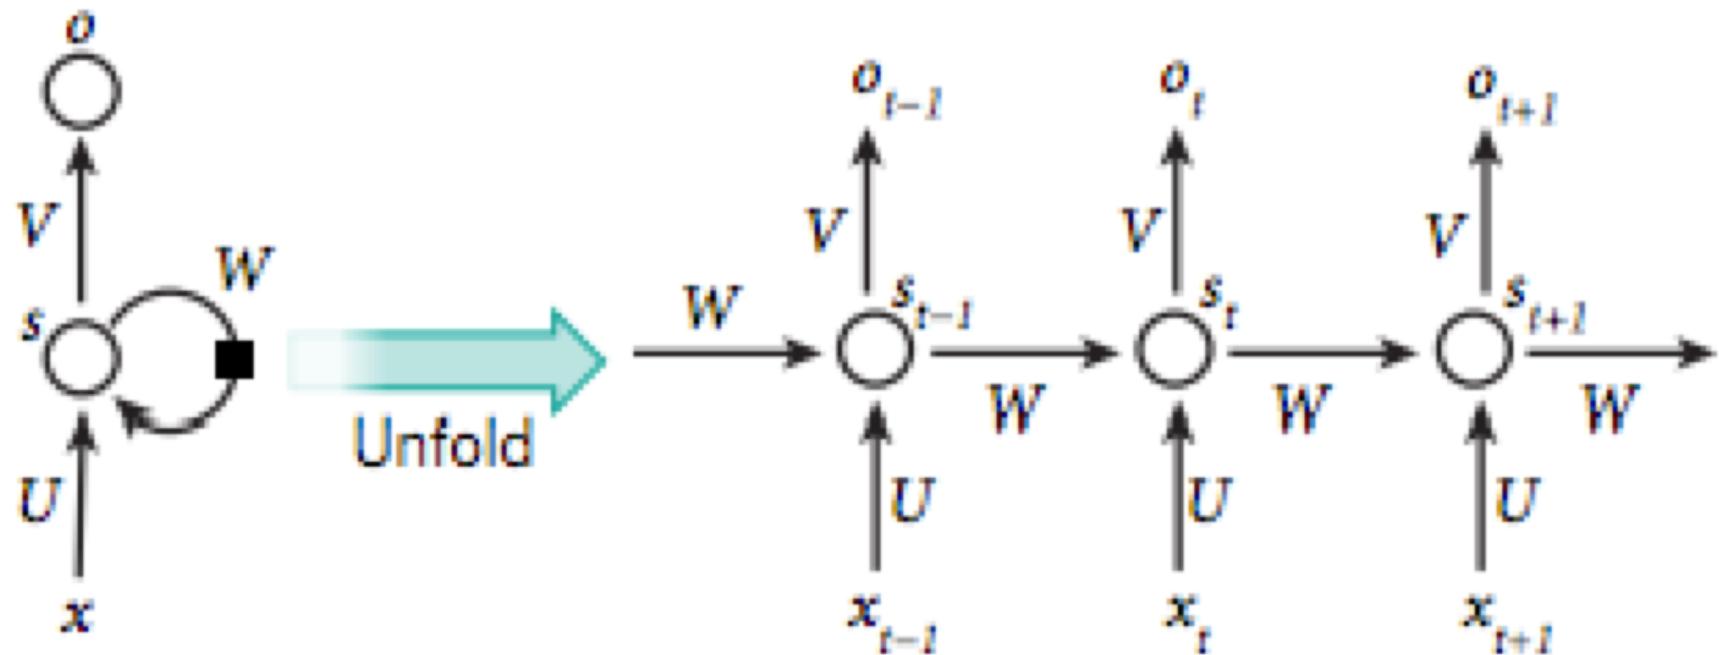
Recurrent Neural Networks (RNN)

Sentiment Analysis

output



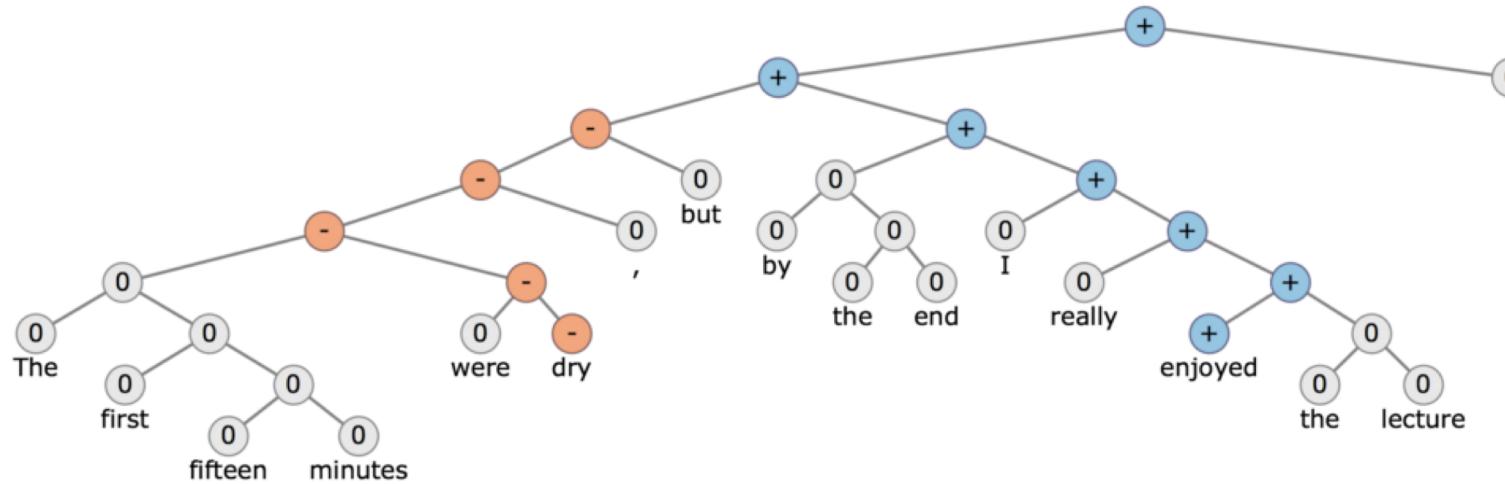
Recurrent Neural Network (RNN)



CS224d: Deep Learning for Natural Language Processing



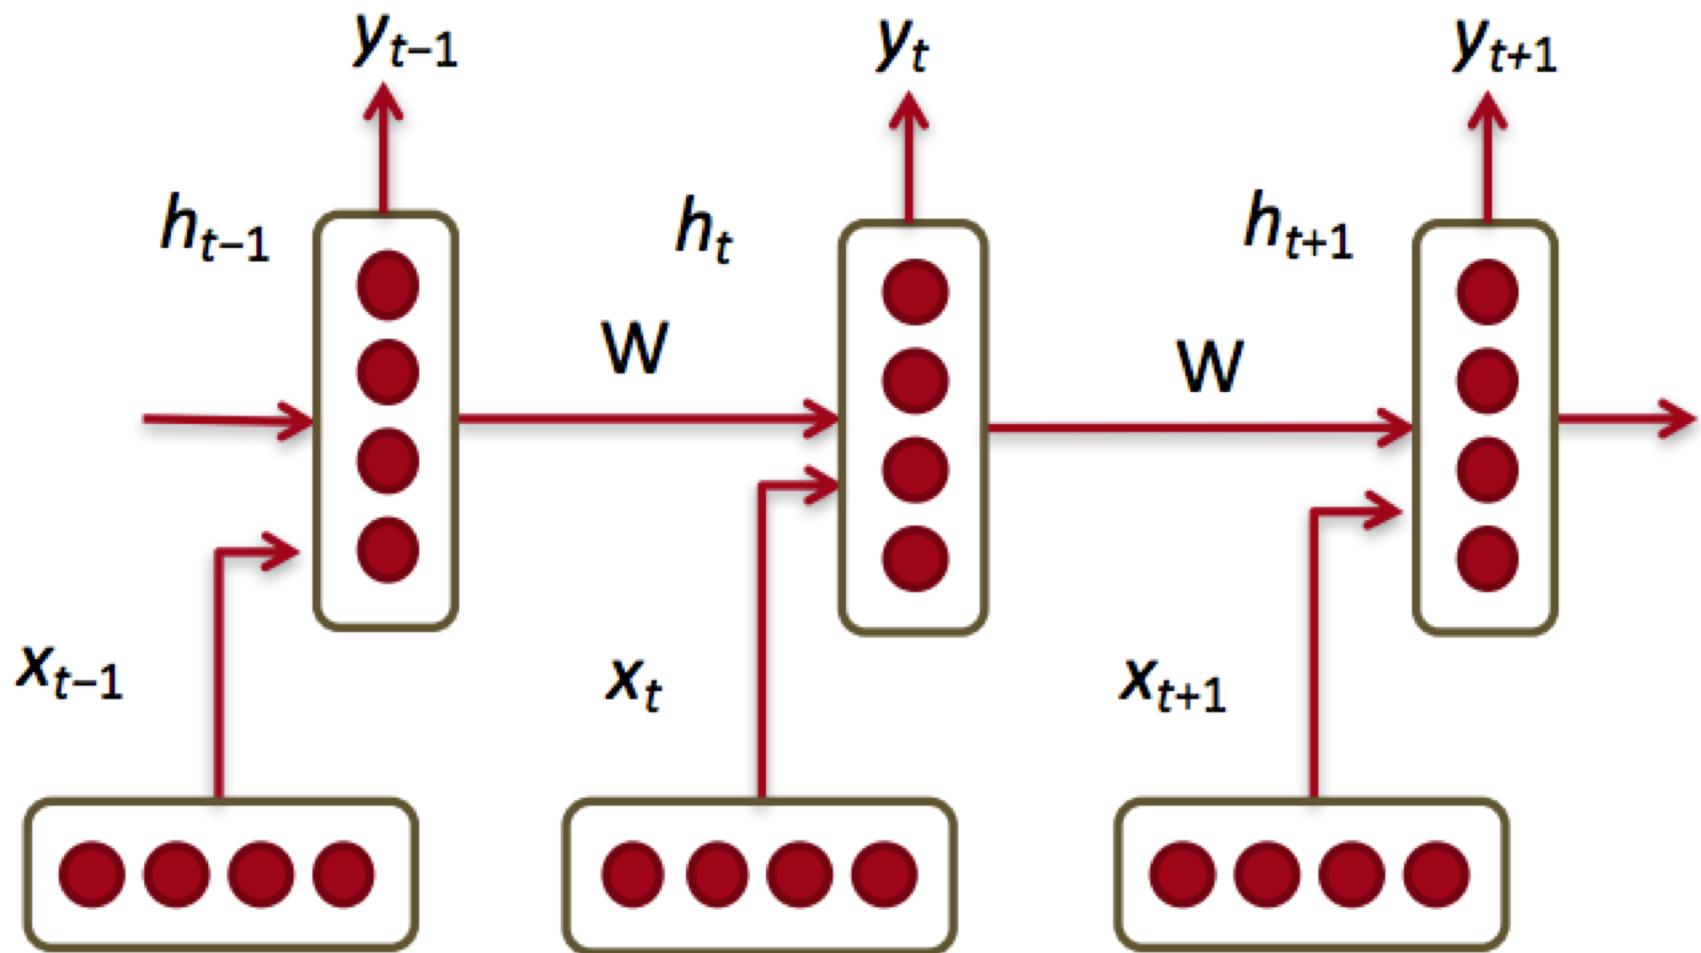
CS224d: Deep Learning for Natural Language Processing



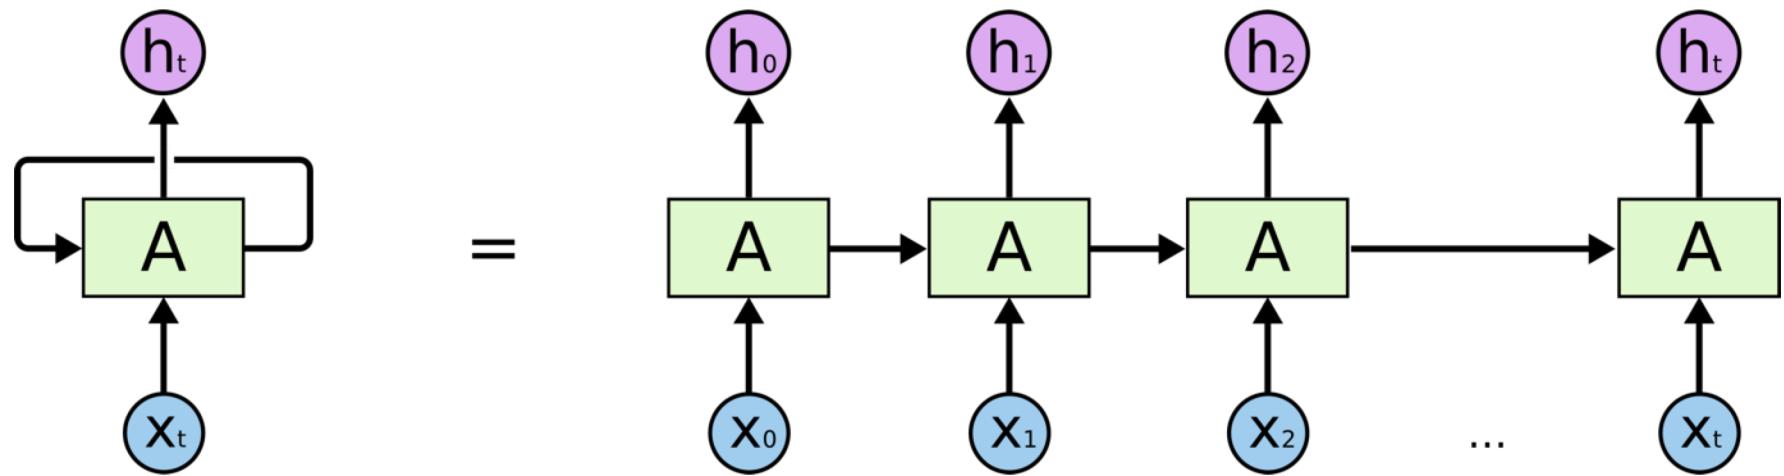
Course Description

Natural language processing (NLP) is one of the most important technologies of the information age. Understanding complex language utterances is also a crucial part of artificial intelligence. Applications of NLP are everywhere because people communicate most everything in language: web search, advertisement, emails, customer service, language translation, radiology reports, etc. There are a large variety of underlying tasks and machine learning models powering NLP applications. Recently, deep learning approaches have obtained very high performance across many different NLP tasks. These models can often be trained with a single end-to-end model and do not require traditional, task-specific feature engineering. In this spring quarter course students will learn to implement, train, debug, visualize and invent their own neural network models. The course provides a deep excursion into cutting-edge research in deep learning applied to NLP. The final project will involve training a complex recurrent neural network and applying it to a large scale NLP problem. On the model side we will cover word vector representations,

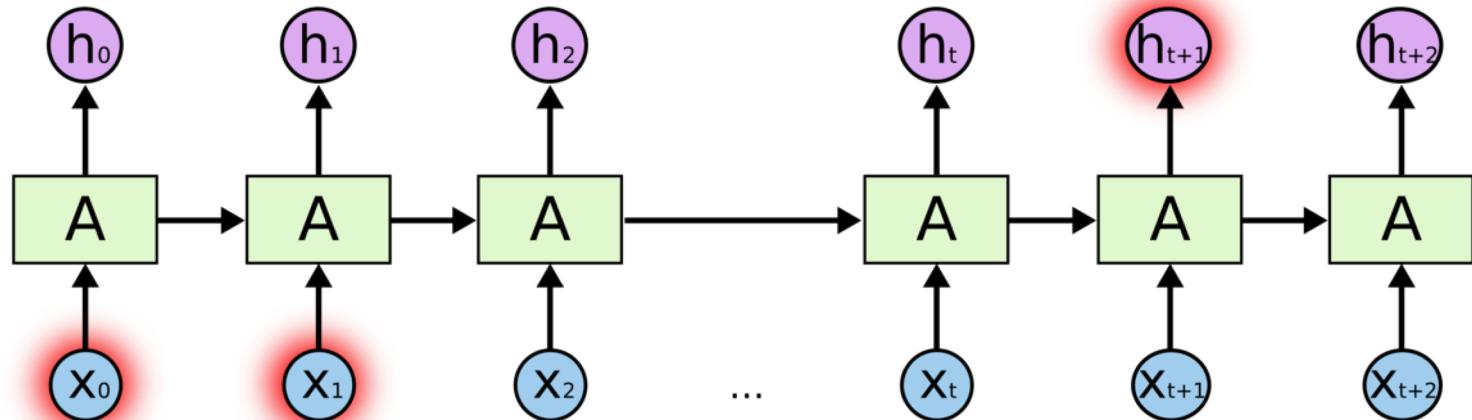
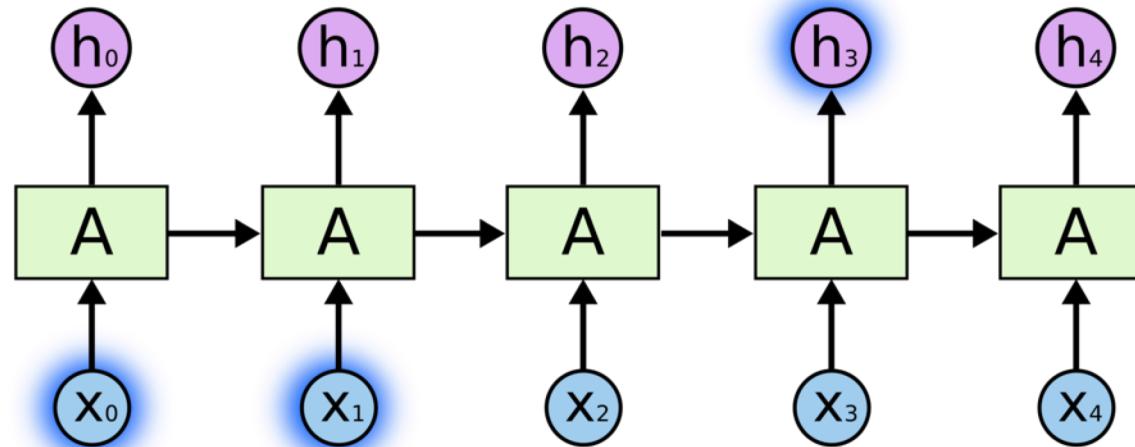
Recurrent Neural Networks (RNNs)



RNN

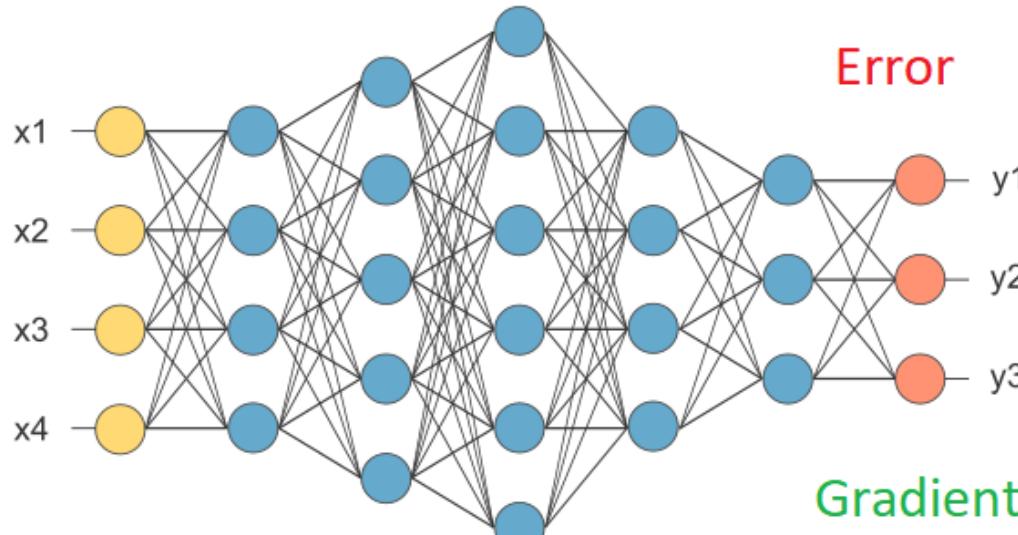


RNN long-term dependencies



I grew up in France... I speak fluent French.

Vanishing Gradient Exploding Gradient

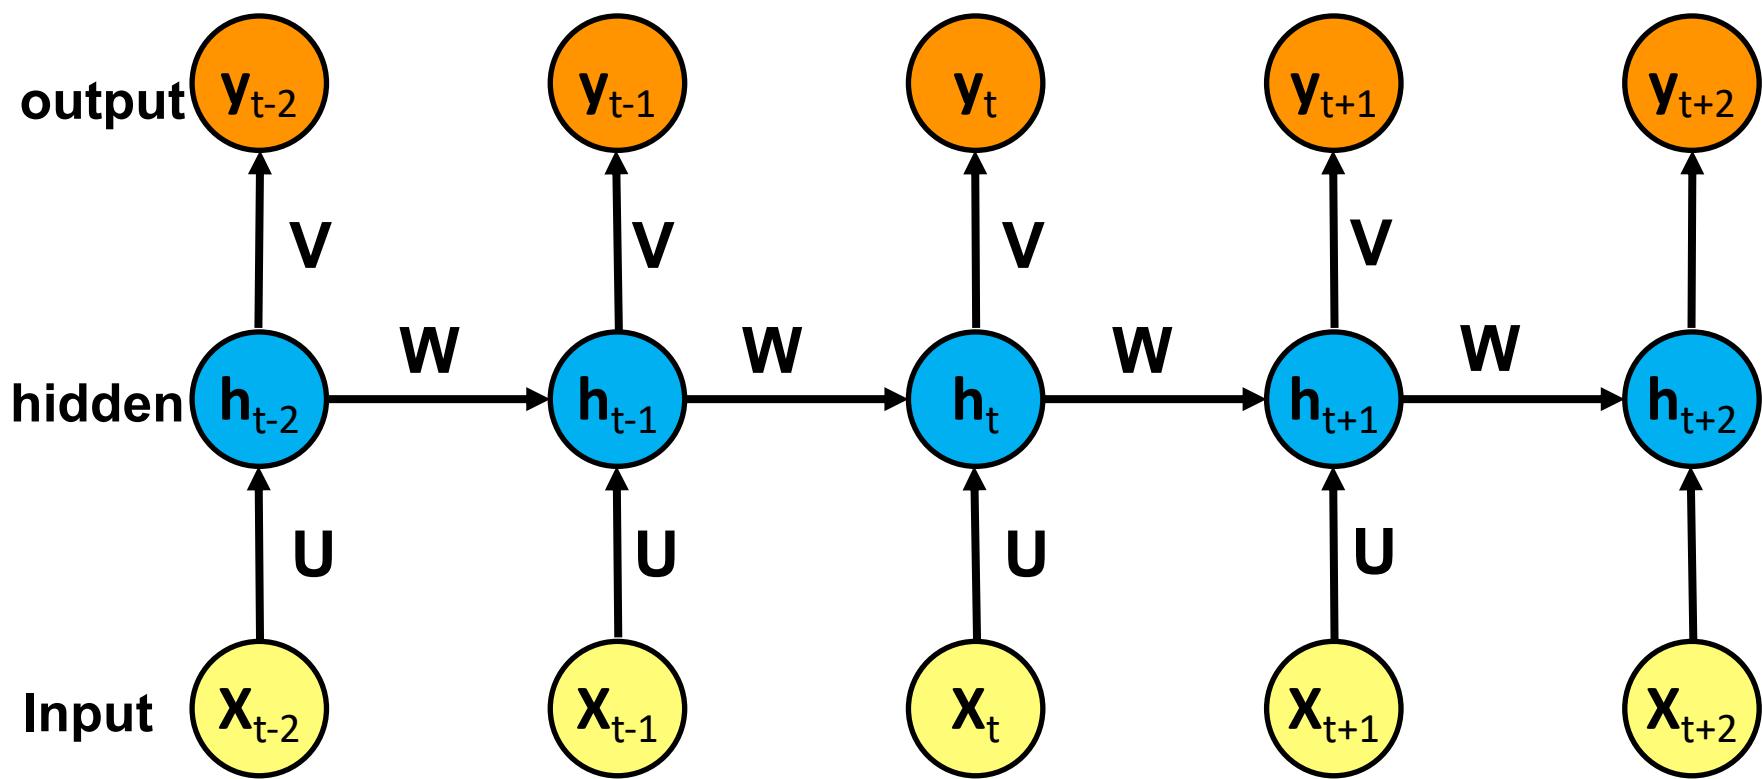


Vanishing Gradient



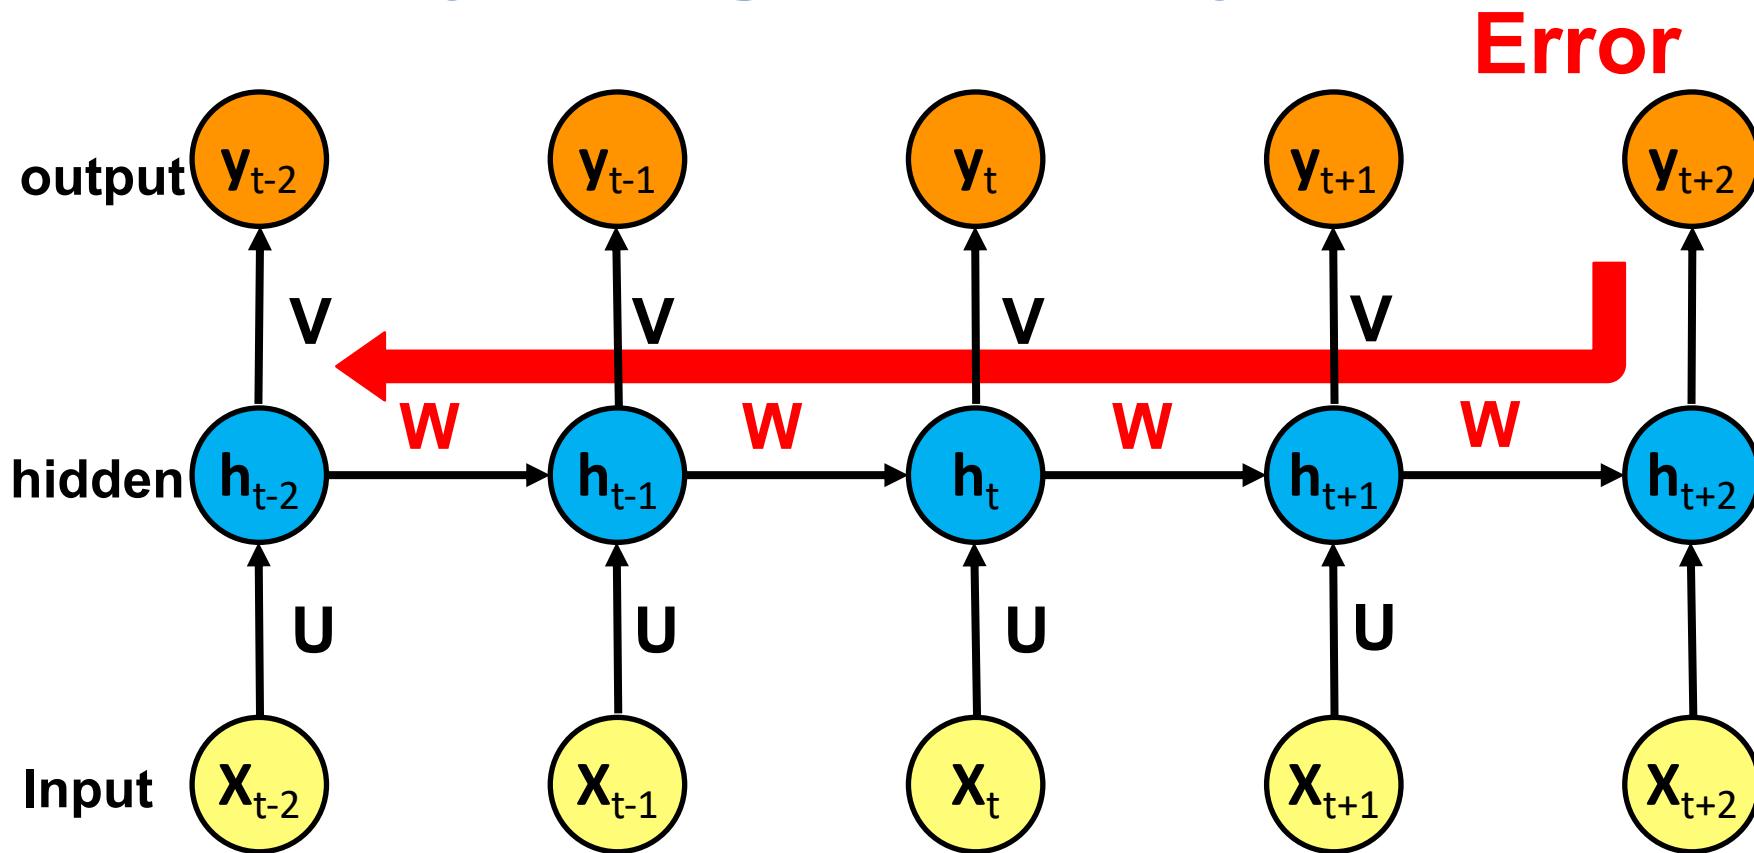
Exploding Gradient

Recurrent Neural Networks (RNN)



RNN

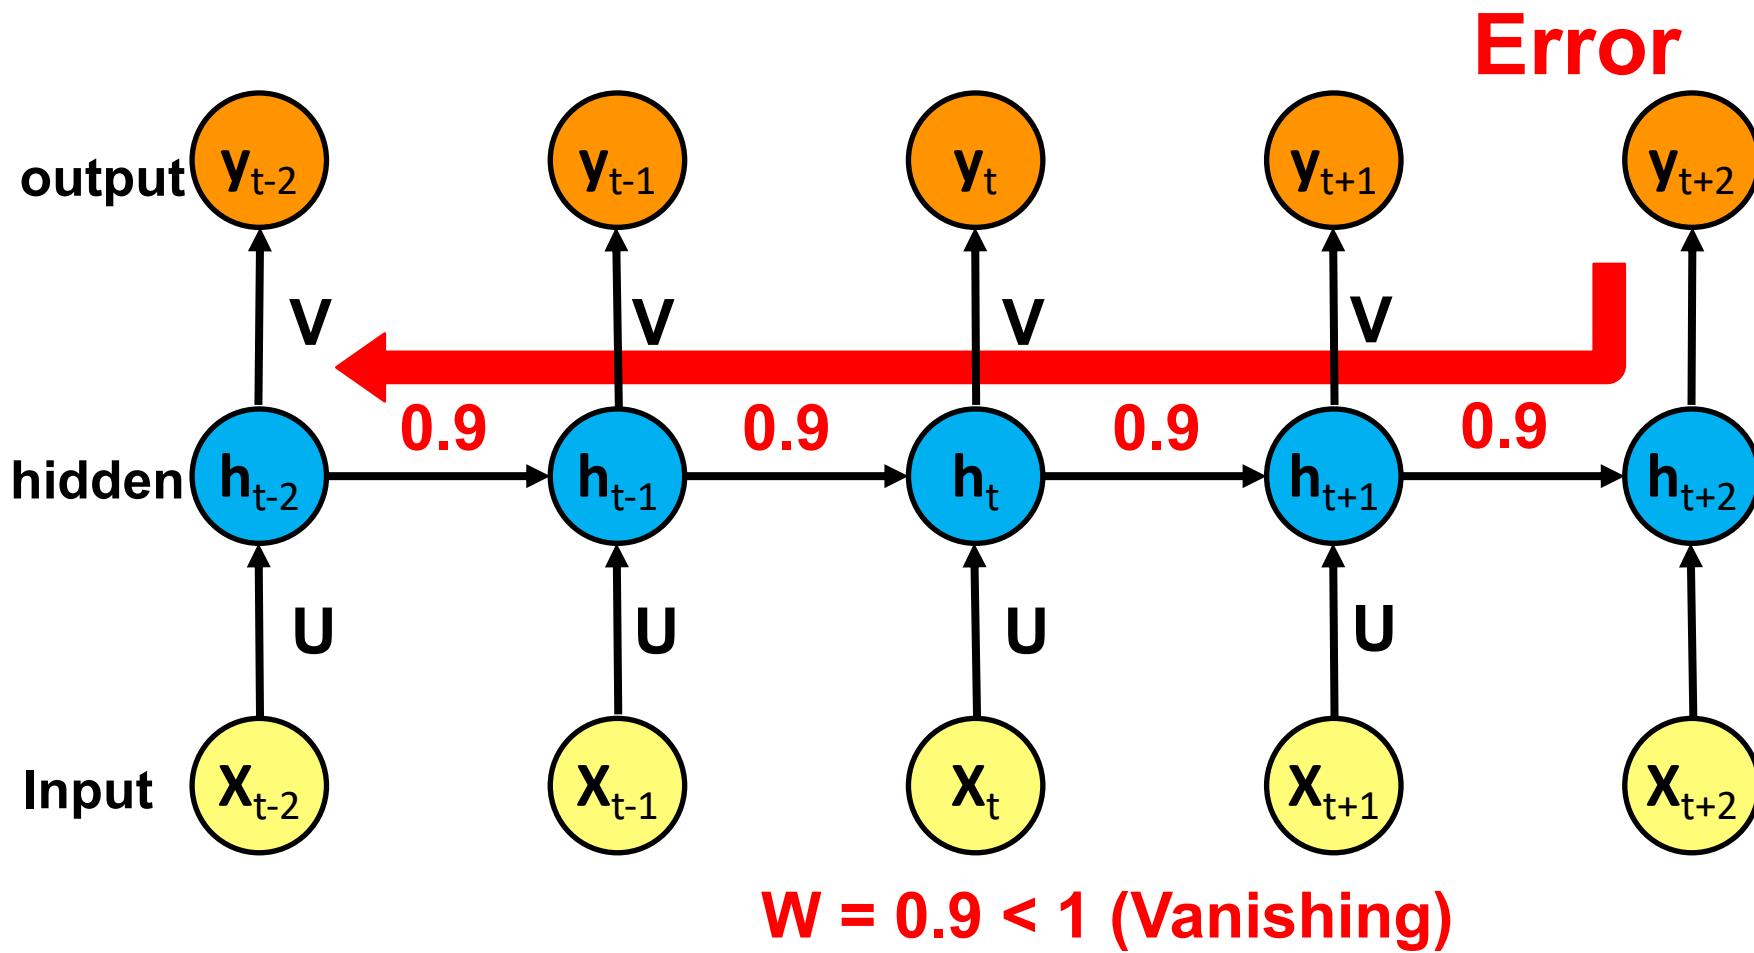
Vanishing Gradient problem Exploding Gradient problem



if $|W| < 1$ (Vanishing)
if $|W| > 1$ (Exploding)

