

金融科技



Tamkang  
University  
淡江大學

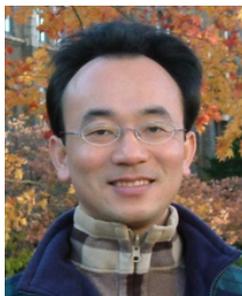
FinTech: Financial Technology

人工智慧與深度學習金融科技  
(Artificial Intelligence and  
Deep Learning for Fintech)

1052FinTech08

MIS EMBA (M2263) (8595)

Fri, 12,13,14 (19:20-22:10) (D409)



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/>

2017-05-05



# 課程大綱 (Syllabus)

週次 (Week)	日期 (Date)	內容 (Subject/Topics)
1	2017/02/17	Fintech 金融科技課程介紹 (Course Orientation for Fintech: Financial Technology)
2	2017/02/24	Fintech 金融科技的演進：貨幣與金融服務 (Evolution of Fintech: Money and Financial Services)
3	2017/03/03	Fintech 金融科技：金融服務科技創新 (Fintech: Technology Innovation in Financial Services)
4	2017/03/10	Fintech 金融科技與金融服務價值鏈 (Fintech and Financial Services Value Chain)
5	2017/03/17	Fintech 金融科技商業模式創新 (Fintech Business Models Innovation)
6	2017/03/24	Fintech 金融科技個案研究 I (Case Study on Fintech I)

# 課程大綱 (Syllabus)

週次 (Week)	日期 (Date)	內容 (Subject/Topics)
7	2017/03/31	金融服務消費者心理與行為 (Consumer Psychology and Behavior on Financial Services)
8	2017/04/07	教學行政觀摩日 (Off-campus study)
9	2017/04/14	區塊鏈技術 (Blockchain Technology) [Invited Speaker: Dr. Raymund Lin, IBM (林俊叡 博士，IBM)]
10	2017/04/21	期中報告 (Midterm Project Report)
11	2017/04/28	Python Pandas財務大數據分析 (Finance Big Data Analytics with Pandas in Python)
12	2017/05/05	人工智慧與深度學習金融科技 (Artificial Intelligence and Deep Learning for Fintech)

# 課程大綱 (Syllabus)

週次 (Week)	日期 (Date)	內容 (Subject/Topics)
13	2017/05/12	Fintech 金融科技個案研究 II (Case Study on Fintech II)
14	2017/05/19	金融科技財富管理：機器人理財顧問 (Robo-Advisors for Wealth Management in Fintech)
15	2017/05/26	投資組合最佳化與程式交易 (Portfolio Optimization and Algorithmic Trading)
16	2017/06/02	金融科技智慧問答系統 (Intelligent Question Answering System for Fintech)
17	2017/06/09	期末報告 I (Final Project Presentation I)
18	2017/06/16	期末報告 II (Final Project Presentation II)

**Artificial Intelligence  
and  
Deep Learning  
for  
Fintech**

**From Algorithmic Trading  
to Personal Finance Bots:  
41 Startups Bringing  
AI to Fintech**

# From Algorithmic Trading To Personal Finance Bots: 41 Startups Bringing AI To Fintech

## AI in Fintech

41 Startups Bringing Artificial Intelligence To Fintech

General Purpose/ Predictive Analytics



Market Research & Sentiment Analysis



Search Engine



Quantitative Trading



Blockchain



Debt Collection



AI Assistants/Bots



Fraud Detection



Credit Scoring



Personal Banking



# Artificial Intelligence (AI) in Fintech

## General Purpose/ Predictive Analytics



## Market Research & Sentiment Analysis



## Search Engine



# Artificial Intelligence (AI) in Fintech

## Quantitative Trading



## Blockchain



## Debt Collection



## AI Assistants/Bots



## Fraud Detection



## Credit Scoring



## Personal Banking

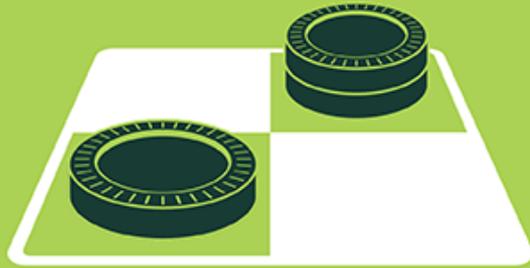


# Artificial Intelligence

## Machine Learning & Deep Learning

### ARTIFICIAL INTELLIGENCE

Early artificial intelligence stirs excitement.



### MACHINE LEARNING

Machine learning begins to flourish.



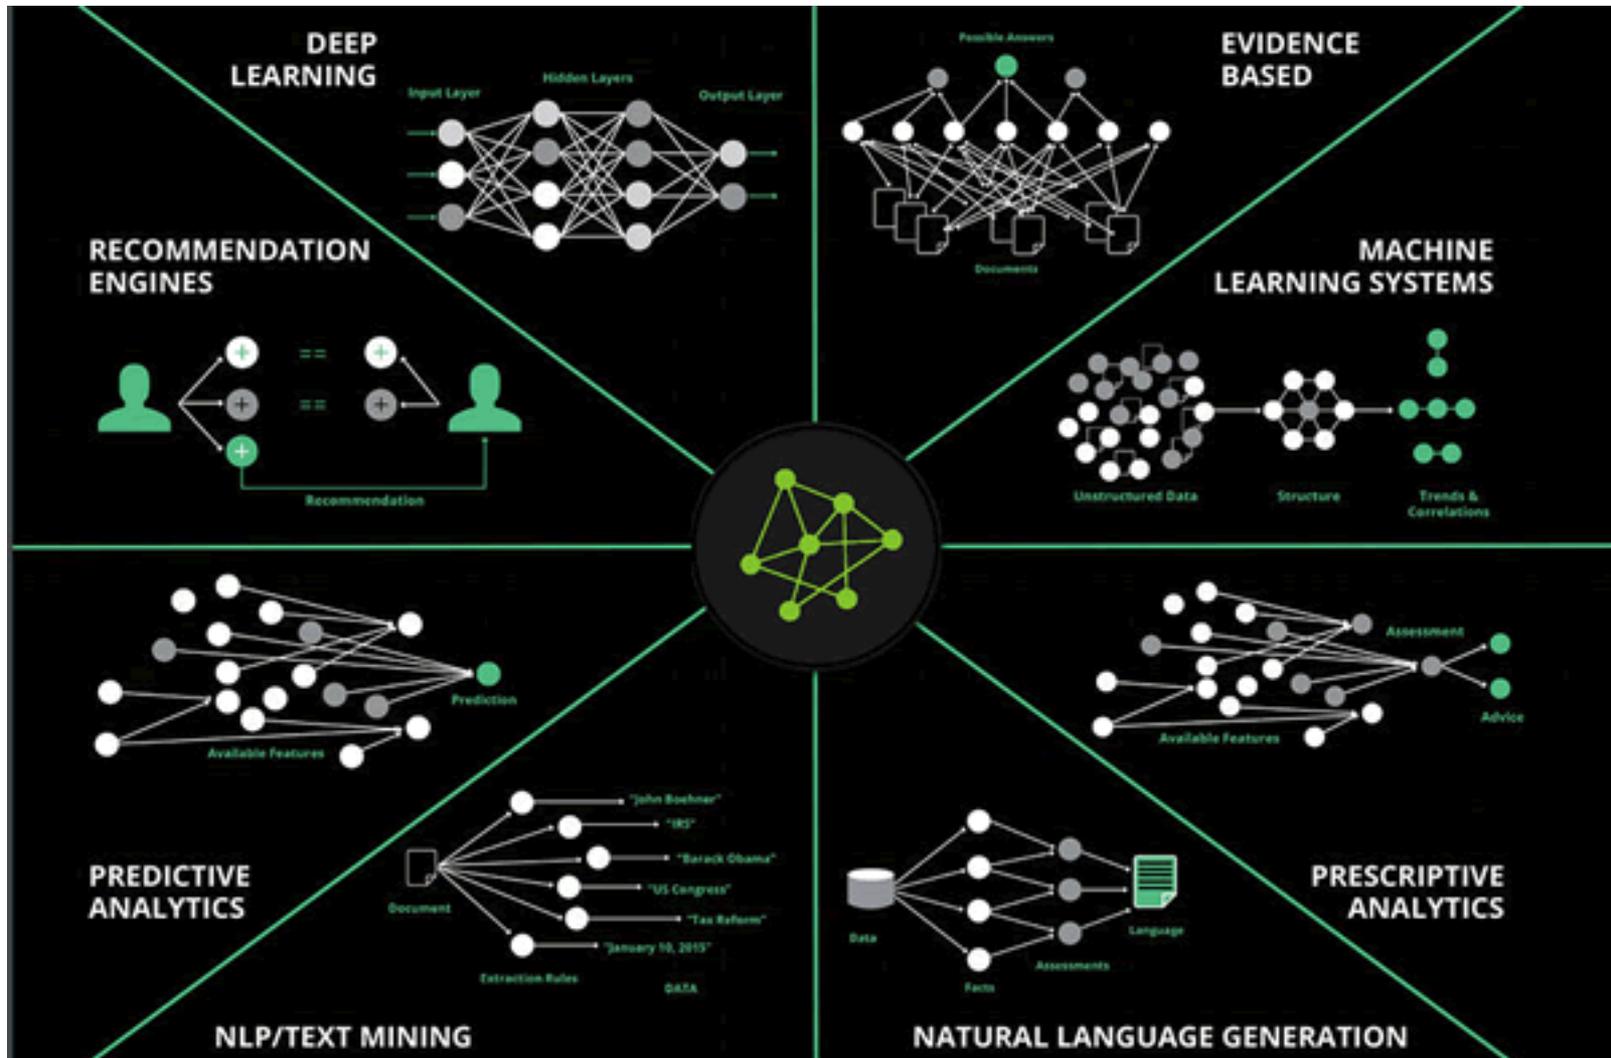
### DEEP LEARNING

Deep learning breakthroughs drive AI boom.



Since an early flush of optimism in the 1950s, smaller subsets of artificial intelligence – first machine learning, then deep learning, a subset of machine learning – have created ever larger disruptions.

# Artificial Intelligence (AI) is many things

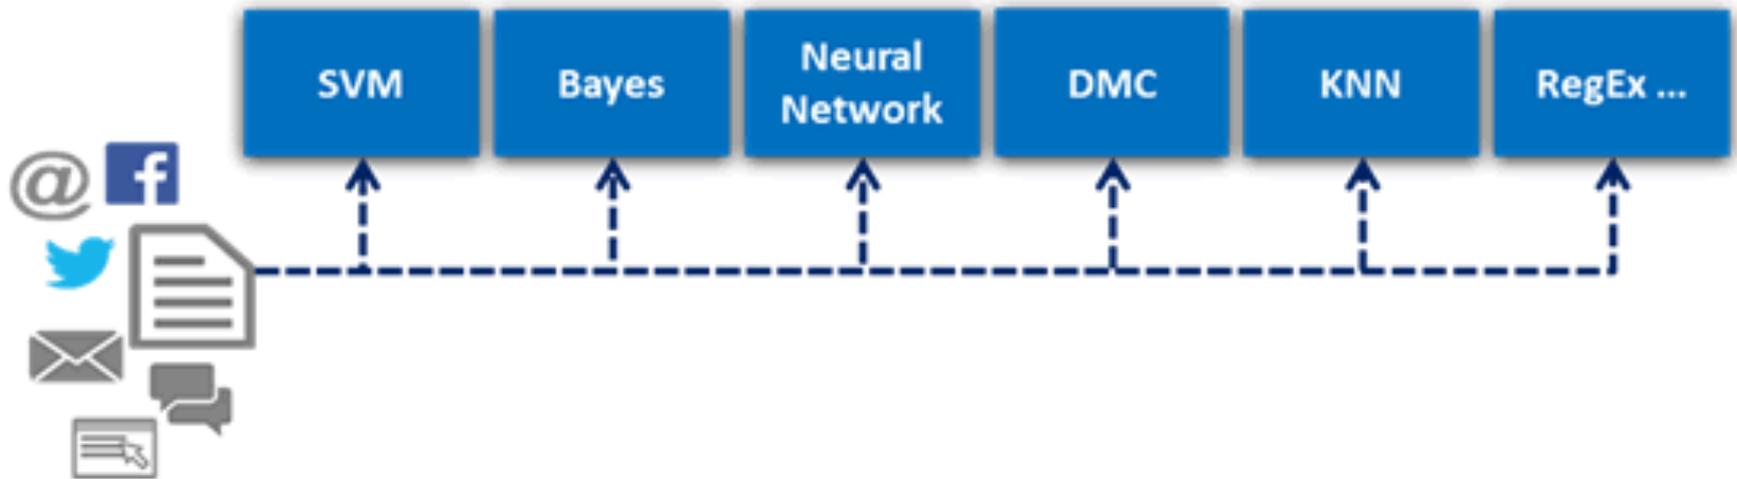


## Ecosystem of AI

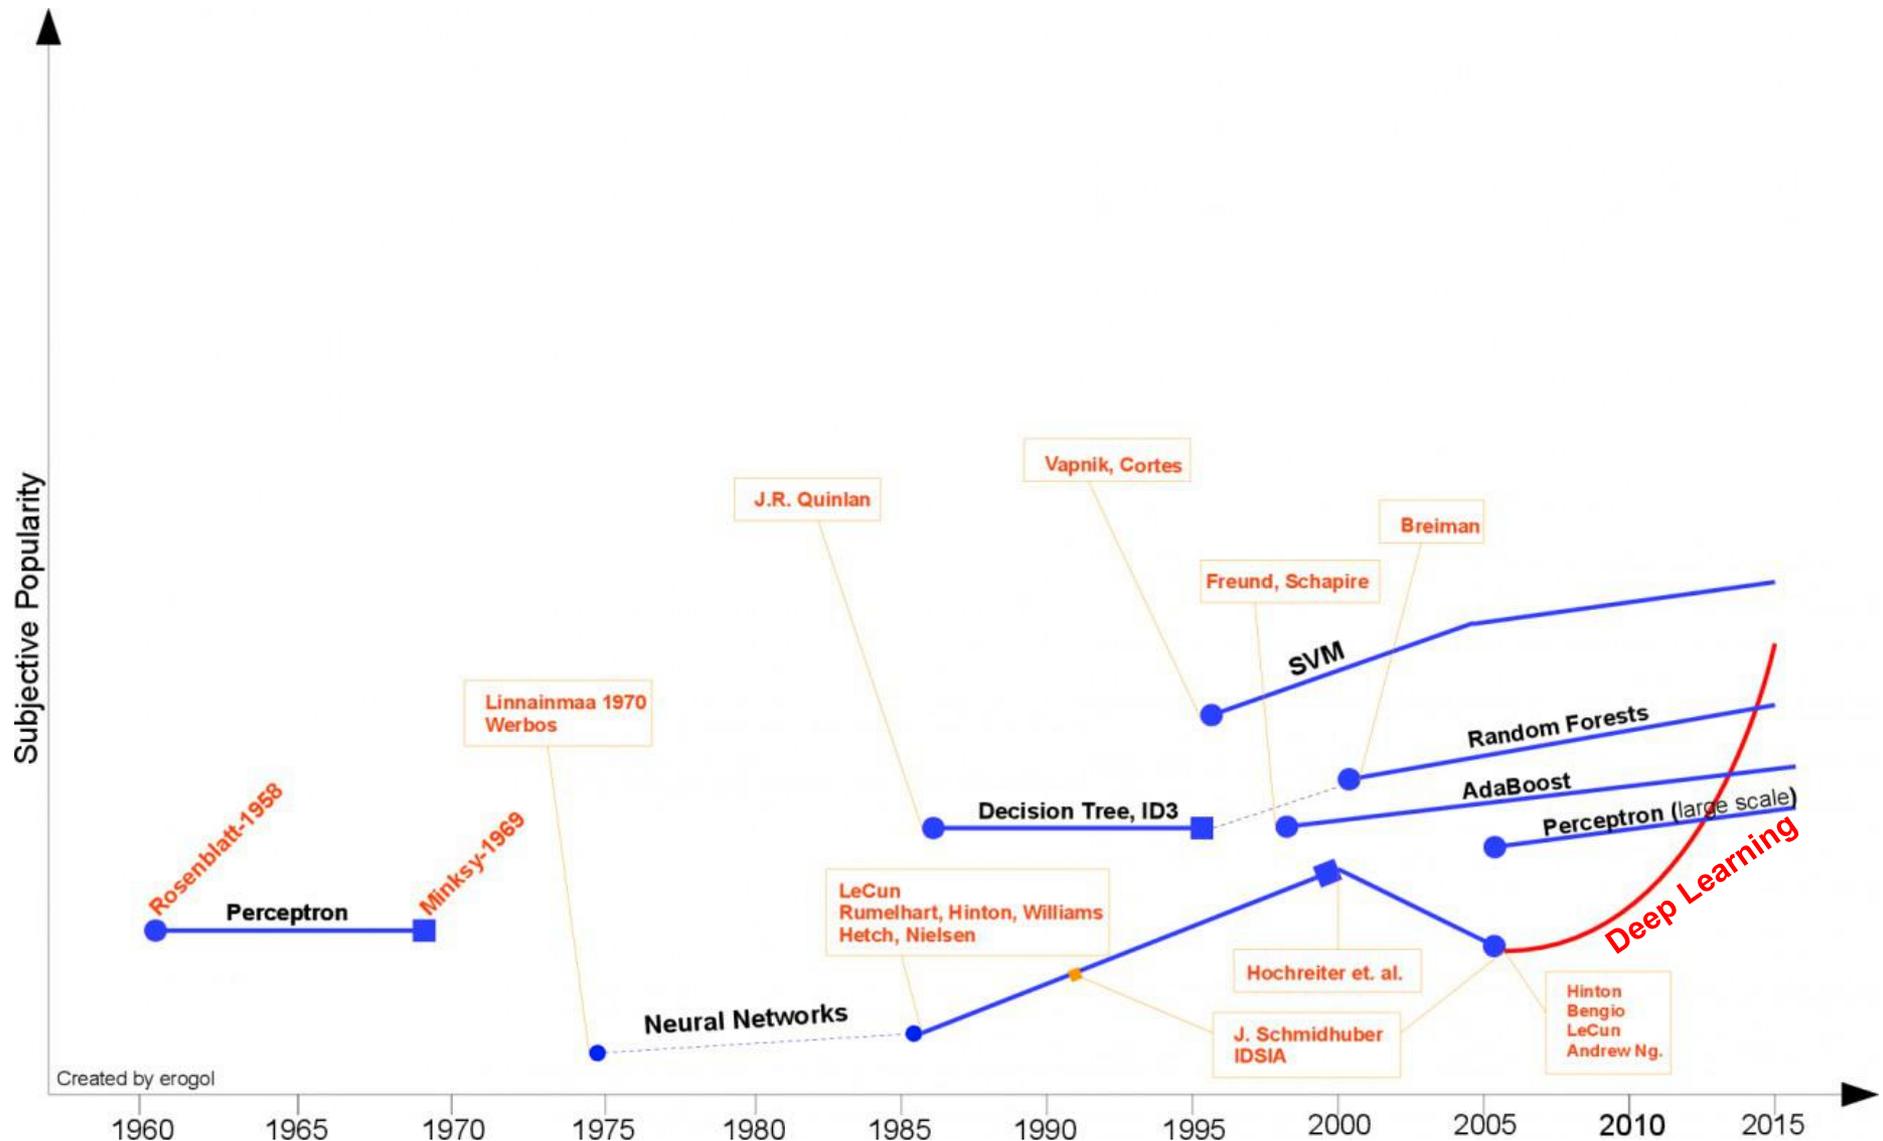
Source: <https://www.i-scoop.eu/artificial-intelligence-cognitive-computing/>

# Artificial Intelligence (AI)

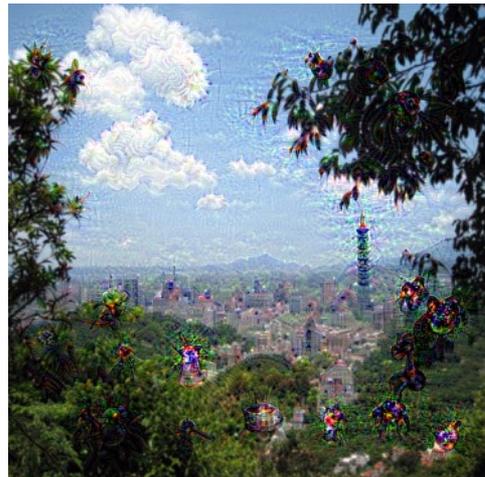
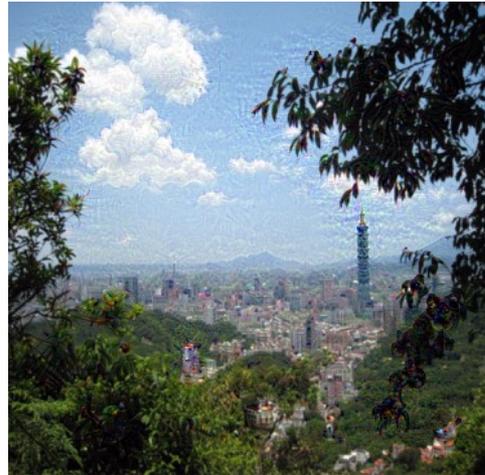
## Intelligent Document Recognition algorithms