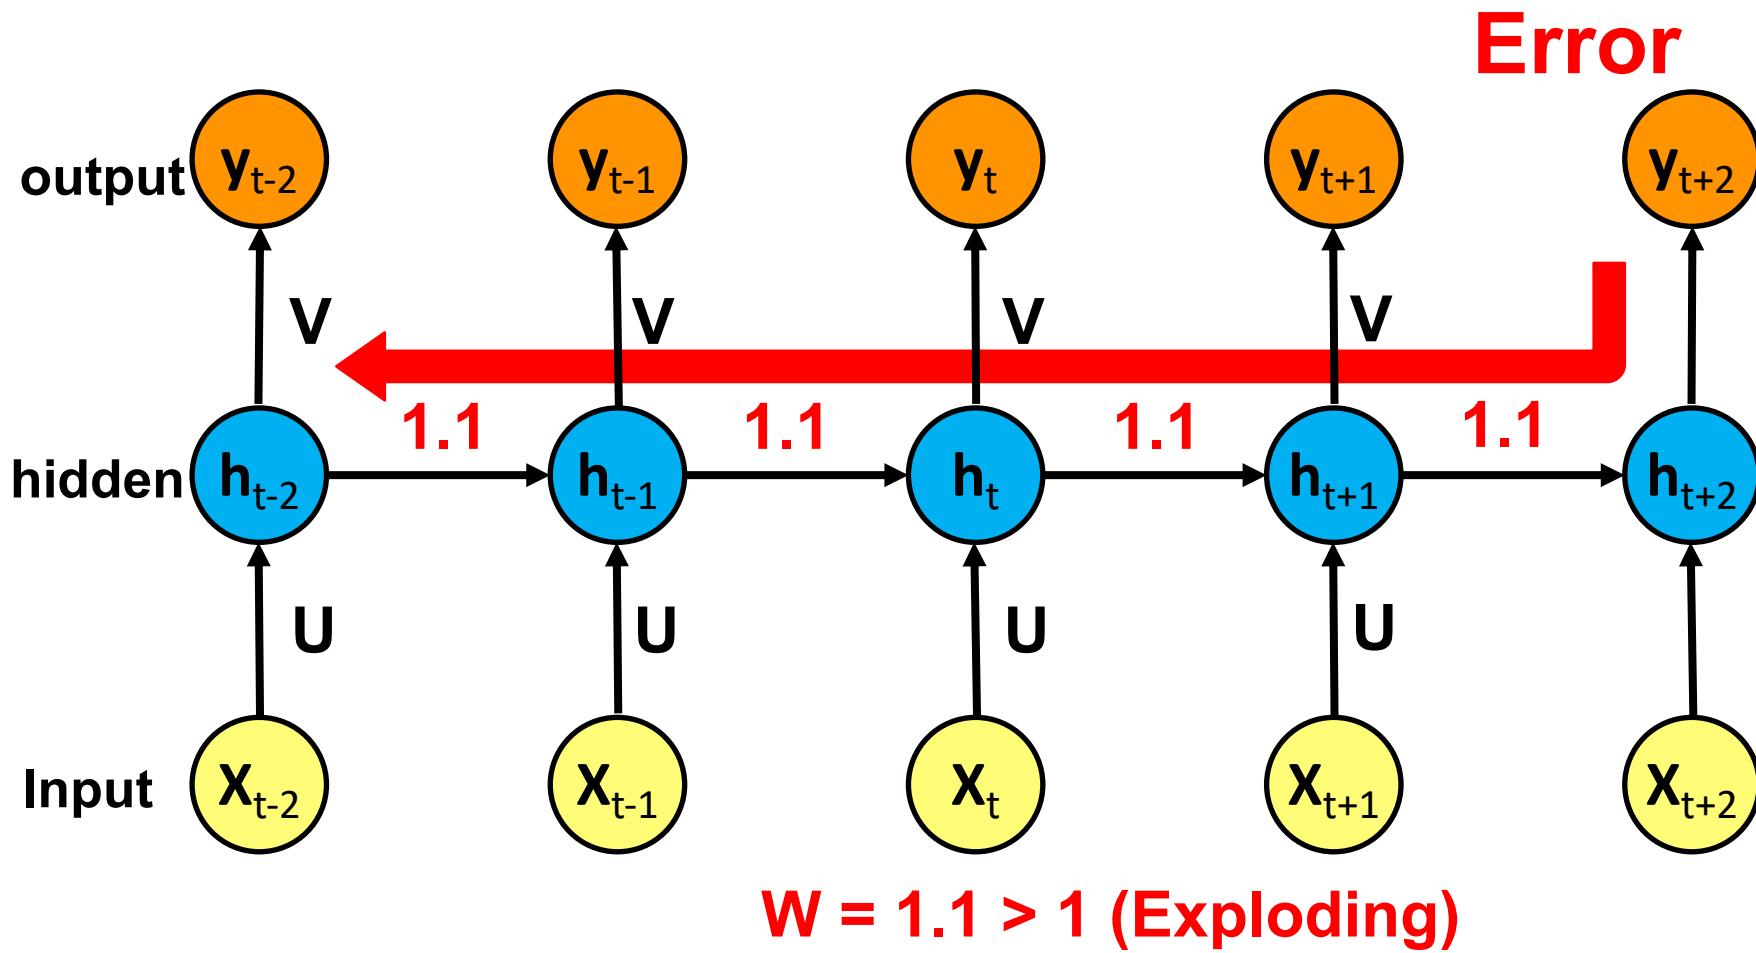
RNN

Vanishing Gradient problem



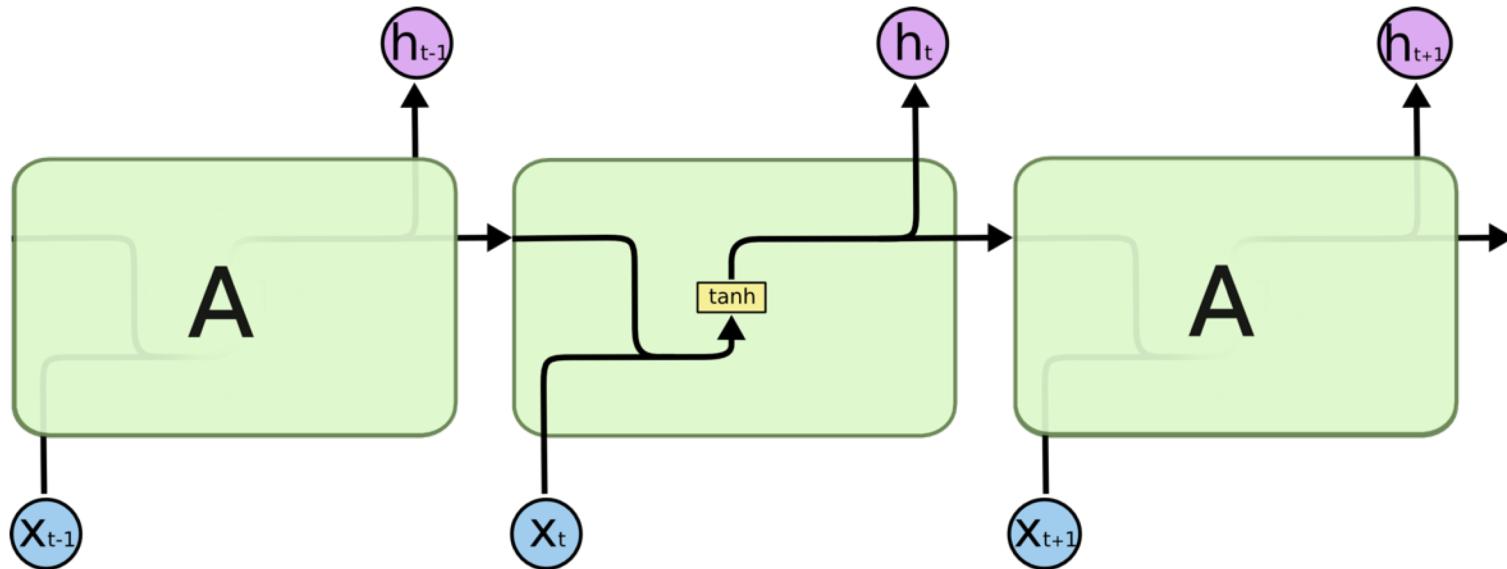
RNN

Exploding Gradient problem

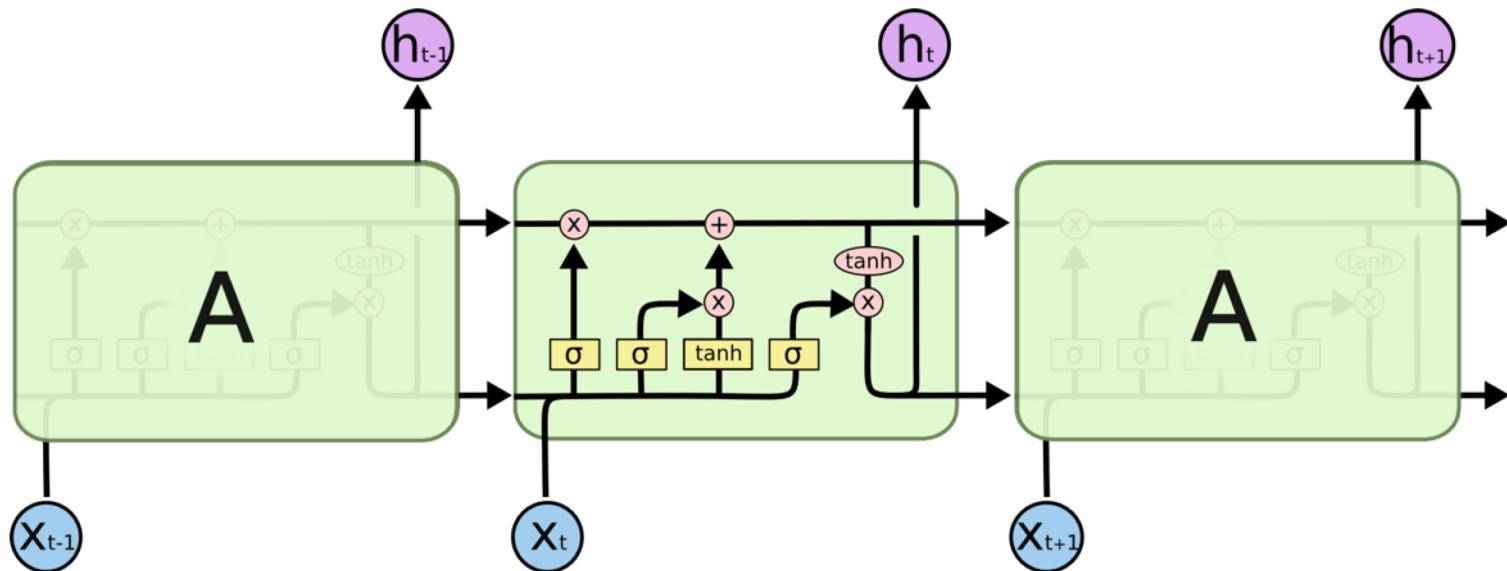


RNN LSTM

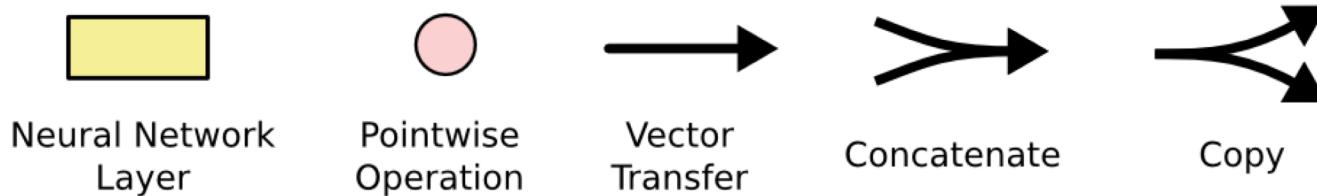
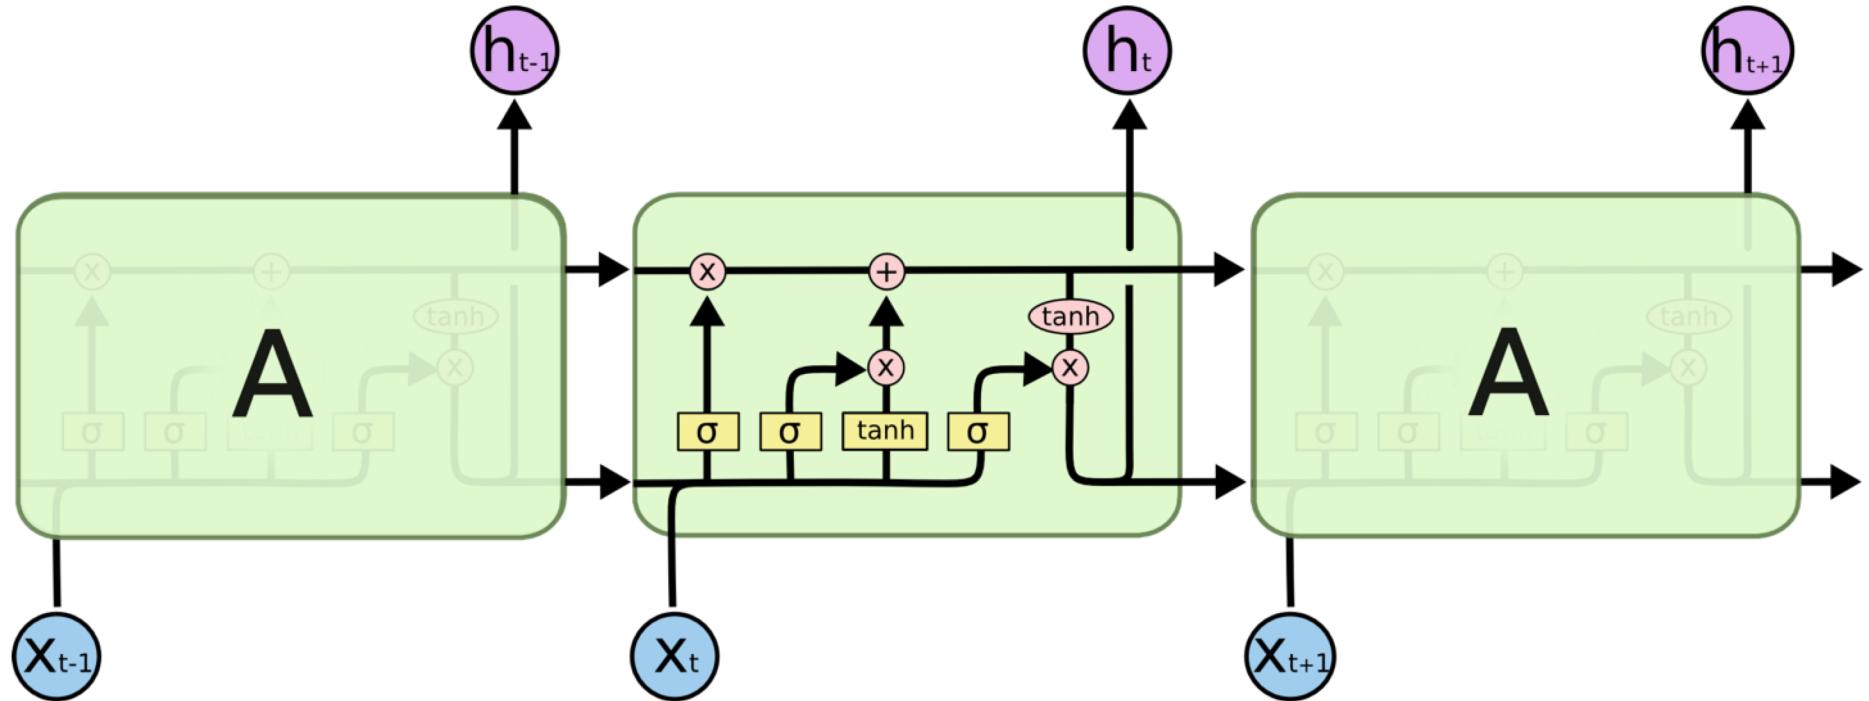
RNN



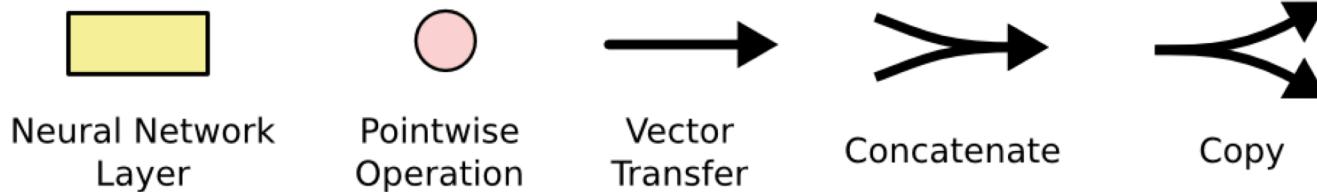
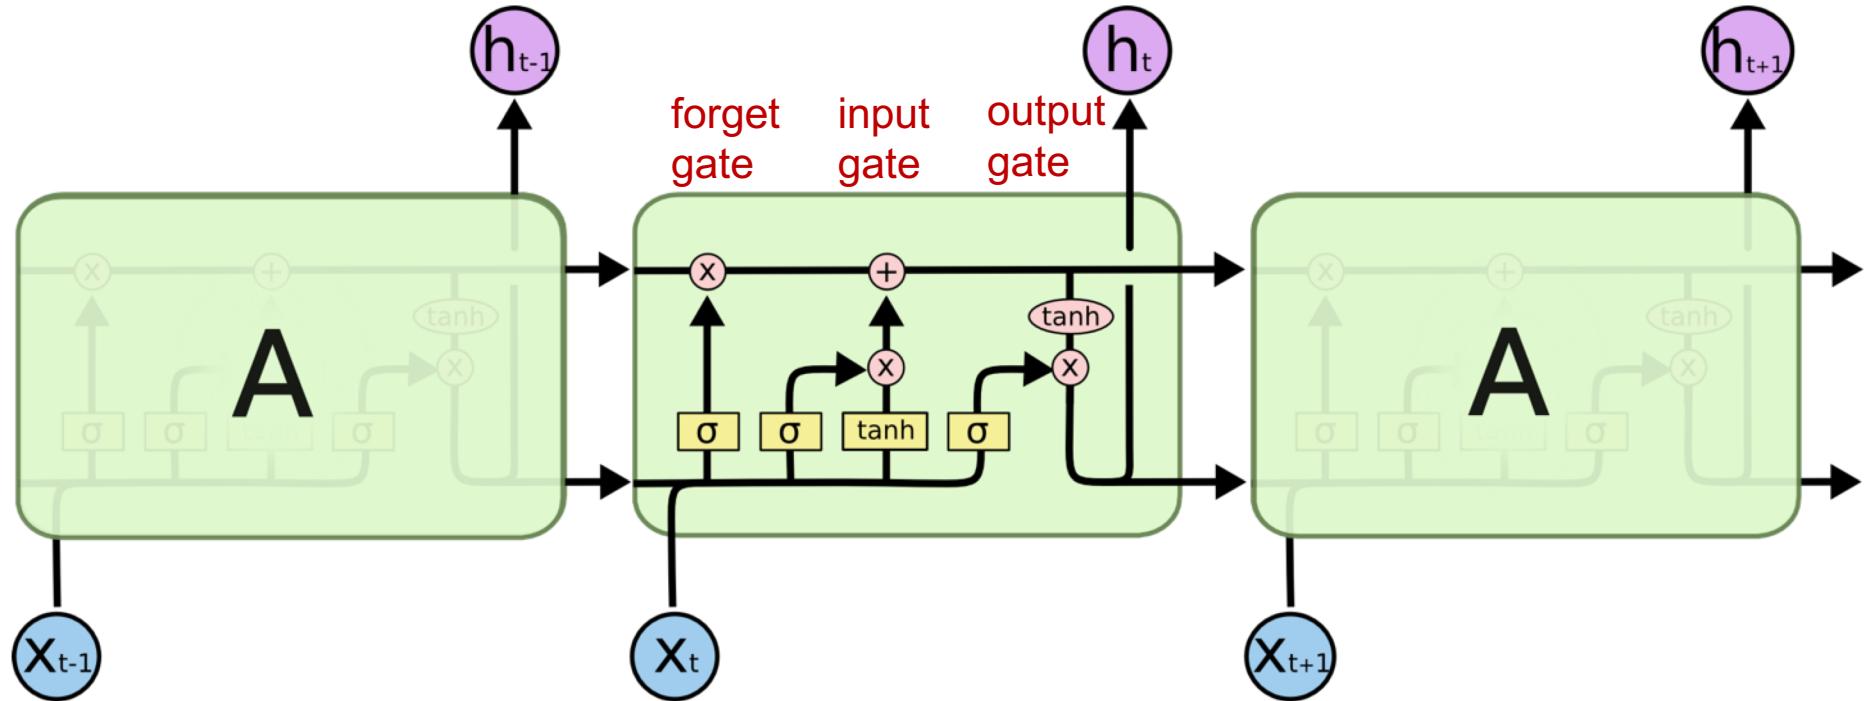
LSTM



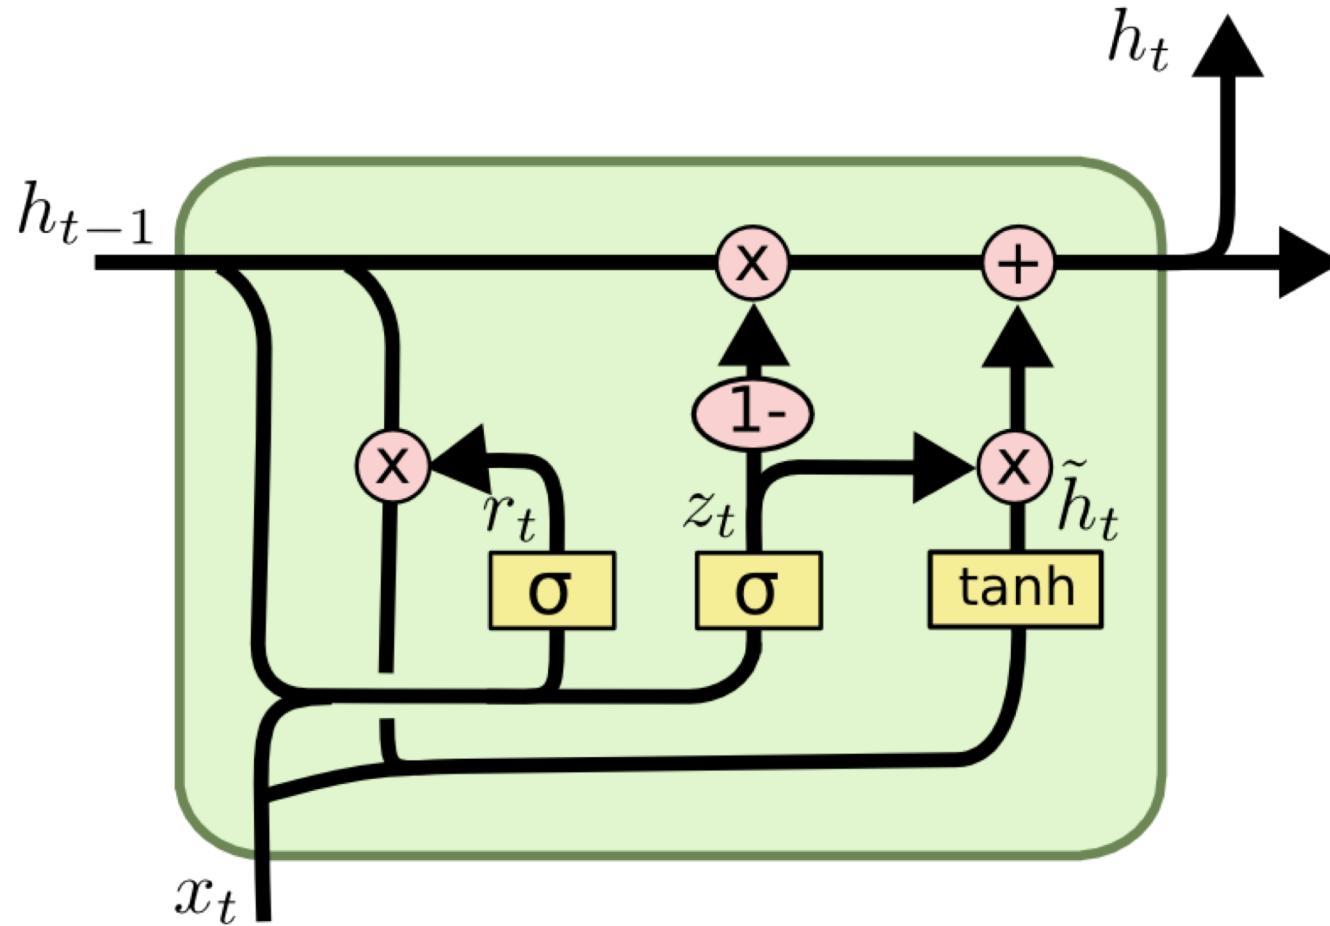
Long Short Term Memory (LSTM)



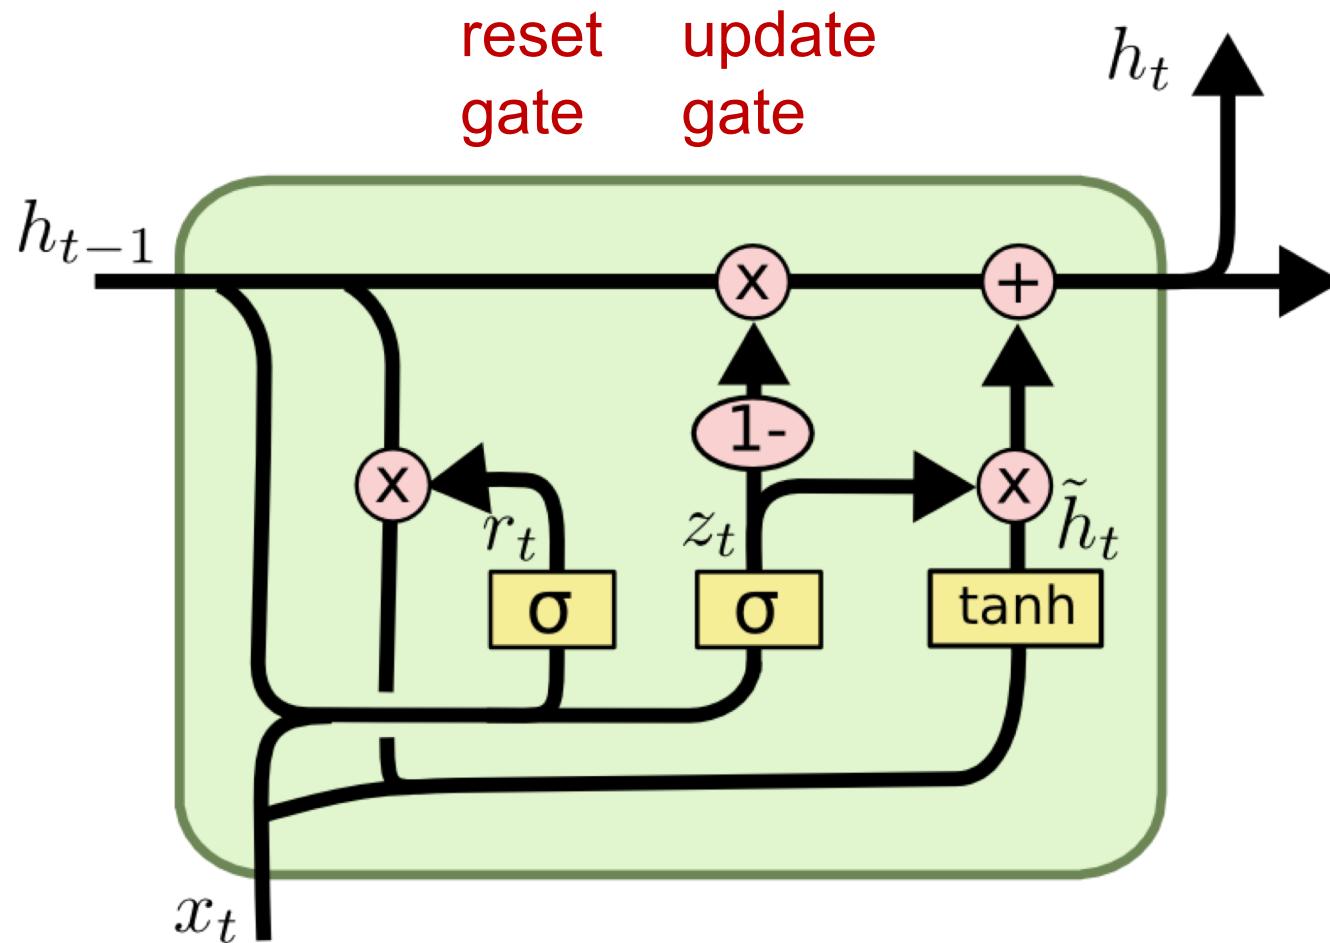
Long Short Term Memory (LSTM)



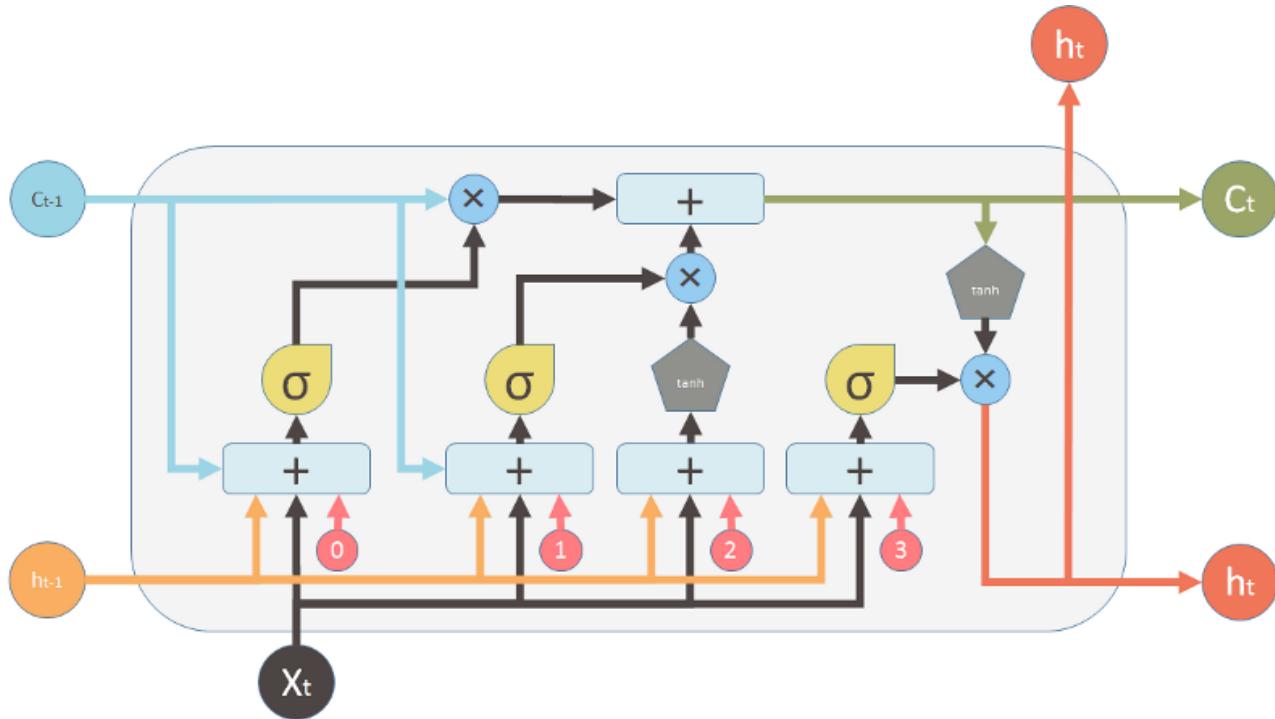
Gated Recurrent Unit (GRU)



Gated Recurrent Unit (GRU)



LSTM



Inputs:



X_t
Input vector



C_{t-1}
Memory from previous block



h_{t-1}
Output of previous block

outputs:



C_t
Memory from current block



h_t
Output of current block

Nonlinearities:



σ
Sigmoid



\tanh
Hyperbolic tangent



0
Bias

Vector operations:

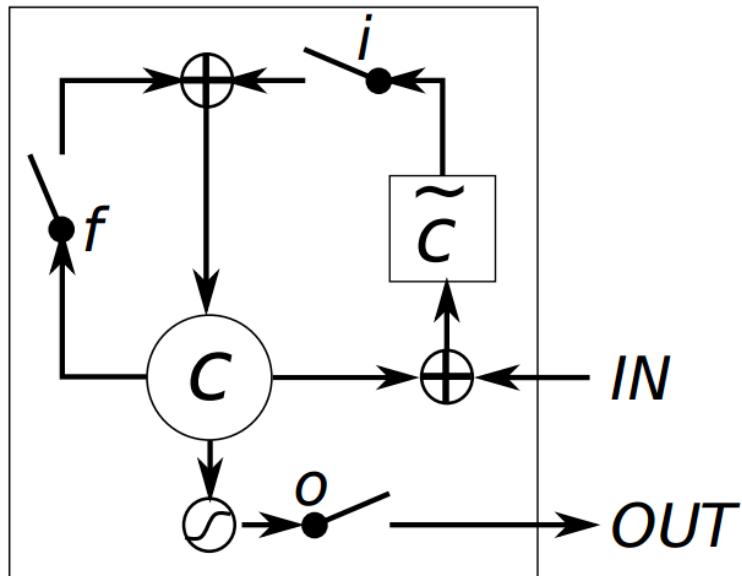


\times
Element-wise multiplication



$+$
Element-wise Summation / Concatenation

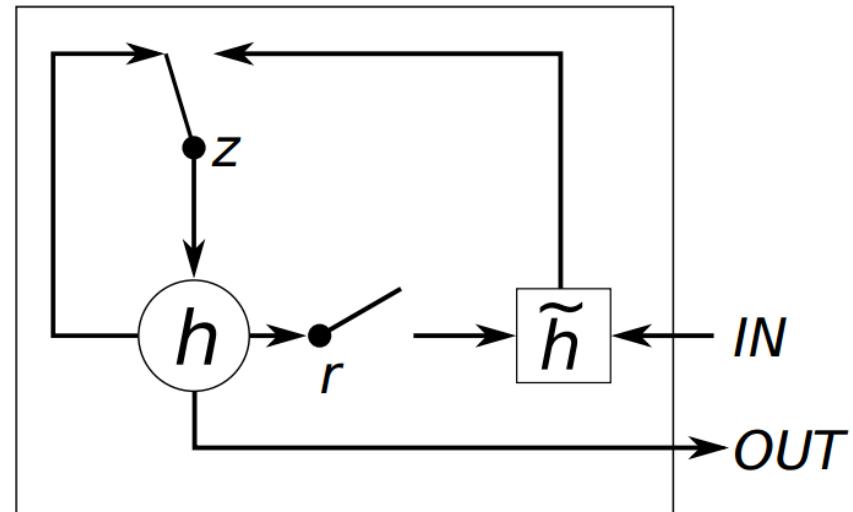
LSTM vs GRU



LSTM

i , f and o are the **input**, **forget** and **output** gates, respectively.

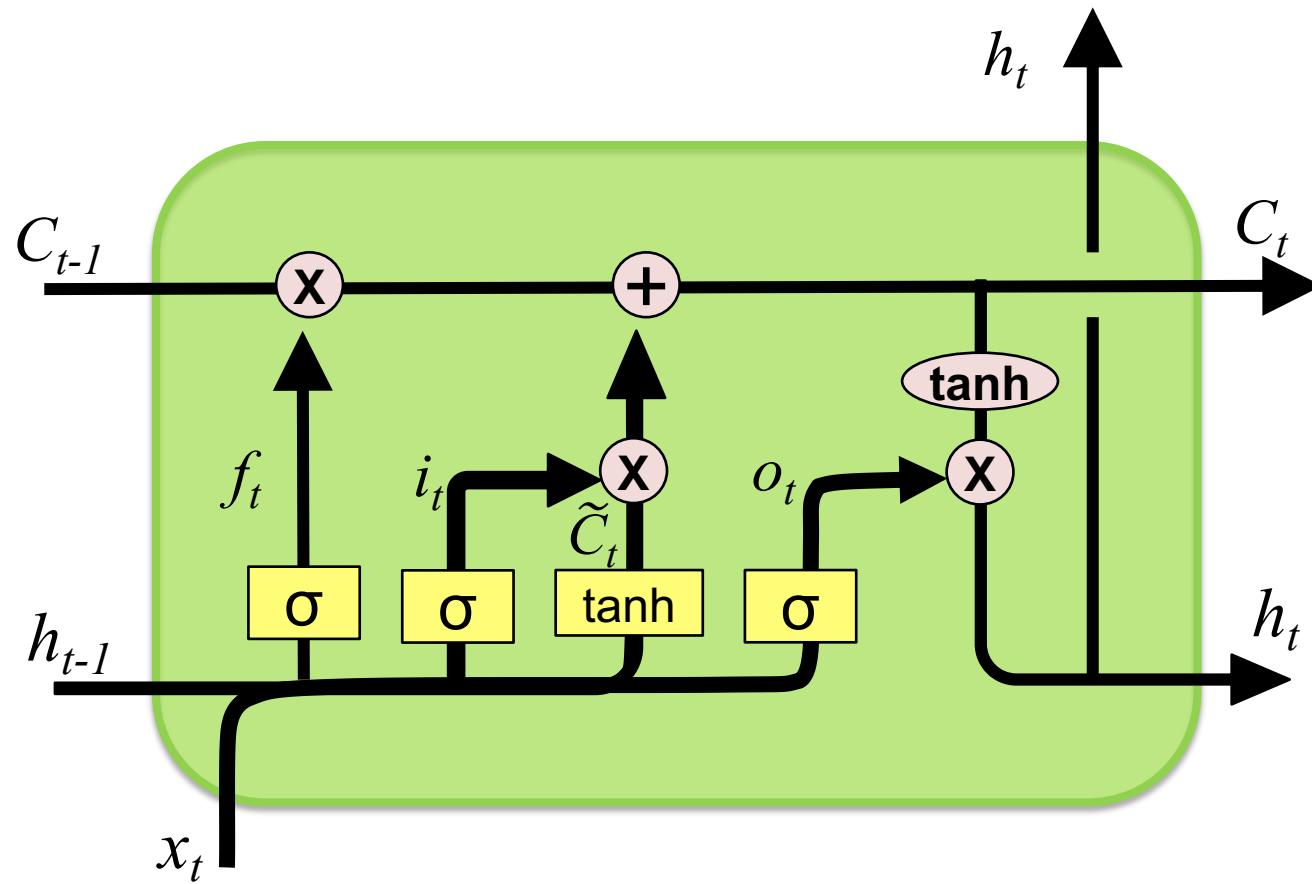
c and \tilde{c} denote the **memory cell** and the **new memory cell content**.



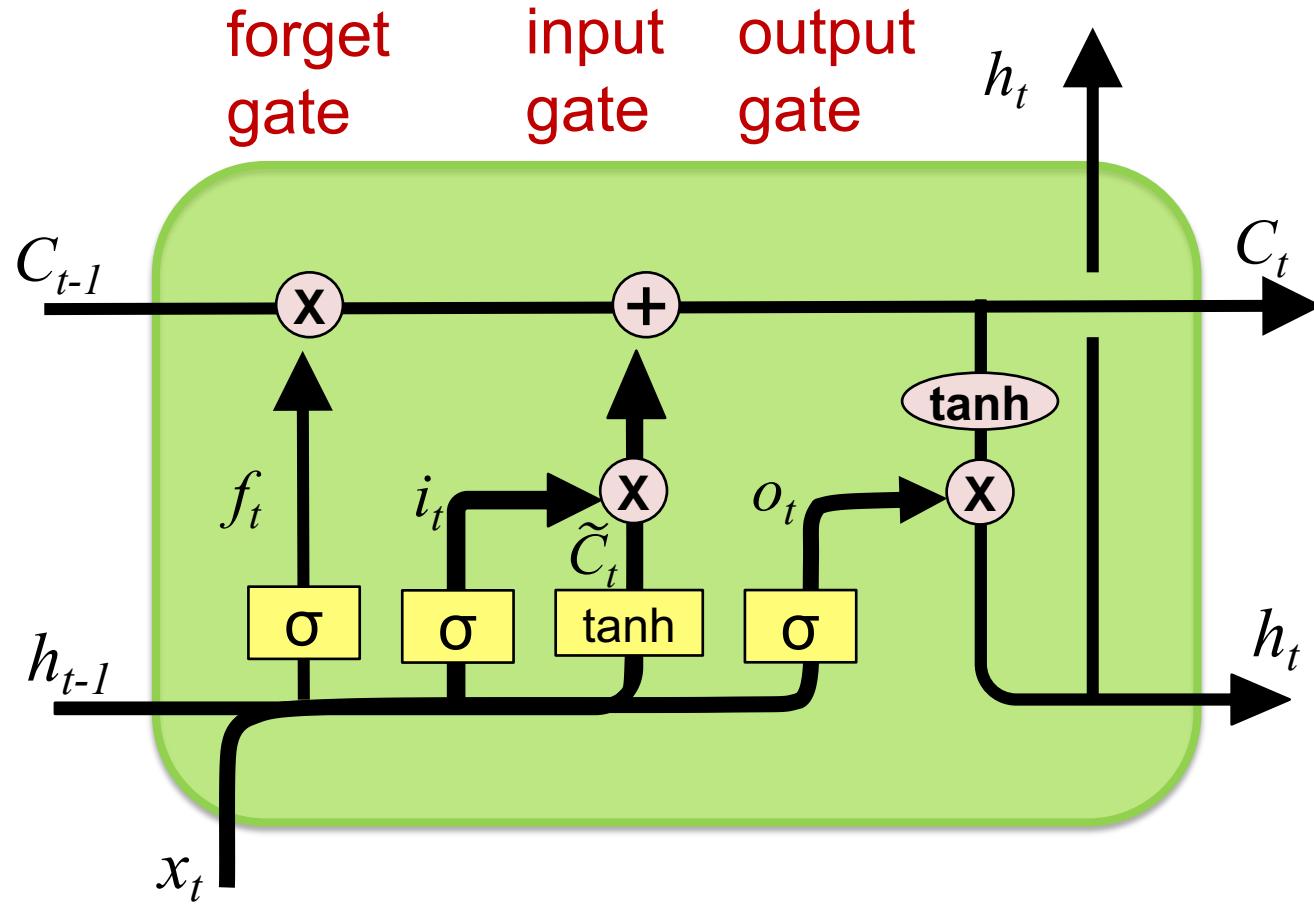
GRU

r and z are the **reset** and **update** gates, and h and \tilde{h} are the **activation** and the **candidate activation**.

Long Short Term Memory (LSTM)

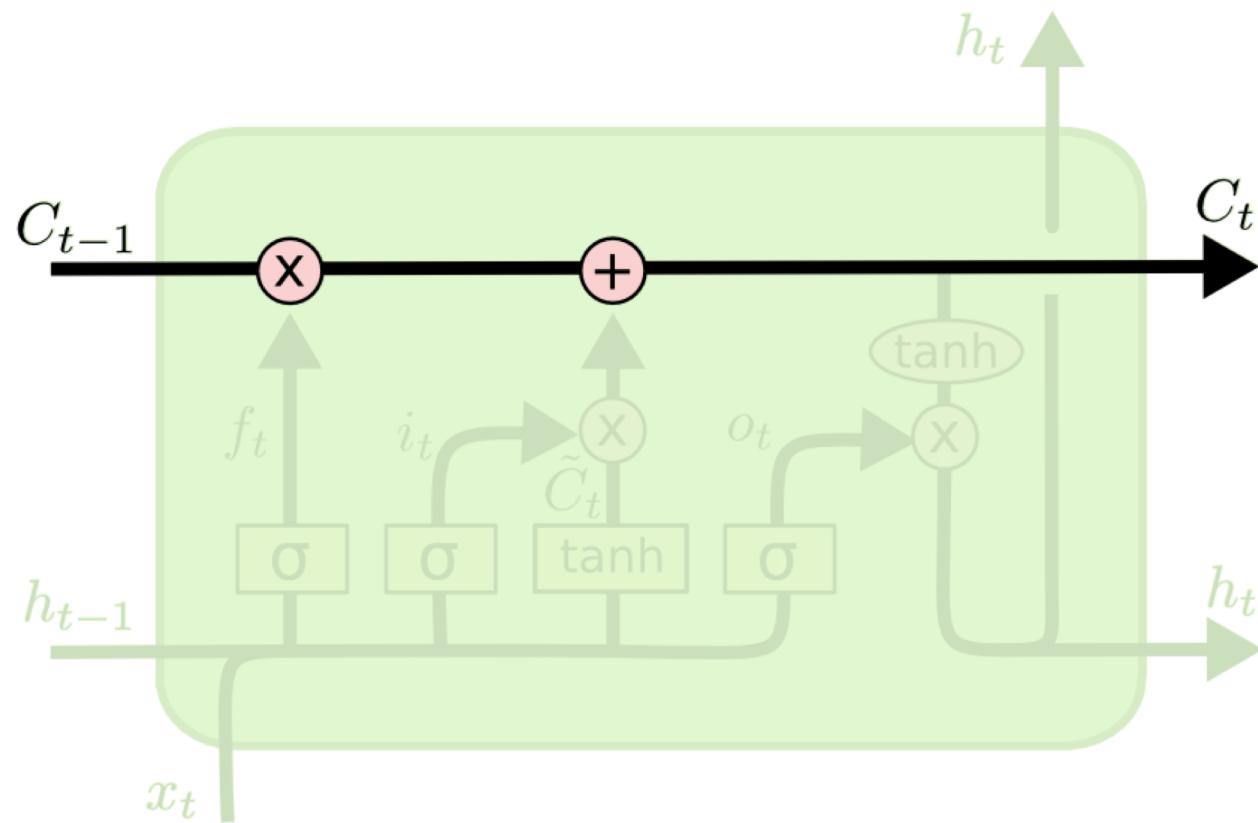


Long Short Term Memory (LSTM)



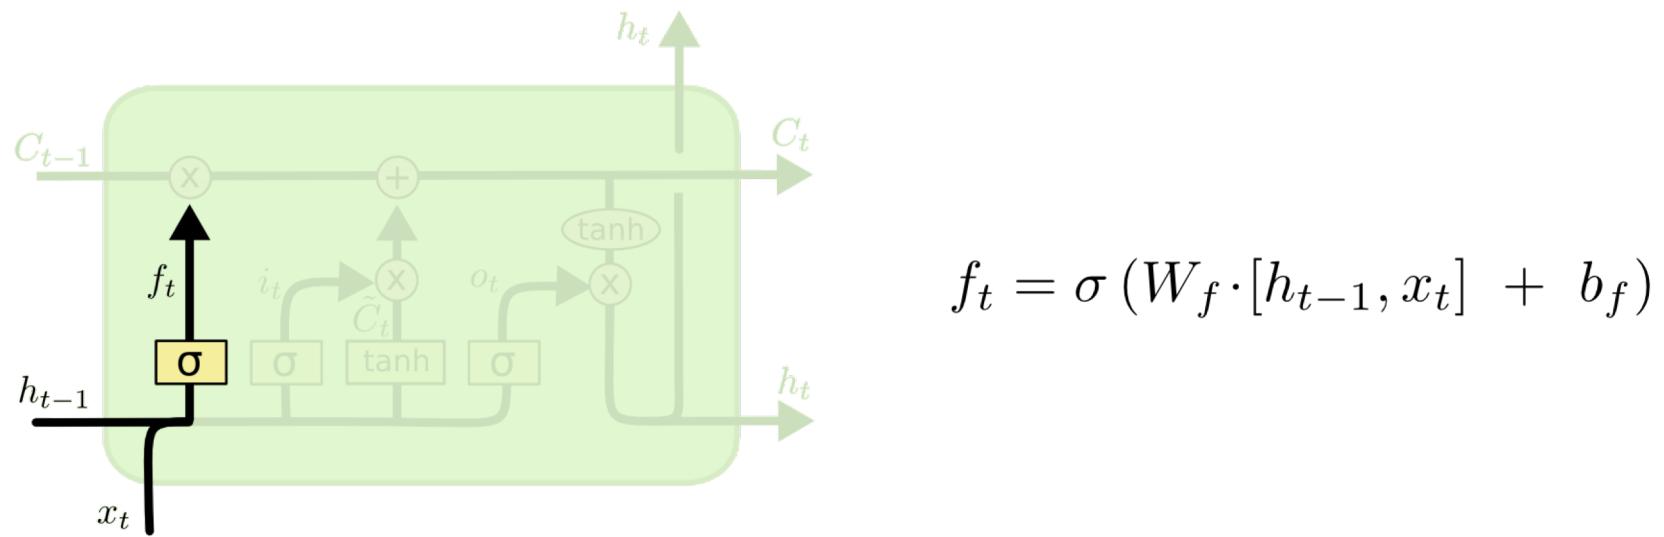
LSTM

Memory state (C)



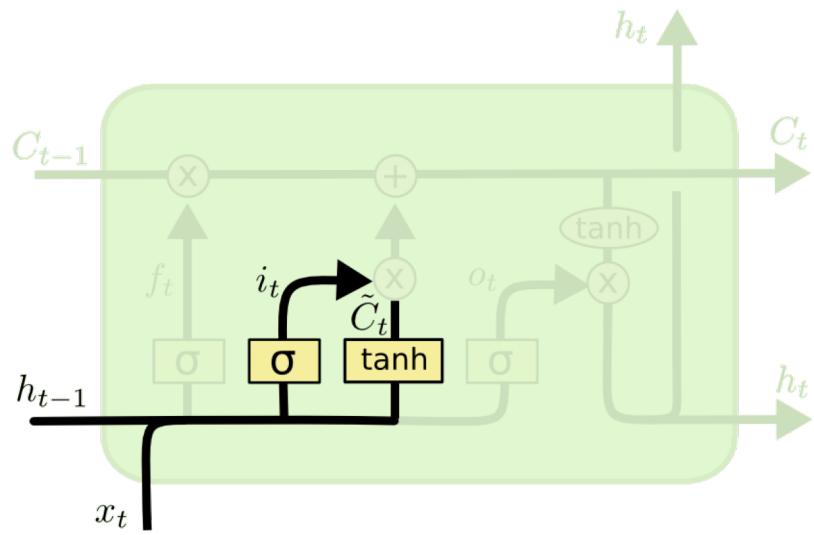
LSTM

forget gate (f)



LSTM

input gate (i)