# Deep Learning Evolution



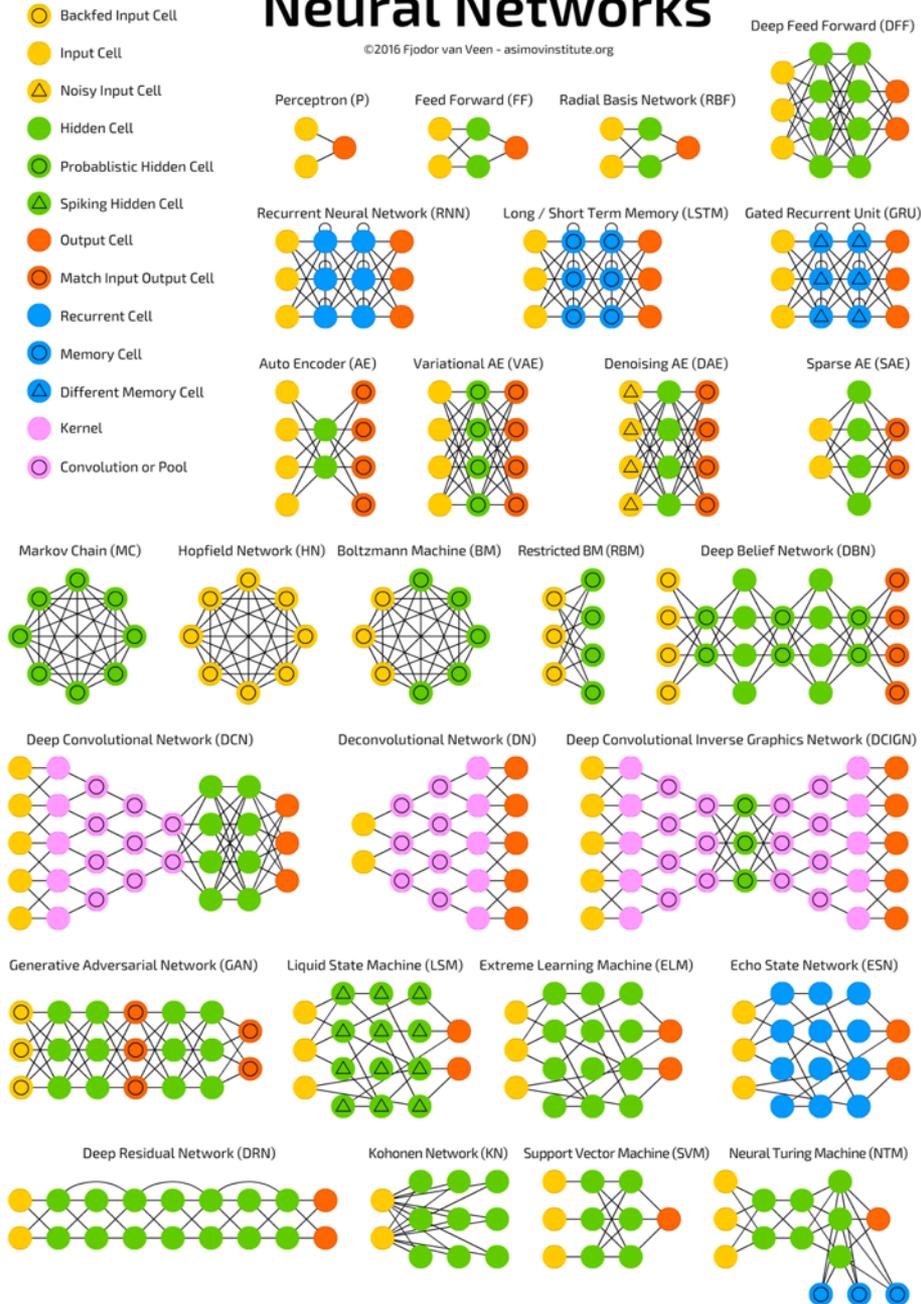
# Deep Dream



# Neural Networks (NN)

## A mostly complete chart of Neural Networks

©2016 Fjodor van Veen - asimovinstitute.org

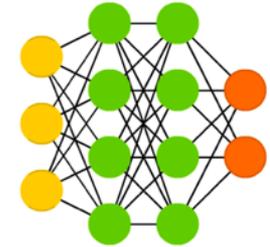


# Neural Networks

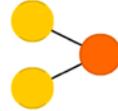
©2016 Fjodor van Veen - asimovinstitute.org

-  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

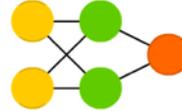
Deep Feed Forward (DFF)



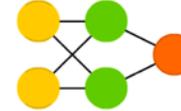
Perceptron (P)



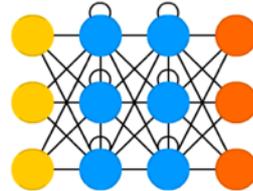
Feed Forward (FF)



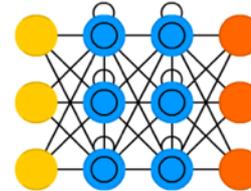
Radial Basis Network (RBF)



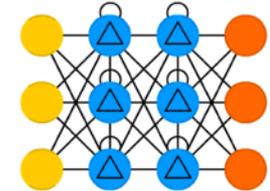
Recurrent Neural Network (RNN)



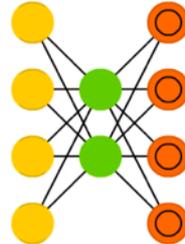
Long / Short Term Memory (LSTM)



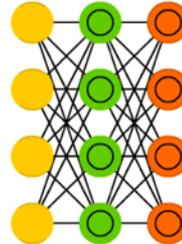
Gated Recurrent Unit (GRU)



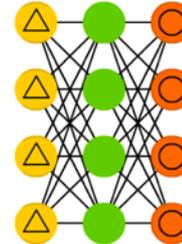
Auto Encoder (AE)



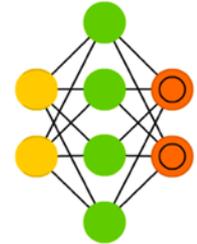
Variational AE (VAE)



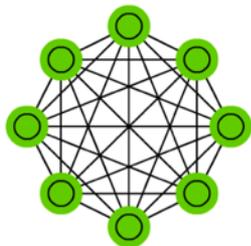
Denoising AE (DAE)



Sparse AE (SAE)



Markov Chain (MC)



Hopfield Network (HN)



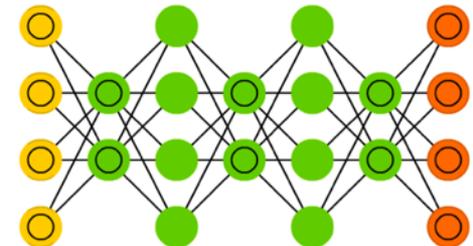
Boltzmann Machine (BM)



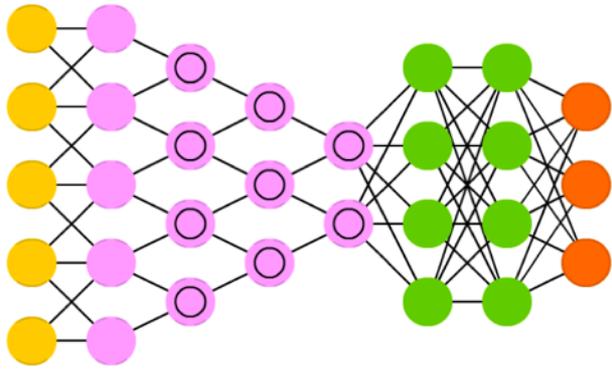
Restricted BM (RBM)



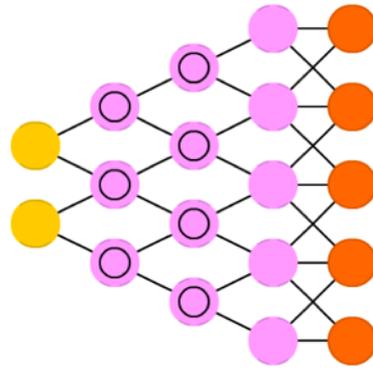
Deep Belief Network (DBN)



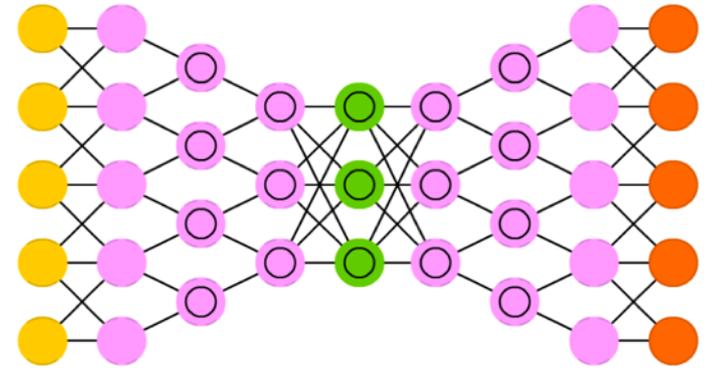
Deep Convolutional Network (DCN)



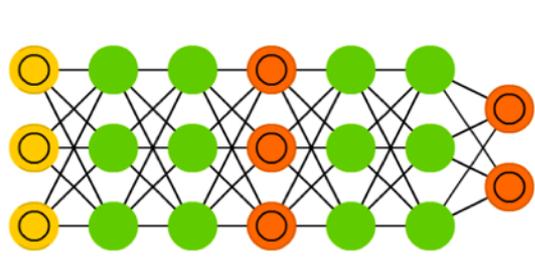
Deconvolutional Network (DN)



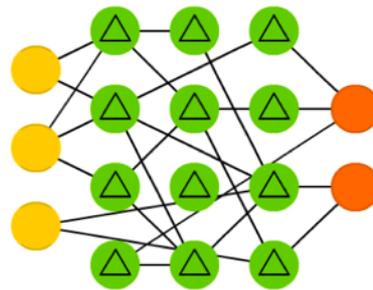
Deep Convolutional Inverse Graphics Network (DCIGN)



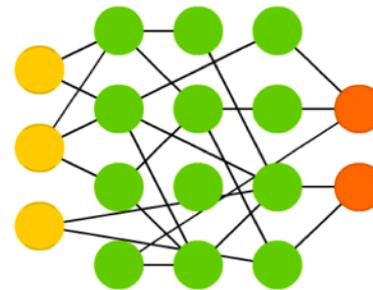
Generative Adversarial Network (GAN)



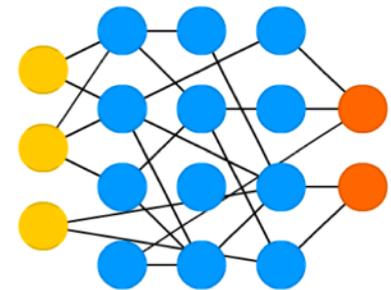
Liquid State Machine (LSM)



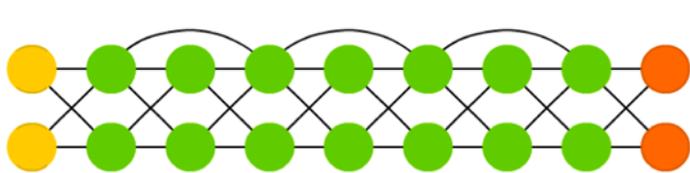
Extreme Learning Machine (ELM)



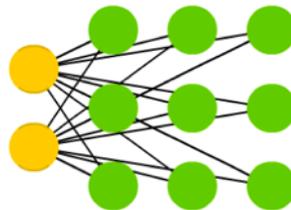
Echo State Network (ESN)



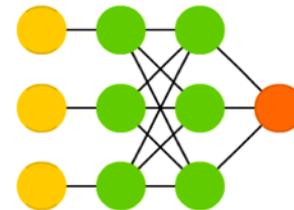
Deep Residual Network (DRN)



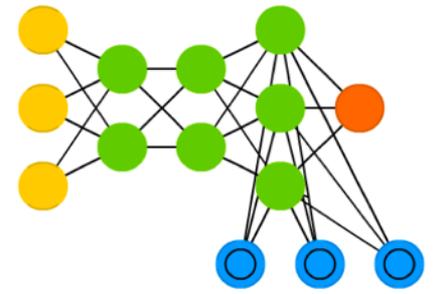
Kohonen Network (KN)



Support Vector Machine (SVM)

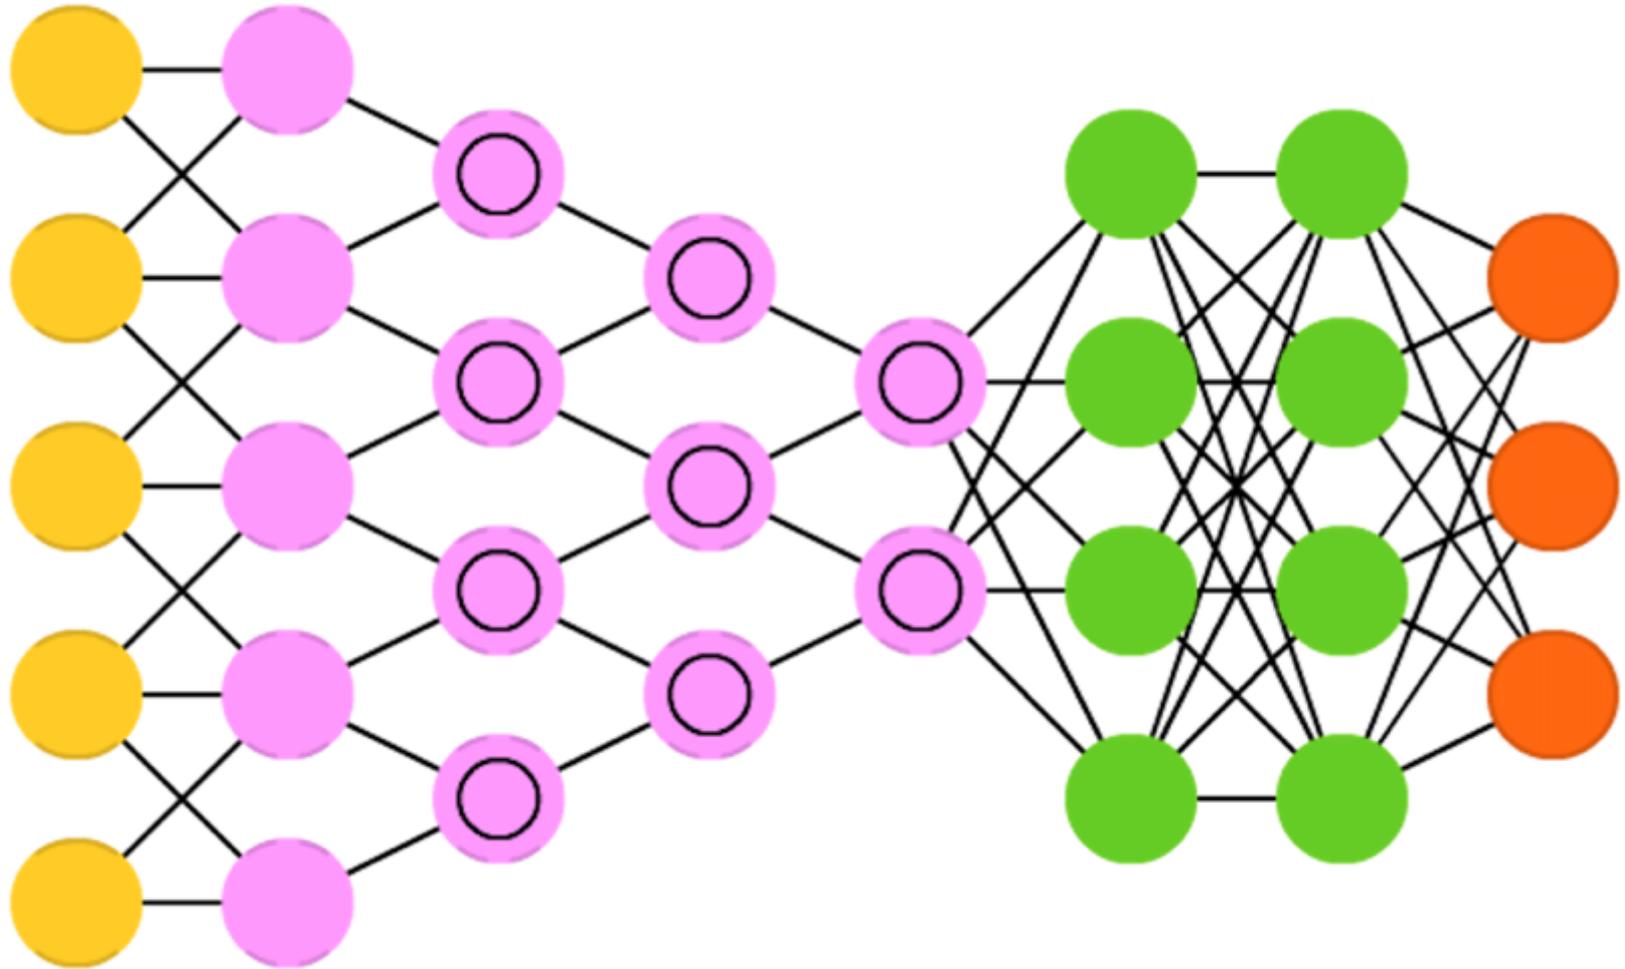


Neural Turing Machine (NTM)

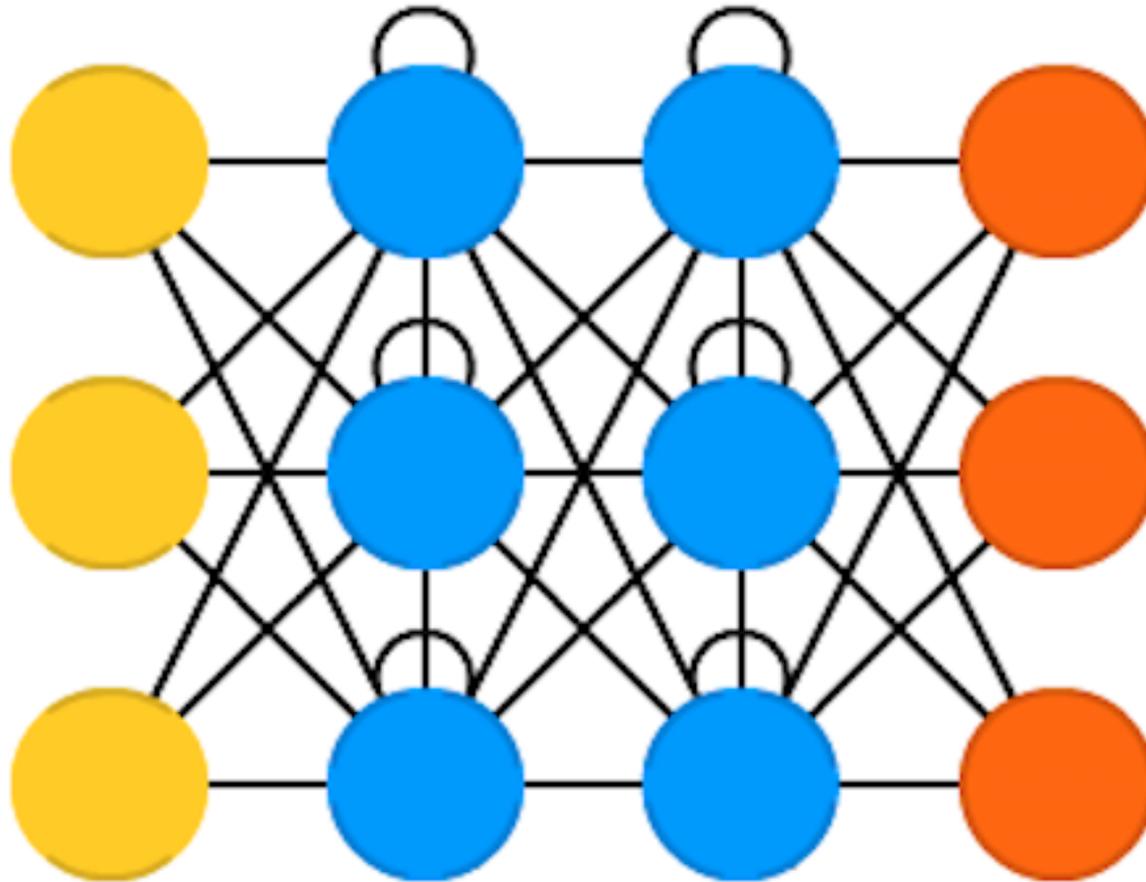


# Convolutional Neural Networks

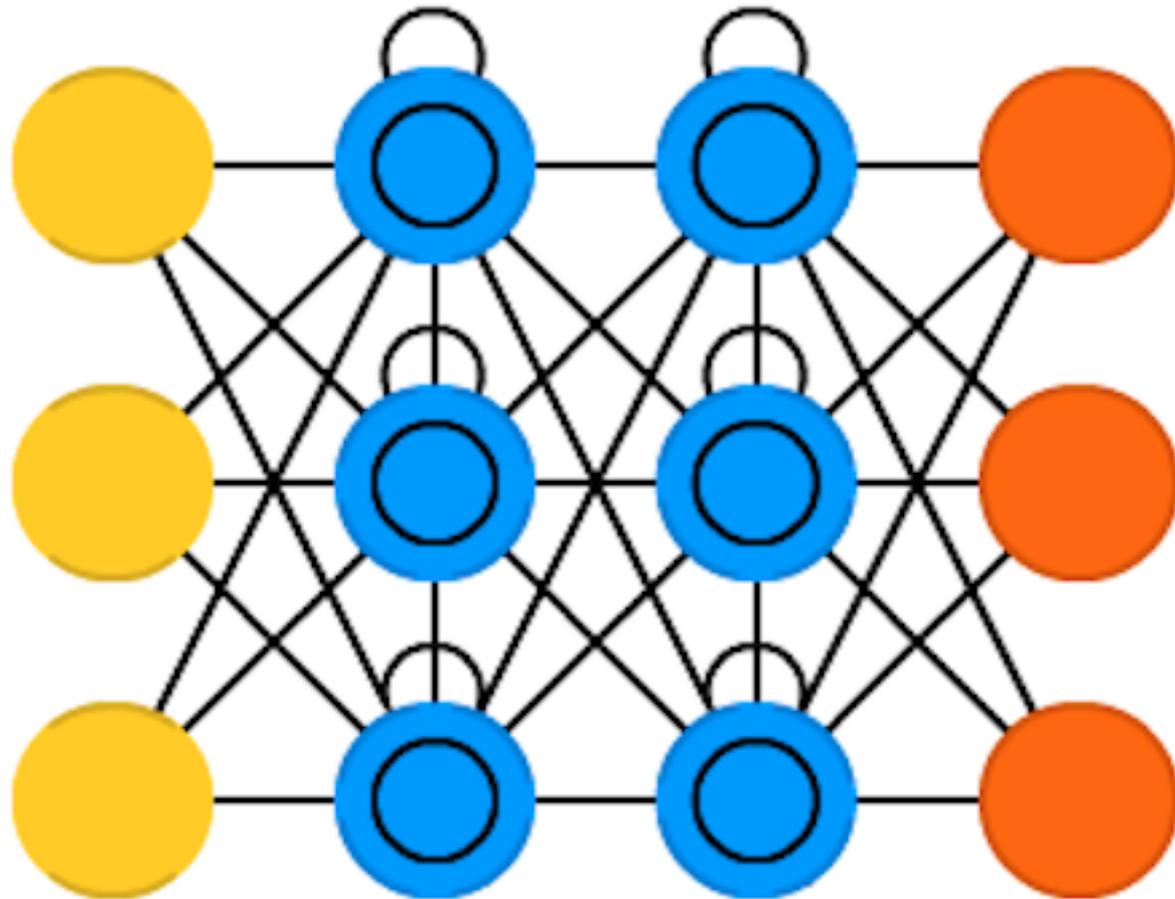
(CNN or Deep Convolutional Neural Networks, DCNN)



# Recurrent Neural Networks (RNN)



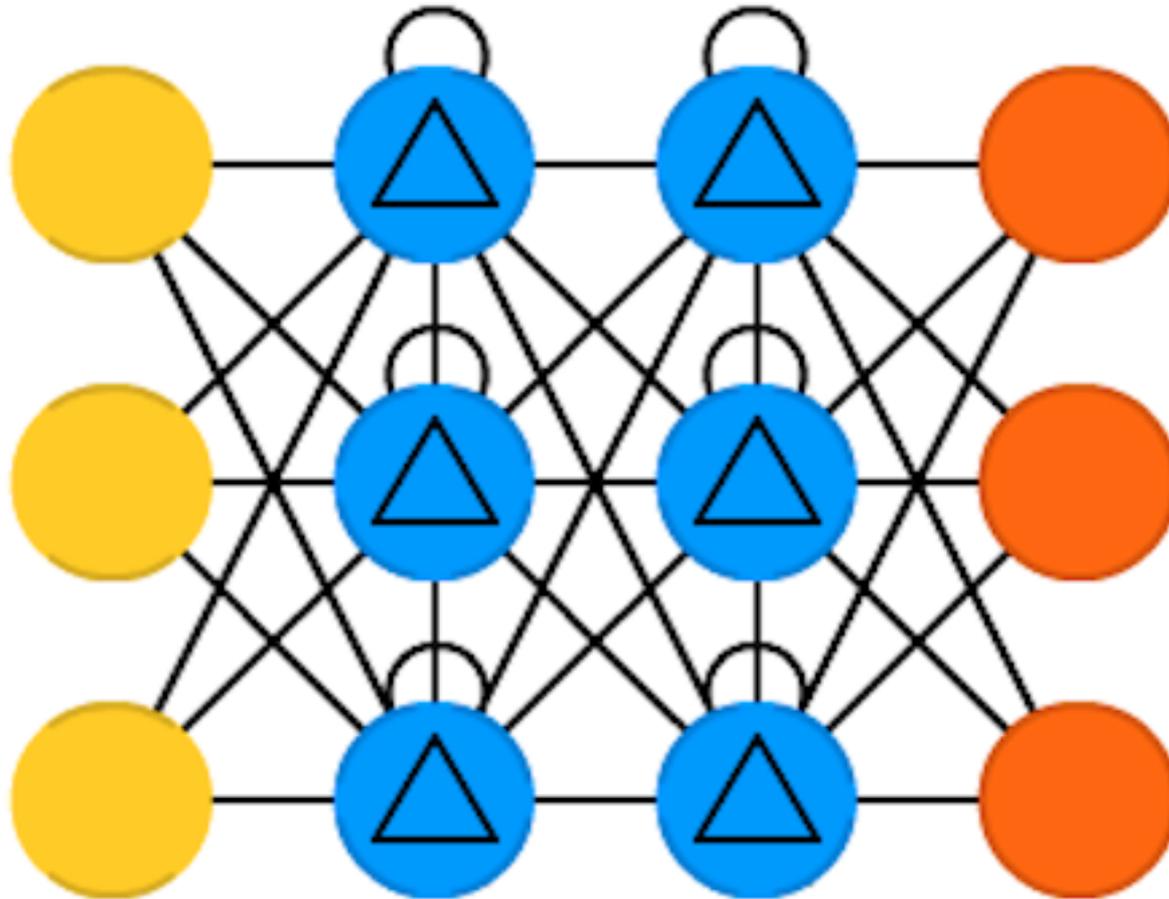
# Long / Short Term Memory (LSTM)