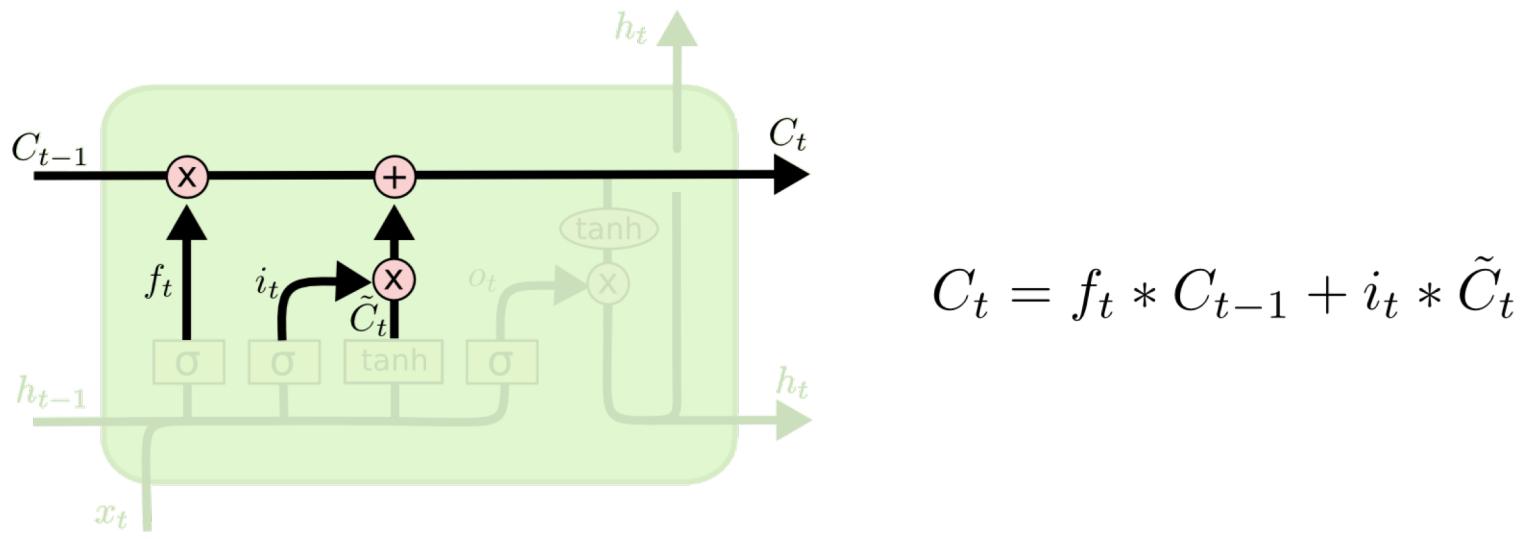


$$i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i)$$

$$\tilde{C}_t = \tanh(W_C \cdot [h_{t-1}, x_t] + b_C)$$

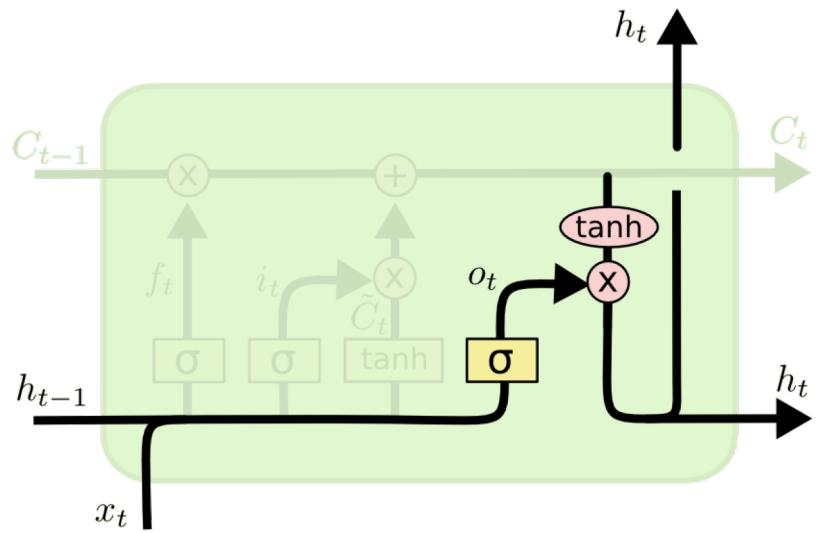
LSTM

Memory state (C)



LSTM

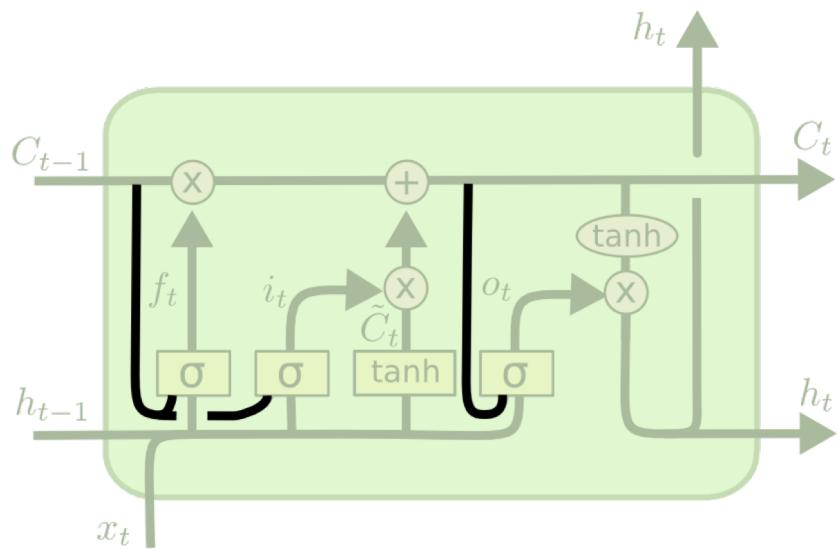
output gate (o)



$$o_t = \sigma (W_o [h_{t-1}, x_t] + b_o)$$
$$h_t = o_t * \tanh (C_t)$$

LSTM

forget (f), input (i), output (o) gates



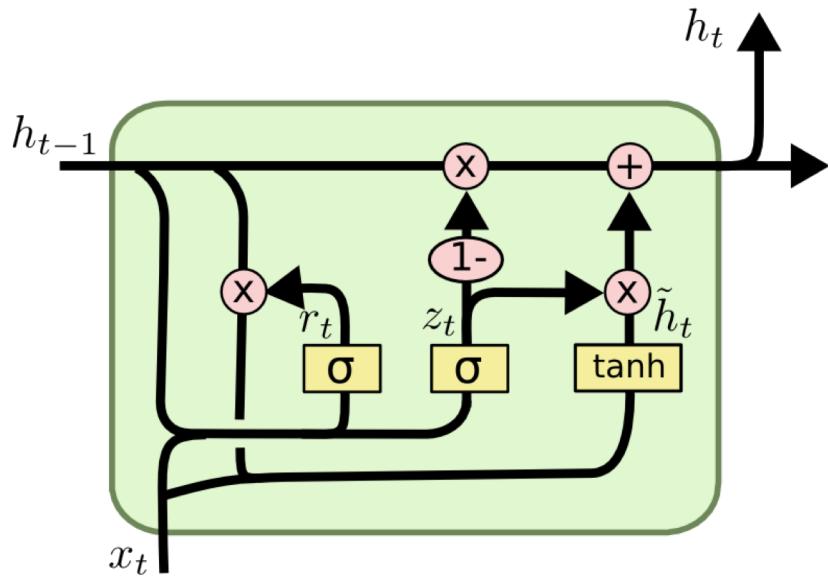
$$f_t = \sigma(W_f \cdot [C_{t-1}, h_{t-1}, x_t] + b_f)$$

$$i_t = \sigma(W_i \cdot [C_{t-1}, h_{t-1}, x_t] + b_i)$$

$$o_t = \sigma(W_o \cdot [C_t, h_{t-1}, x_t] + b_o)$$

Gated Recurrent Unit (GRU)

update (z), reset (r) gates



$$z_t = \sigma (W_z \cdot [h_{t-1}, x_t])$$

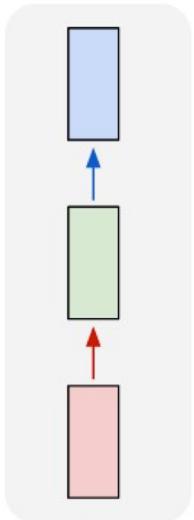
$$r_t = \sigma (W_r \cdot [h_{t-1}, x_t])$$

$$\tilde{h}_t = \tanh (W \cdot [r_t * h_{t-1}, x_t])$$

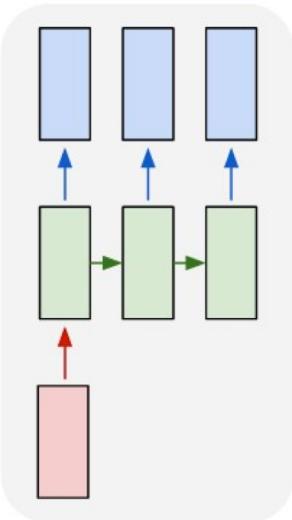
$$h_t = (1 - z_t) * h_{t-1} + z_t * \tilde{h}_t$$

LSTM Recurrent Neural Network

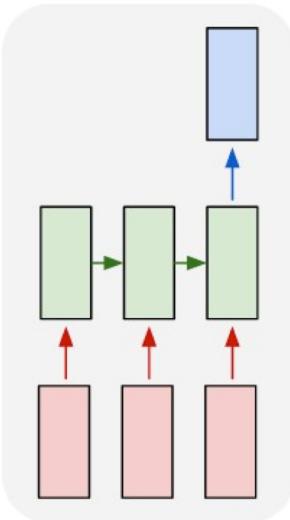
one to one



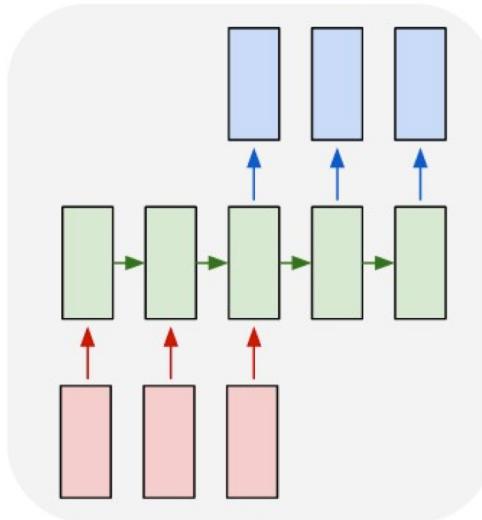
one to many



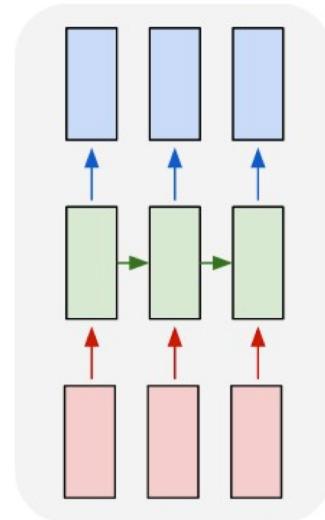
many to one



many to many



many to many



Traditional
Neural
Network

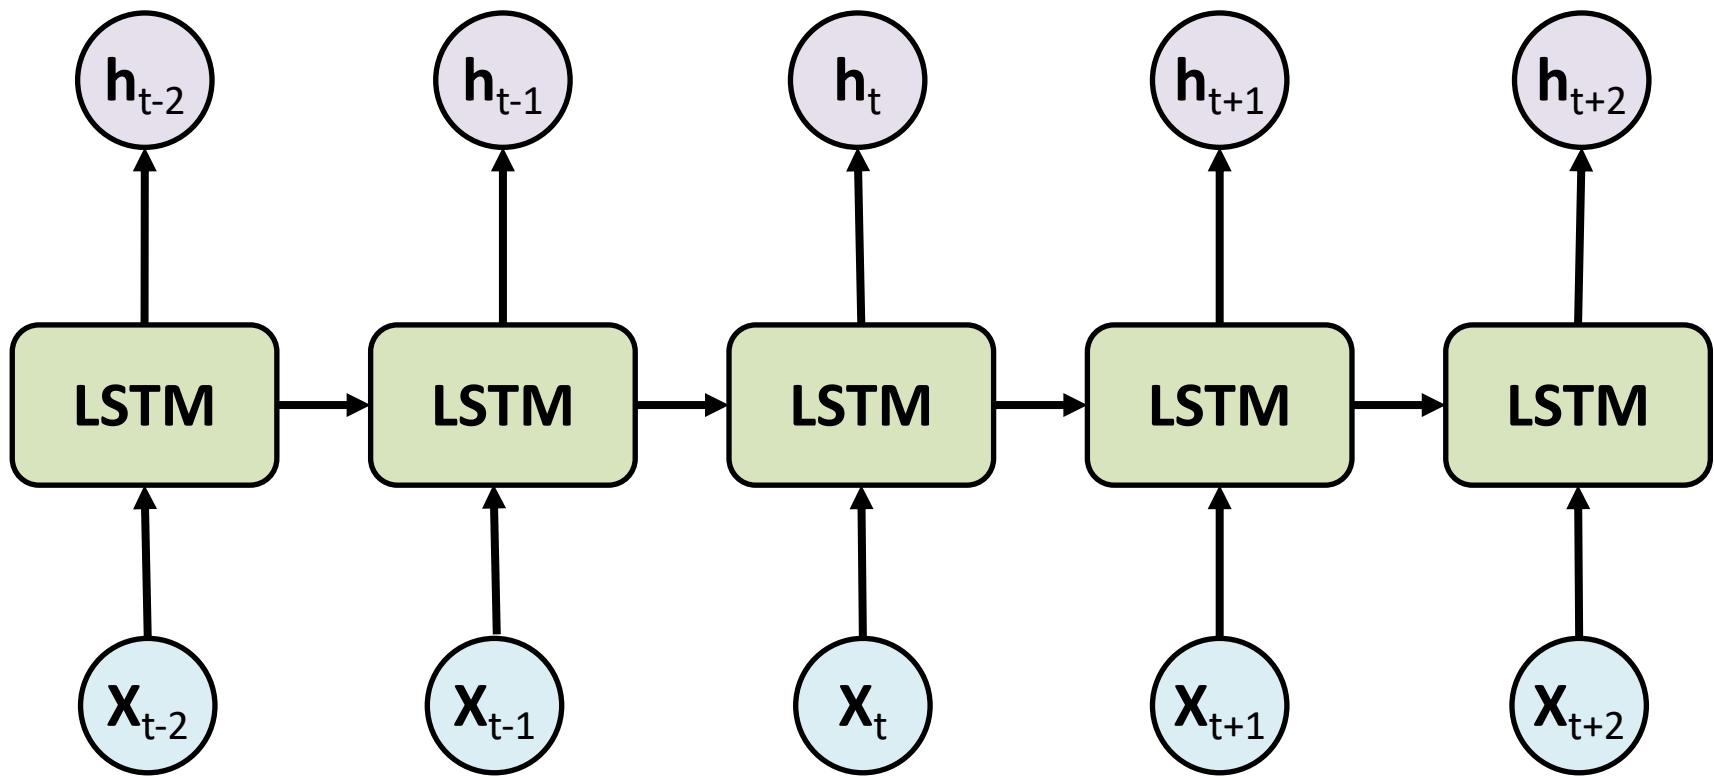
Music
Generation

Sentiment
Classification

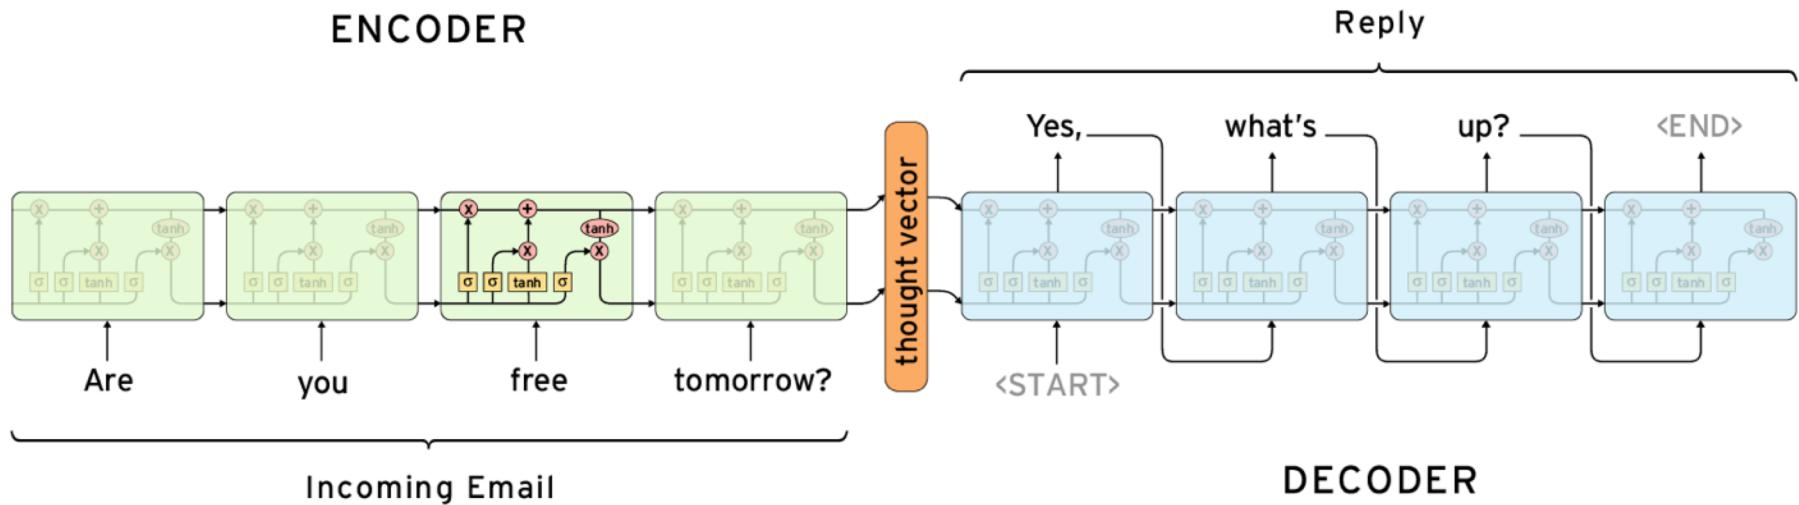
Name
Entity
Recognition

Machine
Translation

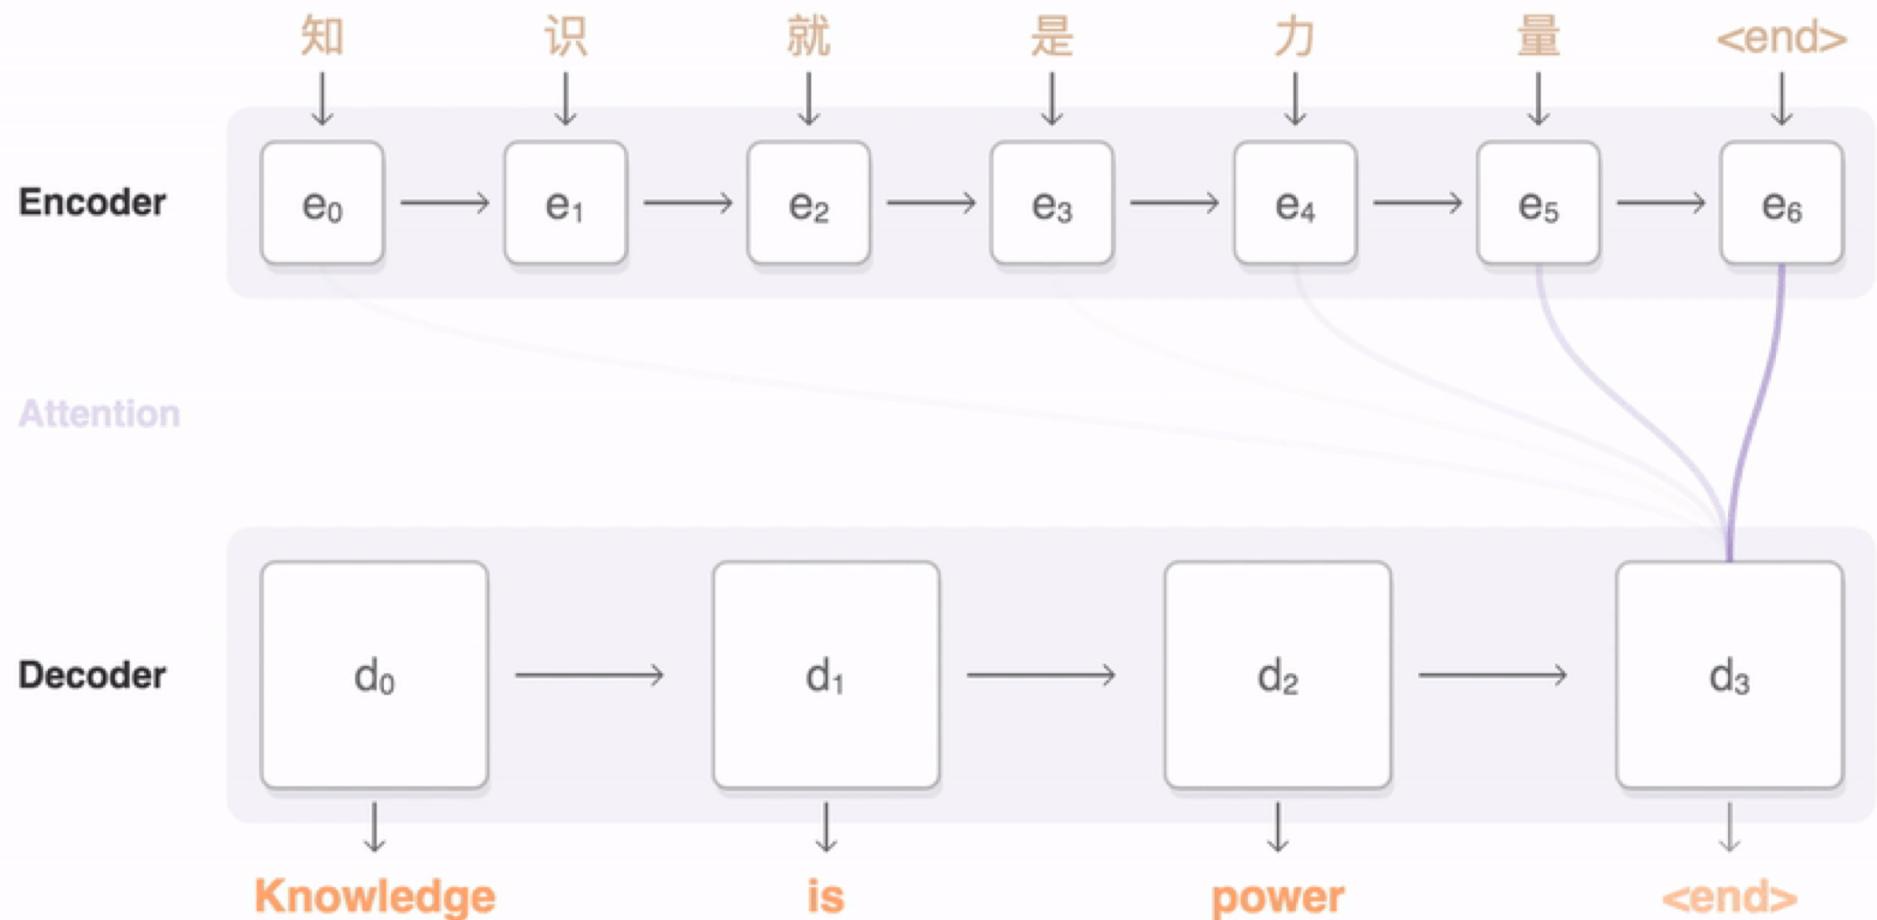
Long Short Term Memory (LSTM) for Time Series Forecasting



The Sequence to Sequence model (seq2seq)

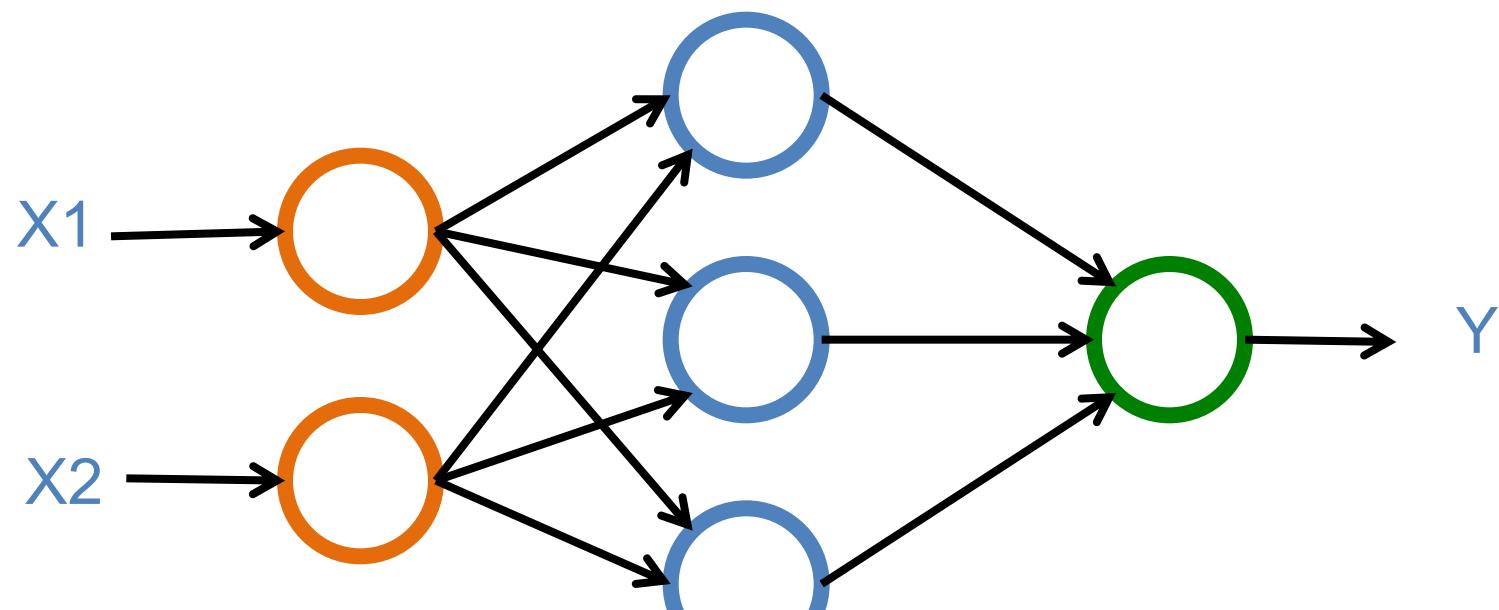


Sequence to Sequence (Seq2Seq)

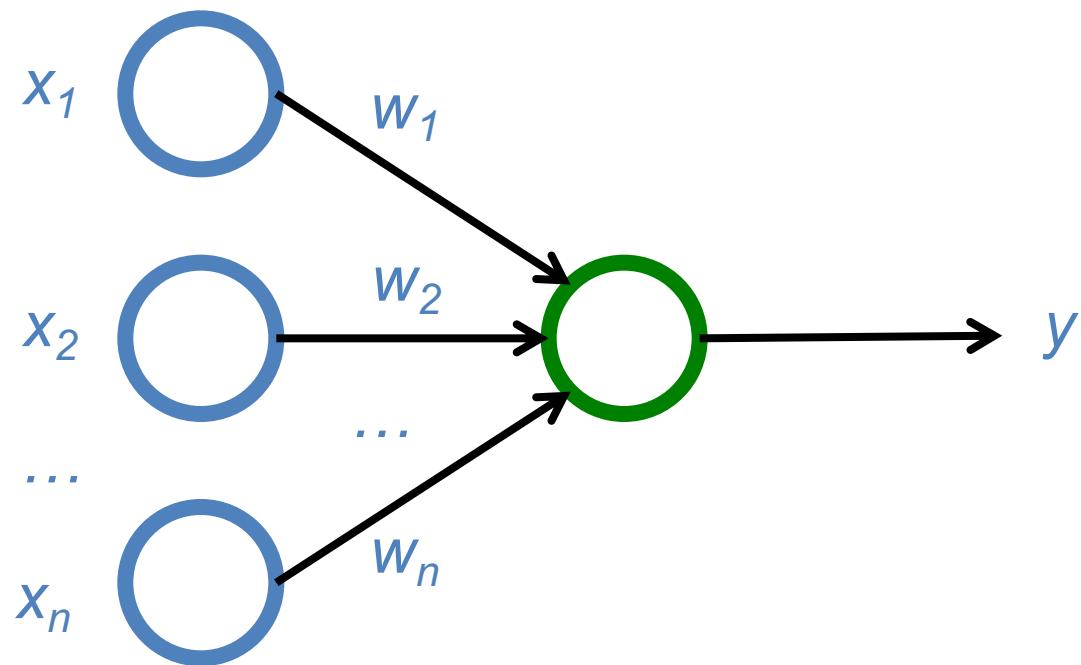


Neural Networks

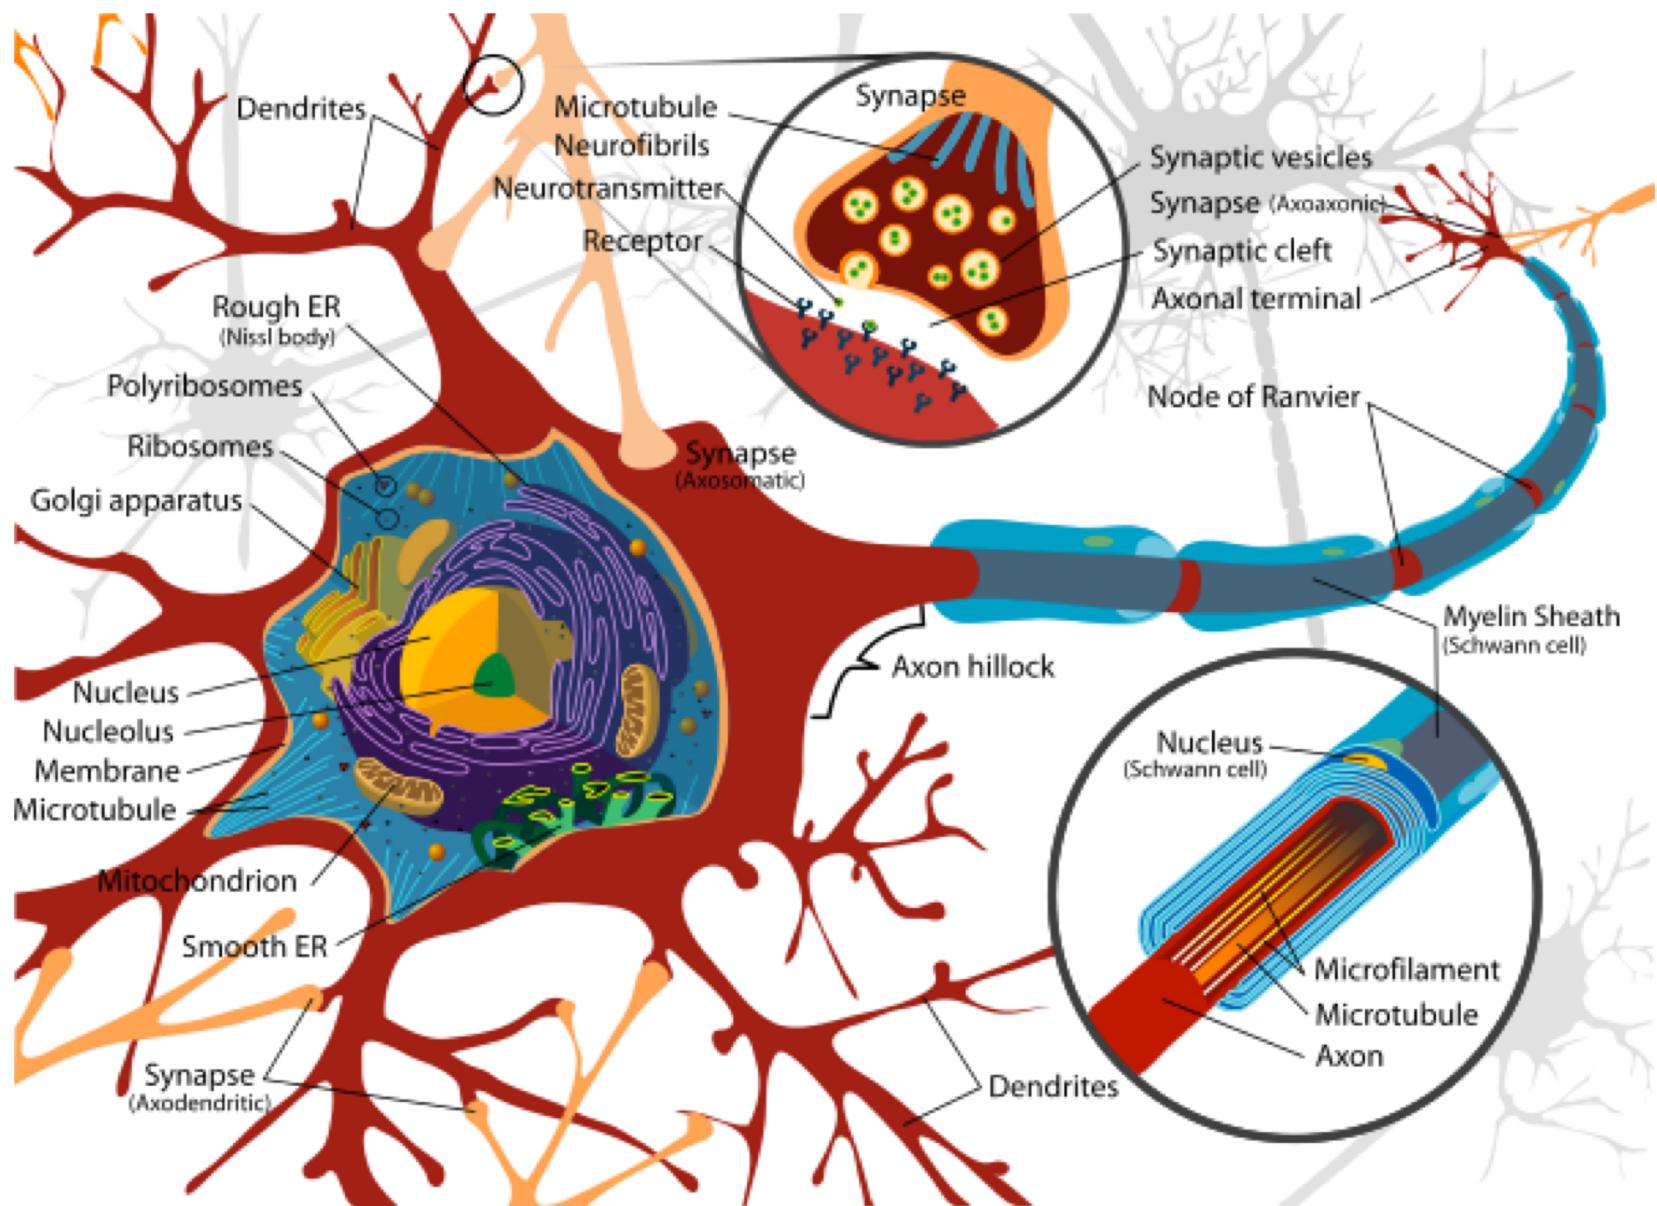
Input Layer **Hidden Layer** **Output Layer**
(X) **(H)** **(Y)**