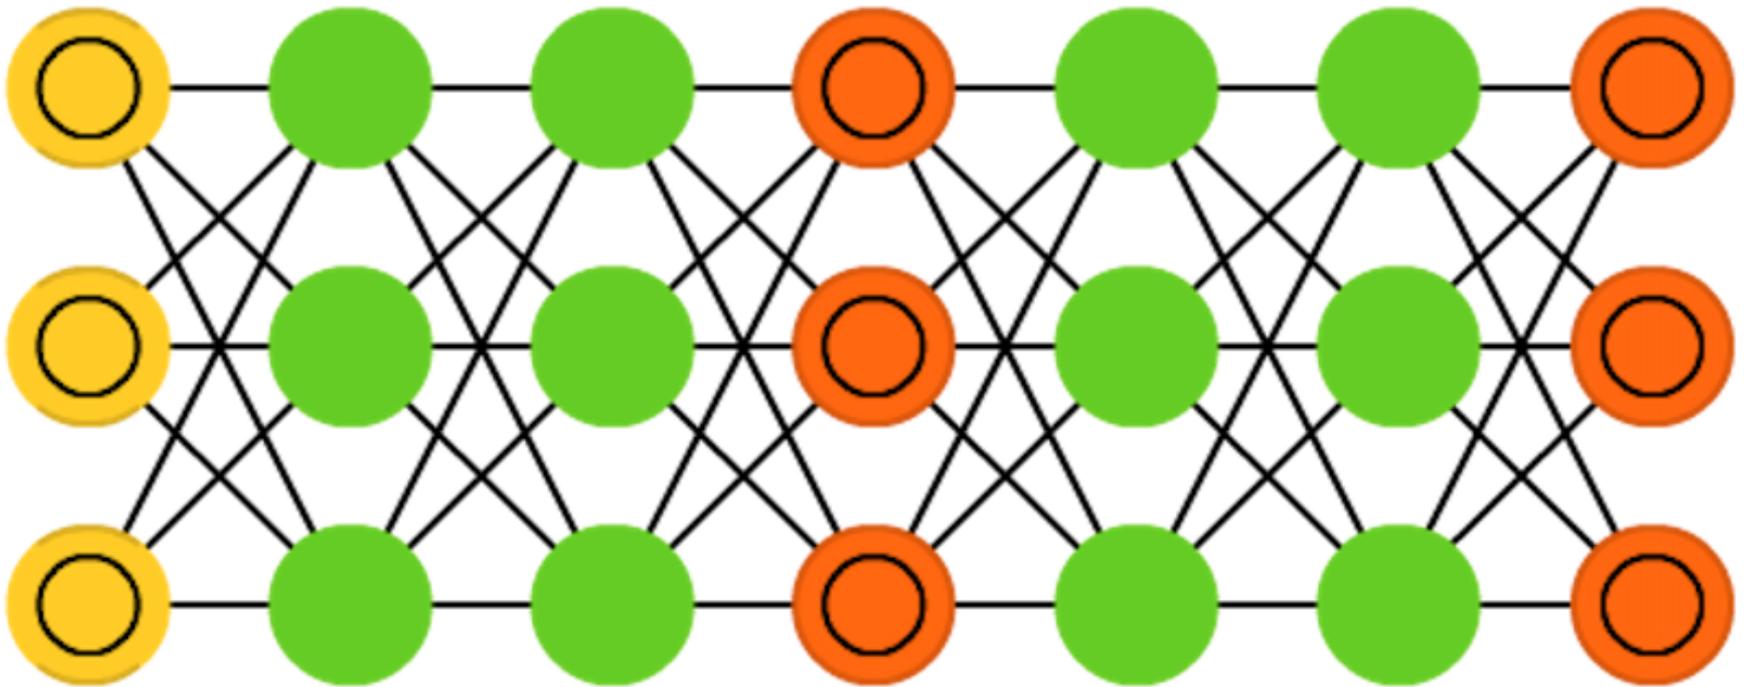
Hochreiter, Sepp, and Jürgen Schmidhuber. "Long short-term memory." *Neural computation* 9.8 (1997): 1735-1780.

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

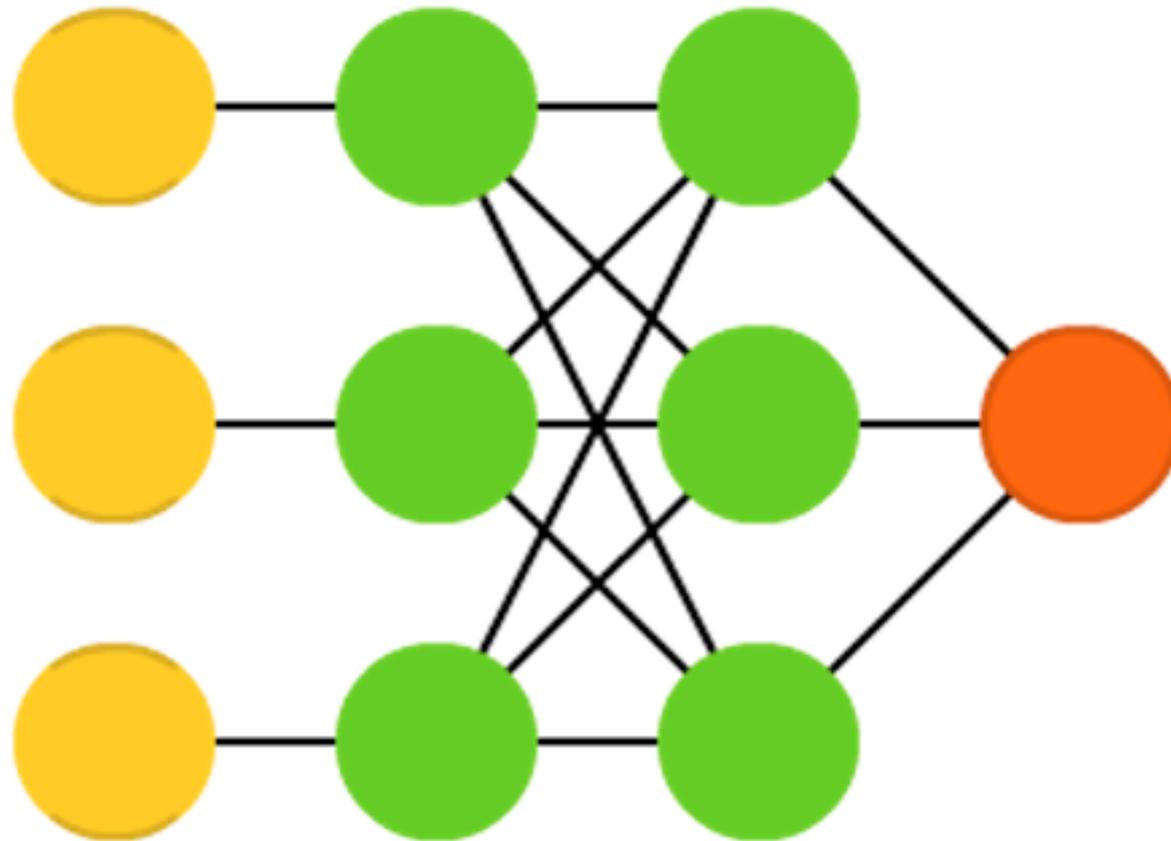
# Gated Recurrent Units (GRU)



# Generative Adversarial Networks (GAN)



# Support Vector Machines (SVM)



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

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

**LeCun, Yann,  
Yoshua Bengio,  
and Geoffrey Hinton.**

**"Deep learning."**

**Nature 521, no. 7553 (2015): 436-  
444.**

## Deep learning

Yann LeCun<sup>1,2</sup>, Yoshua Bengio<sup>3</sup> & Geoffrey Hinton<sup>4,5</sup>

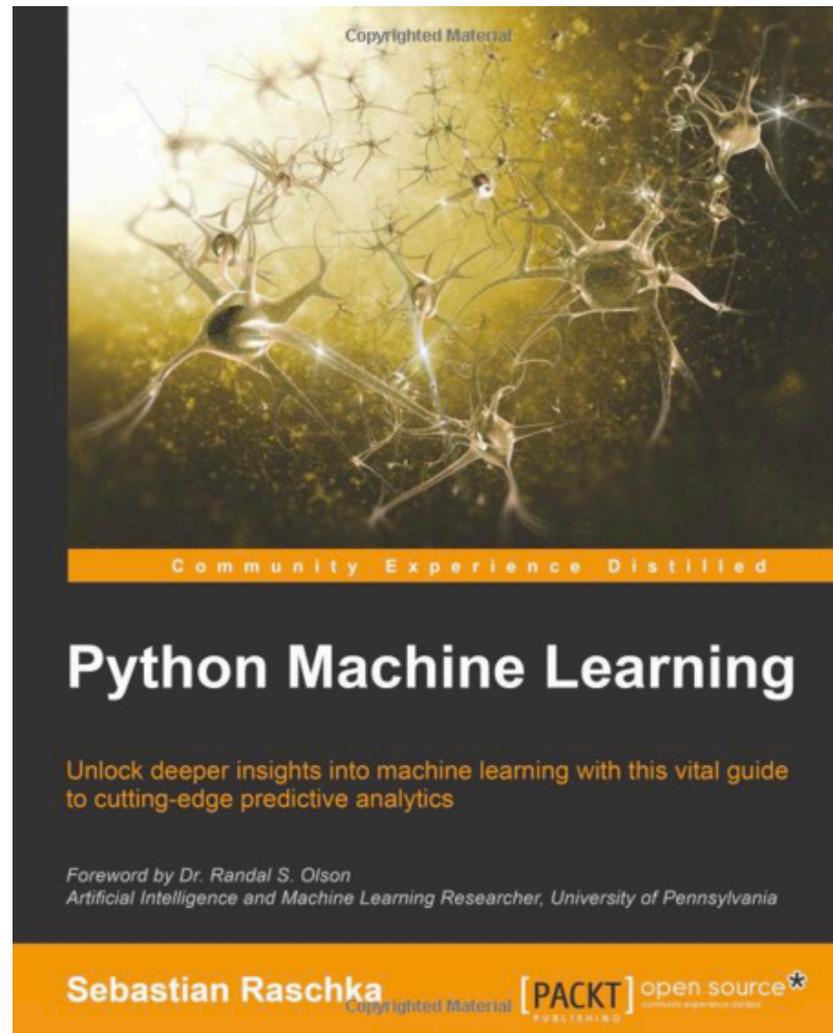
Deep learning allows computational models that are composed of multiple processing layers to learn representations of data with multiple levels of abstraction. These methods have dramatically improved the state-of-the-art in speech recognition, visual object recognition, object detection and many other 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.

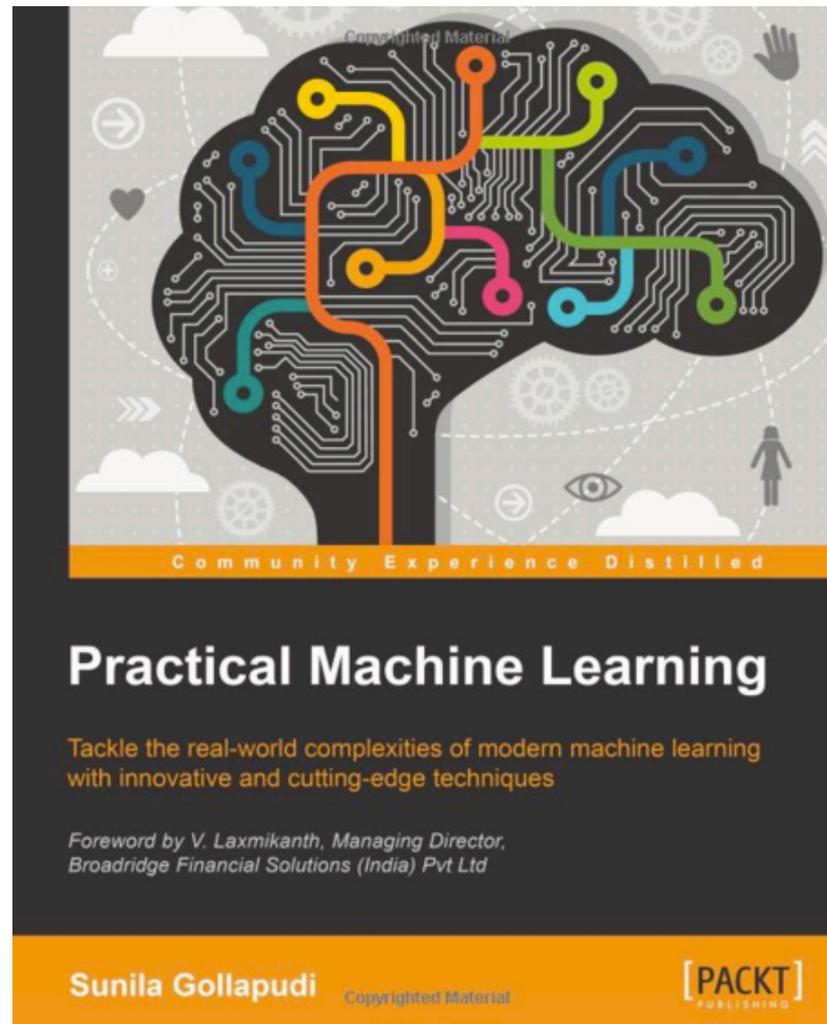
Conventional machine-learning techniques were limited in their ability to process natural data in their raw form. For decades, con-

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<sup>1-4</sup> and speech recognition<sup>5-7</sup>, it has beaten other machine-learning techniques at predicting the activity of potential drug molecules<sup>8</sup>, analysing particle accelerator data<sup>9,10</sup>, reconstructing brain circuits<sup>11</sup>, and predicting the effects of mutations in non-coding DNA on gene expression and disease<sup>12,13</sup>. Perhaps more surprisingly, deep learning has produced extremely promising results for various tasks in natural language understanding<sup>14</sup>, particularly topic classification, sentiment analysis, question answering<sup>15</sup> and language translation<sup>16,17</sup>.

Sebastian Raschka (2015),  
**Python Machine Learning,**  
Packt Publishing



Sunila Gollapudi (2016),  
**Practical Machine Learning,**  
Packt Publishing



# Machine Learning Models

Deep Learning

Association rules

Decision tree

Clustering

Bayesian

Kernel

Ensemble

Dimensionality reduction

Regression Analysis

Instance based

# Neural networks (NN) 1960

# Multilayer Perceptrons (MLP) 1985

# Restricted Boltzmann Machine (RBM) 1986

# 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

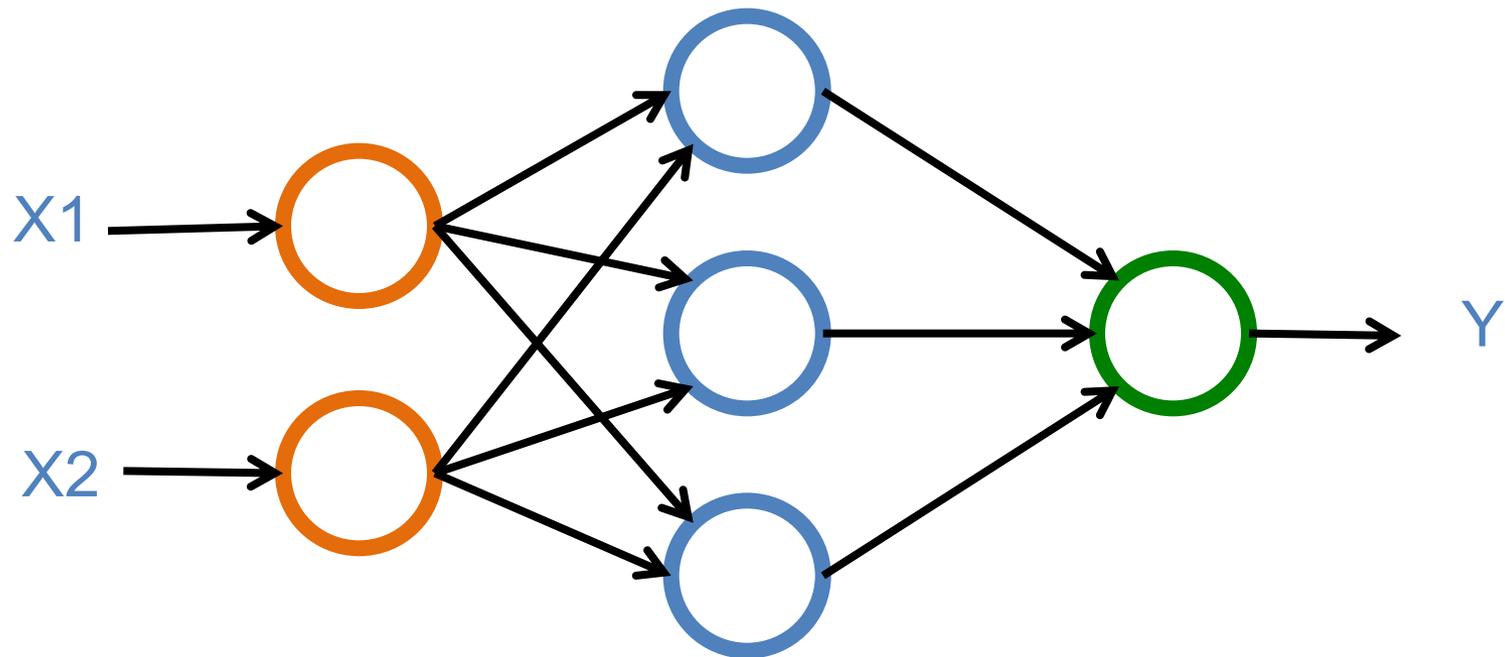
# Max-Pooling CDBN 2011

# Neural Networks

**Input Layer**  
(X)

**Hidden Layer**  
(H)

**Output Layer**  
(Y)



# Deep Learning

**Geoffrey Hinton**

**Yann LeCun**

**Yoshua Bengio**

**Andrew Y. Ng**



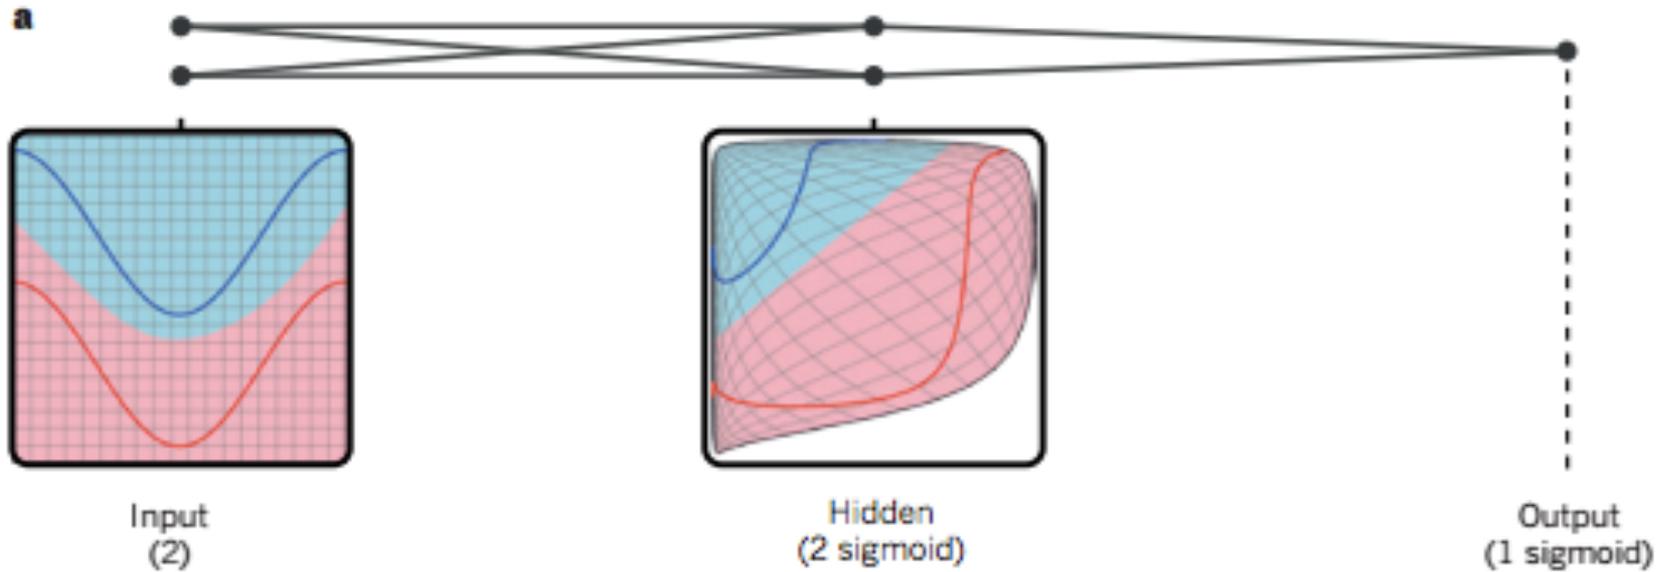
**Geoffrey Hinton**  
**Google**  
**University of Toronto**