The Neuron

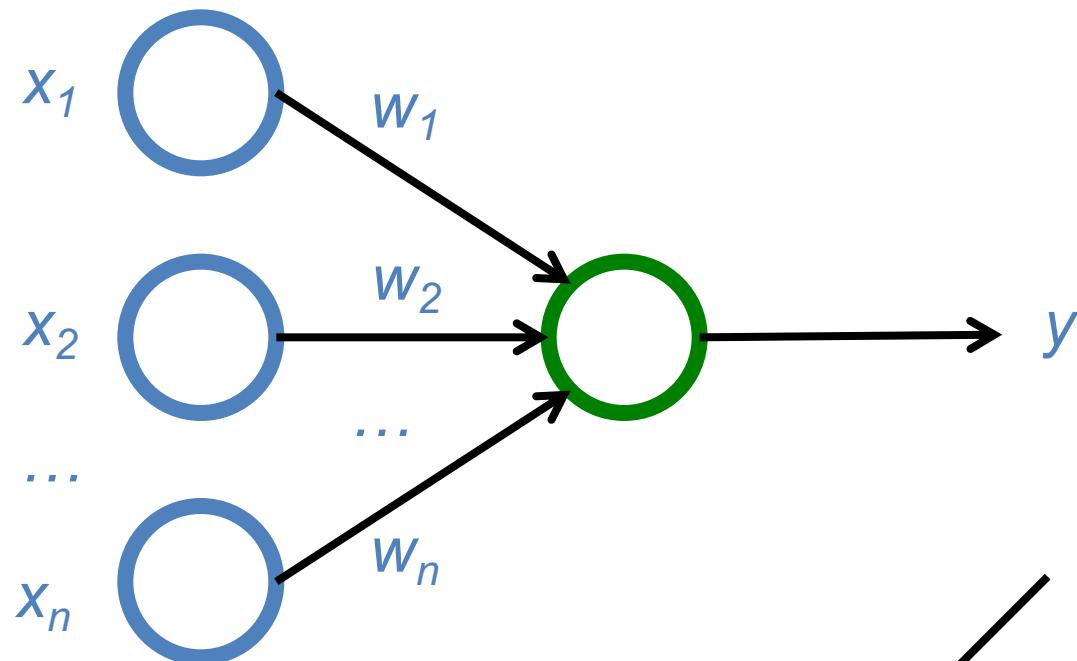


Neuron and Synapse



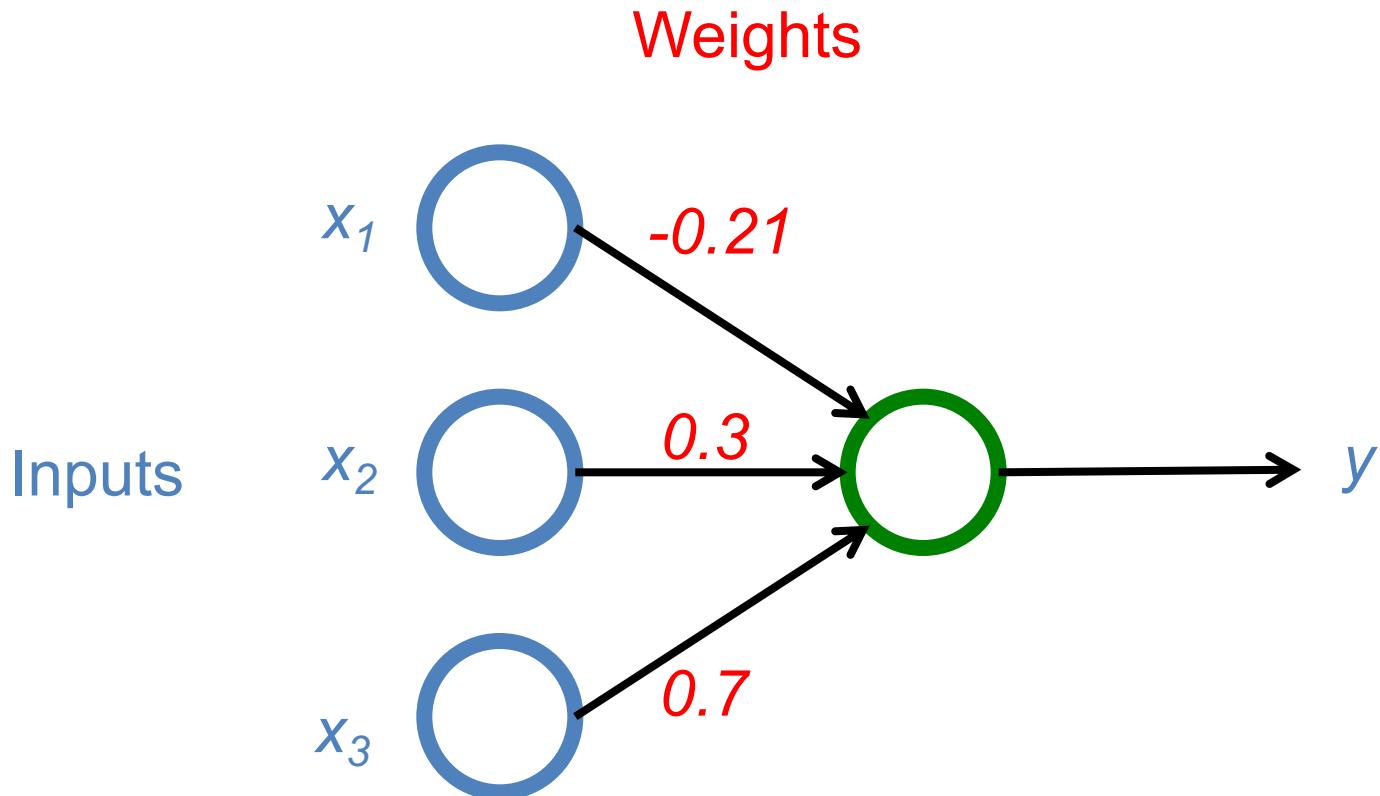
The Neuron

$$y = F\left(\sum_i w_i x_i\right)$$

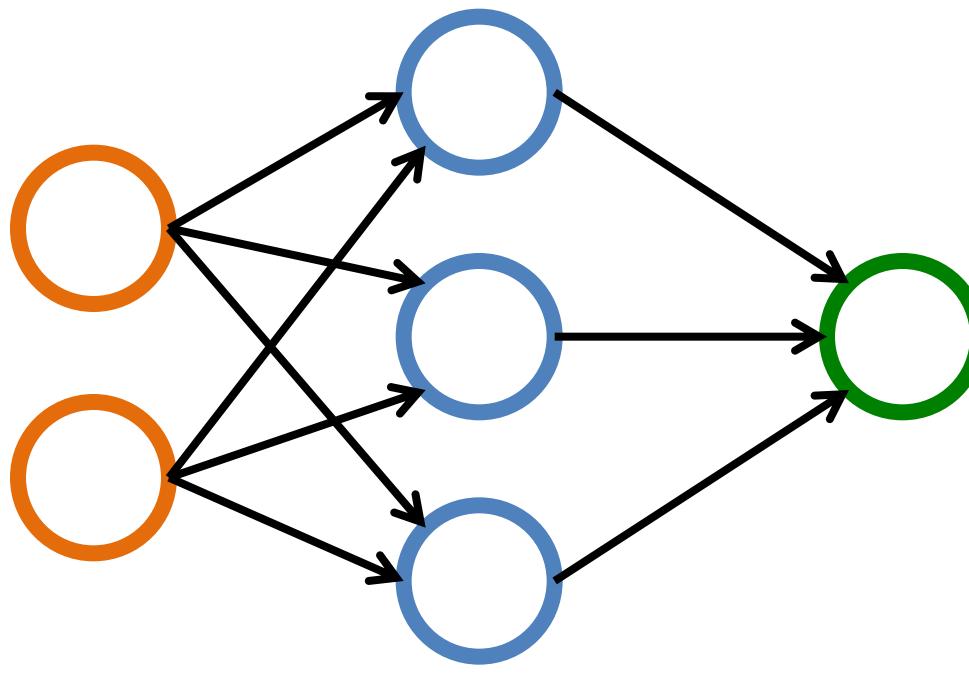


$$F(x) = \max(0, x)$$

$$y = \max(0, -0.21 * x_1 + 0.3 * x_2 + 0.7 * x_3)$$

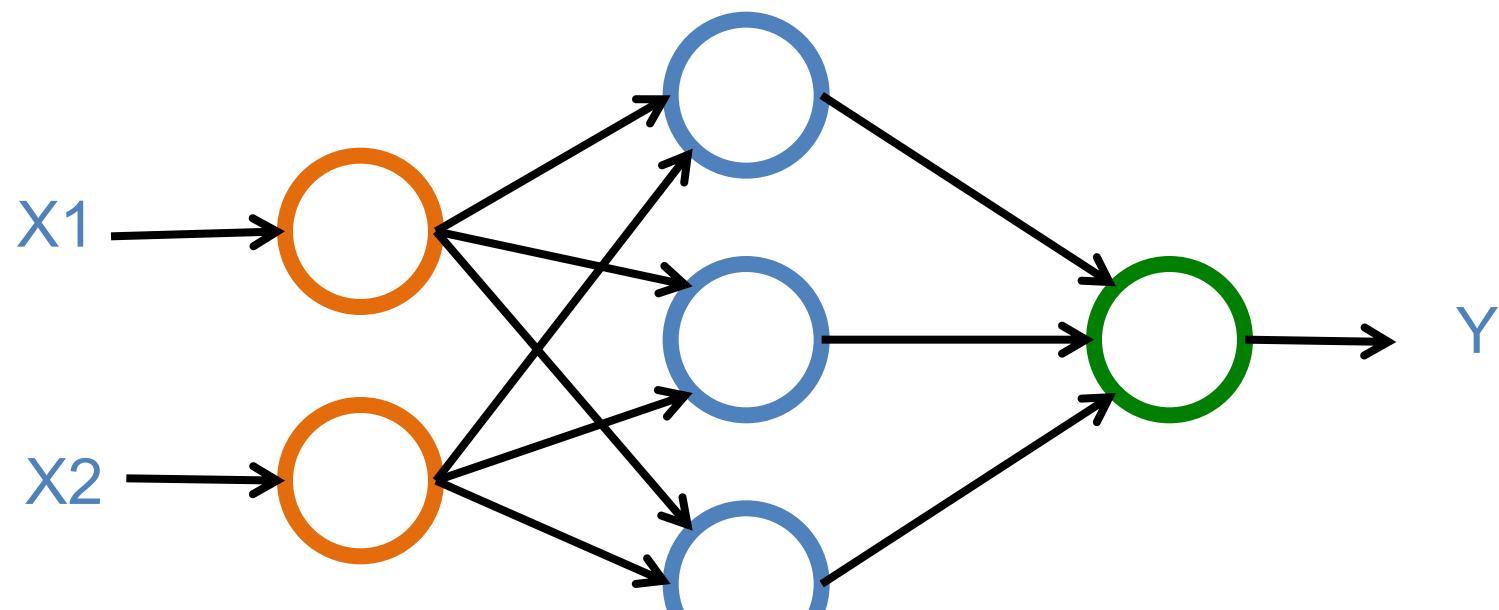


Neural Networks



Neural Networks

Input Layer **Hidden Layer** **Output Layer**
(X) **(H)** **(Y)**



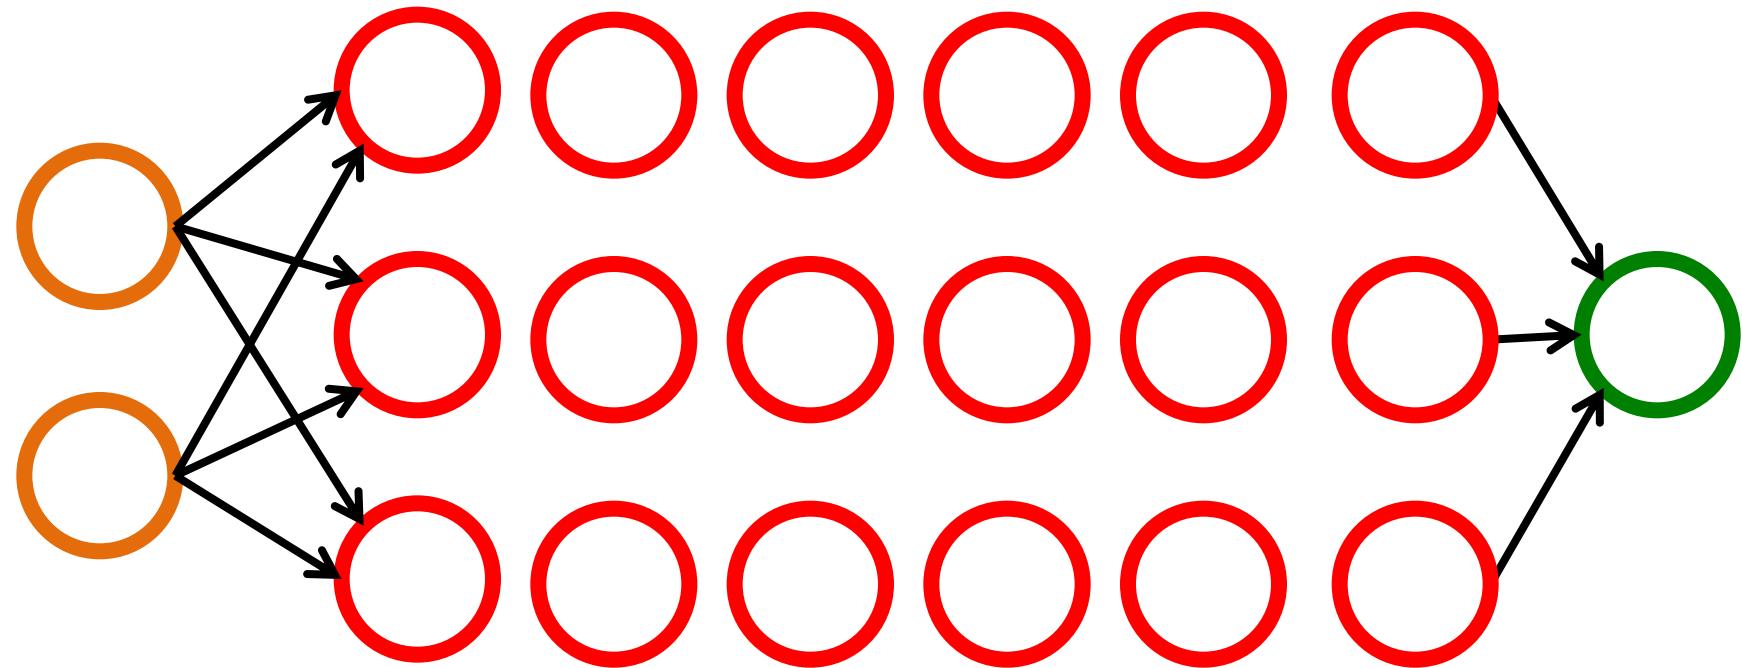
Neural Networks

Input Layer
(X)

Hidden Layers
(H)

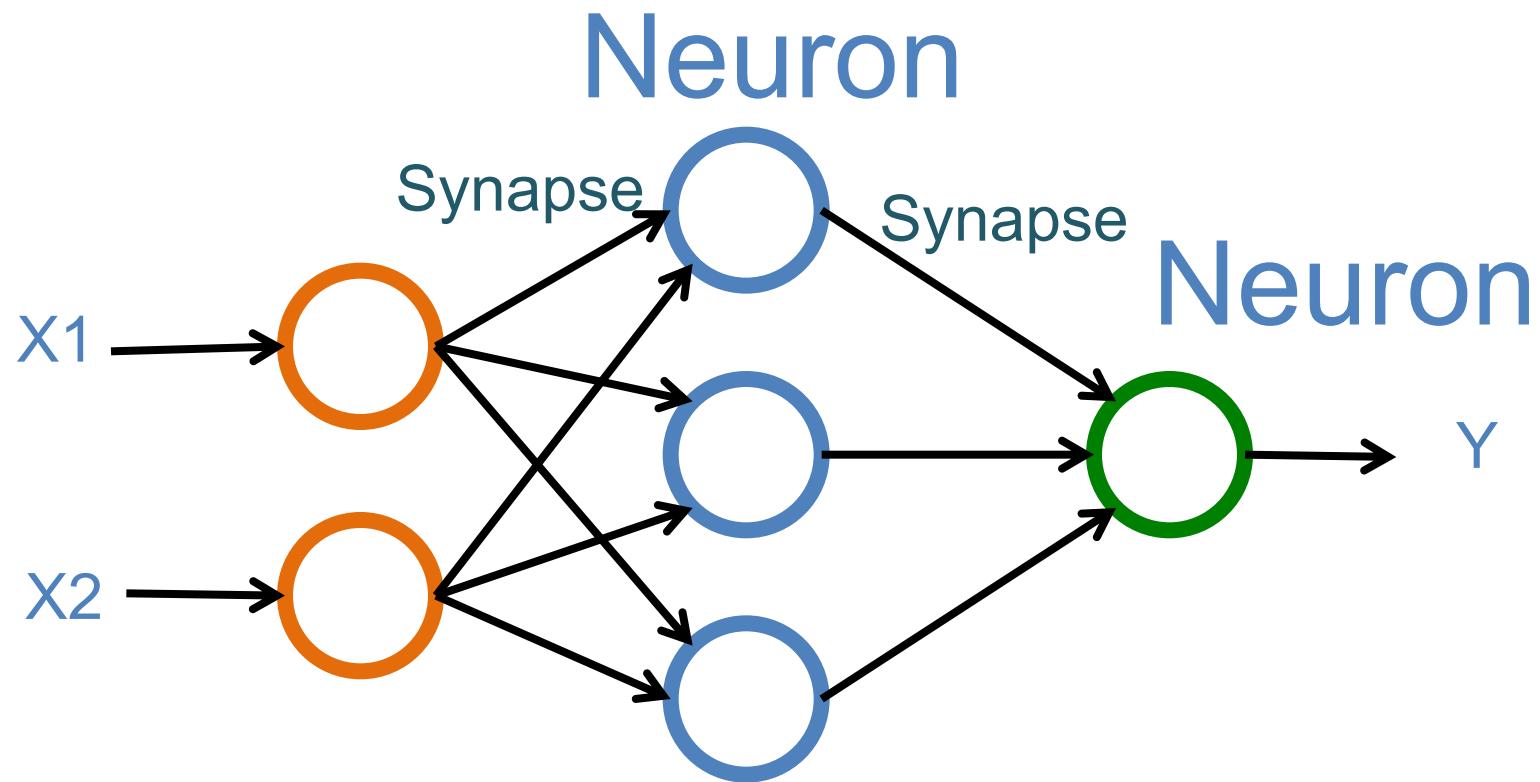
Output Layer
(Y)

Deep Neural Networks
Deep Learning

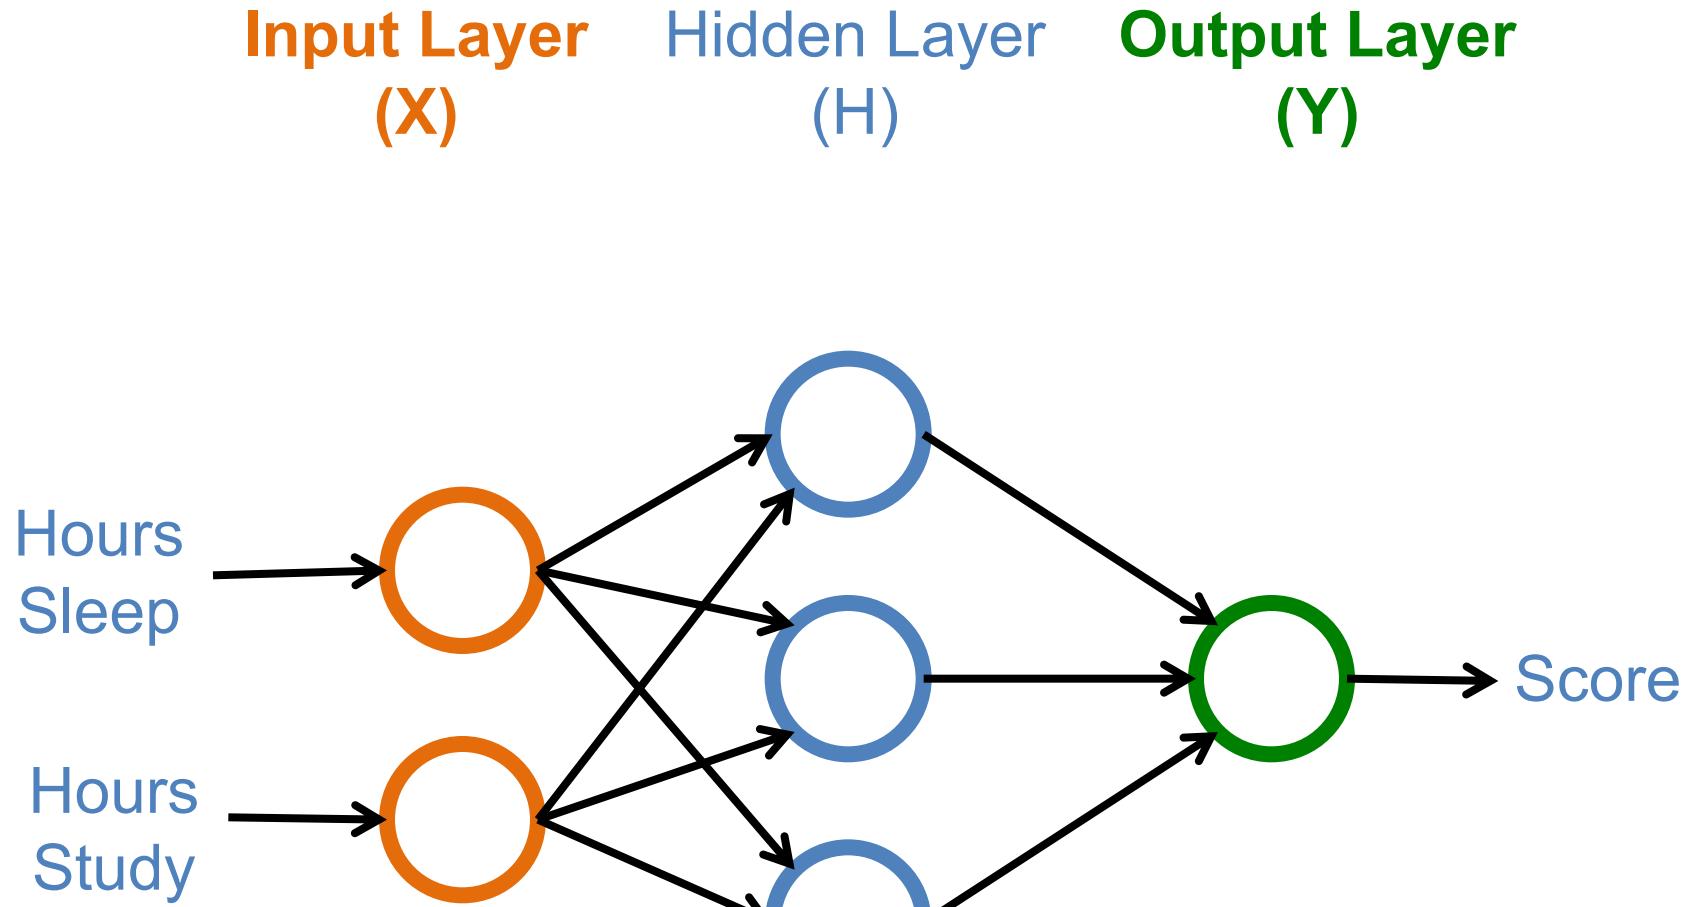


Neural Networks

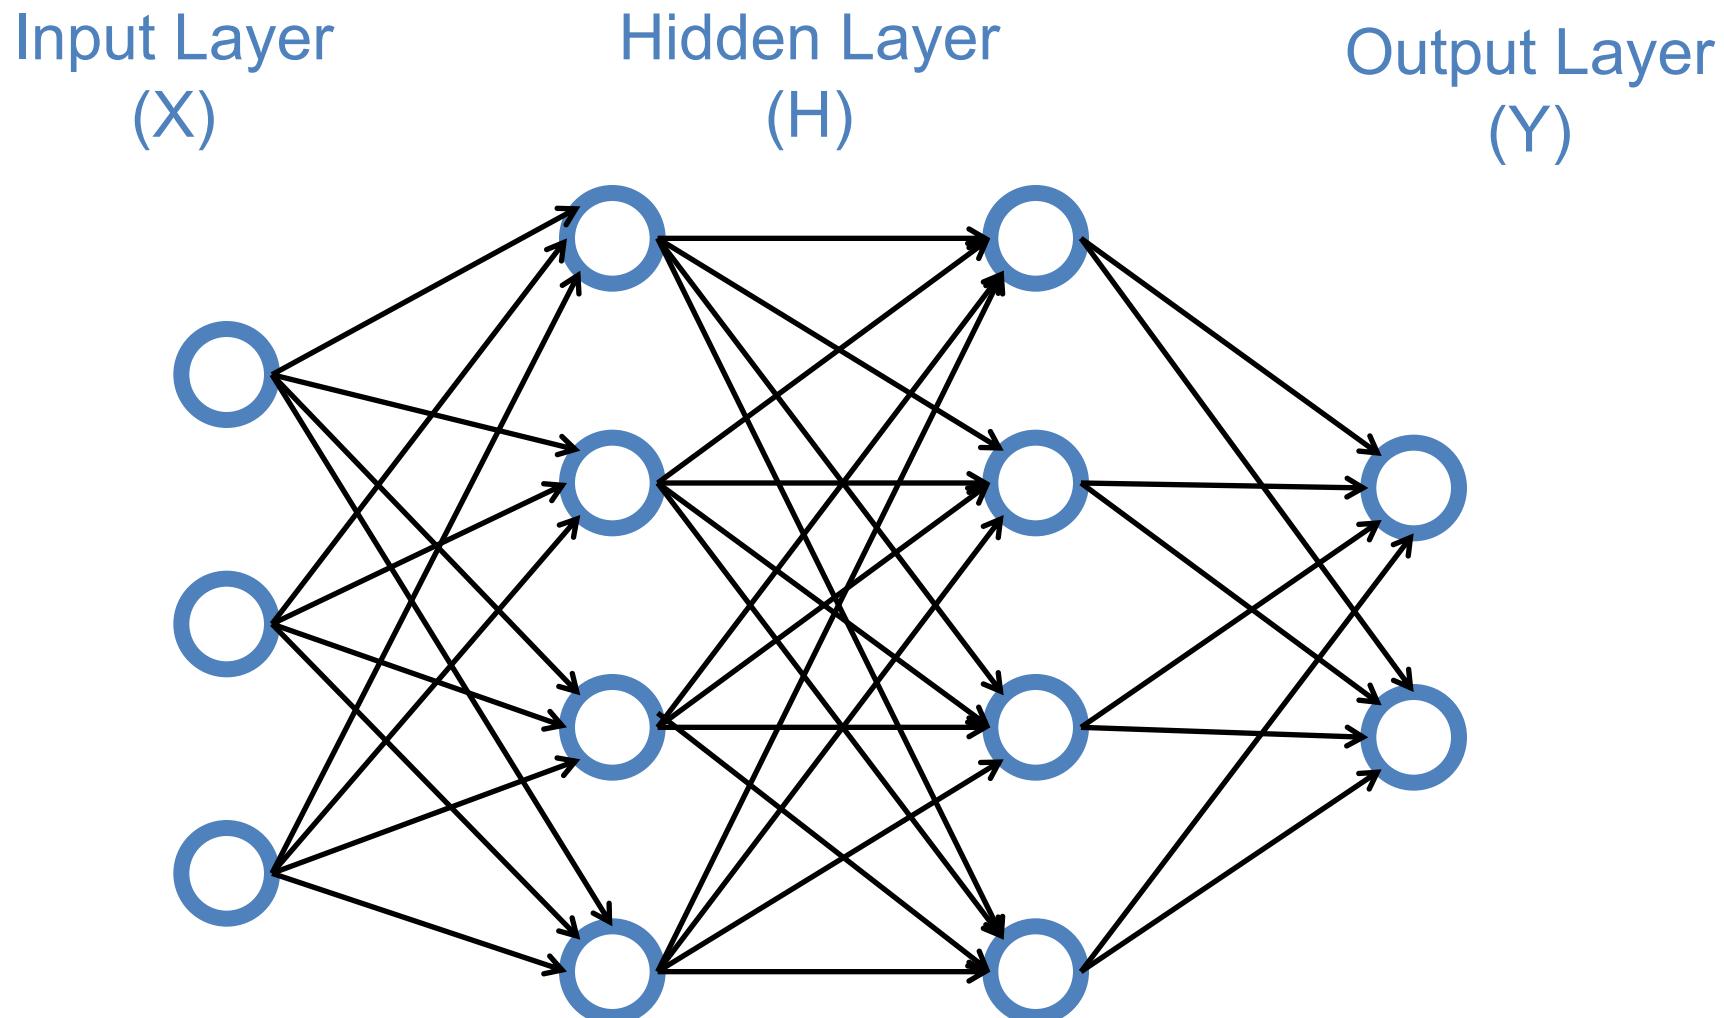
Input Layer Hidden Layer **Output Layer**
(X) (H) **(Y)**



Neural Networks

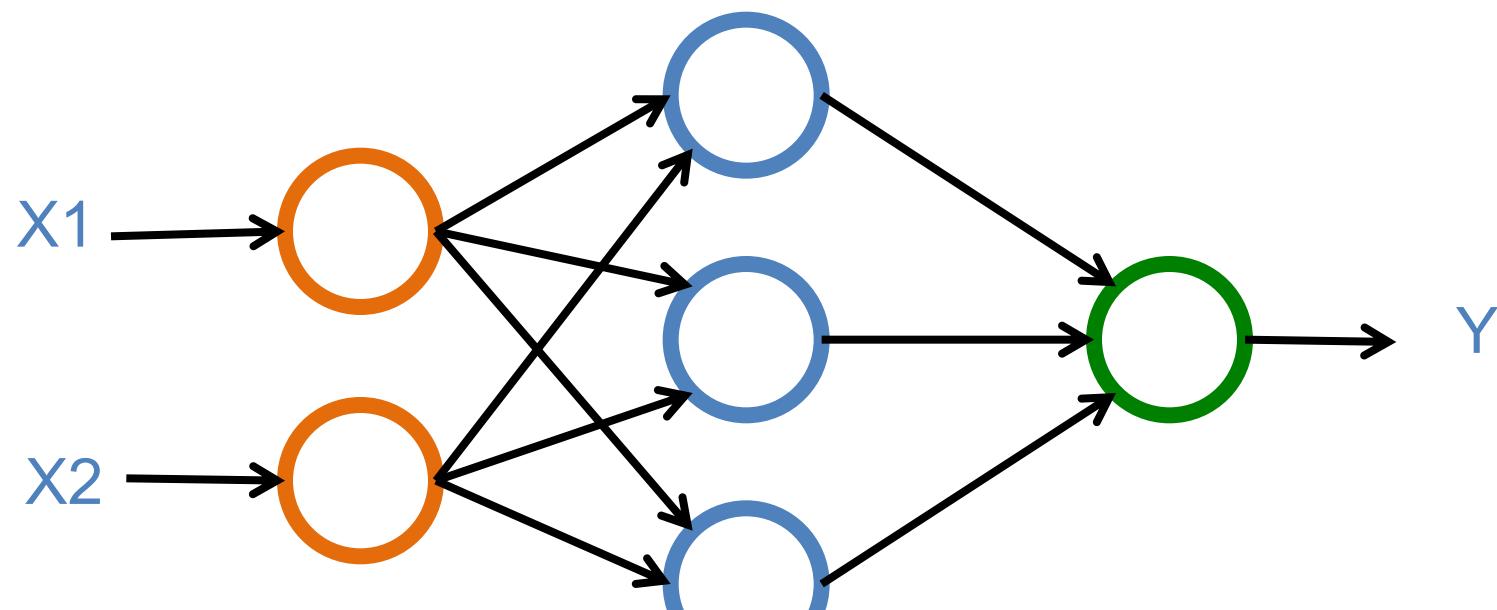


Neural Networks



Neural Networks

Input Layer **Hidden Layer** **Output Layer**
(X) **(H)** **(Y)**



Hours Sleep	Hours Study	Score
3	5	75
5	1	82
10	2	93
8	3	?

A scatter plot diagram illustrating the relationship between Sleep (X-axis), Study (Y-axis), and Score (Z-axis). The X-axis is labeled 'Hours Sleep' and the Y-axis is labeled 'Hours Study'. The Z-axis is labeled 'Score'. The plot shows three data points corresponding to Training and Testing sessions.

	X	Y	Z
	Hours Sleep	Hours Study	Score
Training	3	5	75
Training	5	1	82
Training	10	2	93
Testing	8	3	?

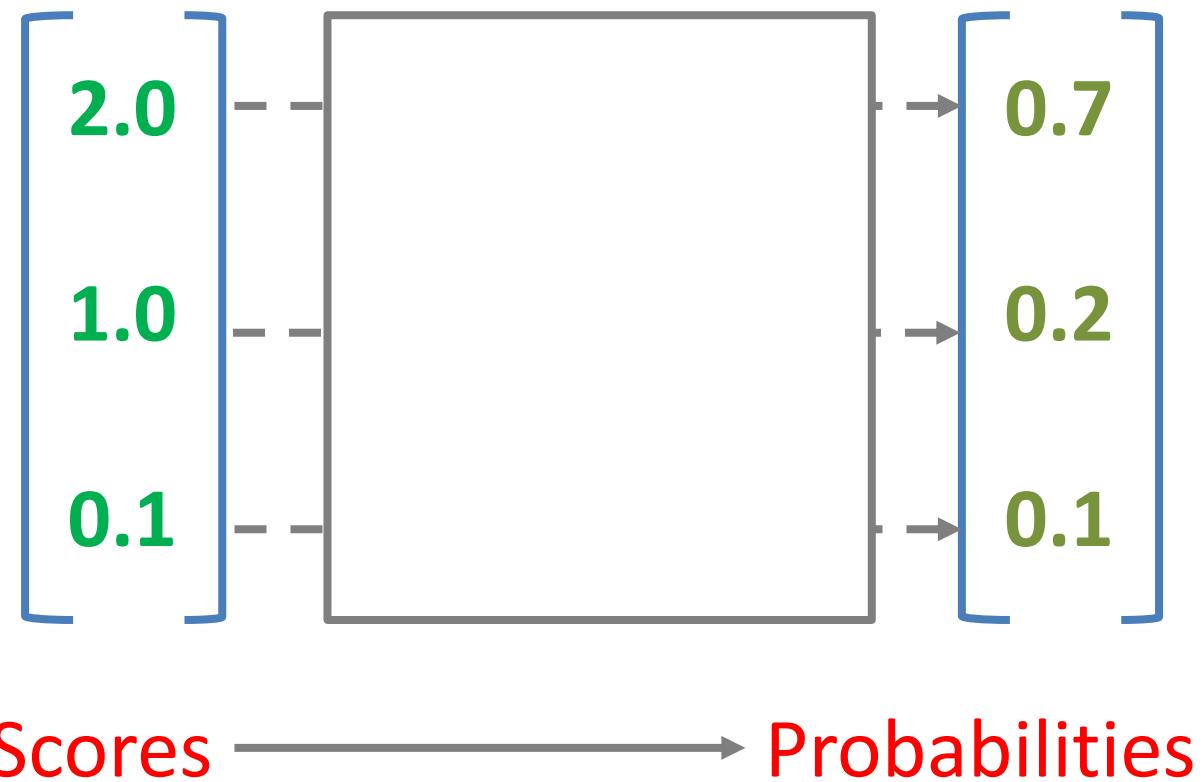
$$Y = wX + b$$

$$Y = W X + b$$

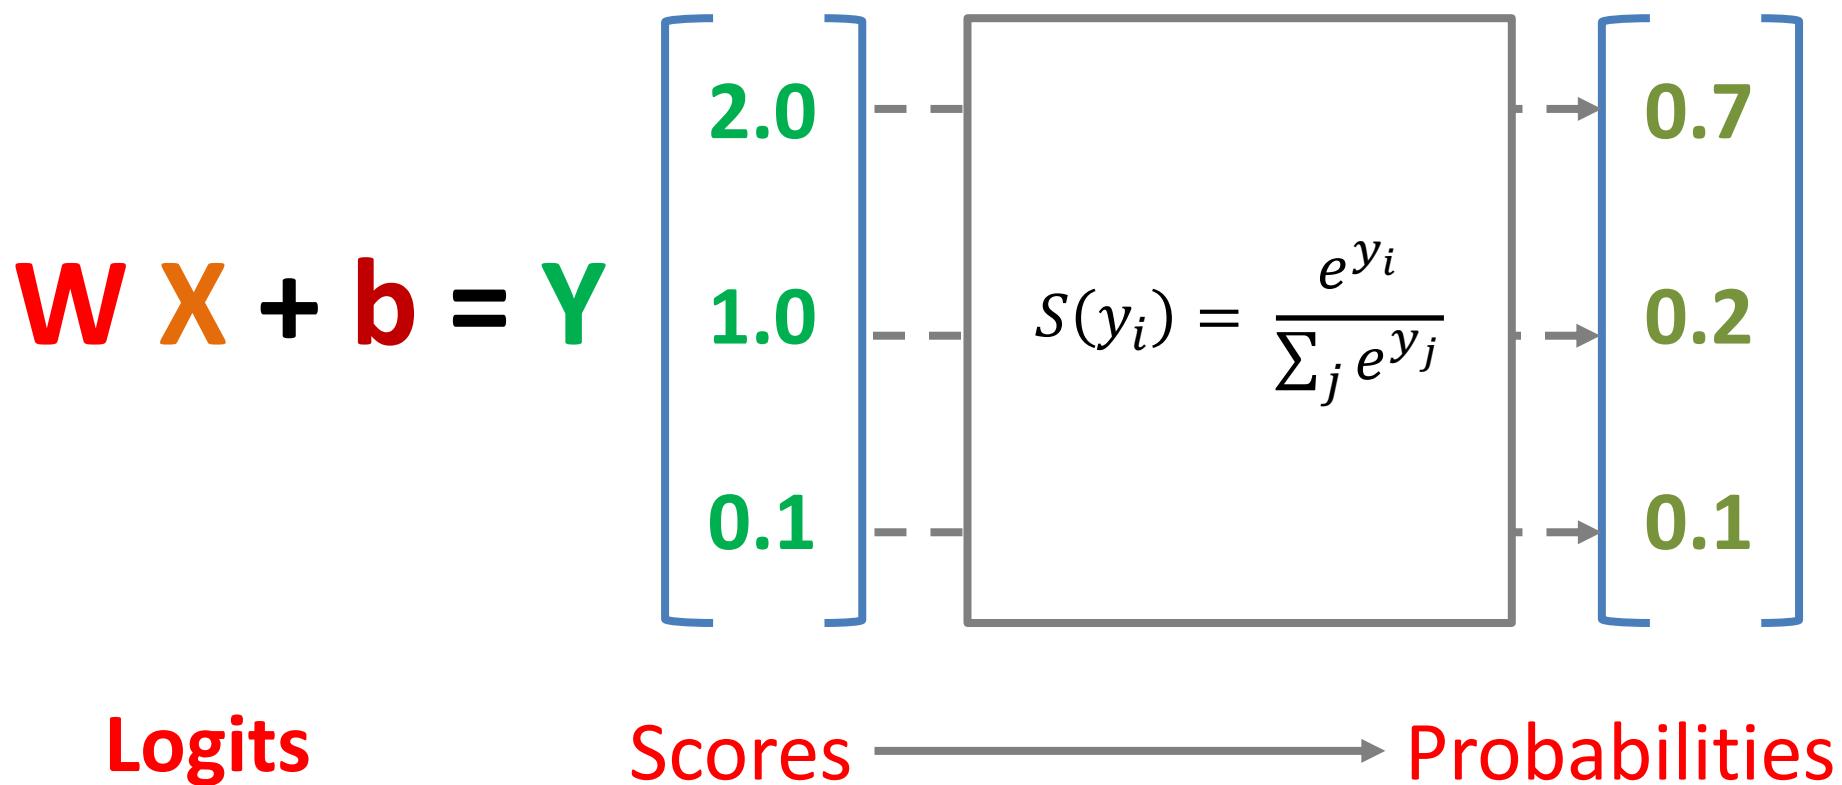
Diagram illustrating the linear equation $Y = W X + b$:

- Output** (green text) points to the variable Y .
- input** (orange text) points to the variable X .
- Weights** (red text) points to the variable W .
- bias** (red text) points to the variable b .
- Trained** (red text) is positioned below the equation, with arrows pointing from it to both W and b , indicating that both weights and bias are learned or "trained" parameters.

$$W X + b = Y$$



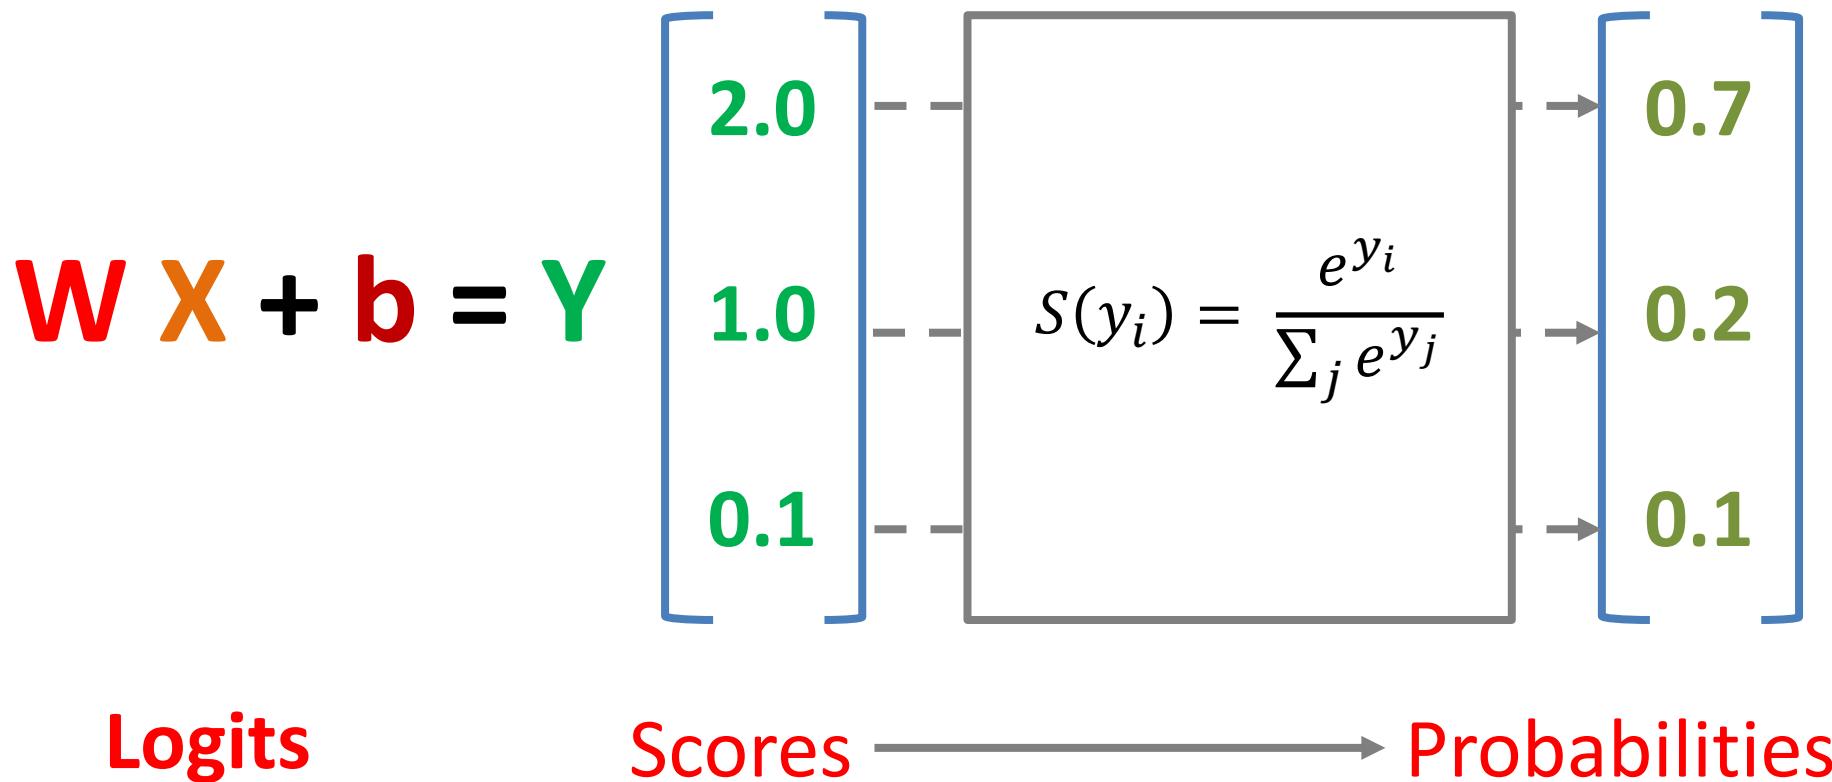
SoftMAX



$$S(y_i) = \frac{e^{y_i}}{\sum_j e^{y_j}} = \frac{e^{2.0}}{e^{2.0} + e^{1.0} + e^{0.1}} = \frac{2.7182^{2.0}}{2.7182^{2.0} + 2.7182^{1.0} + 2.7182^{0.1}} = 0.7$$

$$S(y_i) = \frac{e^{y_i}}{\sum_j e^{y_j}} = \frac{e^{1.0}}{e^{2.0} + e^{1.0} + e^{0.1}} = \frac{2.7182^{1.0}}{2.7182^{2.0} + 2.7182^{1.0} + 2.7182^{0.1}} = 0.2$$

$$S(y_i) = \frac{e^{y_i}}{\sum_j e^{y_j}} = \frac{e^{0.1}}{e^{2.0} + e^{1.0} + e^{0.1}} = \frac{2.7182^{0.1}}{2.7182^{2.0} + 2.7182^{1.0} + 2.7182^{0.1}} = 0.1$$