**LeCun, Yann,  
Yoshua Bengio,  
and Geoffrey Hinton.**

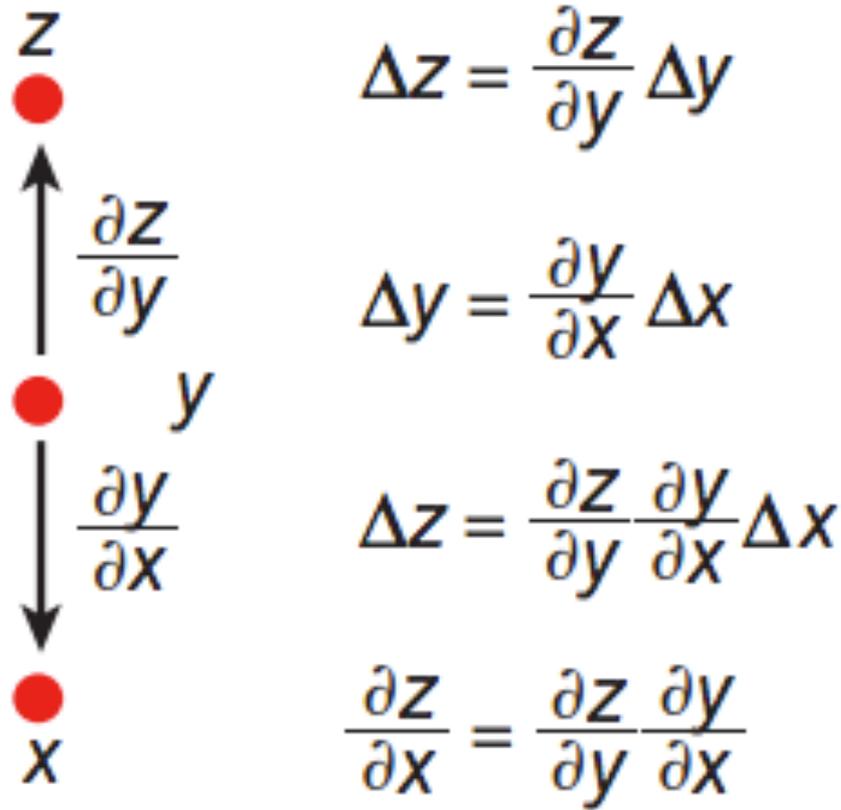
**"Deep learning."**

**Nature 521, no. 7553 (2015): 436-  
444.**

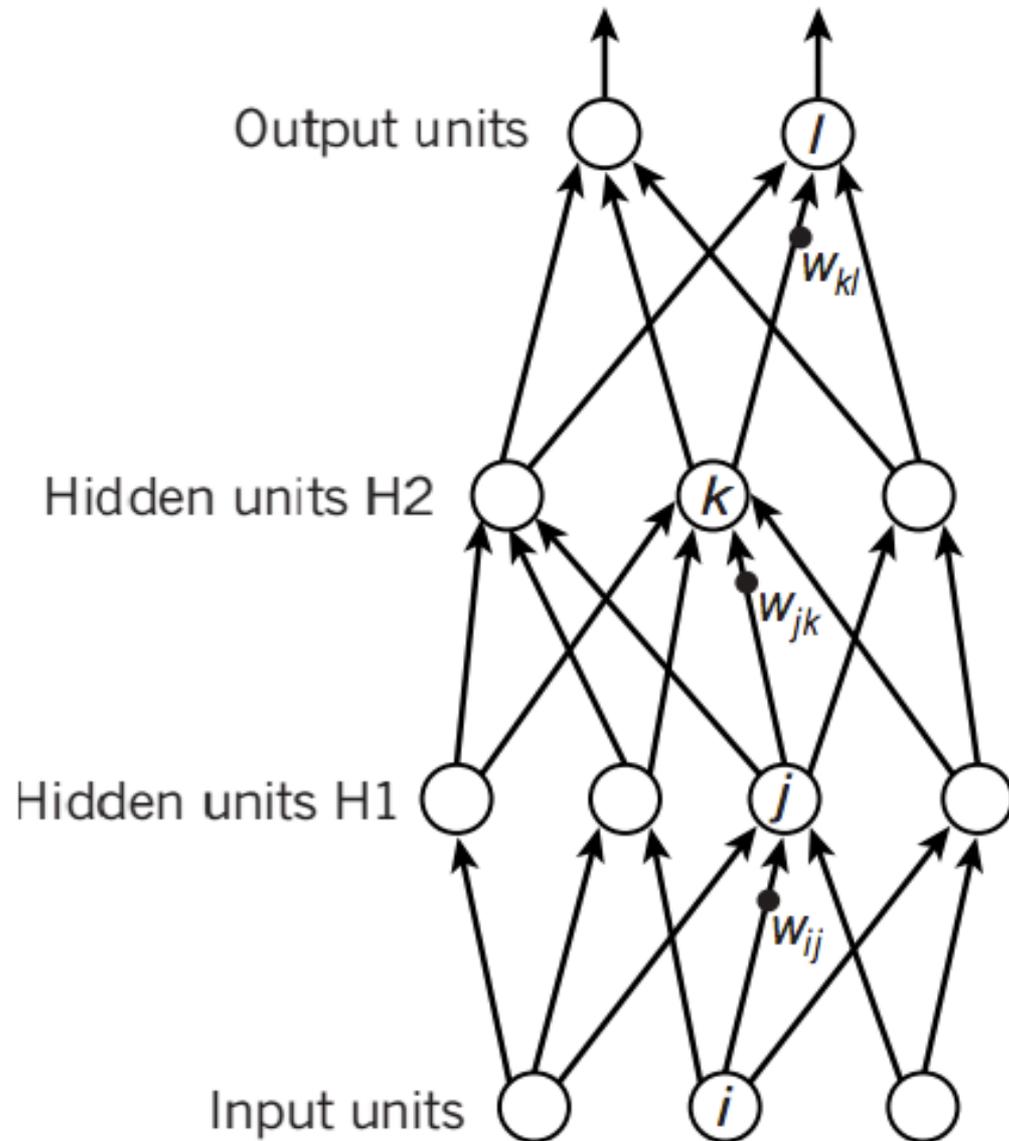
# Deep Learning



# Deep Learning



# Deep Learning



$$y_l = f(z_l)$$

$$z_l = \sum_{k \in H2} w_{kl} y_k$$

$$y_k = f(z_k)$$

$$z_k = \sum_{j \in H1} w_{jk} y_j$$

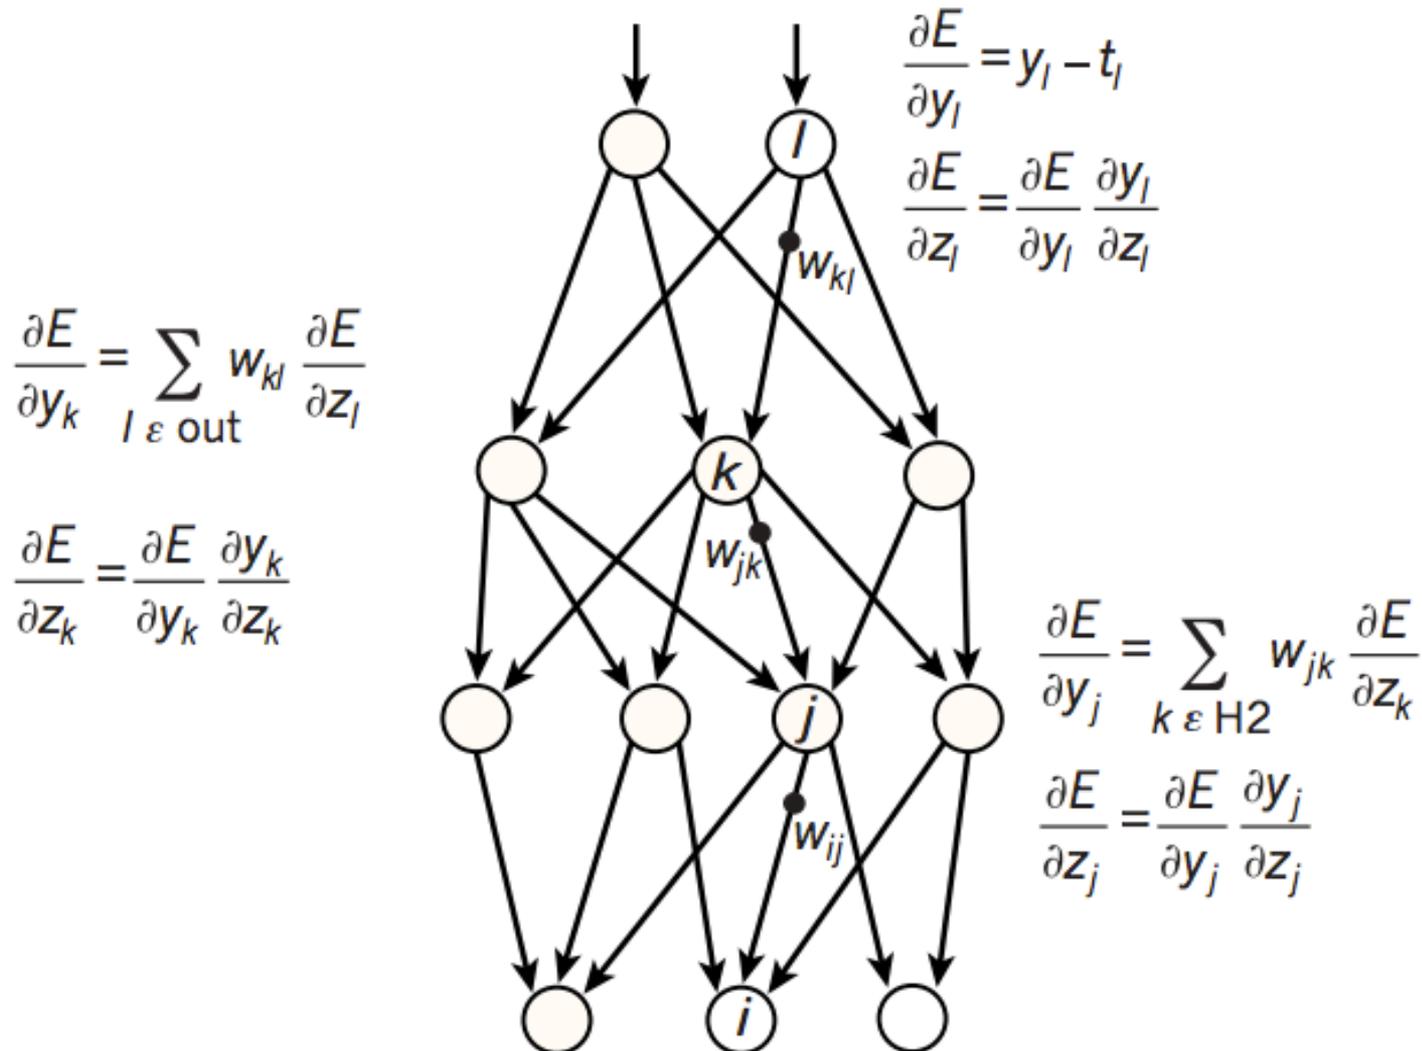
$$y_j = f(z_j)$$

$$z_j = \sum_{i \in \text{Input}} w_{ij} x_i$$

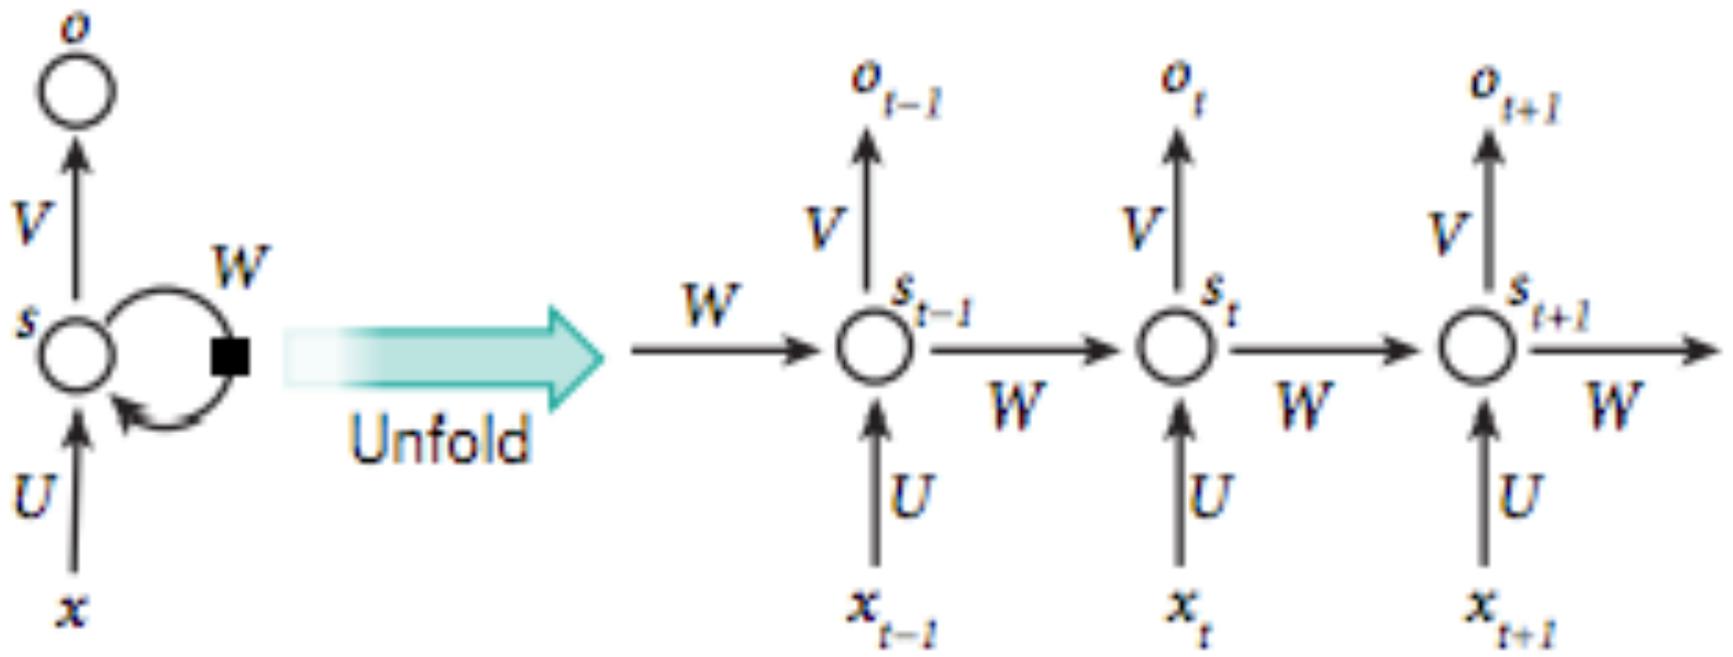
# Deep Learning

**d**

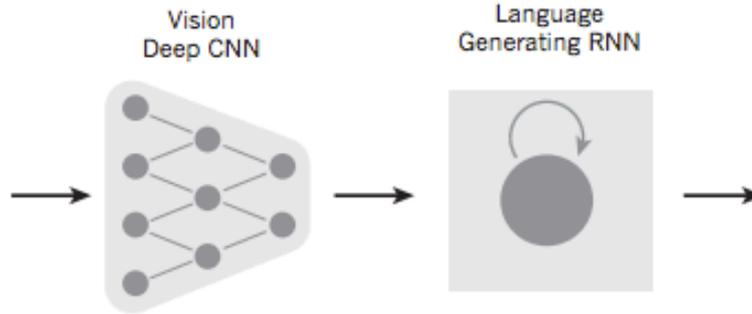
Compare outputs with correct answer to get error derivatives