Training a Network

=

Minimize the Cost Function

Training a Network

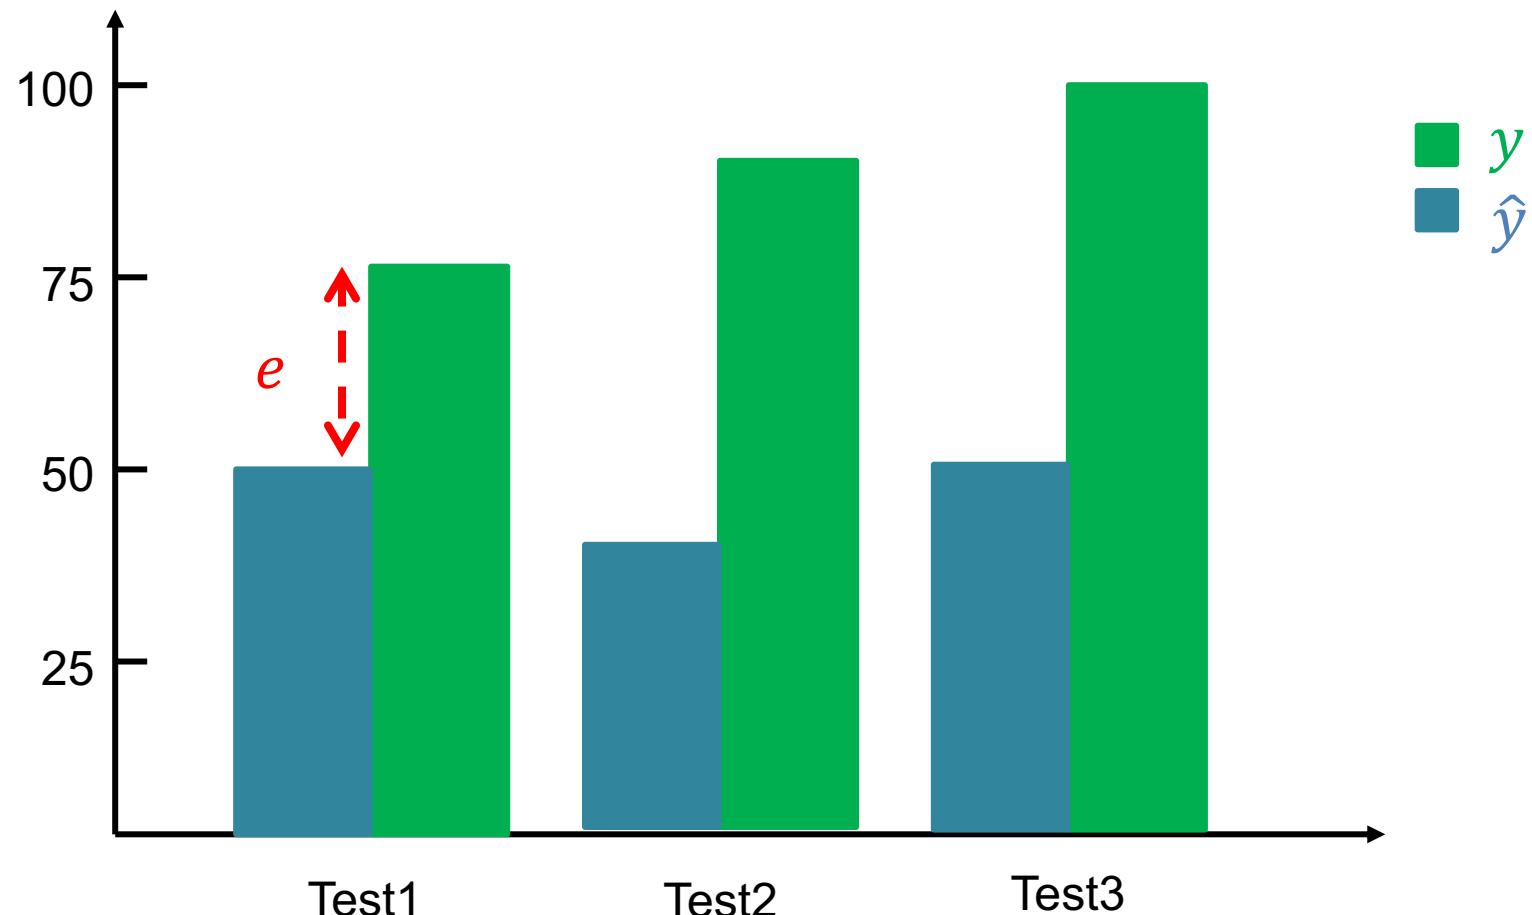
=

Minimize the **Cost Function**

Minimize the **Loss Function**

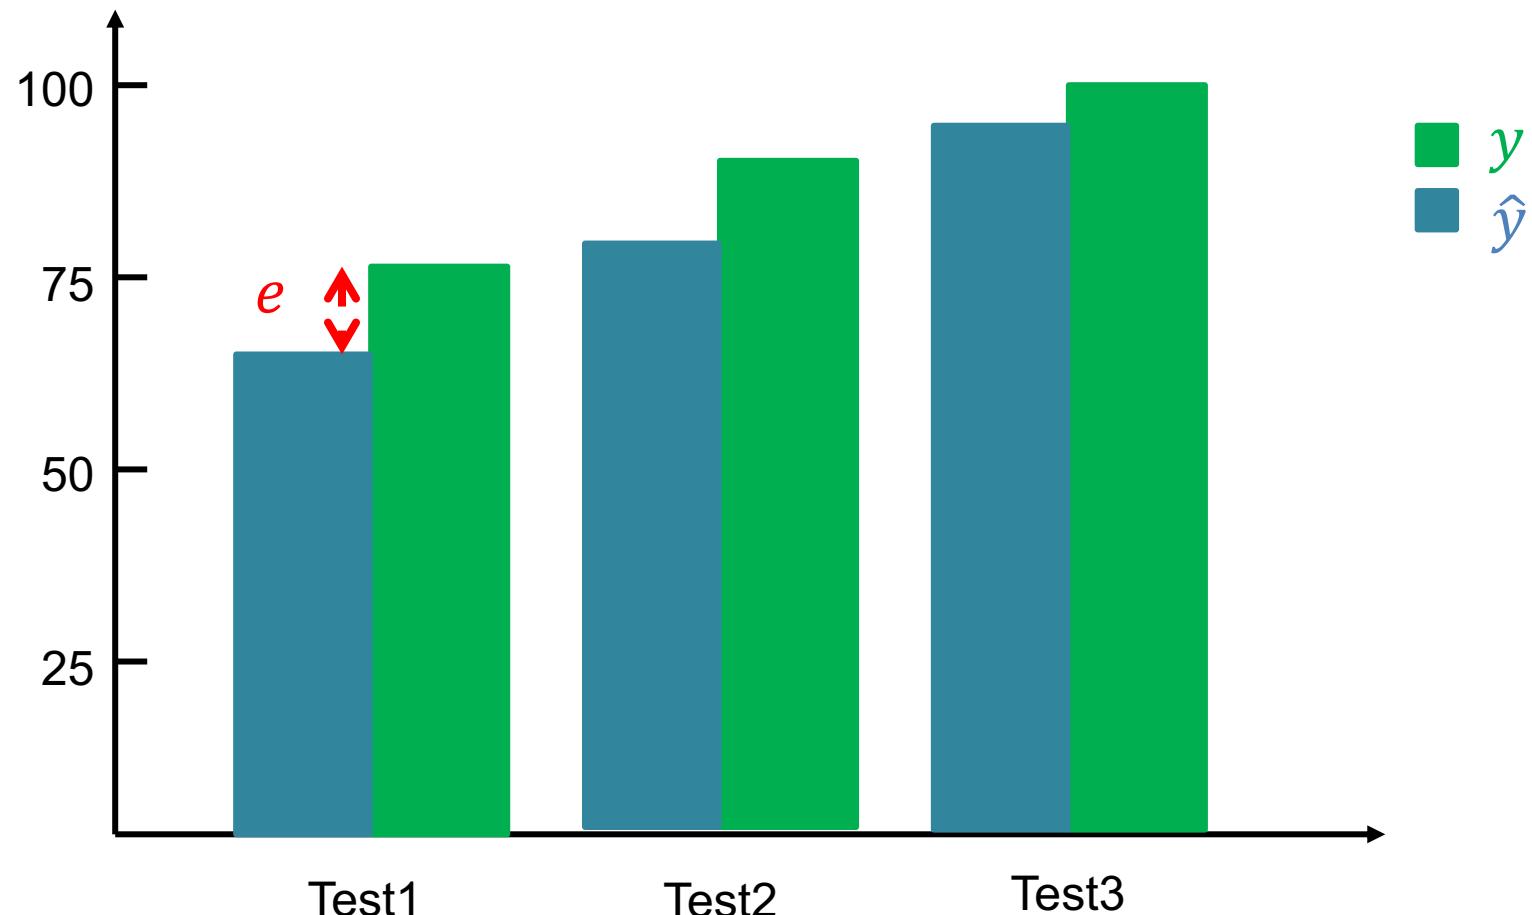
Error = Predict Y - Actual Y

Error : Cost : Loss



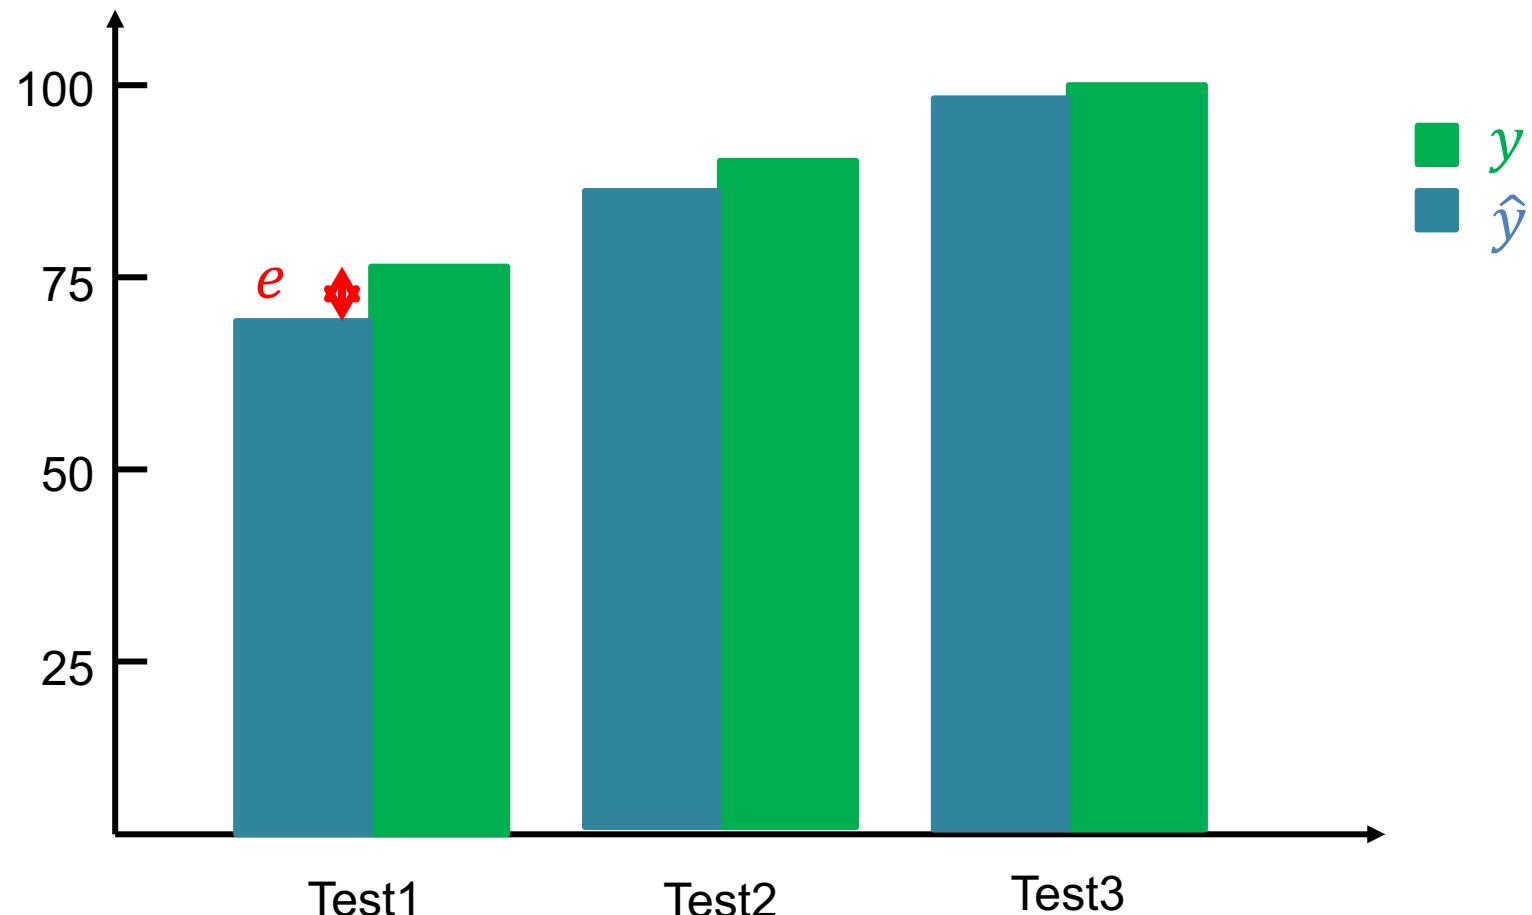
Error = Predict Y - Actual Y

Error : Cost : Loss



Error = Predict Y - Actual Y

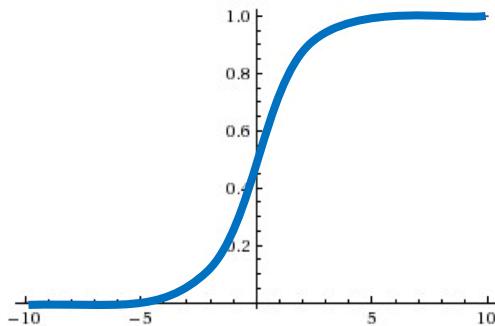
Error : Cost : Loss



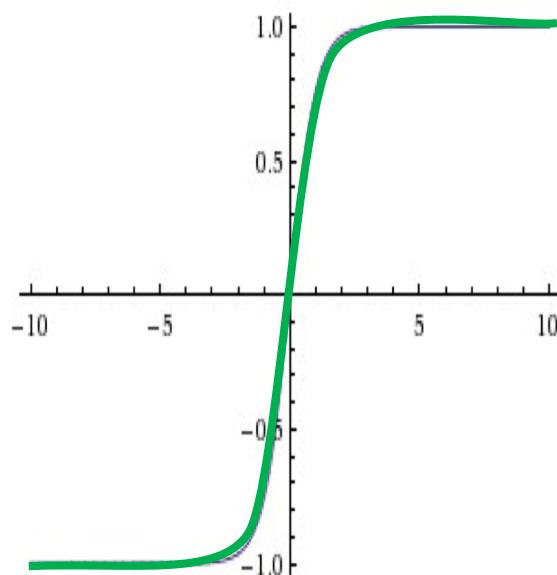
Activation Functions

Activation Functions

Sigmoid

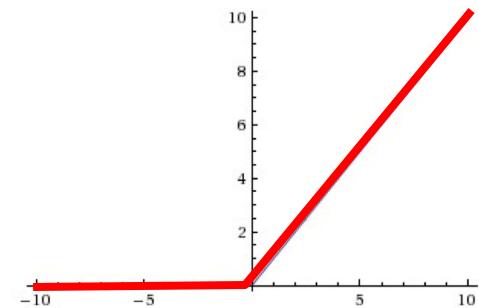


TanH



ReLU

(Rectified Linear Unit)



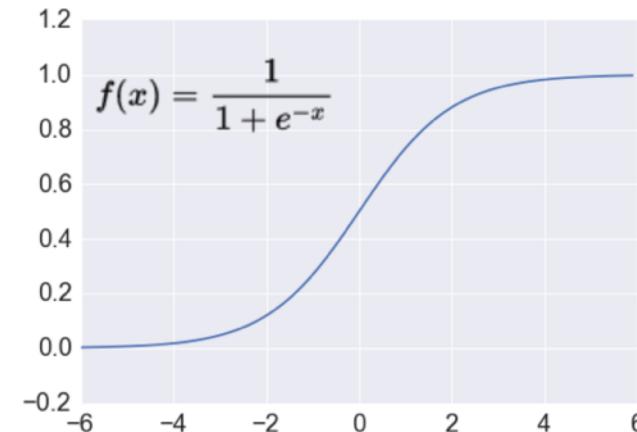
[0, 1]

[-1, 1]

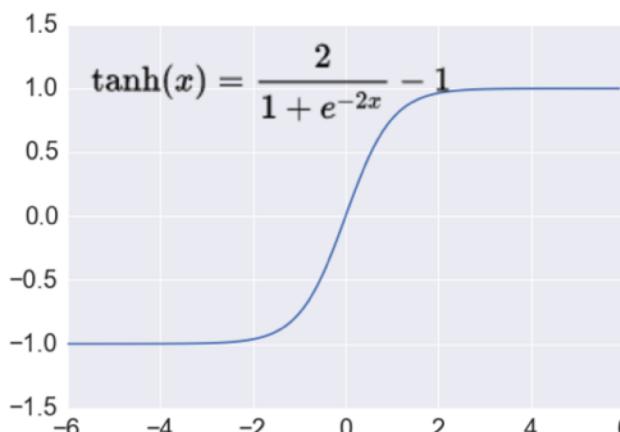
$$f(x) = \max(0, x)$$

Activation Functions

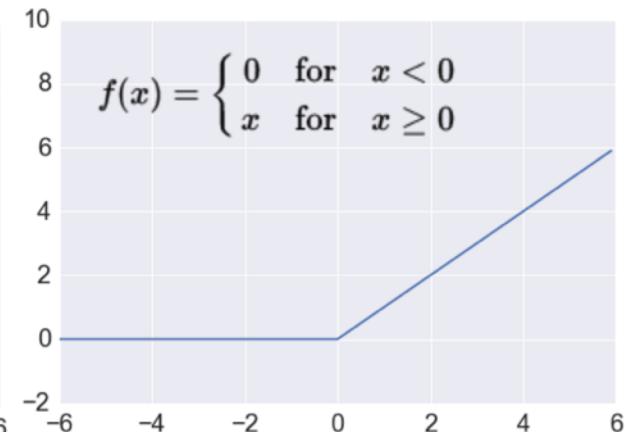
Sigmoid



TanH



ReLU



Loss Function

Binary Classification: 2 Class

Activation Function:
Sigmoid

Loss Function:
Binary Cross-Entropy

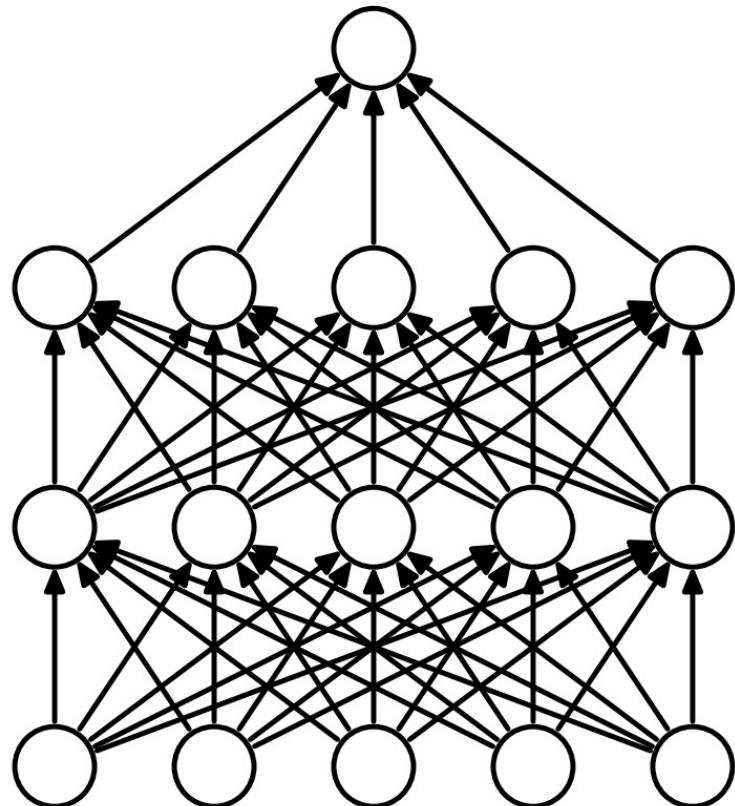
Multiple Classification: 10 Class

Activation Function:
SoftMAX

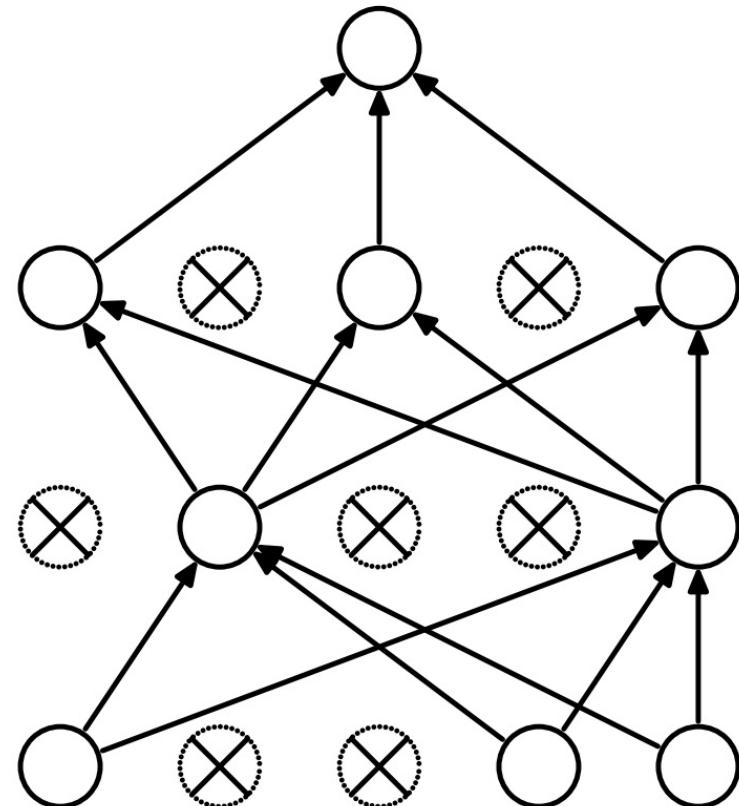
Loss Function:
Categorical Cross-Entropy

Dropout

Dropout: a simple way to prevent neural networks from overfitting



(a) Standard Neural Net



(b) After applying dropout.

Source: Srivastava, Nitish, Geoffrey E. Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov.

"Dropout: a simple way to prevent neural networks from overfitting." *Journal of machine learning research* 15, no. 1 (2014): 1929-1958.

Learning Algorithm

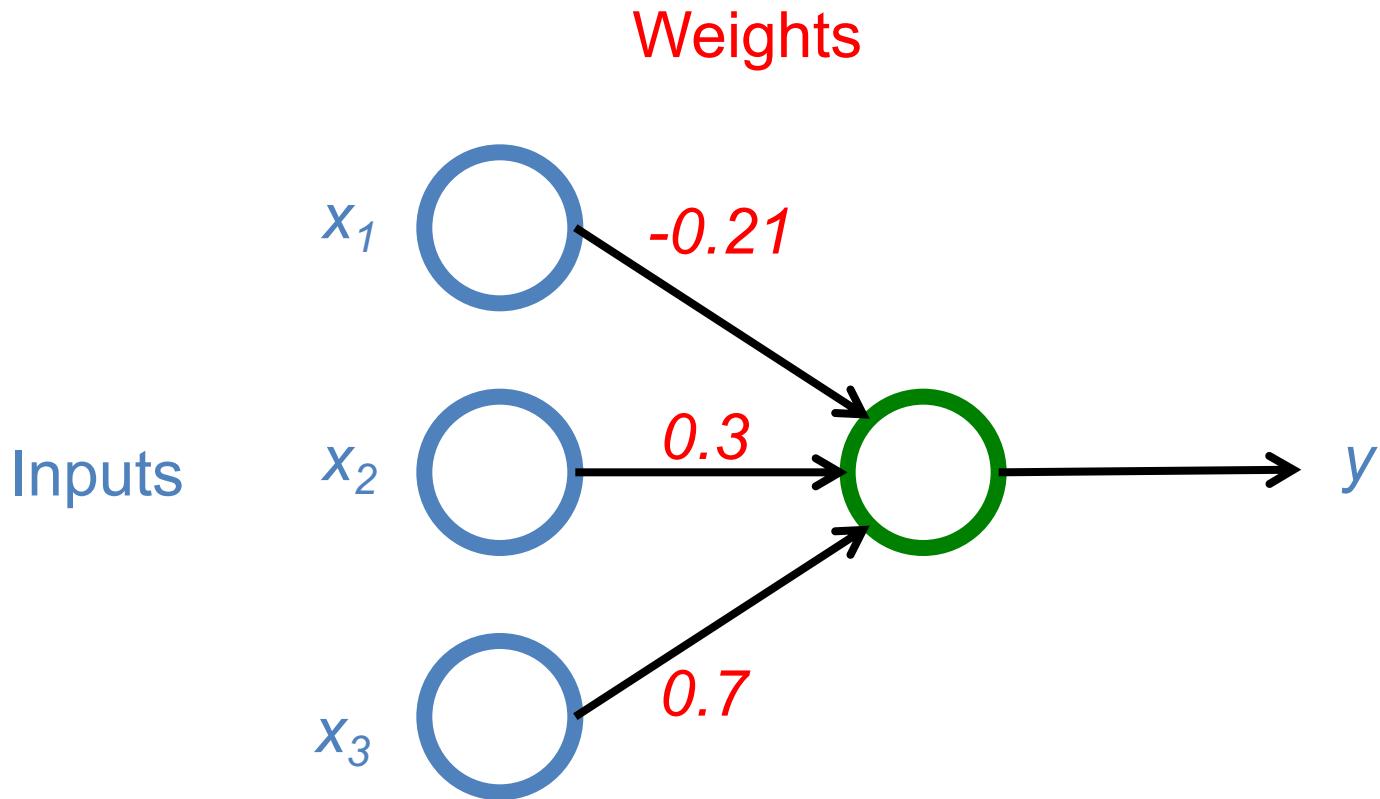
While not done:

Pick a random training example “(input, label)”

Run neural network on “input”

Adjust weights on edges to make output closer to “label”

$$y = \max(0, -0.21 * x_1 + 0.3 * x_2 + 0.7 * x_3)$$

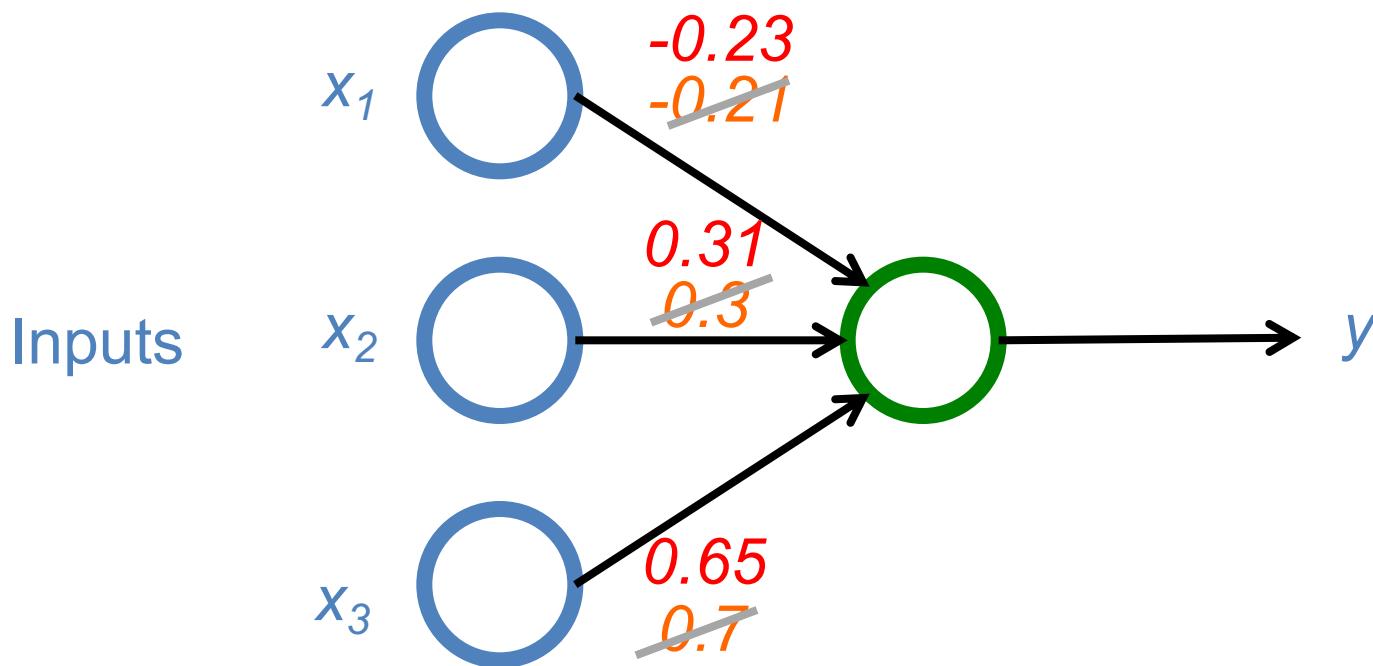


Next time:

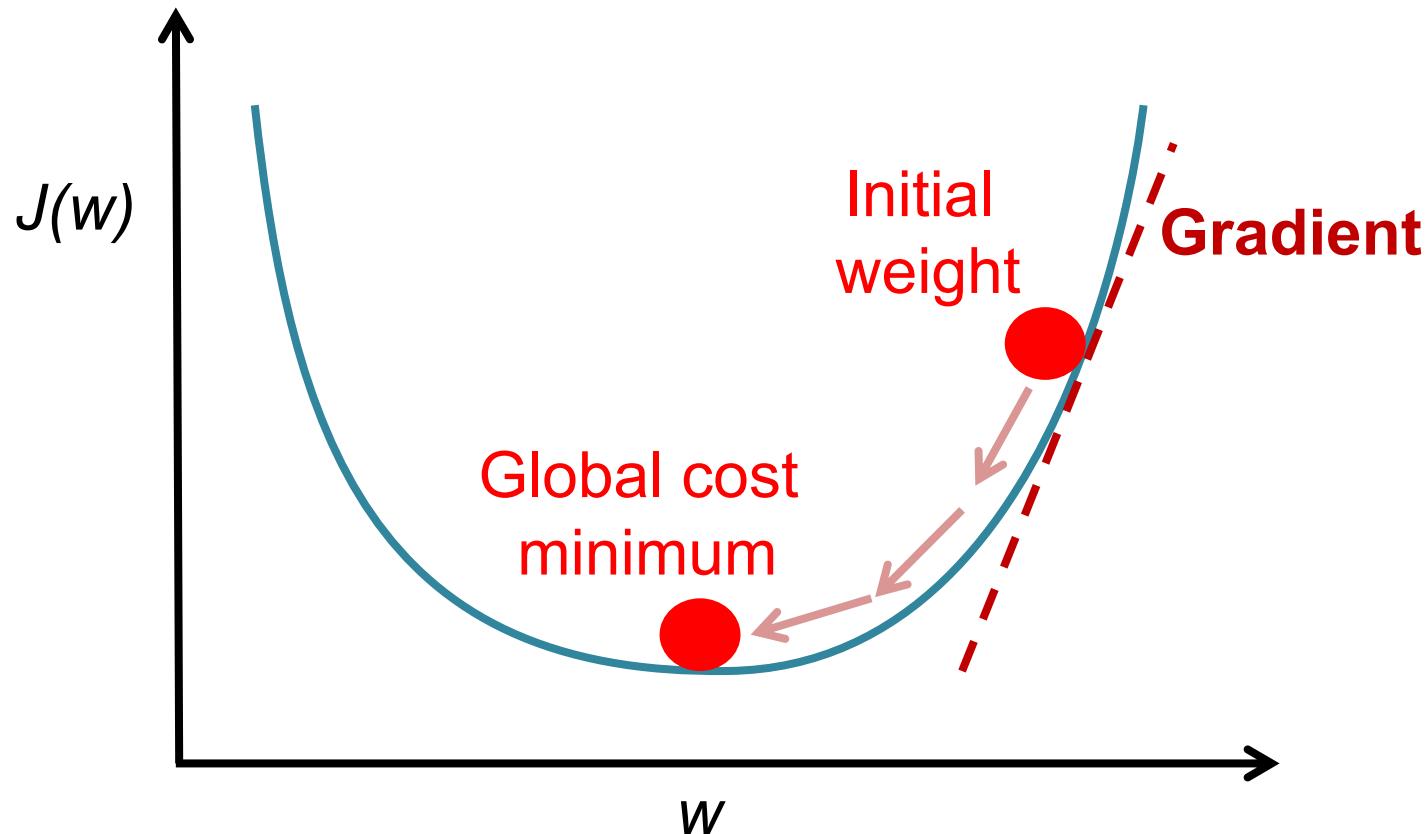
$$y = \max(0, -0.23 * x_1 + 0.31 * x_2 + 0.65 * x_3)$$

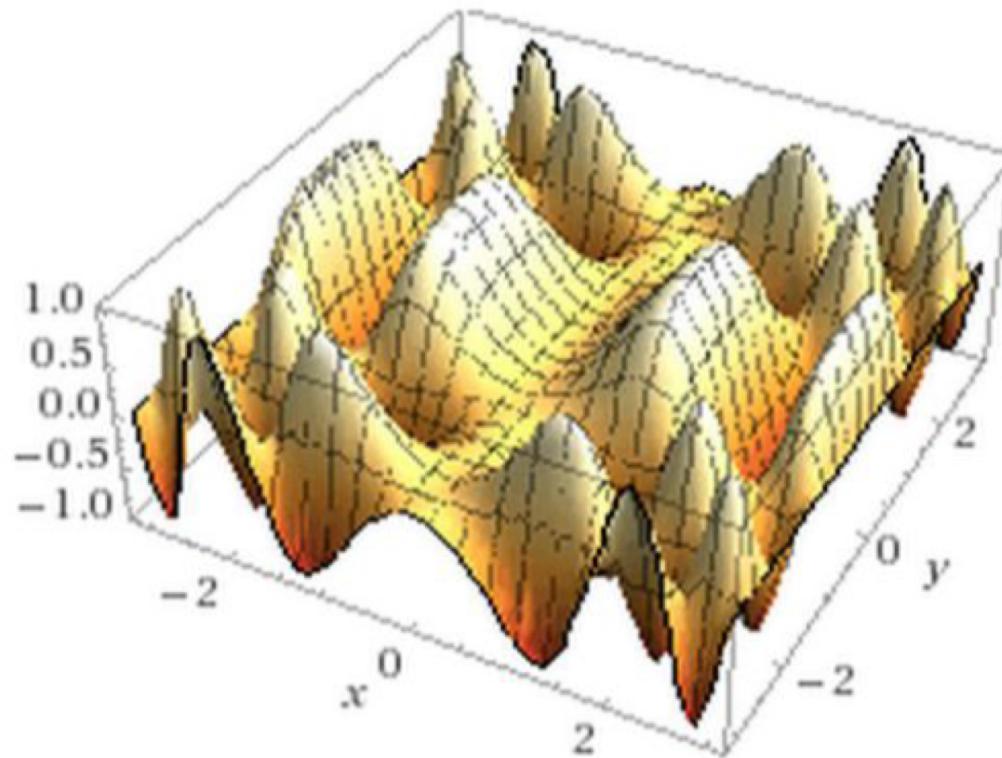
$$y = \max(0, -0.21 * x_1 + 0.3 * x_2 + 0.7 * x_3)$$

Weights



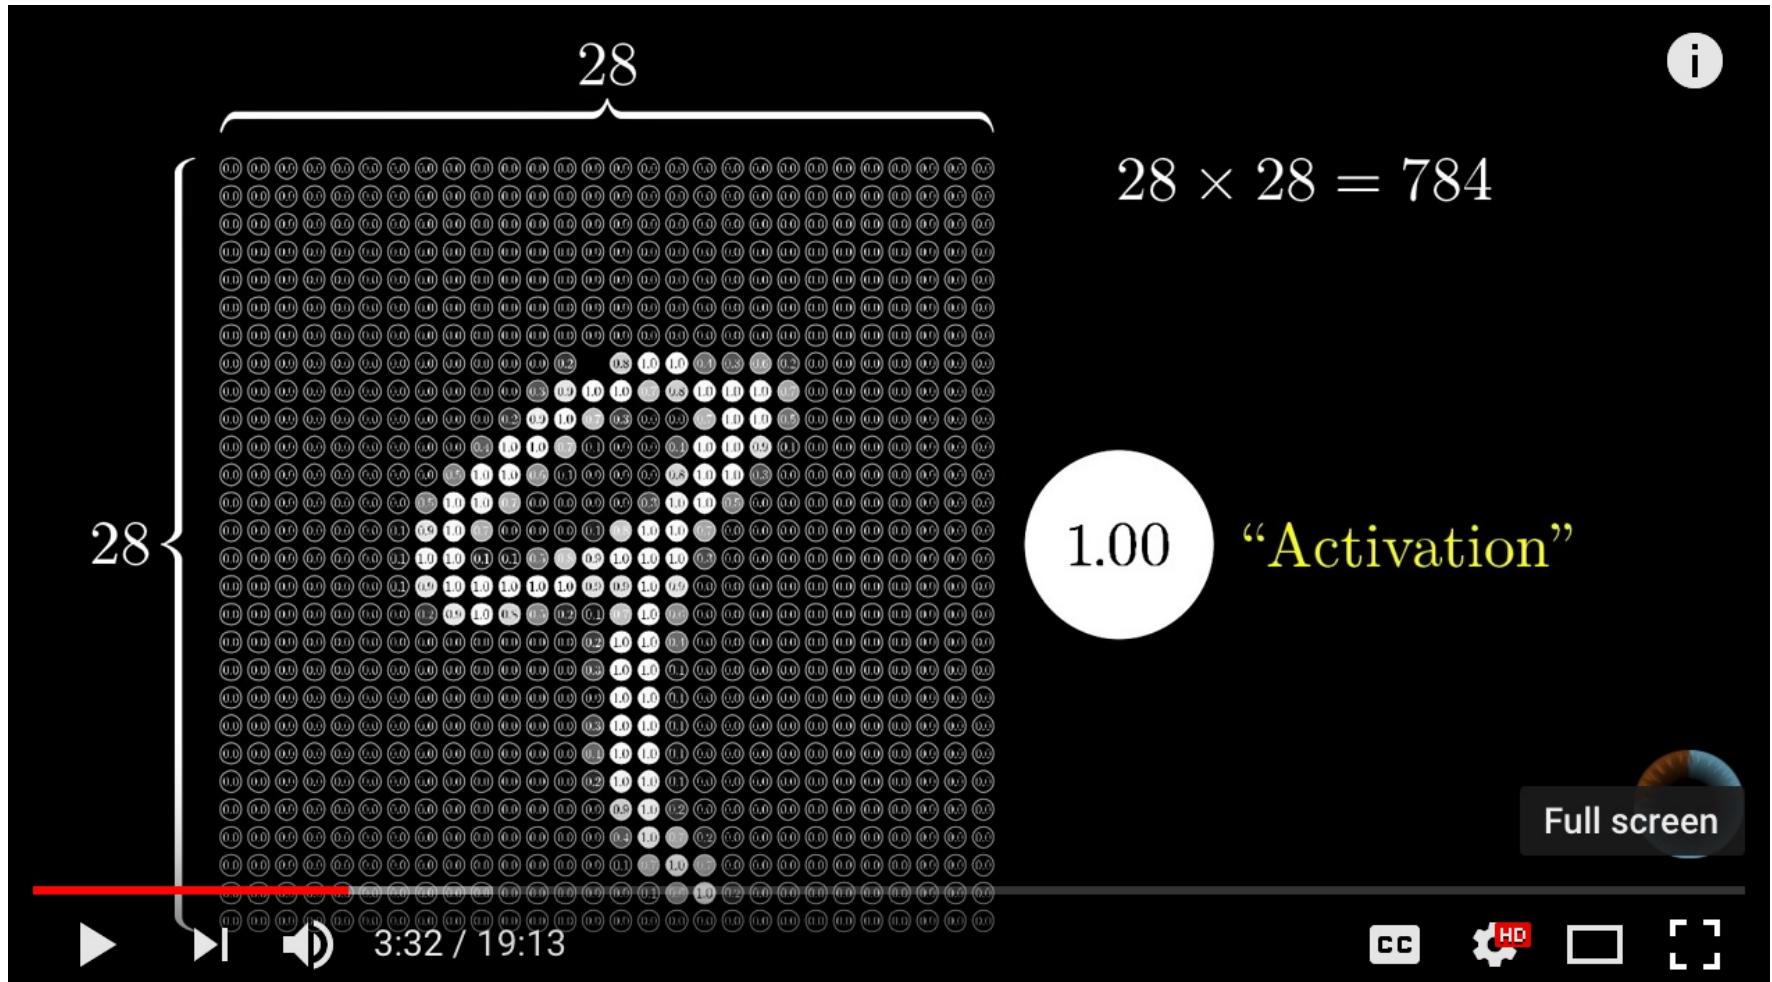
Optimizer: Stochastic Gradient Descent (SGD)





This shows a function of 2 variables: real neural nets are functions of hundreds of millions of variables!

Neural Network and Deep Learning



Gradient Descent

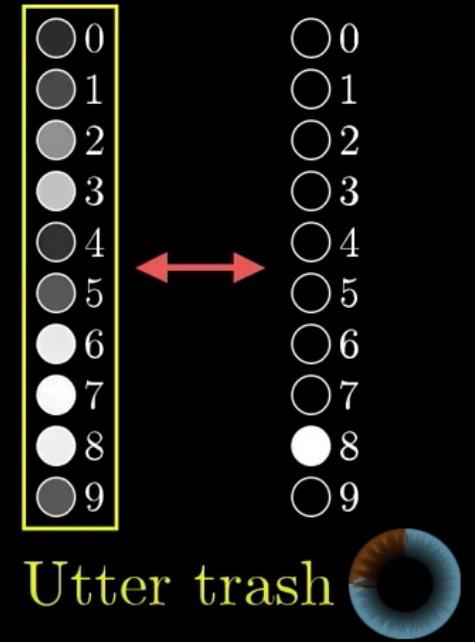
how neural networks learn

Average cost of
all training data...

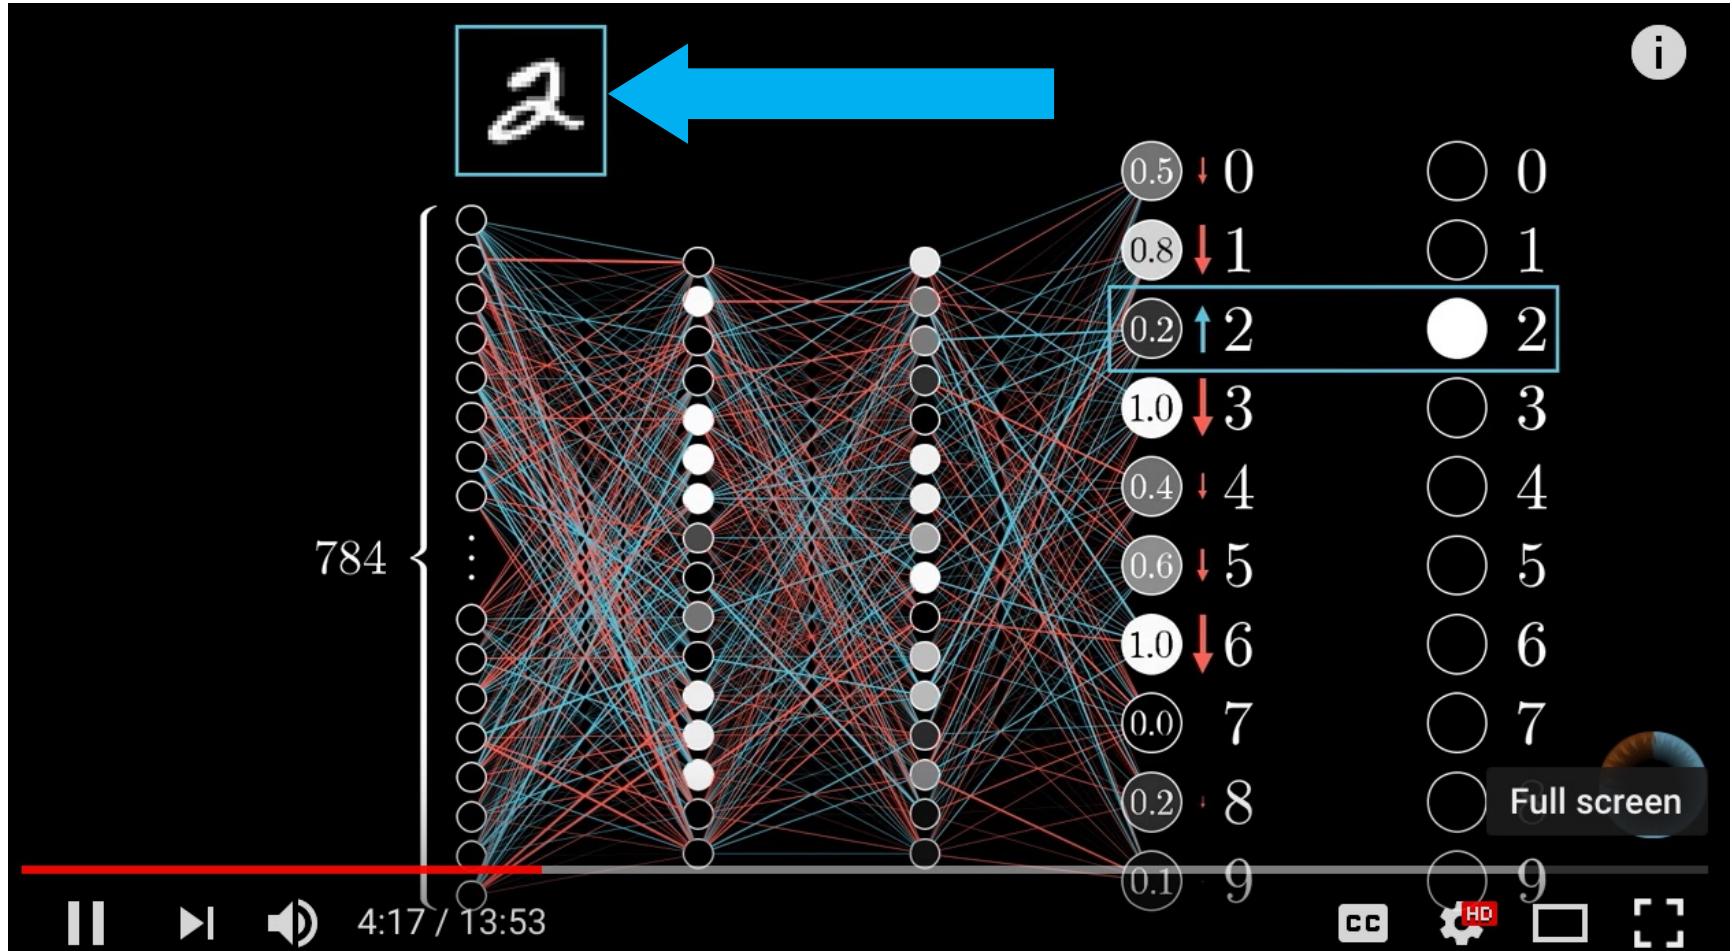
Cost of \hat{y}

$$\left\{ \begin{array}{l} (0.18 - 0.00)^2 + \\ (0.29 - 0.00)^2 + \\ (0.58 - 0.00)^2 + \\ (0.77 - 0.00)^2 + \\ (0.20 - 0.00)^2 + \\ (0.36 - 0.00)^2 + \\ (0.93 - 0.00)^2 + \\ (1.00 - 0.00)^2 + \\ (0.95 - 1.00)^2 + \\ (0.35 - 0.00)^2 \end{array} \right.$$

What's the “cost”
of this difference?



Backpropagation



Source: 3Blue1Brown (2017), What is backpropagation really doing? | Chapter 3, deep learning,
<https://www.youtube.com/watch?v=Ilg3gGewQ5U>

Learning Algorithm

While not done:

Pick a random training example “(input, label)”

Run neural network on “input”

Adjust weights on edges to make output closer to “label”

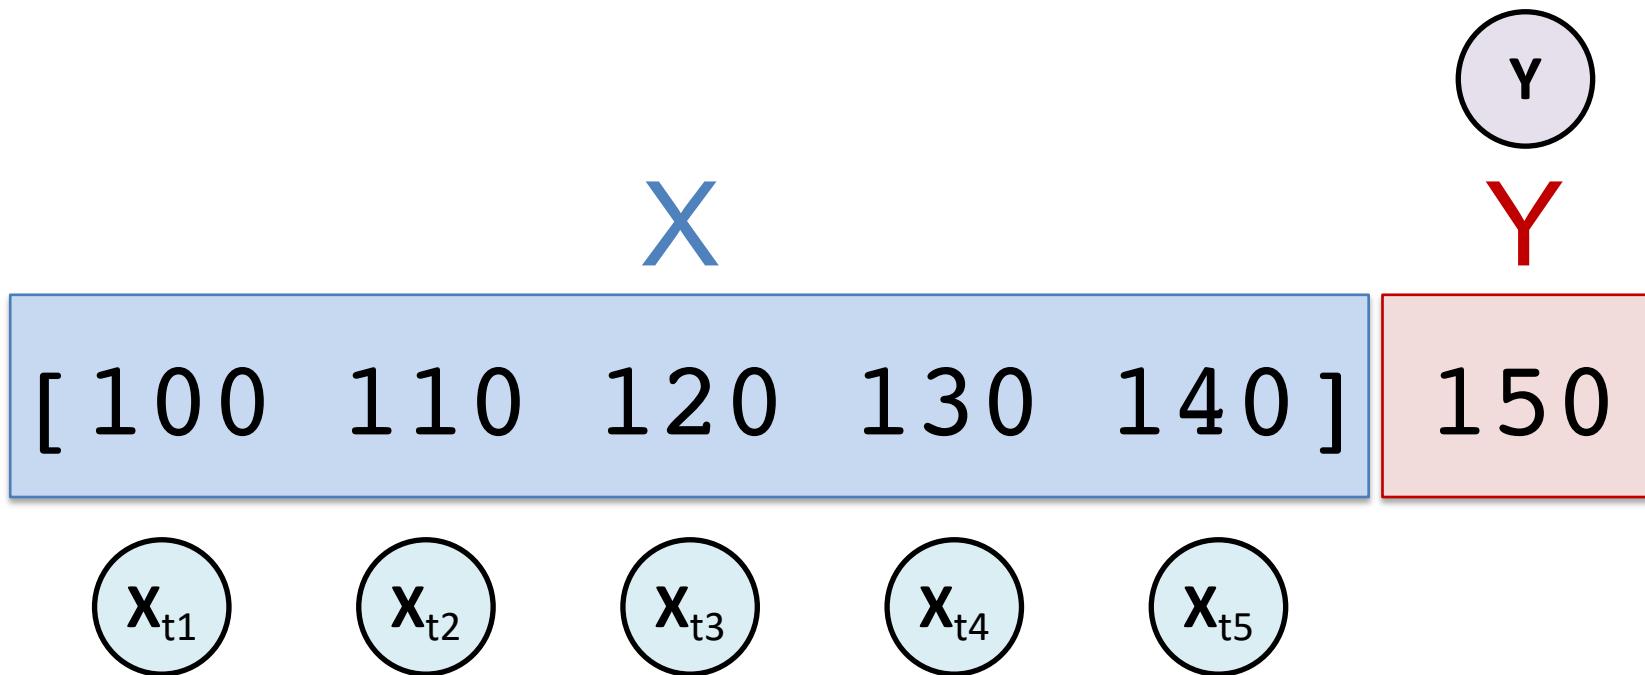
Financial Time Series Forecasting

Time Series Data

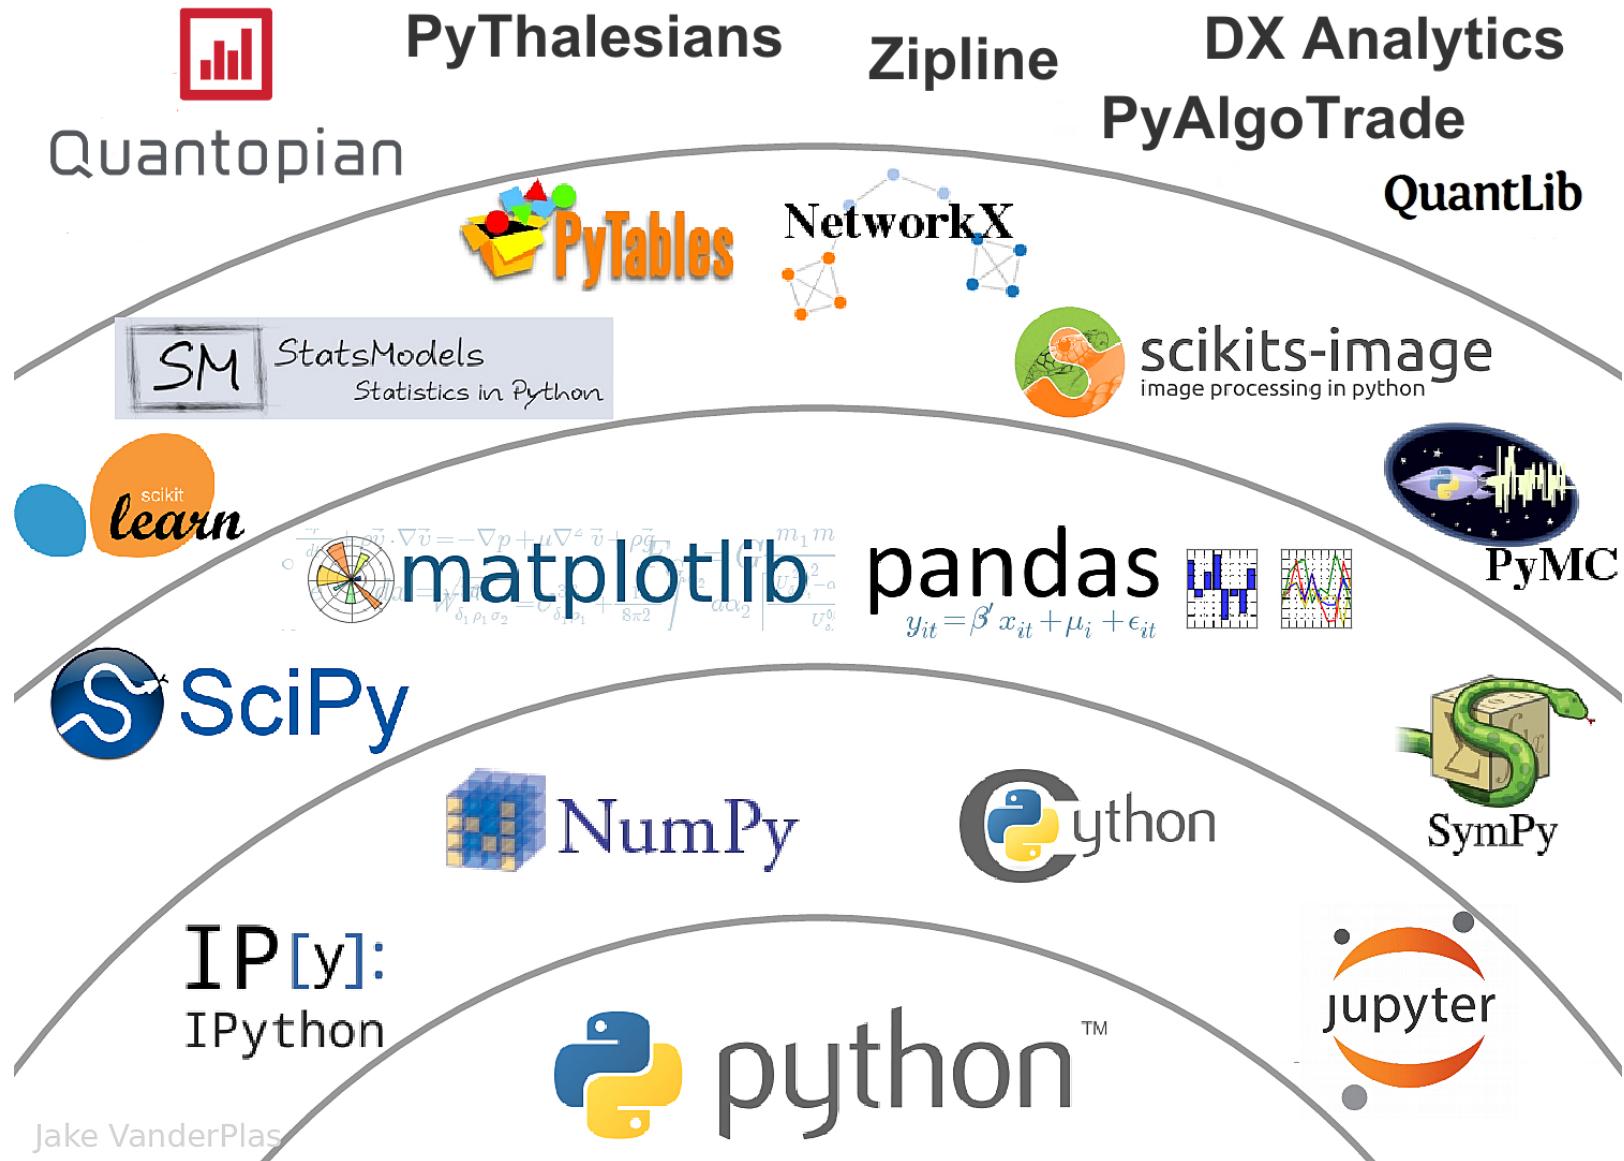


Time Series Data

[100, 110, 120, 130, 140, 150]



The Quant Finance PyData Stack



Python in Google Colab

<https://colab.research.google.com/drive/1FEG6DnGwvfUbeo4zJ1zTunjMqf2RkCrT>

CO python101.ipynb ☆

File Edit View Insert Runtime Tools Help

CODE TEXT CELL CELL COMMENT SHARE A

```
1 # !pip install pandas_datareader
2 import pandas as pd
3 import pandas_datareader.data as web
4 import matplotlib.pyplot as plt
5 import seaborn as sns
6 import datetime as dt
7 %matplotlib inline
8
9 #Read Stock Data from Yahoo Finance
10 end = dt.datetime.now()
11 #start = dt.datetime(end.year-2, end.month, end.day)
12 start = dt.datetime(2016, 1, 1)
13 df = web.DataReader("AAPL", 'yahoo', start, end)
14 df.to_csv('AAPL.csv')
15 df.from_csv('AAPL.csv')
16 df.tail()
17
18 df['Adj Close'].plot(legend=True, figsize=(12, 8), title='AAPL', label='Adj Close')
19 plt.figure(figsize=(12,9))
20 top = plt.subplot2grid((12,9), (0, 0), rowspan=10, colspan=9)
21 bottom = plt.subplot2grid((12,9), (10,0), rowspan=2, colspan=9)
22 top.plot(df.index, df['Adj Close'], color='blue') #df.index gives the dates
23 bottom.bar(df.index, df['Volume'])
24
25 # set the labels
26 top.axes.get_xaxis().set_visible(False)
27 top.set_title('AAPL')
28 top.set_ylabel('Adj Close')
29 bottom.set_ylabel('Volume')
30
31 plt.figure(figsize=(12,9))
32 sns.distplot(df['Adj Close'].dropna(), bins=50, color='purple')
33
34 # simple moving averages
35 df['MA05'] = df['Adj Close'].rolling(5).mean() #5 days
36 df['MA20'] = df['Adj Close'].rolling(20).mean() #20 days
37 df['MA60'] = df['Adj Close'].rolling(60).mean() #60 days
38 df2 = pd.DataFrame({'Adj Close': df['Adj Close'], 'MA05': df['MA05'], 'MA20': df['MA20'], 'MA60': df['MA60']})
39 df2.plot(figsize=(12, 9), legend=True, title='AAPL')
40 df2.to_csv('AAPL_MA.csv')
41 fig = plt.gcf()
42 fig.set_size_inches(12, 9)
43 fig.savefig('AAPL_plot.png', dpi=300)
```

pandas

Python Data Analysis

Library

**providing high-performance, easy-to-use
data structures and data analysis tools
for the Python programming language.**

pandas Ecosystem

- Statistics and Machine Learning
 - [Statsmodels](#)
 - [sklearn-pandas](#)
- Visualization
 - [Bokeh](#)
 - [yhat/ggplot](#)
 - [Seaborn](#)
 - [Vincent](#)
 - [IPython Vega](#)
 - [Plotly](#)
 - [Pandas-Qt](#)
- IDE
 - [IPython](#)
 - [quantopian/qgrid](#)
 - [Spyder](#)
- API
 - [pandas-datareader](#)
 - [quandl/Python](#)
 - [pydatastream](#)
 - [pandaSDMX](#)
 - [fredapi](#)
- Domain Specific
 - [Geopandas](#)
 - [xarray](#)
- Out-of-core
 - [Dask](#)
 - [Blaze](#)
 - [Odo](#)

pandas-datareader

pandas-datareader
latest

Search docs

What's New

- Remote Data Access
- Caching queries
- Other Data Sources
- Data Readers

Take the [Python User's Survey](#) and let us know more about your usage of Python.

Sponsored · Ads served ethically

Read the Docs v: latest ▾

pandas-datareader

Up to date remote data access for pandas, works for multiple versions of pandas.

[pypi v0.7.0](#) [build failing](#) [coverage 79%](#) [docs passing](#) [health 95%](#)

⚠ Warning

As of v0.6.0 Yahoo!, Google Options, Google Quotes and EDGAR have been immediately deprecated due to large changes in their API and no stable replacement.

⚠ Note

As of v0.6.0 Google finance is still functioning for historical price data, although there are frequent reports of failures. Failure is frequently encountered when bulk downloading historical price data.

Usage

Starting in 0.19.0, pandas no longer supports `pandas.io.data` or `pandas.io.wb`, so you must replace your imports from `pandas.io` with those from `pandas_datareader`:

```
from pandas.io import data, wb # becomes
from pandas_datareader import data, wb
```

Many functions from the data module have been included in the top level API.

<https://pandas-datareader.readthedocs.io/en/latest/>

Quandl

ALTERNATIVE DATA FINANCIAL DATA

Quandl DOCS & HELP BLOG LOG IN SIGN UP

Analysis Tools ALL TOOLS API R PYTHON EXCEL



Get Financial Data Directly into Python

Get millions of financial and economic datasets from hundreds of publishers directly into Python.

[Load Quandl Data Directly Into Python](#)

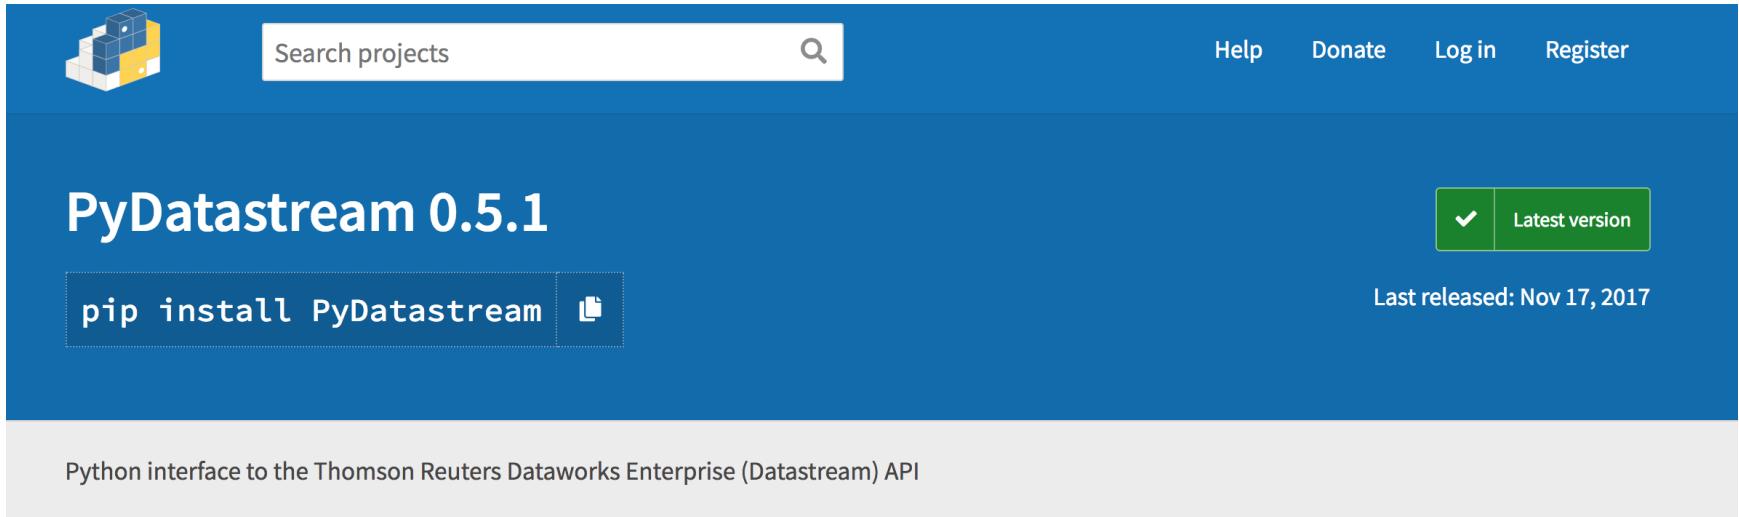
All the Data You Want

Quandl unifies financial and economic datasets from hundreds of publishers on a single user-friendly platform.

[Directly Into Python](#)

<https://www.quandl.com/tools/python>

PyDatastream



The screenshot shows the PyDatastream project page on the Python Package Index (pypi.org). The top navigation bar includes a search icon, a search bar labeled "Search projects", and links for "Help", "Donate", "Log in", and "Register". A logo consisting of three 3D cubes is visible. The main title "PyDatastream 0.5.1" is displayed prominently. Below it, a "Latest version" button is shown with a green checkmark icon. A "Last released: Nov 17, 2017" message is also present. On the left, there's a code snippet for installing the package via pip: "pip install PyDatastream" followed by a download icon. The page content area describes the project as a Python interface to the Thomson Reuters Dataworks Enterprise (Datastream) API.

PyDatastream 0.5.1

Last released: Nov 17, 2017

pip install PyDatastream

Python interface to the Thomson Reuters Dataworks Enterprise (Datastream) API

Navigation

Project description

Release history

Download files

Project description

PyDatastream is a Python interface to the **Thomson Dataworks Enterprise (DWE)** SOAP API (non free), with some convenience functions for retrieving Datastream data specifically. This package requires valid credentials for this API.

For the documentation please refer to README.md inside the package or on the GitHub (<https://github.com/vfilimonov/pydatastream/blob/master/README.md>).

<https://pypi.org/project/PyDatastream/>

pandasSDMX

Table Of Contents

[pandaSDMX: Statistical Data and Metadata eXchange in Python](#)

- [Supported data providers](#)
- [Main features](#)
- [Example](#)
- [Quick install](#)
- [pandaSDMX Links](#)
- [Table of contents](#)

[Indices and tables](#)

This Page

[Show Source](#)

Quick search

 Go

Take the [Python User's Survey](#) and let us know more about your usage of Python.

Sponsored · Ads served ethically

pandaSDMX: Statistical Data and Metadata eXchange in Python

pandaSDMX is an Apache 2.0-licensed [Python](#) client to retrieve and acquire statistical data and metadata disseminated in [SDMX](#) 2.1, an ISO-standard widely used by institutions such as statistics offices, central banks, and international organisations. pandaSDMX exposes datasets and related structural metadata including dataflows, codelists, and datastructure definitions as [pandas](#) Series or multi-indexed DataFrames. Many other output formats and storage backends are available thanks to [Odo](#).

Supported data providers

pandaSDMX ships with built-in support for the following agencies (others may be configured by the user):

- [Australian Bureau of Statistics \(ABS\)](#)
- [European Central Bank \(ECB\)](#)
- [Eurostat](#)
- [French National Institute for Statistics \(INSEE\)](#)
- [Instituto Nacional de la Estadística y Geografía - INEGI \(Mexico\)](#)
- [International Monetary Fund \(IMF\) - SDMX Central only](#)
- [International Labour Organization \(ILO\)](#)
- [Italian statistics Office \(ISTAT\)](#)
- [Norges Bank \(Norway\)](#)
- [Organisation for Economic Cooperation and Development \(OECD\)](#)
- [United Nations Statistics Division \(UNSD\)](#)
- [UNESCO \(free registration required\)](#)
- [World Bank - World Integrated Trade Solution \(WITS\)](#)

Fred API



ECONOMIC RESEARCH
FEDERAL RESERVE BANK *of* ST. LOUIS

REGISTER | SIGN IN

Search

FRED® Economic Data Information Services Publications Working Papers Economists About St. Louis Fed Home

Home > Developer API > FRED API

[API Keys](#) | [Terms of Use](#)

FRED® API

[General Documentation](#) | [API](#) | [Toolkits](#)

The FRED® API is a web service that allows developers to write programs and build applications that retrieve economic data from the [FRED®](#) and [ALFRED®](#) websites hosted by the [Economic Research Division](#) of the [Federal Reserve Bank of St. Louis](#). Requests can be customized according to data source, release, category, series, and other preferences.

General Documentation

- [Overview](#)
- [What is FRED®?](#)
- [What is ALFRED®?](#)
- [FRED® versus ALFRED®](#)
- [Real-Time Periods](#)
- [Errors](#)

API

Categories

- [fred/category](#) – Get a category.
- [fred/category/children](#) – Get the child categories for a specified parent category.
- [fred/category/related](#) – Get the related categories for a category.
- [fred/category/series](#) – Get the series in a category.
- [fred/category/tags](#) – Get the tags for a category.
- [fred/category/related_tags](#) – Get the related tags for a category.

<https://research.stlouisfed.org/docs/api/fred/>

Python Pandas for Finance

! pip install pandas_datareader

```
1 | ! pip install pandas_datareader

Collecting pandas_datareader
  Downloading https://files.pythonhosted.org/packages/cc/5c/ea5b6dcfd0f55c5fb1e37fb45335ec01cceca199b8a79339137f5ed269e0/pandas_datarea
  100% |██████████| 112kB 2.7MB/s
Collecting lxml (from pandas_datareader)
  Downloading https://files.pythonhosted.org/packages/03/a4/9eea8035fc7c7670e5eab97f34ff2ef0ddd78a491bf96df5accedb0e63f5/lxml-4.2.5-cp3
  100% |██████████| 5.8MB 7.5MB/s
Requirement already satisfied: pandas>=0.19.2 in /usr/local/lib/python3.6/dist-packages (from pandas_datareader) (0.22.0)
Requirement already satisfied: requests>=2.3.0 in /usr/local/lib/python3.6/dist-packages (from pandas_datareader) (2.18.4)
Requirement already satisfied: wrapt in /usr/local/lib/python3.6/dist-packages (from pandas_datareader) (1.10.11)
Requirement already satisfied: python-dateutil>=2 in /usr/local/lib/python3.6/dist-packages (from pandas>=0.19.2->pandas_datareader) (2
Requirement already satisfied: numpy>=1.9.0 in /usr/local/lib/python3.6/dist-packages (from pandas>=0.19.2->pandas_datareader) (1.14.6)
Requirement already satisfied: pytz>=2011k in /usr/local/lib/python3.6/dist-packages (from pandas>=0.19.2->pandas_datareader) (2018.5)
Requirement already satisfied: idna<2.7,>=2.5 in /usr/local/lib/python3.6/dist-packages (from requests>=2.3.0->pandas_datareader) (2.6)
Requirement already satisfied: chardet<3.1.0,>=3.0.2 in /usr/local/lib/python3.6/dist-packages (from requests>=2.3.0->pandas_datareader)
Requirement already satisfied: certifi>=2017.4.17 in /usr/local/lib/python3.6/dist-packages (from requests>=2.3.0->pandas_datareader) (
Requirement already satisfied: urllib3<1.23,>=1.21.1 in /usr/local/lib/python3.6/dist-packages (from requests>=2.3.0->pandas_datareader)
Requirement already satisfied: six>=1.5 in /usr/local/lib/python3.6/dist-packages (from python-dateutil>=2->pandas>=0.19.2->pandas_data
Installing collected packages: lxml, pandas-datareader
Successfully installed lxml-4.2.5 pandas-datareader-0.7.0
```

conda install pandas-datareader

```
iMyday—bash—80x24
[iMyday-MacBook-Pro:~ imyday$ conda install pandas-datareader
Fetching package metadata .....
Solving package specifications: .

Package plan for installation in environment /Users/imyday/anaconda:

The following NEW packages will be INSTALLED:

pandas-datareader: 0.2.1-py36_0
requests-file:    1.4.1-py36_0

Proceed ([y]/n)? y

requests-file- 100% |#####
pandas-datarea 100% |#####
[iMyday-MacBook-Pro:~ imyday$ conda list
# packages in environment at /Users/imyday/anaconda:
#
_license          1.1                  py36_1
alabaster         0.7.9                py36_0
anaconda          4.3.1                np111py36_0
anaconda-client   1.6.0                py36_0
anaconda-navigator 1.5.0                py36_0
anaconda-project  0.4.1                py36_0
```

AAPL



Finance Data from Yahoo Finance

```
# !pip install pandas_datareader
import pandas_datareader.data as web
import datetime as dt
#Read Stock Data from Yahoo Finance
end = dt.datetime(2017, 12, 31)
start = dt.datetime(2016, 1, 1)
df = web.DataReader("AAPL", 'yahoo', start, end)
df.to_csv('AAPL.csv')
df.from_csv('AAPL.csv')
df.tail()
```

```
# !pip install pandas_datareader
import pandas as pd
import pandas_datareader.data as web
import matplotlib.pyplot as plt
import seaborn as sns
import datetime as dt
%matplotlib inline

#Read Stock Data from Yahoo Finance
end = dt.datetime.now()
#start = dt.datetime(end.year-2, end.month,
end.day)
start = dt.datetime(2016, 1, 1)
df = web.DataReader("AAPL", 'yahoo', start, end)
df.to_csv('AAPL.csv')
df.from_csv('AAPL.csv')
df.tail()
```

```
df[ 'Adj Close' ].plot(legend=True, figsize=(12, 8), title='AAPL', label='Adj Close')
```



```
plt.figure(figsize=(12,9))
top = plt.subplot2grid((12,9), (0, 0),
rowspan=10, colspan=9)
bottom = plt.subplot2grid((12,9), (10,0),
rowspan=2, colspan=9)
top.plot(df.index, df[ 'Adj Close'],
color='blue') #df.index gives the dates
bottom.bar(df.index, df[ 'Volume' ])
```

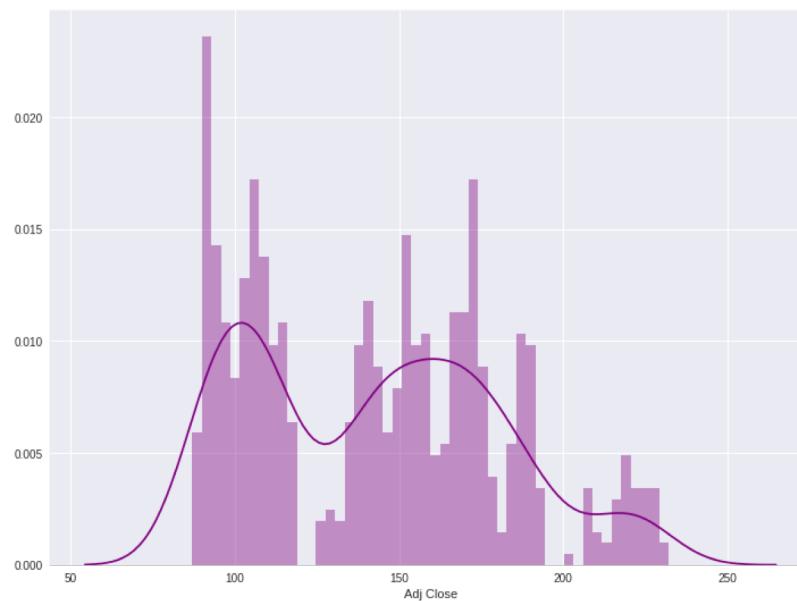


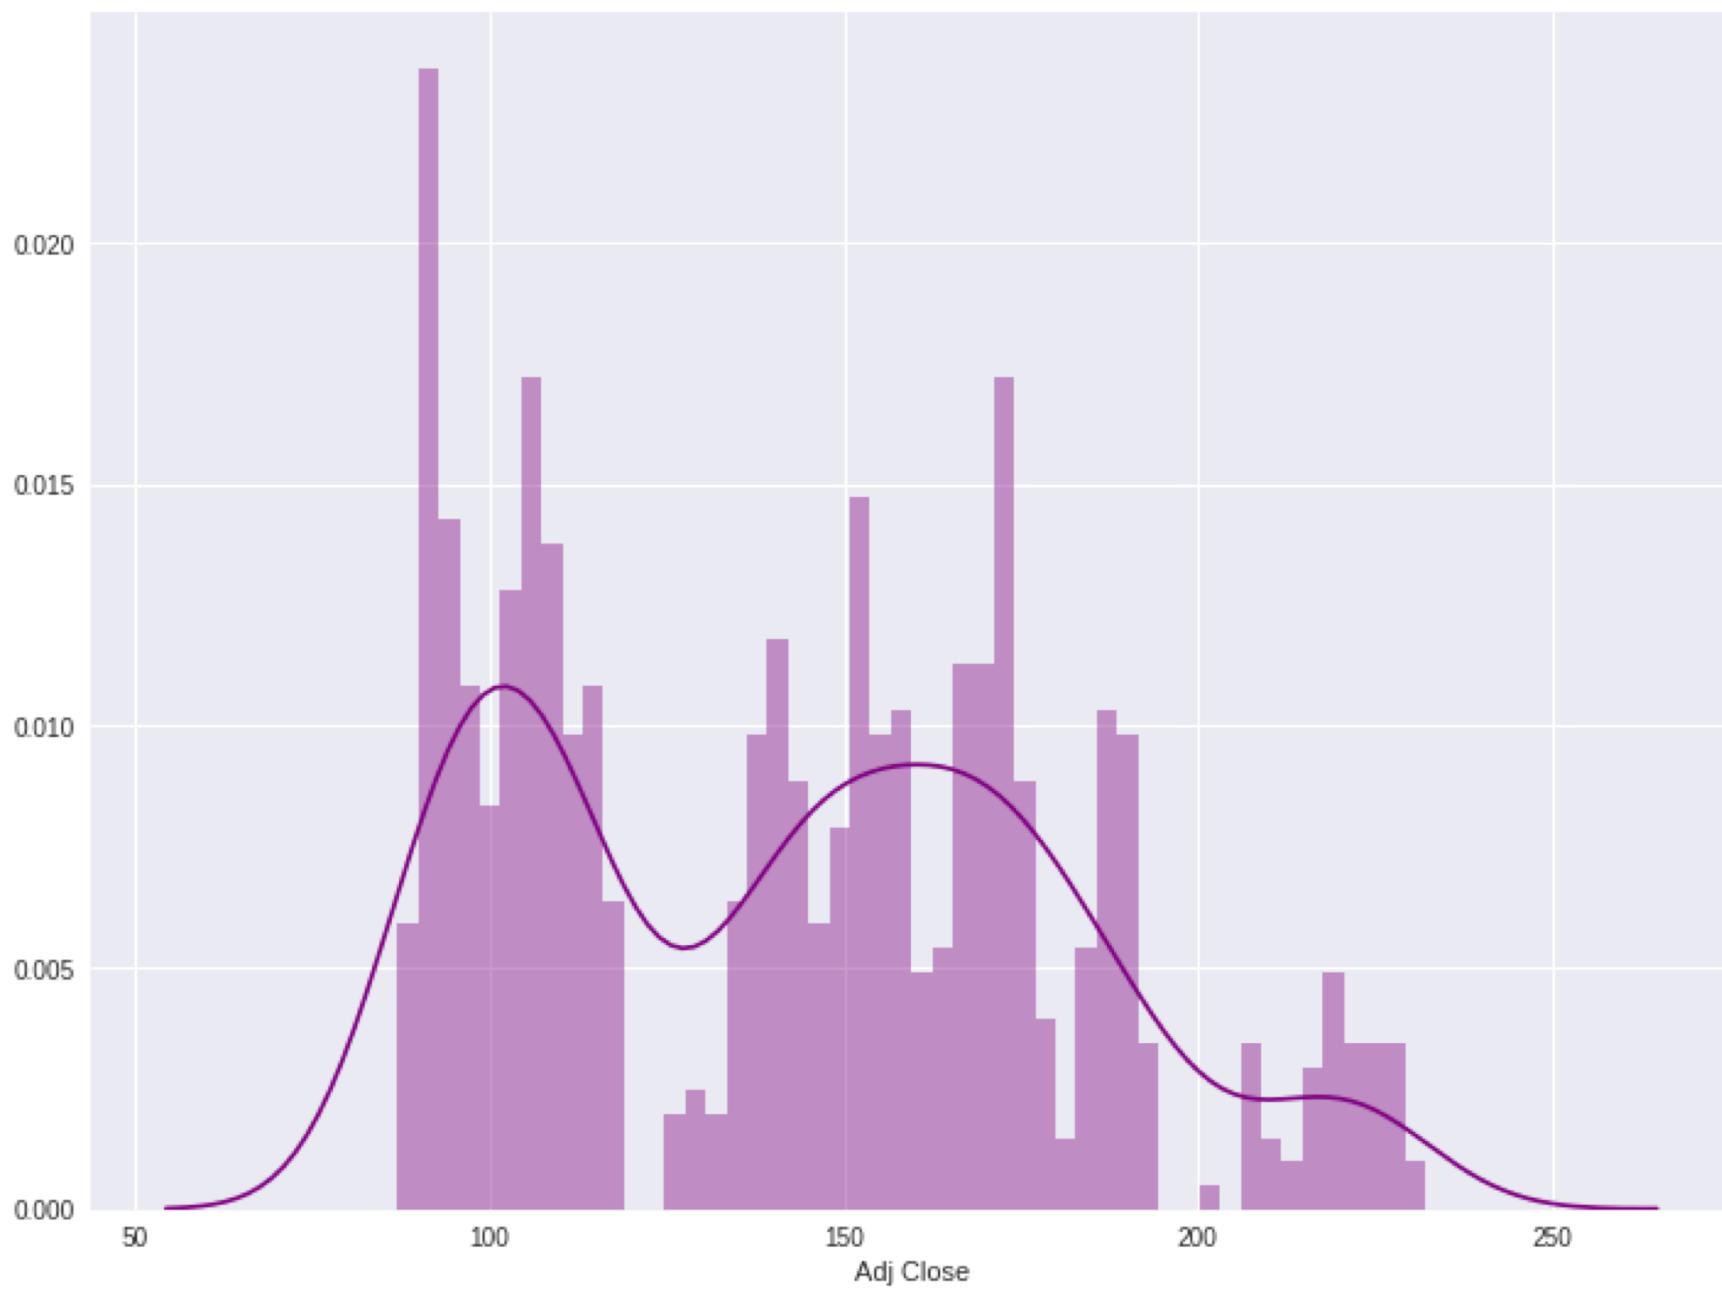
AAPL



```
# set the labels
top.axes.get_xaxis().set_visible(False)
top.set_title('AAPL')
top.set_ylabel('Adj Close')
bottom.set_ylabel('Volume')

plt.figure(figsize=(12,9))
sns.distplot(df['Adj Close'].dropna(),
bins=50, color='purple')
```





```
# simple moving averages
df[ 'MA05' ] = df[ 'Adj Close' ].rolling(5).mean()
#5 days
df[ 'MA20' ] = df[ 'Adj
Close' ].rolling(20).mean() #20 days
df[ 'MA60' ] = df[ 'Adj
Close' ].rolling(60).mean() #60 days
df2 = pd.DataFrame({ 'Adj Close': df[ 'Adj
Close' ], 'MA05': df[ 'MA05' ], 'MA20': df[ 'MA20' ],
'MA60': df[ 'MA60' ]})
df2.plot(figsize=(12, 9), legend=True,
title='AAPL')
df2.to_csv('AAPL_MA.csv')
fig = plt.gcf()
fig.set_size_inches(12, 9)
fig.savefig('AAPL_plot.png', dpi=300)
plt.show()
```

AAPL



```

# !pip install pandas_datareader
import pandas as pd
import pandas_datareader.data as web
import matplotlib.pyplot as plt
import seaborn as sns
import datetime as dt
%matplotlib inline

#Read Stock Data from Yahoo Finance
end = dt.datetime.now()
#start = dt.datetime(end.year-2, end.month, end.day)
start = dt.datetime(2016, 1, 1)
df = web.DataReader("AAPL", 'yahoo', start, end)
df.to_csv('AAPL.csv')
df.from_csv('AAPL.csv')
df.tail()

df['Adj Close'].plot(legend=True, figsize=(12, 8), title='AAPL', label='Adj Close')
plt.figure(figsize=(12,9))
top = plt.subplot2grid((12,9), (0, 0), rowspan=10, colspan=9)
bottom = plt.subplot2grid((12,9), (10,0), rowspan=2, colspan=9)
top.plot(df.index, df['Adj Close'], color='blue') #df.index gives the dates
bottom.bar(df.index, df['Volume'])

# set the labels
top.axes.get_xaxis().set_visible(False)
top.set_title('AAPL')
top.set_ylabel('Adj Close')
bottom.set_ylabel('Volume')

plt.figure(figsize=(12,9))
sns.distplot(df['Adj Close'].dropna(), bins=50, color='purple')

# simple moving averages
df['MA05'] = df['Adj Close'].rolling(5).mean() #5 days
df['MA20'] = df['Adj Close'].rolling(20).mean() #20 days
df['MA60'] = df['Adj Close'].rolling(60).mean() #60 days
df2 = pd.DataFrame({'Adj Close': df['Adj Close'], 'MA05': df['MA05'], 'MA20': df['MA20'], 'MA60': df['MA60']})
df2.plot(figsize=(12, 9), legend=True, title='AAPL')
df2.to_csv('AAPL_MA.csv')
fig = plt.gcf()
fig.set_size_inches(12, 9)
fig.savefig('AAPL_plot.png', dpi=300)
plt.show()

```

```

1 # !pip install pandas_datareader
2 import pandas as pd
3 import pandas_datareader.data as web
4 import matplotlib.pyplot as plt
5 import seaborn as sns
6 import datetime as dt
7 %matplotlib inline
8
9 #Read Stock Data from Yahoo Finance
10 end = dt.datetime.now()
11 #start = dt.datetime(end.year-2, end.month, end.day)
12 start = dt.datetime(2016, 1, 1)
13 df = web.DataReader("AAPL", 'yahoo', start, end)
14 df.to_csv('AAPL.csv')
15 df.from_csv('AAPL.csv')
16 df.tail()
17
18 df['Adj Close'].plot(legend=True, figsize=(12, 8), title='AAPL', label='Adj Close')
19 plt.figure(figsize=(12,9))
20 top = plt.subplot2grid((12,9), (0, 0), rowspan=10, colspan=9)
21 bottom = plt.subplot2grid((12,9), (10,0), rowspan=2, colspan=9)
22 top.plot(df.index, df['Adj Close'], color='blue') #df.index gives the dates
23 bottom.bar(df.index, df['Volume'])
24
25 # set the labels
26 top.axes.get_xaxis().set_visible(False)
27 top.set_title('AAPL')
28 top.set_ylabel('Adj Close')
29 bottom.set_ylabel('Volume')
30
31 plt.figure(figsize=(12,9))
32 sns.distplot(df['Adj Close'].dropna(), bins=50, color='purple')
33
34 # simple moving averages
35 df['MA05'] = df['Adj Close'].rolling(5).mean() #5 days
36 df['MA20'] = df['Adj Close'].rolling(20).mean() #20 days
37 df['MA60'] = df['Adj Close'].rolling(60).mean() #60 days
38 df2 = pd.DataFrame({'Adj Close': df['Adj Close'], 'MA05': df['MA05'], 'MA20': df['MA20'], 'MA60': df['MA60']})
39 df2.plot(figsize=(12, 9), legend=True, title='AAPL')
40 df2.to_csv('AAPL_MA.csv')
41 fig = plt.gcf()
42 fig.set_size_inches(12, 9)
43 fig.savefig('AAPL_plot.png', dpi=300)
44 plt.show()

```

Finance Data from Quandl

```
# ! pip install quandl
import quandl
# quandl.ApiConfig.api_key = "YOURAPIKEY"
df = quandl.get("WIKI/AAPL", start_date="2016-01-01", end_date="2017-12-31")
df.to_csv('AAPL.csv')
df.from_csv('AAPL.csv')
df.tail()
```

```
1 # ! pip install quandl
2 import quandl
3 # quandl.ApiConfig.api_key = "YOURAPIKEY"
4 df = quandl.get("WIKI/AAPL", start_date="2016-01-01", end_date="2017-12-31")
5 df.to_csv('AAPL.csv')
6 df.from_csv('AAPL.csv')
7 df.tail()
```

```
/usr/local/lib/python3.6/dist-packages/ipykernel_launcher.py:5: FutureWarning: from_csv is deprecated. Please use read_csv(...) instead
"""
```

Date	Open	High	Low	Close	Volume	Ex-Dividend	Split Ratio	Adj. Open	Adj. High	Adj. Low	Adj. Close	Adj. Volume
2017-12-22	174.68	175.424	174.500	175.01	16052615.0	0.0	1.0	174.68	175.424	174.500	175.01	16052615.0
2017-12-26	170.80	171.470	169.679	170.57	32968167.0	0.0	1.0	170.80	171.470	169.679	170.57	32968167.0
2017-12-27	170.10	170.780	169.710	170.60	21672062.0	0.0	1.0	170.10	170.780	169.710	170.60	21672062.0
2017-12-28	171.00	171.850	170.480	171.08	15997739.0	0.0	1.0	171.00	171.850	170.480	171.08	15997739.0
2017-12-29	170.52	170.590	169.220	169.23	25643711.0	0.0	1.0	170.52	170.590	169.220	169.23	25643711.0

Deep Learning with TensorFlow

TensorFlow

TensorFlow™

Install

Develop

Community

API ▾

Resources ▾



Search

GITHUB

An open source machine learning
framework for everyone

GET STARTED



Get started with
TensorFlow

There are new tutorials to get started
with Tensorflow using tf.keras and
eager execution. Run the Colab
notebooks directly in the browser.



TensorFlow 1.12 is here!

TensorFlow 1.12 is available, see the
release notes for the latest updates.



Announcing TensorFlow.js

Learn about our JavaScript library for
machine learning in the browser.

<https://www.tensorflow.org/>

Deep Learning Software

- **TensorFlow**
 - TensorFlow™ is an open source software library for high performance numerical computation.
- **Keras**
 - Deep Learning library for **TensorFlow**, **CNTK**
- **PyTorch**
 - An open source deep learning platform that provides a seamless path from research prototyping to production deployment.
- **CNTK**
 - Computational Network Toolkit by Microsoft Research

tf.keras

Keras:

High-level API

for TensorFlow

Keras

Search docs

Home

Keras: The Python Deep Learning library

You have just found Keras.

Guiding principles

Getting started: 30 seconds to Keras

Installation

Configuring your Keras backend

Support

Why this name, Keras?

Activations

Applications

Backend

Callbacks

Constraints

Contributing

Datasets

Initializers

Losses

Metrics

Docs » Home

 Edit on GitHub

Keras: The Python Deep Learning library



Keras

You have just found Keras.

Keras is a high-level neural networks API, written in Python and capable of running on top of [TensorFlow](#), [CNTK](#), or [Theano](#). It was developed with a focus on enabling fast experimentation. *Being able to go from idea to result with the least possible delay is key to doing good research.*

Use Keras if you need a deep learning library that:

- Allows for easy and fast prototyping (through user friendliness, modularity, and extensibility).
- Supports both convolutional networks and recurrent networks, as well as combinations of the two.
- Runs seamlessly on CPU and GPU.

Read the documentation at [Keras.io](#).

Keras is compatible with: **Python 2.7-3.6**.

<http://keras.io/>

PyTorch

[Get Started](#)[Features](#)[Ecosystem](#)[Blog](#)[Tutorials](#)[Docs](#)[Resources](#)[GitHub](#)

FROM RESEARCH TO PRODUCTION

An open source deep learning platform that provides a seamless path from research prototyping to production deployment.

[Get Started >](#)

KEY FEATURES &
CAPABILITIES

[See all Features >](#)

<http://pytorch.org/>



Keras



Keras

- Keras is a **high-level neural networks API**
- Written in Python and capable of running on top of **TensorFlow**, **CNTK**, or **Theano**.
- It was developed with a focus on enabling fast experimentation.
- Being able to go from idea to result with the least possible delay is key to doing good research.



TensorFlow

Google TensorFlow

TensorFlow™

GET STARTED TUTORIALS HOW TO API RESOURCES ABOUT

Fork me on GitHub

TensorFlow is an Open Source Software
Library for Machine Intelligence

GET STARTED

About TensorFlow

TensorFlow™ is an open source software library for numerical computation using data flow graphs. Nodes in the graph represent mathematical operations, while the graph edges represent the multidimensional data arrays (tensors) communicated between them. The flexible architecture allows you to deploy computation to one or more CPUs or GPUs in a desktop, server, or mobile device with a single API.



<https://www.tensorflow.org/>

TensorFlow Playground

Tinker With a **Neural Network** Right Here in Your Browser.
Don't Worry, You Can't Break It. We Promise.

Iterations
000,582Learning rate
0.03Activation
TanhRegularization
NoneRegularization rate
0Problem type
Classification

DATA

Which dataset do you want to use?



Ratio of training to test data: 50%

Noise: 0

Batch size: 10

INPUT

Which properties do you want to feed in?

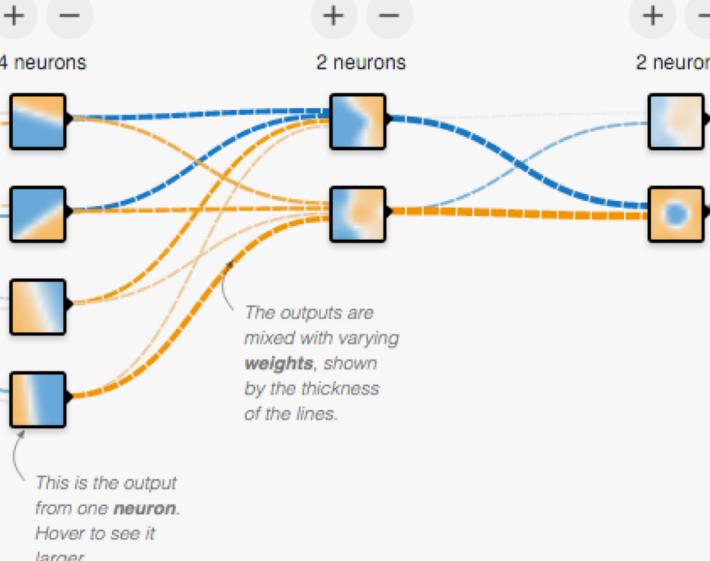


3 HIDDEN LAYERS

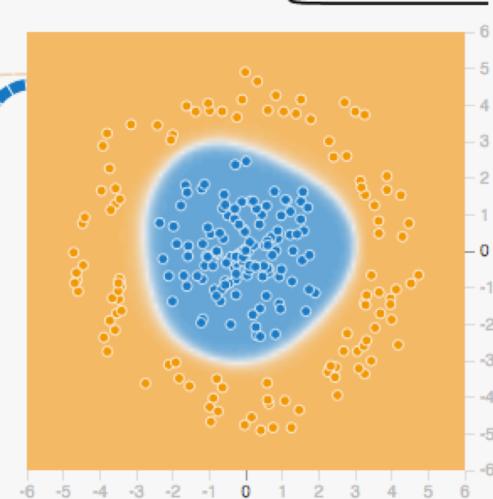
+ - 4 neurons

+ - 2 neurons

+ - 2 neurons



OUTPUT

Test loss 0.000
Training loss 0.000

TensorFlow

is an

Open Source

Software Library

for

Machine Intelligence

numerical computation using data flow graphs

Tensor

- 3
 - # a rank 0 tensor; this is a **scalar** with shape []
- [1., 2., 3.]
 - # a rank 1 tensor; this is a **vector** with shape [3]
- [[1., 2., 3.], [4., 5., 6.]]
 - # a rank 2 tensor; a **matrix** with shape [2, 3]
- [[[1., 2., 3.]], [[7., 8., 9.]]]
 - # a rank 3 **tensor** with shape [2, 1, 3]

Scalar

80

Vector

[50 60 70]

Matrix

$$\begin{bmatrix} 50 & 60 & 70 \\ 55 & 65 & 75 \end{bmatrix}$$

Tensor

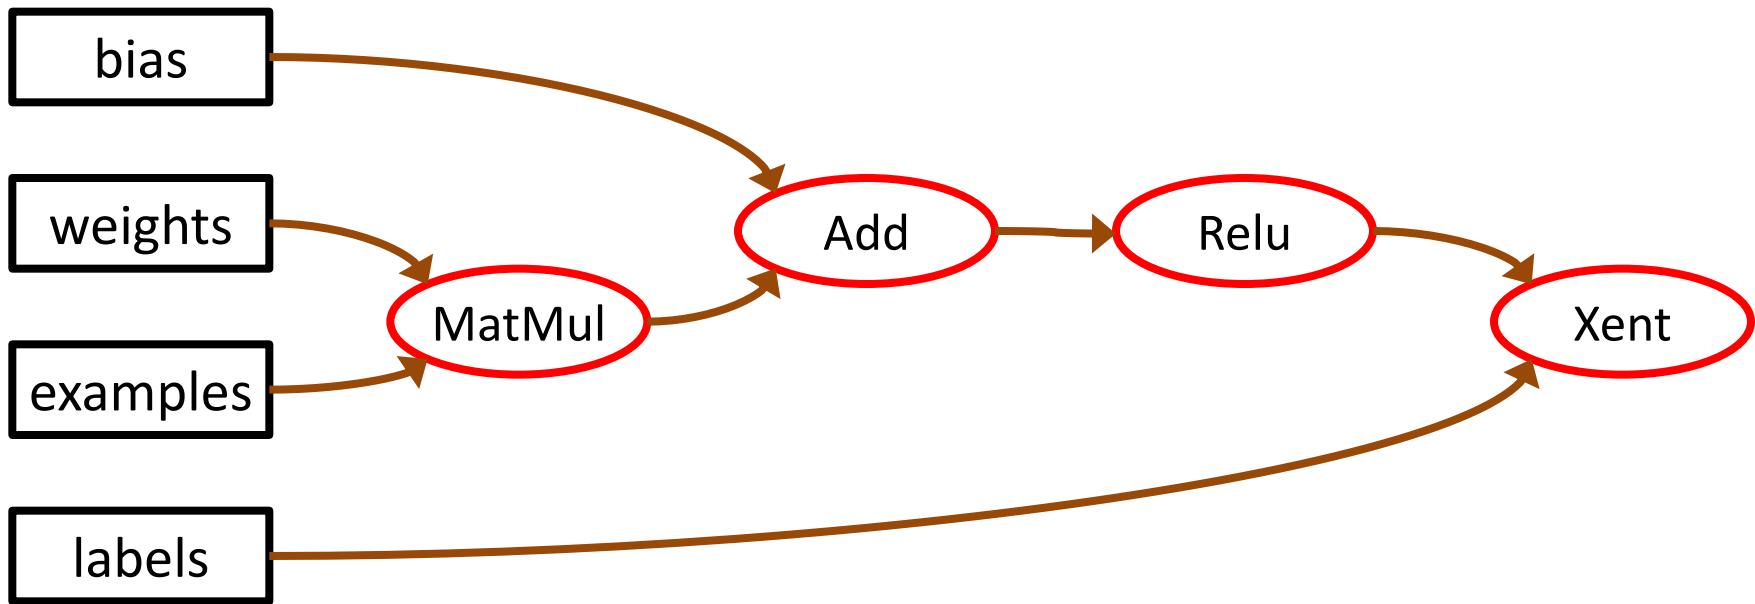
$$\begin{bmatrix} [50 60 70] & [70 80 90] \\ [55 65 75] & [75 85 95] \end{bmatrix}$$

Nodes:
mathematical operations

edges:
multidimensional data arrays
(tensors)
communicated between nodes

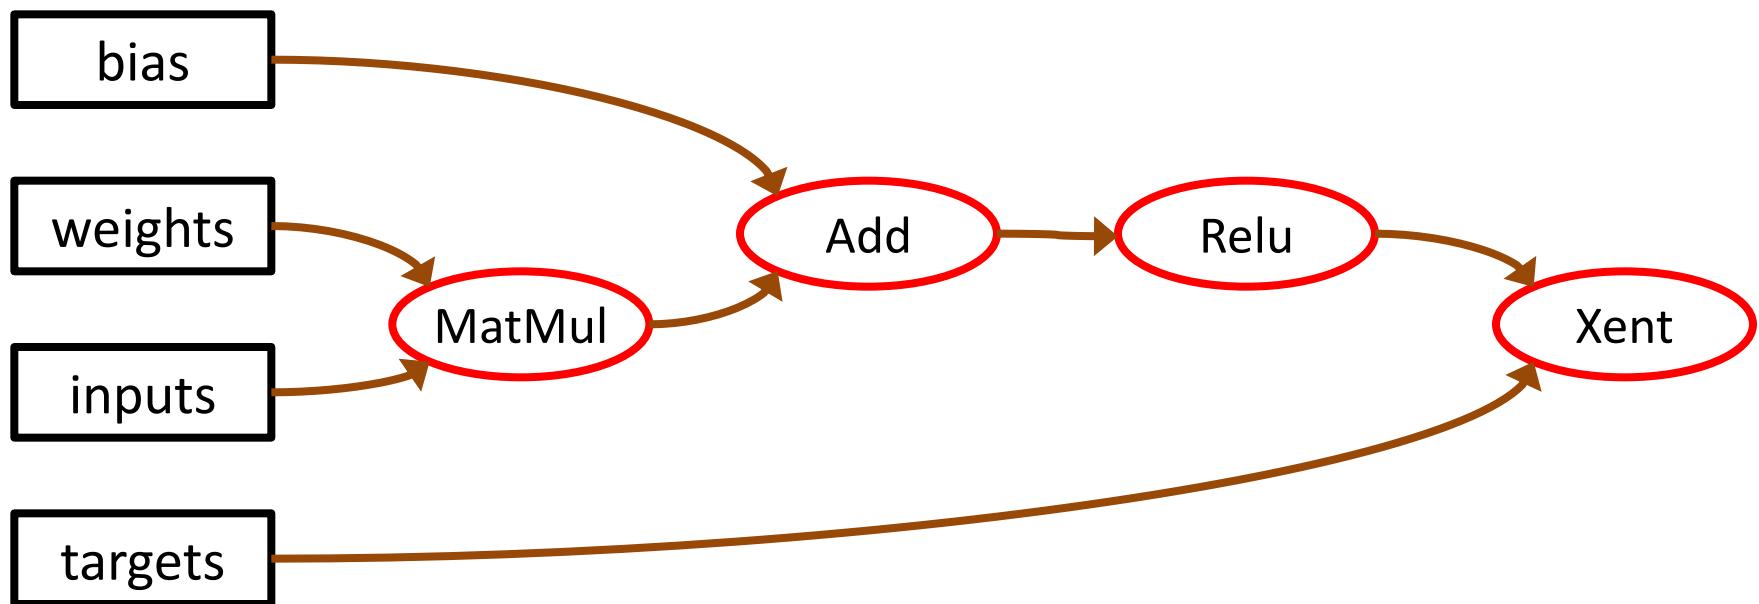
Computation is a Dataflow Graph

Graph of Nodes,
also called Operations or ops.

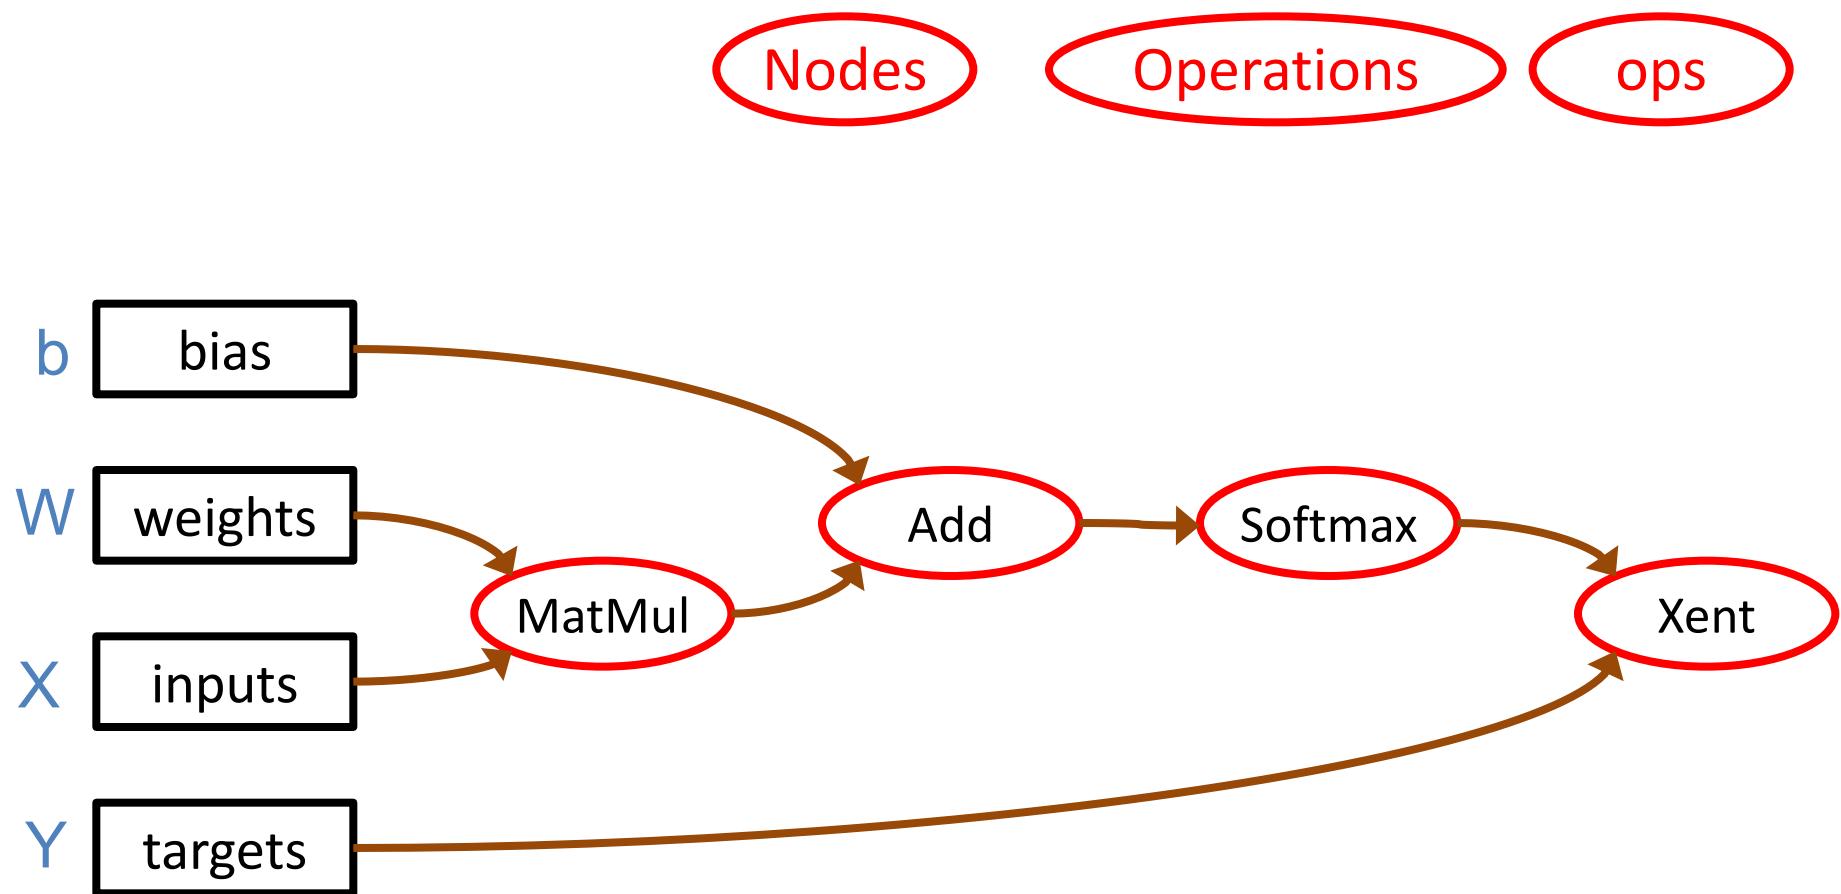


Computation is a Dataflow Graph

Edges are N-dimensional arrays: **Tensors**



Logistic Regression as Dataflow Graph

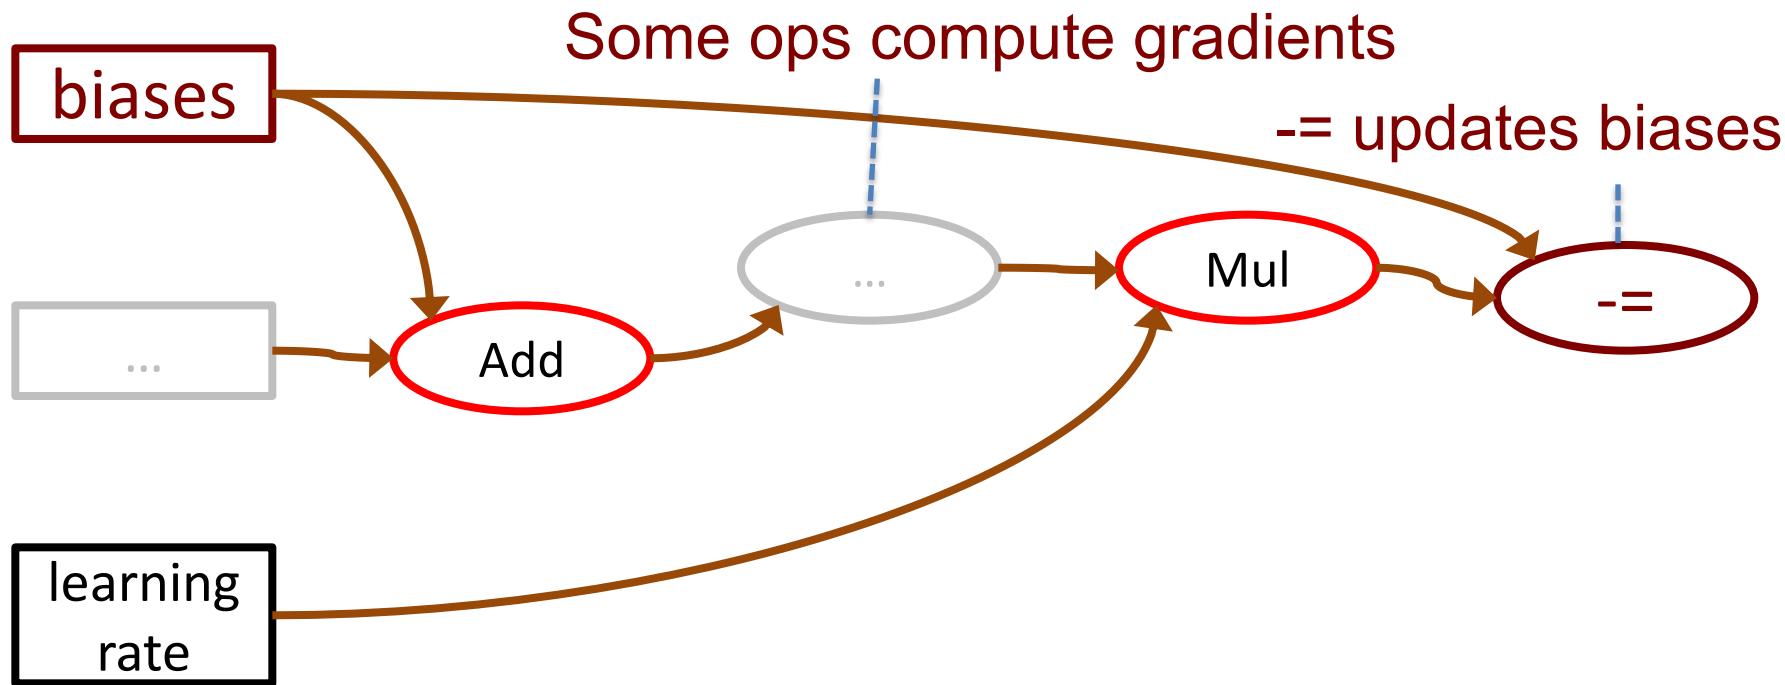


Edges are N-dimensional arrays: **Tensors**

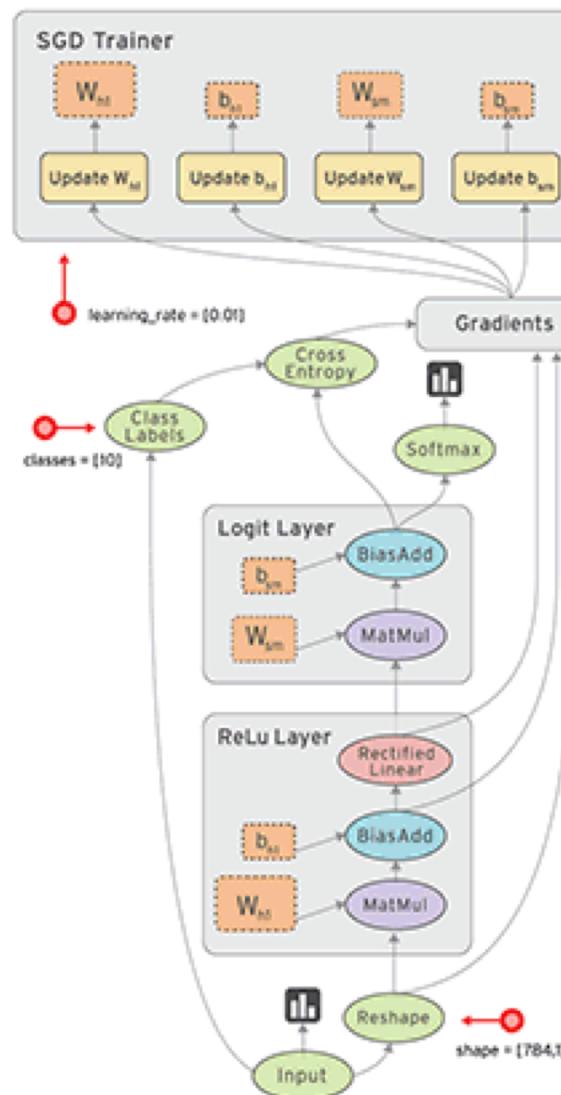
Computation is a Dataflow Graph

with state

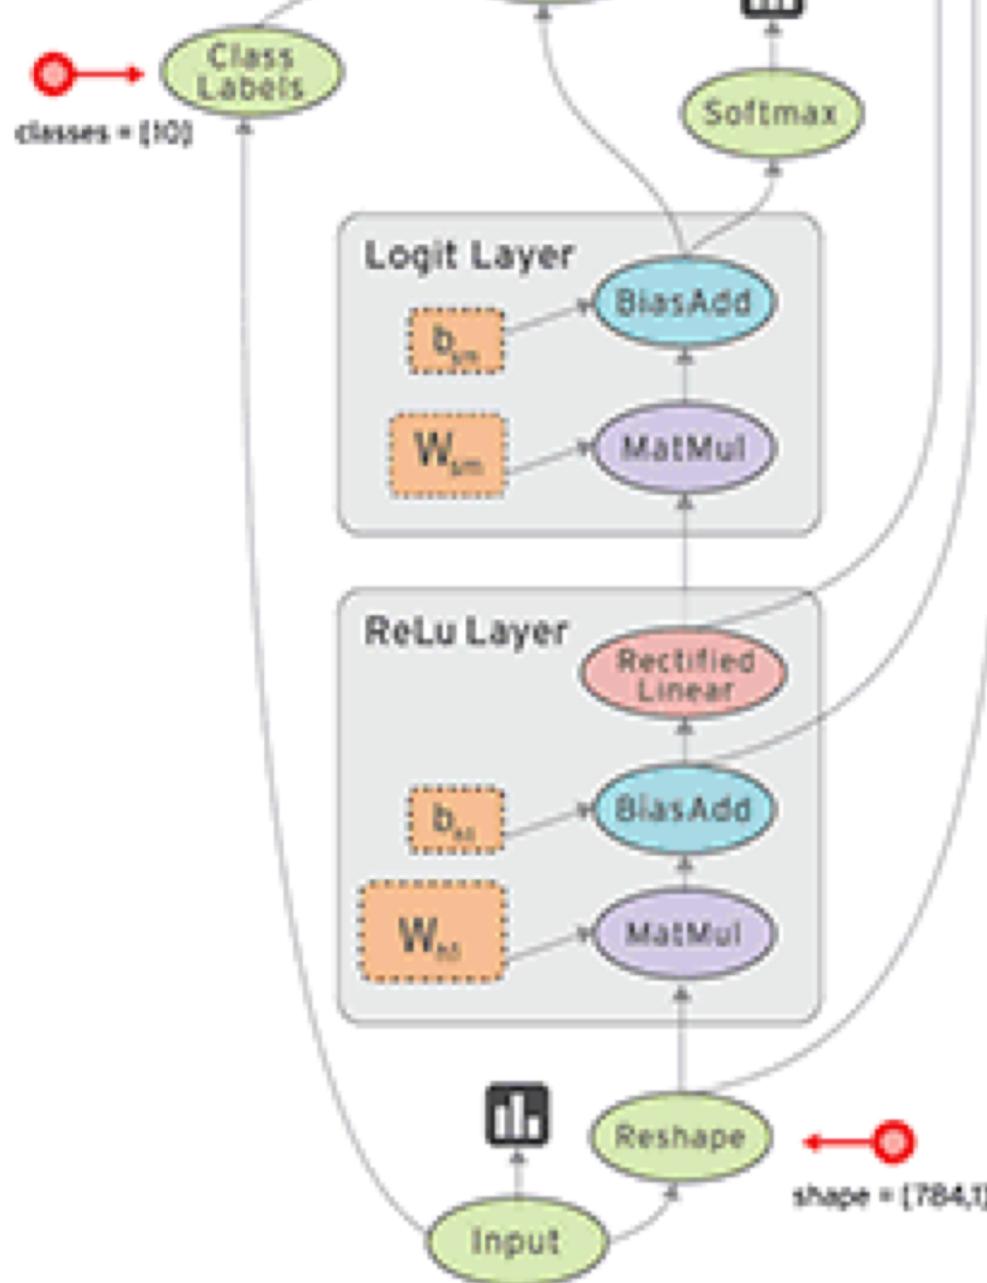
'Biases' is a variable



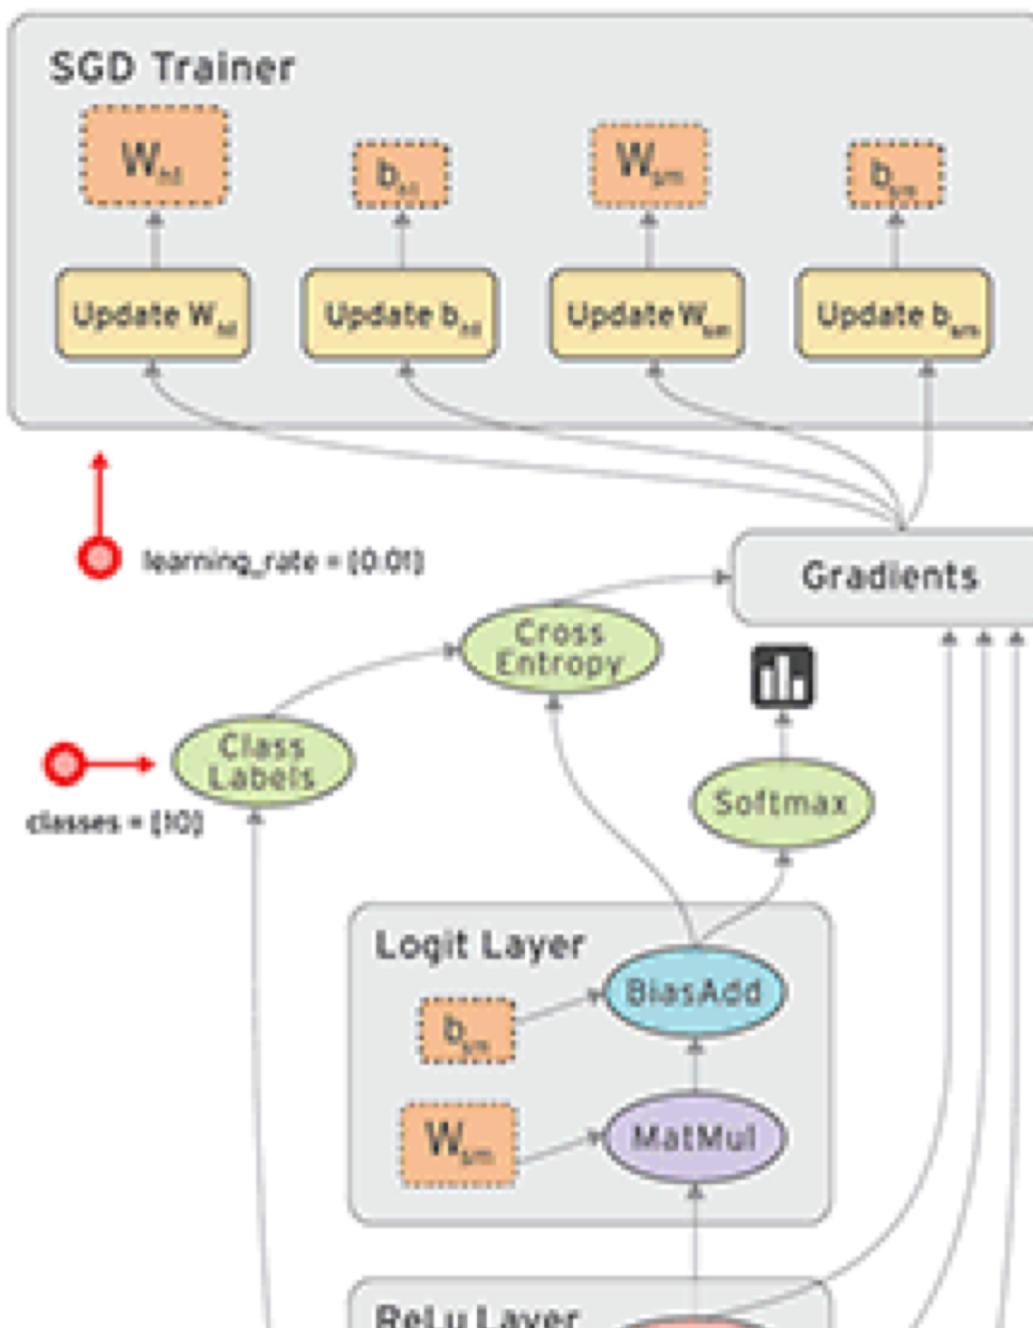
Data Flow Graph



Data Flow Graph



Data Flow Graph



TensorFlow

TensorBoard

TensorBoard

EVENTS

IMAGES

GRAPHS

HISTOGRAMS



Fit to screen

Download PNG

Run (1)

Session runs (0)

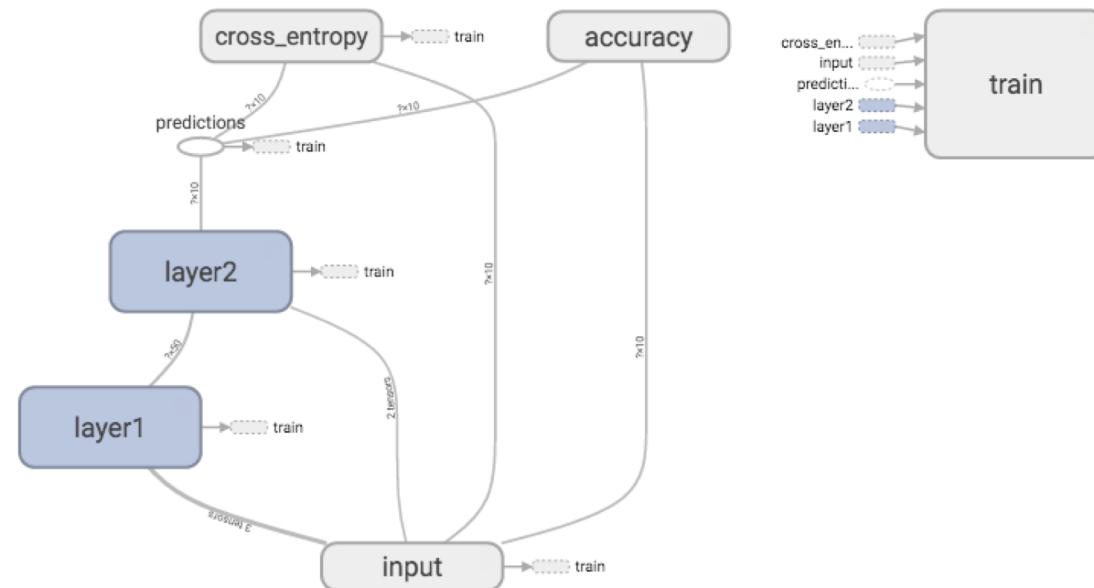
Upload

Color Structure Device
color: same substructure gray: unique substructure

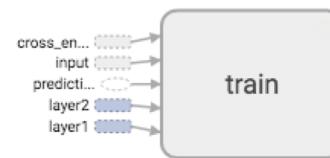
Graph (* = expandable)

- Namespace*
- OpNode
- Unconnected series*
- Connected series*
- Constant
- Summary
- Dataflow edge
- Control dependency edge
- Reference edge

Main Graph



Auxiliary nodes



Getting Started with TensorFlow

TensorFlow™

Install

Develop

API r1.4

Deploy

Extend

Community

Versions

>



Search

GITHUB

Develop

GET STARTED

PROGRAMMER'S GUIDE

TUTORIALS

PERFORMANCE

MOBILE

Getting Started

[Getting Started With TensorFlow](#)

MNIST For ML Beginners

Deep MNIST for Experts

TensorFlow Mechanics 101

tf.estimator Quickstart

Building Input Functions with tf.
estimator

TensorBoard: Visualizing Learning

TensorBoard: Graph Visualization

TensorBoard Histogram Dashboard

TensorFlow Versions

Getting Started With TensorFlow

This guide gets you started programming in TensorFlow. Before using this guide, [install TensorFlow](#).

To get the most out of this guide, you should know the following:

- How to program in Python.
- At least a little bit about arrays.
- Ideally, something about machine learning. However, if you know little or nothing about machine learning, then this is still the first guide you should read.

TensorFlow provides multiple APIs. The lowest level API--TensorFlow Core-- provides you with complete programming control. We recommend TensorFlow Core for machine learning researchers and others who require fine levels of control over their models. The higher level APIs are built on top of TensorFlow Core. These higher level APIs are typically easier to learn and use than TensorFlow Core. In addition, the higher level APIs make repetitive tasks easier and more consistent between different users. A high-level API like tf.estimator helps you manage data sets, estimators, training and inference.

This guide begins with a tutorial on TensorFlow Core. Later, we demonstrate how to implement the same model in tf.estimator. Knowing TensorFlow Core principles will give you a great mental model of how things are working internally when you use the more compact higher level API.

Contents

TensorFlow Core tutorial

Importing TensorFlow

The Computational
Graph

tf.train API

Complete program

tf.estimator

Basic usage

A custom model

Next steps

Try your first TensorFlow

```
$ python
```

```
>>> import tensorflow as tf
>>> hello = tf.constant('Hello, TensorFlow!')
>>> sess = tf.Session()
>>> sess.run(hello)
'Hello, TensorFlow!'
>>> a = tf.constant(10)
>>> b = tf.constant(32)
>>> sess.run(a+b)
42
>>>
```

Hello TensorFlow

```
import tensorflow as tf
hello = tf.constant('Hello, TensorFlow!')
sess = tf.Session()
sess.run(hello)
```

b'Hello, TensorFlow!'

`tf.Session()`
`sess.run()`

```
import tensorflow as tf
sess = tf.Session()
a = tf.constant(10)
b = tf.constant(32)
sess.run(a+b)
```

42

Linear Regression Model

```
import tensorflow as tf

# Model parameters
W = tf.Variable([.3], dtype=tf.float32)
b = tf.Variable([- .3], dtype=tf.float32)
# Model input and output
x = tf.placeholder(tf.float32)
linear_model = W*x + b
y = tf.placeholder(tf.float32)

# loss
loss = tf.reduce_sum(tf.square(linear_model - y)) # sum of the squares
# optimizer
optimizer = tf.train.GradientDescentOptimizer(0.01)
train = optimizer.minimize(loss)

# training data
x_train = [1, 2, 3, 4]
y_train = [0, -1, -2, -3]
# training loop
init = tf.global_variables_initializer()
sess = tf.Session()
sess.run(init) # reset values to wrong
for i in range(1000):
    sess.run(train, {x: x_train, y: y_train})

# evaluate training accuracy
curr_W, curr_b, curr_loss = sess.run([W, b, loss], {x: x_train, y: y_train})
print("W: %s b: %s loss: %s" % (curr_W, curr_b, curr_loss))
```

W: [-0.9999969] b: [0.99999082] loss: 5.69997e-11

tf.estimator

```
import numpy as np
import tensorflow as tf

feature_columns = [tf.feature_column.numeric_column("x", shape=[1])]

estimator = tf.estimator.LinearRegressor(feature_columns=feature_columns)

x_train = np.array([1., 2., 3., 4.])
y_train = np.array([0., -1., -2., -3.])
x_eval = np.array([2., 5., 8., 1.])
y_eval = np.array([-1.01, -4.1, -7, 0.])
input_fn = tf.estimator.inputs.numpy_input_fn(
    {"x": x_train}, y_train, batch_size=4, num_epochs=None, shuffle=True)
train_input_fn = tf.estimator.inputs.numpy_input_fn(
    {"x": x_train}, y_train, batch_size=4, num_epochs=1000, shuffle=False)
eval_input_fn = tf.estimator.inputs.numpy_input_fn(
    {"x": x_eval}, y_eval, batch_size=4, num_epochs=1000, shuffle=False)

estimator.train(input_fn=input_fn, steps=1000)

train_metrics = estimator.evaluate(input_fn=train_input_fn)
eval_metrics = estimator.evaluate(input_fn=eval_input_fn)
print("train metrics: %r" % train_metrics)
print("eval metrics: %r" % eval_metrics)
```

```
train metrics: {'average_loss': 2.7210228e-07, 'loss': 1.0884091e-06, 'global_step': 1000}
eval metrics: {'average_loss': 0.0025725411, 'loss': 0.010290165, 'global_step': 1000}
```

Deep Learning for Financial Time Series Forecasting

Deep Learning for Financial Market Prediction

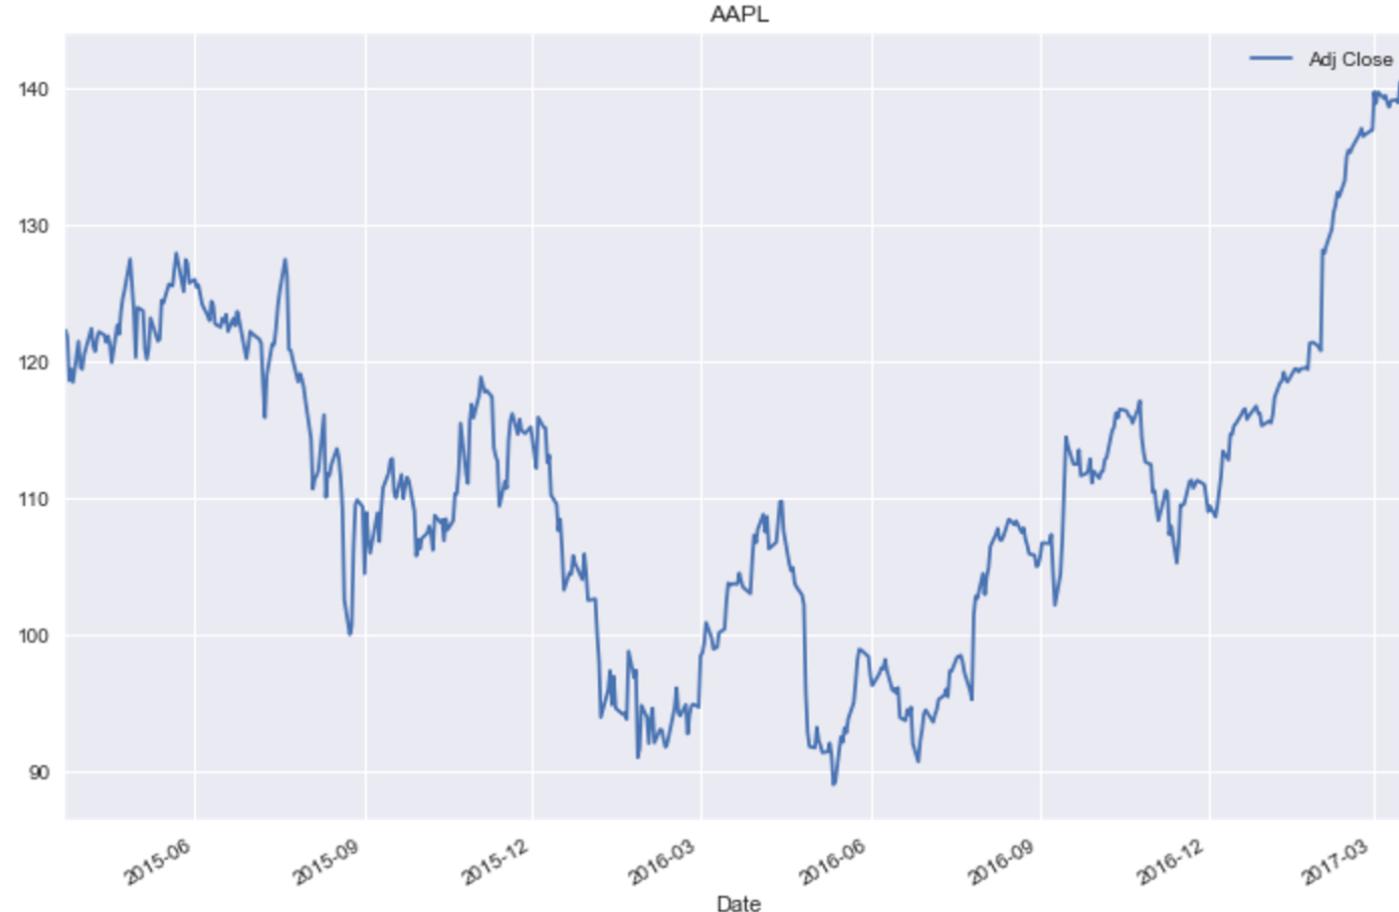
Stock Market Prediction

Stock Price Prediction

Time Series Prediction

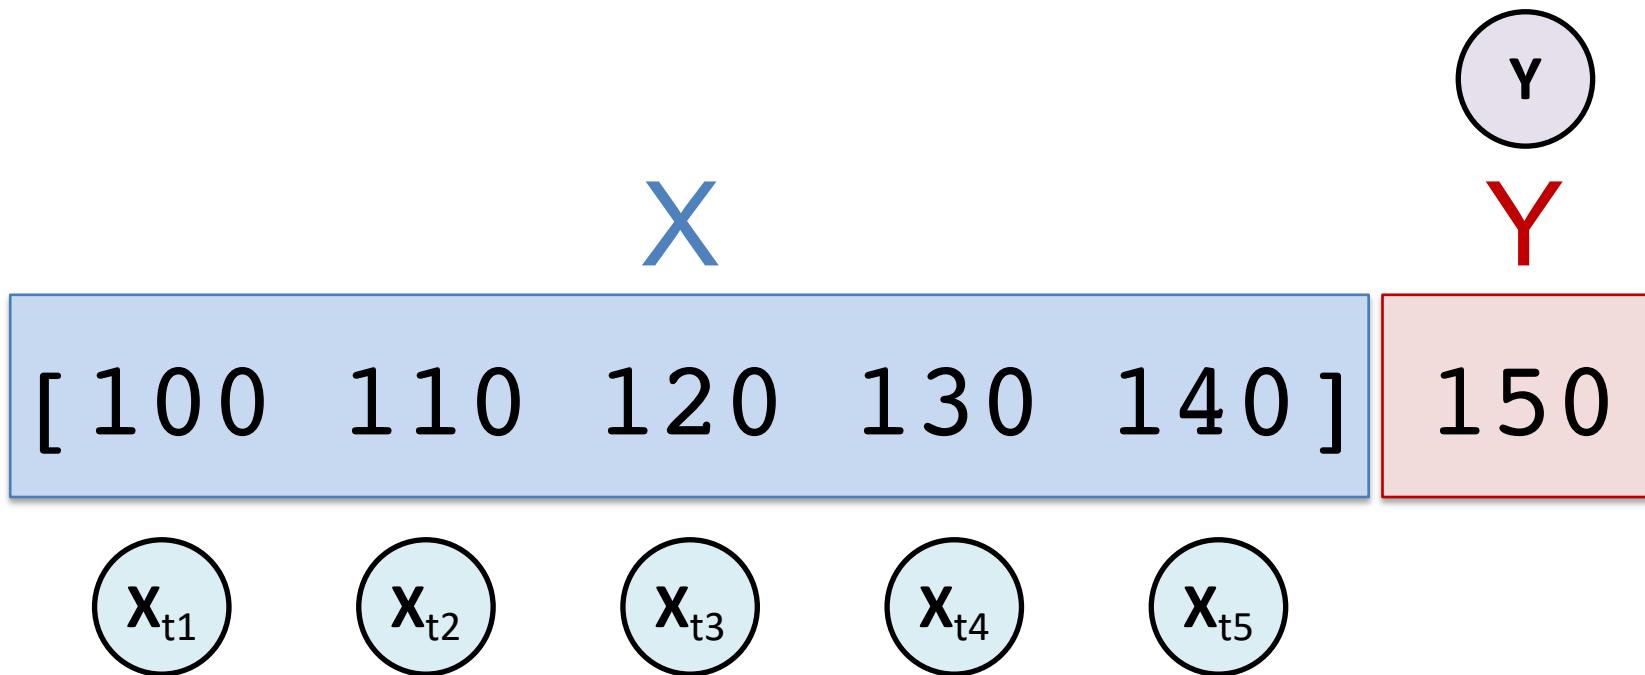
Time Series Data

```
df['Adj Close'].plot(legend=True, figsize=(12, 8), title='AAPL', label='Adj Close')  
<matplotlib.axes._subplots.AxesSubplot at 0x1150bac88>
```

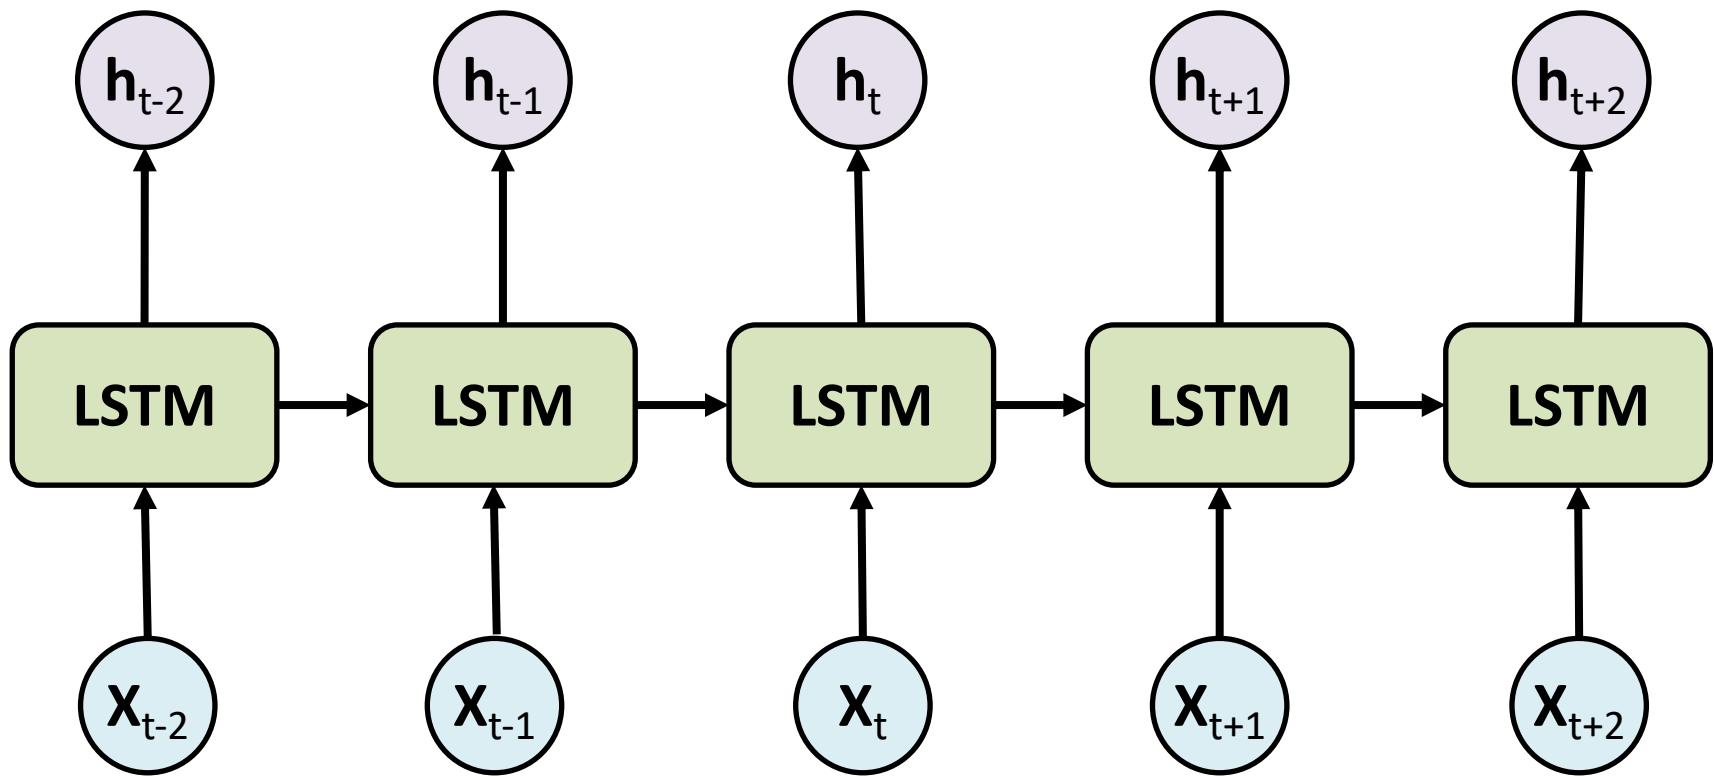


Time Series Data

[100, 110, 120, 130, 140, 150]



Long Short Term Memory (LSTM) for Time Series Forecasting



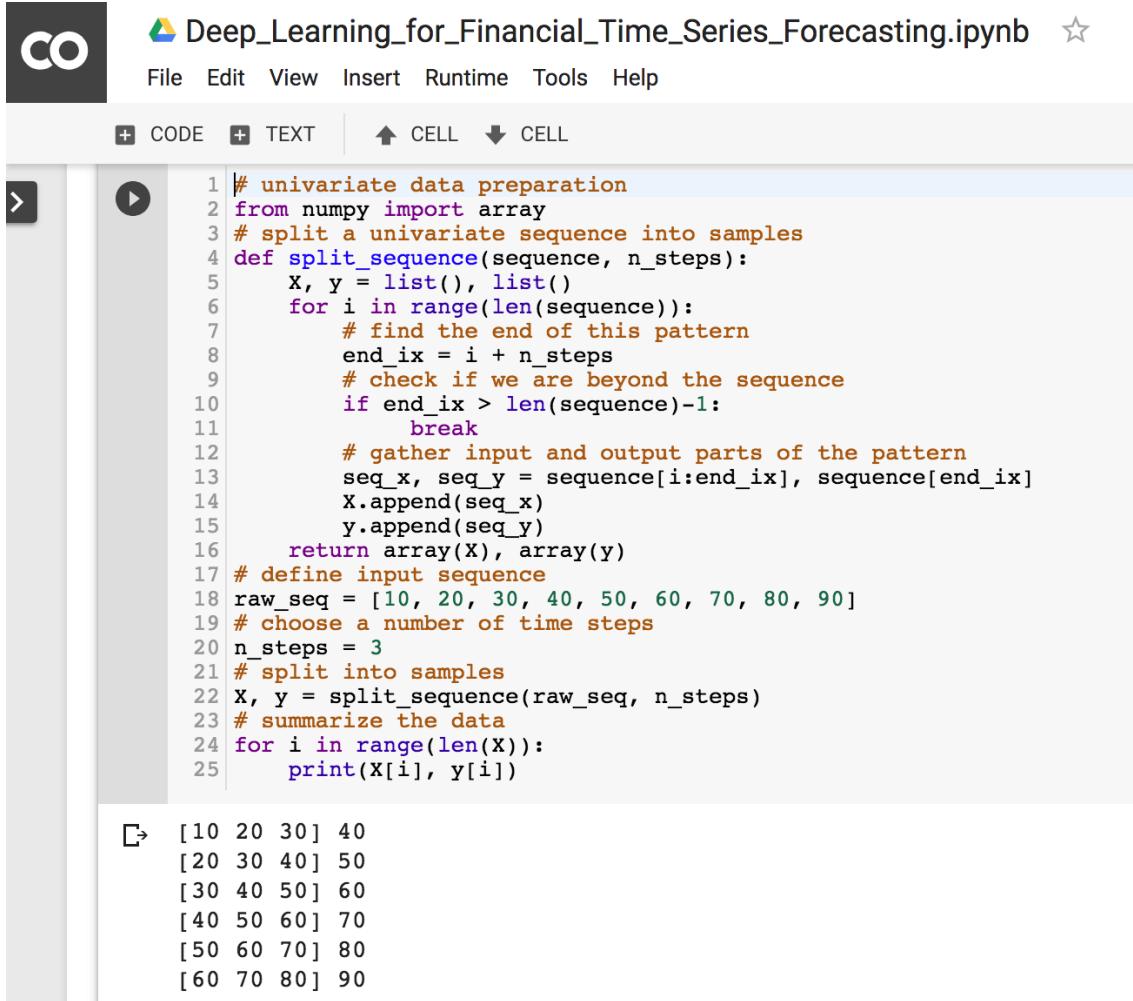
Time Series Data

[10, 20, 30, 40, 50, 60, 70, 80, 90]

X	Y
[10 20 30]	40
[20 30 40]	50
[30 40 50]	60
[40 50 60]	70
[50 60 70]	80
[60 70 80]	90

Deep Learning for Financial Time Series Forecasting

<https://colab.research.google.com/drive/1aEK0eSev8Q-Y0nNY32geFk7CB8pVgSQM>



The screenshot shows a Google Colab notebook titled "Deep_Learning_for_Financial_Time_Series_Forecasting.ipynb". The code cell contains Python code for preparing univariate time series data. It defines a function to split a sequence into samples and then applies it to a raw sequence of values [10, 20, 30, 40, 50, 60, 70, 80, 90] with a step size of 3. The output is a list of samples: [[10, 20, 30], 40], [[20, 30, 40], 50], [[30, 40, 50], 60], [[40, 50, 60], 70], [[50, 60, 70], 80], and [[60, 70, 80], 90].

```
# univariate data preparation
from numpy import array
# split a univariate sequence into samples
def split_sequence(sequence, n_steps):
    X, y = list(), list()
    for i in range(len(sequence)):
        # find the end of this pattern
        end_ix = i + n_steps
        # check if we are beyond the sequence
        if end_ix > len(sequence)-1:
            break
        # gather input and output parts of the pattern
        seq_x, seq_y = sequence[i:end_ix], sequence[end_ix]
        X.append(seq_x)
        y.append(seq_y)
    return array(X), array(y)
# define input sequence
raw_seq = [10, 20, 30, 40, 50, 60, 70, 80, 90]
# choose a number of time steps
n_steps = 3
# split into samples
X, y = split_sequence(raw_seq, n_steps)
# summarize the data
for i in range(len(X)):
    print(X[i], y[i])
```

[10 20 30] 40
[20 30 40] 50
[30 40 50] 60
[40 50 60] 70
[50 60 70] 80
[60 70 80] 90

Deep Learning for Financial Time Series Forecasting

<https://colab.research.google.com/drive/1aEK0eSev8Q-Y0nNY32geFk7CB8pVgSQM>

Deep_Learning_for_Financial_Time_Series_Forecasting.ipynb

File Edit View Insert Runtime Tools Help

CODE TEXT CELL CELL

CONNECTED EDITING

LSTM for Time Series Forecasting

```
# univariate lstm example
from numpy import array
from keras.models import Sequential
from keras.layers import LSTM
from keras.layers import Dense
import matplotlib.pyplot as plt
%matplotlib inline

# define dataset
X = array([[100, 110, 120], [110, 120, 130], [120, 130, 140], [130, 140, 150], [140, 150, 160]])
y = array([130, 140, 150, 160, 170])
# reshape from [samples, timesteps] into [samples, timesteps, features]
X = X.reshape((X.shape[0], X.shape[1], 1))
# define model
model = Sequential()
model.add(LSTM(50, activation='relu', input_shape=(3, 1)))
model.add(Dense(1))
model.compile(optimizer='adam', loss='mse')
# fit model
history = model.fit(X, y, epochs=2000, verbose=0)
# demonstrate prediction
x_input = array([150, 160, 170])
x_input = x_input.reshape((1, 3, 1))
yhat = model.predict(x_input, verbose=0)
print('yhat', yhat)
print(model.summary())
# list all data in history
print(history.history.keys())
# summarize history for loss
print('loss:', '%f' % history.history['loss'][-1])
print('loss:', history.history['loss'][-1])
plt.plot(history.history['loss'])
plt.title('model loss')
plt.ylabel('loss')
plt.xlabel('epoch')
plt.show()

yhat [[181.34615]]
```

Deep Learning for Financial Time Series Forecasting

<https://colab.research.google.com/drive/1aEK0eSev8Q-Y0nNY32geFk7CB8pVgSQM>

The screenshot shows a Google Colab notebook titled "Deep_Learning_for_Financial_Time_Series_Forecasting.ipynb". The code cell contains a Python script for univariate LSTM time series forecasting. The script includes imports for numpy, keras, matplotlib, and tensorflow, defines a function to split a sequence into samples, and demonstrates the use of an LSTM model to predict the next value in a sequence. The code is numbered from 1 to 44.

```
1 # univariate lstm example
2 from numpy import array
3 from keras.models import Sequential
4 from keras.layers import LSTM
5 from keras.layers import Dense
6 import matplotlib.pyplot as plt
7 %matplotlib inline
8 # split a univariate sequence into samples
9 def split_sequence(sequence, n_steps):
10     X, y = list(), list()
11     for i in range(len(sequence)):
12         # find the end of this pattern
13         end_ix = i + n_steps
14         # check if we are beyond the sequence
15         if end_ix > len(sequence)-1:
16             break
17         # gather input and output parts of the pattern
18         seq_x, seq_y = sequence[i:end_ix], sequence[end_ix]
19         X.append(seq_x)
20         y.append(seq_y)
21     return array(X), array(y)
22 # define input sequence
23 raw_seq = [10, 20, 30, 40, 50, 60, 70, 80, 90]
24 # choose a number of time steps
25 n_steps = 3
26 # split into samples
27 X, y = split_sequence(raw_seq, n_steps)
28 # reshape from [samples, timesteps] into [samples, timesteps, features]
29 n_features = 1
30 X = X.reshape((X.shape[0], X.shape[1], n_features))
31 # define model
32 model = Sequential()
33 model.add(LSTM(50, activation='relu', input_shape=(n_steps, n_features)))
34 model.add(Dense(1))
35 model.compile(optimizer='adam', loss='mse')
36 # fit model
37 history = model.fit(X, y, epochs=500, verbose=0)
38 # demonstrate prediction
39 x_input = array([70, 80, 90])
40 x_input = x_input.reshape((1, n_steps, n_features))
41 yhat = model.predict(x_input, verbose=0)
42 print(yhat)
43 print('yhat', yhat)
44 print(model.summary())
```

Deep Learning for Financial Time Series Forecasting

<https://colab.research.google.com/drive/1aEK0eSev8Q-Y0nNY32geFk7CB8pVgSQM>

```
Using TensorFlow backend.
```

```
[[102.31296]]
```

```
yhat [[102.31296]]
```

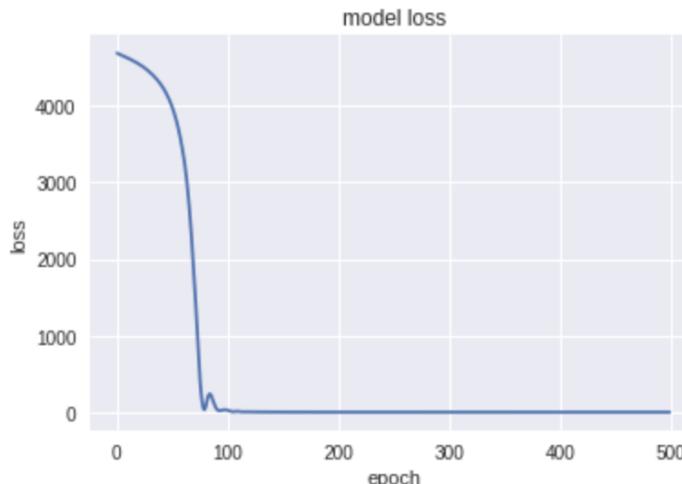
Layer (type)	Output Shape	Param #
<hr/>		
lstm_1 (LSTM)	(None, 50)	10400
dense_1 (Dense)	(None, 1)	51
<hr/>		
Total params: 10,451		
Trainable params: 10,451		
Non-trainable params: 0		

```
None
```

```
dict_keys(['loss'])
```

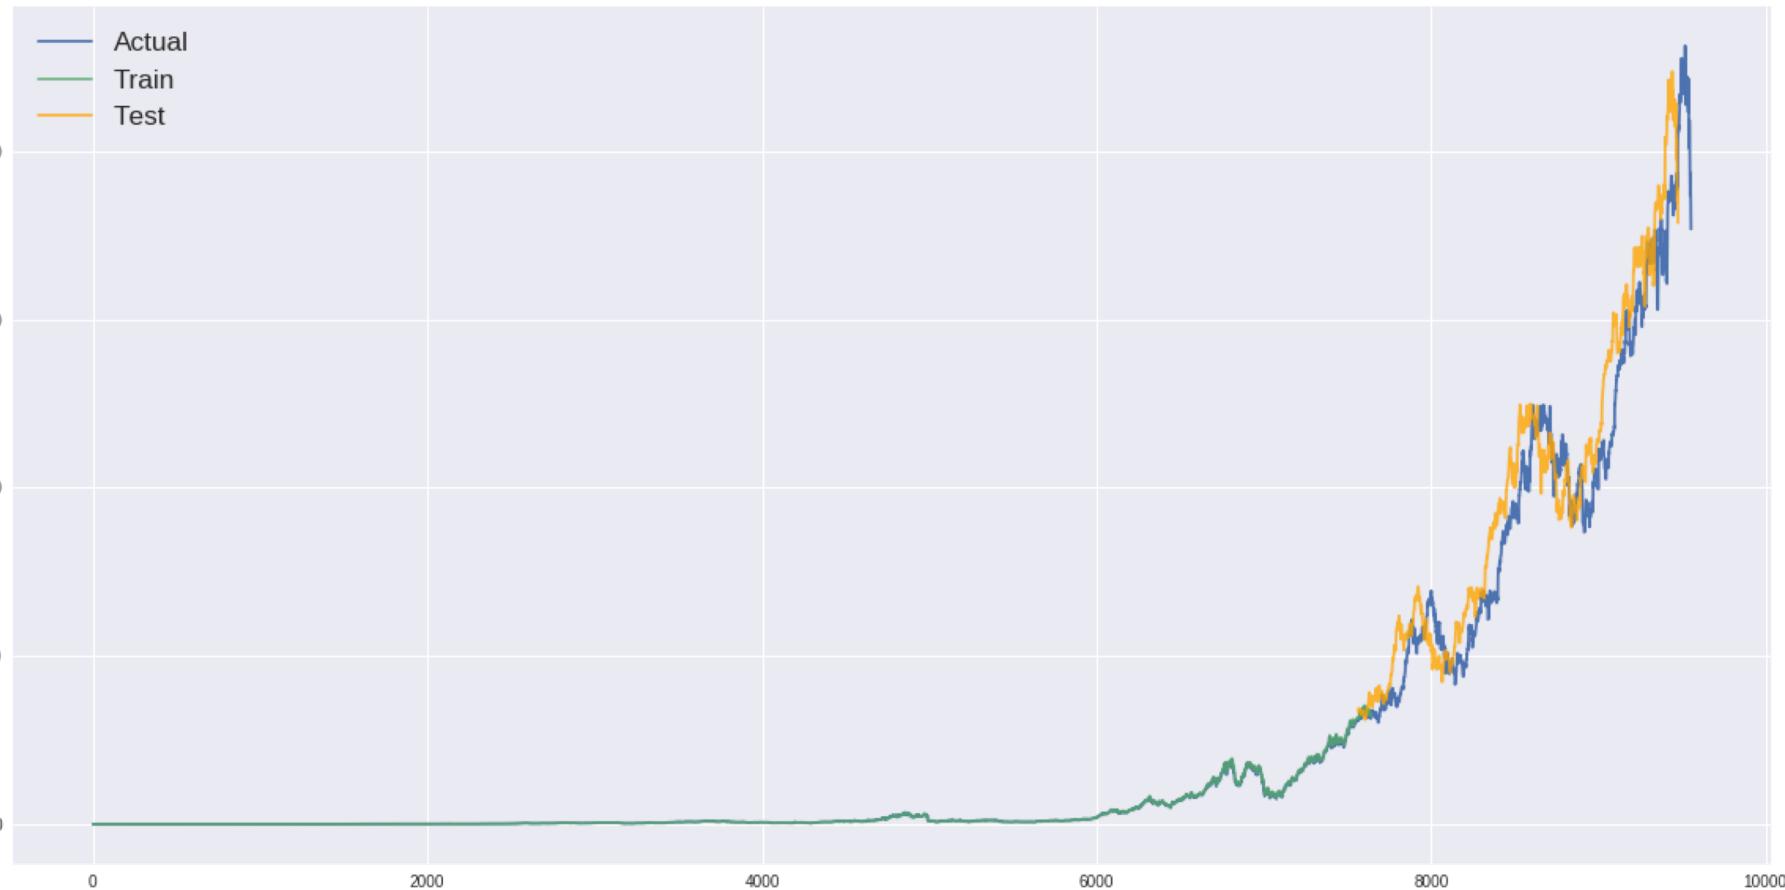
```
loss: 0.000000
```

```
loss: 1.2578432517784677e-07
```

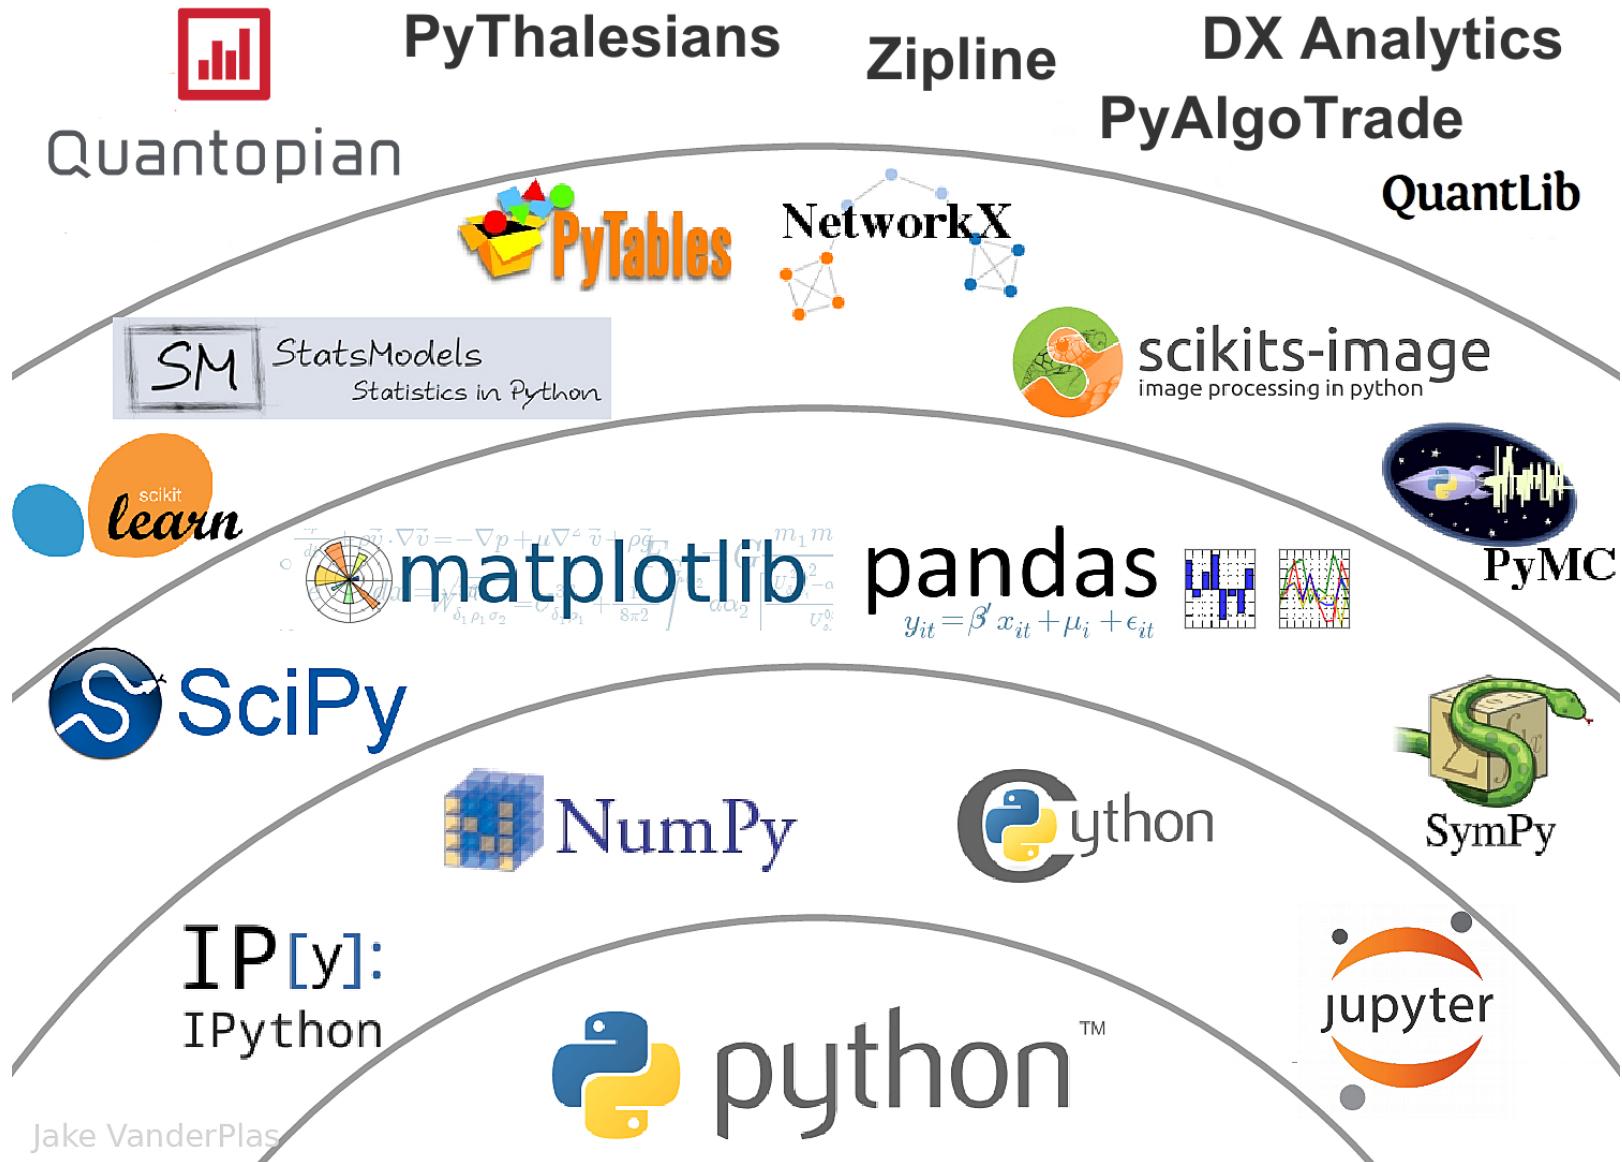


Deep Learning for Financial Time Series Forecasting

<https://colab.research.google.com/drive/1aEK0eSev8Q-Y0nNY32geFk7CB8pVgSQM>



The Quant Finance PyData Stack



AI + VDI

POC

AI + VDI POS

TensorFlow Models

- M1: Basic Classification (Image Classification) (65 Seconds)
 - https://colab.research.google.com/github/tensorflow/docs/blob/master/site/en/tutorials/keras/basic_classification.ipynb
- M2: Basic Text Classification (Text Classification) (46 Seconds)
 - https://colab.research.google.com/github/tensorflow/docs/blob/master/site/en/tutorials/keras/basic_text_classification.ipynb
- M3: Basic Regression (Predict House Prices) (43 Seconds)
 - https://colab.research.google.com/github/tensorflow/docs/blob/master/site/en/tutorials/keras/basic_regression.ipynb
- M4: Pix2Pix Eager (Option) (7-8 Hours)
 - https://colab.research.google.com/github/tensorflow/tensorflow/blob/master/tensorflow/contrib/eager/python/examples/pix2pix/pix2pix_eager.ipynb
- M5. NMT with Attention (Option) (20-30 minutes)
 - https://colab.research.google.com/github/tensorflow/tensorflow/blob/master/tensorflow/contrib/eager/python/examples/nmt_with_attention/nmt_with_attention.ipynb

Basic Classification

Fashion MNIST Image Classification

<https://colab.research.google.com/drive/19PJOJi1vn1kjcutlzNHjRSLbeVI4kd5z>

tf01_basic_classification.ipynb ★

File Edit View Insert Runtime Tools Help

COMMENT SHARE A

+ CODE + TEXT ↑ CELL ↓ CELL CONNECT EDITING

Table of contents Code snippets Files X

Copyright 2018 The TensorFlow Authors.

Licensed under the Apache License, Version 2.0 (the "License");

MIT License

Train your first neural network: basic classification

Import the Fashion MNIST dataset

Explore the data

Preprocess the data

Build the model

Setup the layers

Compile the model

Train the model

Evaluate accuracy

Make predictions

Copyright 2018 The TensorFlow Authors.

↳ 2 cells hidden

▶ **Train your first neural network: basic classification**

 [View on TensorFlow.org](#)  [Run in Google Colab](#)  [View source on GitHub](#)

This guide trains a neural network model to classify images of clothing, like sneakers and shirts. It's okay if you don't understand all the details, this is a fast-paced overview of a complete TensorFlow program with the details explained as we go.

This guide uses [tf.keras](#), a high-level API to build and train models in TensorFlow.

```
1 # memory footprint support libraries/code
2 !ln -sf /opt/bin/nvidia-smi /usr/bin/nvidia-smi
3 !pip install gputil
4 !pip install psutil
5 !pip install humanize
6 import psutil
7 import humanize
8 import os
9 import GPUtil as GPU
10 GPUs = GPU.getGPUs()
11 gpu = GPUs[0]
12 def printm():
13     process = psutil.Process(os.getpid())
14     print("Gen RAM Free: " + humanize.naturalsize( psutil.virtual_memory().available ), " | Pro")
15     print("GPU RAM Free: {0:.0f}MB | Used: {1:.0f}MB | Util {2:.0f}% | Total {3:.0f}MB".format
16 printm()
```

SECTION

Text Classification

IMDB Movie Reviews

https://colab.research.google.com/drive/1x16h1GhHsLlrLYtPCvCHaoO1W-i_gror

co tf02_basic-text-classification.ipynb ★

File Edit View Insert Runtime Tools Help

CONNECT ▾ EDITING ▾

Table of contents Code snippets Files X

Copyright 2018 The TensorFlow Authors.

Licensed under the Apache License, Version 2.0 (the "License");

MIT License

Text classification with movie reviews

Download the IMDB dataset

Explore the data

Convert the integers back to words

Prepare the data

Build the model

Hidden units

Loss function and optimizer

Create a validation set

Train the model

Evaluate the model

▶ Copyright 2018 The TensorFlow Authors.
↳ 2 cells hidden

Text classification with movie reviews

 [View on TensorFlow.org](#)  [Run in Google Colab](#)  [View source on GitHub](#)

This notebook classifies movie reviews as *positive* or *negative* using the text of the review. This is an example of *binary*—or two-class—classification, an important and widely applicable kind of machine learning problem.

We'll use the [IMDB dataset](#) that contains the text of 50,000 movie reviews from the [Internet Movie Database](#). These are split into 25,000 reviews for training and 25,000 reviews for testing. The training and testing sets are *balanced*, meaning they contain an equal number of positive and negative reviews.

This notebook uses `tf.keras`, a high-level API to build and train models in TensorFlow. For a more advanced text classification tutorial using `tf.keras`, see the [MLCC Text Classification Guide](#).

```
# memory footprint support libraries/code
!ln -sf /opt/bin/nvidia-smi /usr/bin/nvidia-smi
!pip install gputil
!pip install psutil
!pip install humanize
import psutil
import humanize
import os
import GPUUtil as GPU
GPUs = GPU.getGPUs()
gpu = GPUs[0]
def printm():
    processes = psutil.Process(os.getpid())
    print("GPU Utilization: " + str(gpu.utilization) + "%")
```

Basic Regression

Predict House Prices

https://colab.research.google.com/drive/1v4c8ZHTnRtgd2_25K_AURjR6SCVBRdlj

CO tf03_basic-regression.ipynb 

File Edit View Insert Runtime Tools Help

CONNECT EDITING

Table of contents Code snippets Files X

Copyright 2018 The TensorFlow Authors.

Predict house prices: regression

The Boston Housing Prices dataset

Examples and features

Labels

Normalize features

Create the model

Train the model

Predict

Conclusion

+ SECTION

▶ Copyright 2018 The TensorFlow Authors.
↳ 2 cells hidden

View on TensorFlow.org Run in Google Colab View source on GitHub

In a *regression* problem, we aim to predict the output of a continuous value, like a price or a probability. Contrast this with a *classification* problem, where we aim to predict a discrete label (for example, where a picture contains an apple or an orange).

This notebook builds a model to predict the median price of homes in a Boston suburb during the mid-1970s. To do this, we'll provide the model with some data points about the suburb, such as the crime rate and the local property tax rate.

This example uses the `tf.keras` API, see [this guide](#) for details.

```
# memory footprint support libraries/code
!ln -sf /opt/bin/nvidia-smi /usr/bin/nvidia-smi
!pip install gputil
!pip install psutil
!pip install humanize
import psutil
import humanize
import os
import GPUUtil as GPU
GPUs = GPU.getGPUs()
gpu = GPUs[0]
def printm():
    process = psutil.Process(os.getpid())
    print("Gen RAM Free: " + humanize.naturalsize( psutil.virtual_memory().available ), " | Proc size: "
    print("GPU RAM Free: {0:.0f}MB | Used: {1:.0f}MB | Util {2:.0f}% | Total {3:.0f}MB".format(gpu.memo
```

AI+VDI POC

ISAC+TKU Test

- AI+VDI POC Folder (3+1 ipynb) (v3.0.20181120)
 - <https://drive.google.com/open?id=1qHOemktbEmUz-ot8eFxIKbGwJvXlrjtc>
- run3models.ipynb
 - https://colab.research.google.com/drive/1HQ1GrIqQUUPCct7_AVgoMwMrh0UqMm0f

Summary

- Deep Learning for Financial Time Series Forecasting with TensorFlow
 - Deep Learning
 - Financial Time Series Forecasting
 - TensorFlow

References

- Wes McKinney (2017), "Python for Data Analysis: Data Wrangling with Pandas, NumPy, and IPython", 2nd Edition, O'Reilly Media.
<https://github.com/wesm/pydata-book>
- Avinash Jain (2017), Introduction To Python Programming, Udemy,
<https://www.udemy.com/pythonforbeginnersintro/>
- Yves Hilpisch (2014), Python for Finance: Analyze Big Financial Data, O'Reilly
- Michael Heydt (2015) , Mastering Pandas for Finance, Packt Publishing
- Michael Heydt (2015), Learning Pandas - Python Data Discovery and Analysis Made Easy, Packt Publishing
- James Ma Weiming (2015), Mastering Python for Finance, Packt Publishing
- Fabio Nelli (2015), Python Data Analytics: Data Analysis and Science using PANDAs, matplotlib and the Python Programming Language, Apress
- Jake VanderPlas (2016), Python Data Science Handbook: Essential Tools for Working with Data, O'Reilly Media.
- Jason Brownlee (2018), How to Develop LSTM Models for Time Series Forecasting,
<https://machinelearningmastery.com/how-to-develop-lstm-models-for-time-series-forecasting/>
- Deep Learning Basics: Neural Networks Demystified,
<https://www.youtube.com/playlist?list=PLiaHhY2iBX9hdHaRr6b7XevZtgZRa1PoU>
- Deep Learning SIMPLIFIED,
<https://www.youtube.com/playlist?list=PLjJh1vISEYgvGod9wWiydumYI8hOXixNu>
- 3Blue1Brown (2017), But what *is* a Neural Network? | Chapter 1, deep learning,
<https://www.youtube.com/watch?v=aircAruvnKk>
- 3Blue1Brown (2017), Gradient descent, how neural networks learn | Chapter 2, deep learning,
<https://www.youtube.com/watch?v=IHZwWFHWa-w>
- 3Blue1Brown (2017), What is backpropagation really doing? | Chapter 3, deep learning,
<https://www.youtube.com/watch?v=Ilg3gGewQ5U>
- TensorFlow: <https://www.tensorflow.org/>
- Keras: <http://keras.io/>
- Udacity, Deep Learning, https://www.youtube.com/playlist?list=PLAwxTw4SYaPn_OWPFT9ulXLuQrlmzHfOV
- <http://p.migdal.pl/2017/04/30/teaching-deep-learning.html>
- <https://github.com/leriomaggio/deep-learning-keras-tensorflow>