# Recurrent Neural Network (RNN)



# From image to text



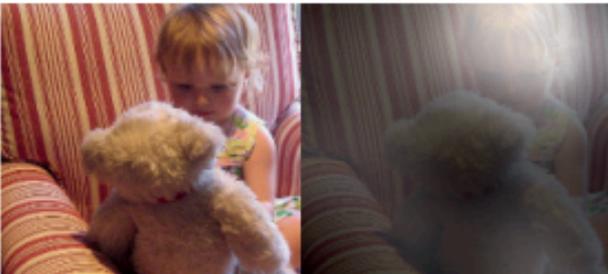
A woman is throwing a **frisbee** in a park.



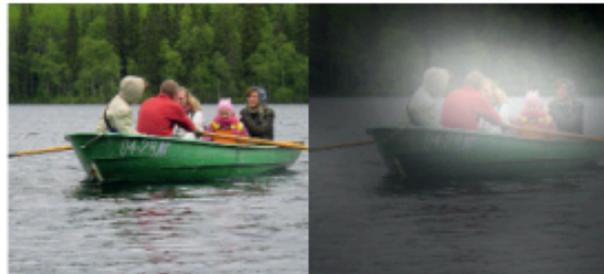
A **dog** is standing on a hardwood floor.



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)

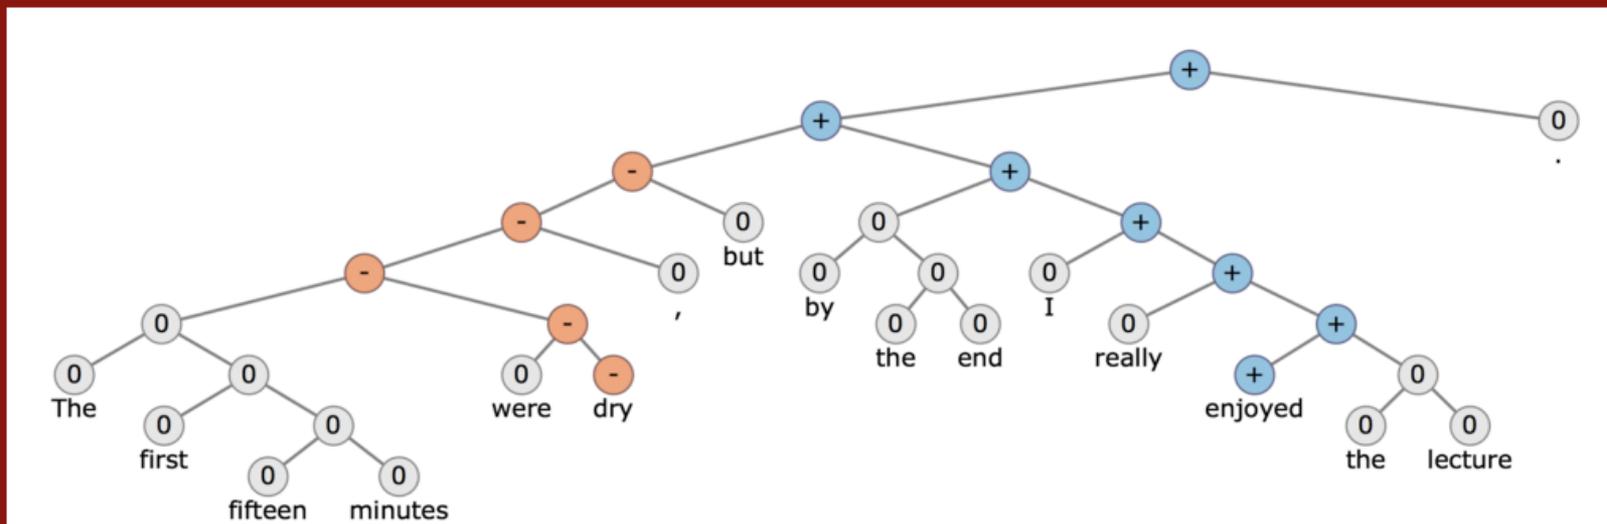
Text: recurrent neural network (RNN)



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

# CS224d: Deep Learning for Natural Language Processing

CS224d: Deep Learning for Natural Language Processing

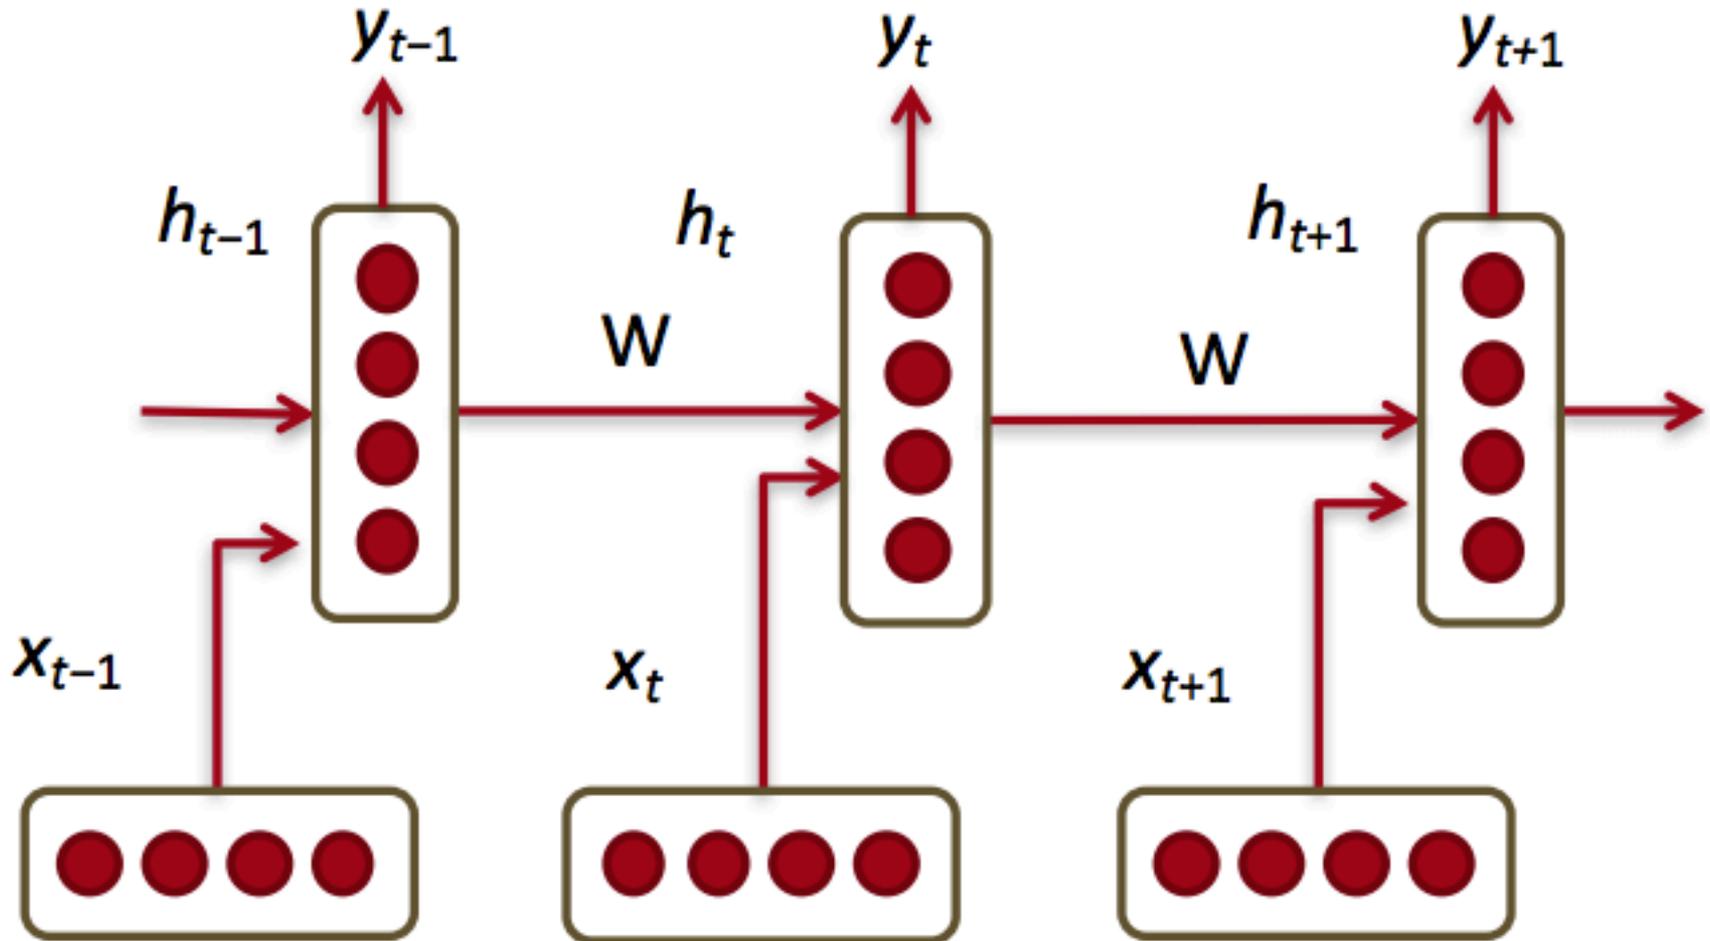


## Course Description

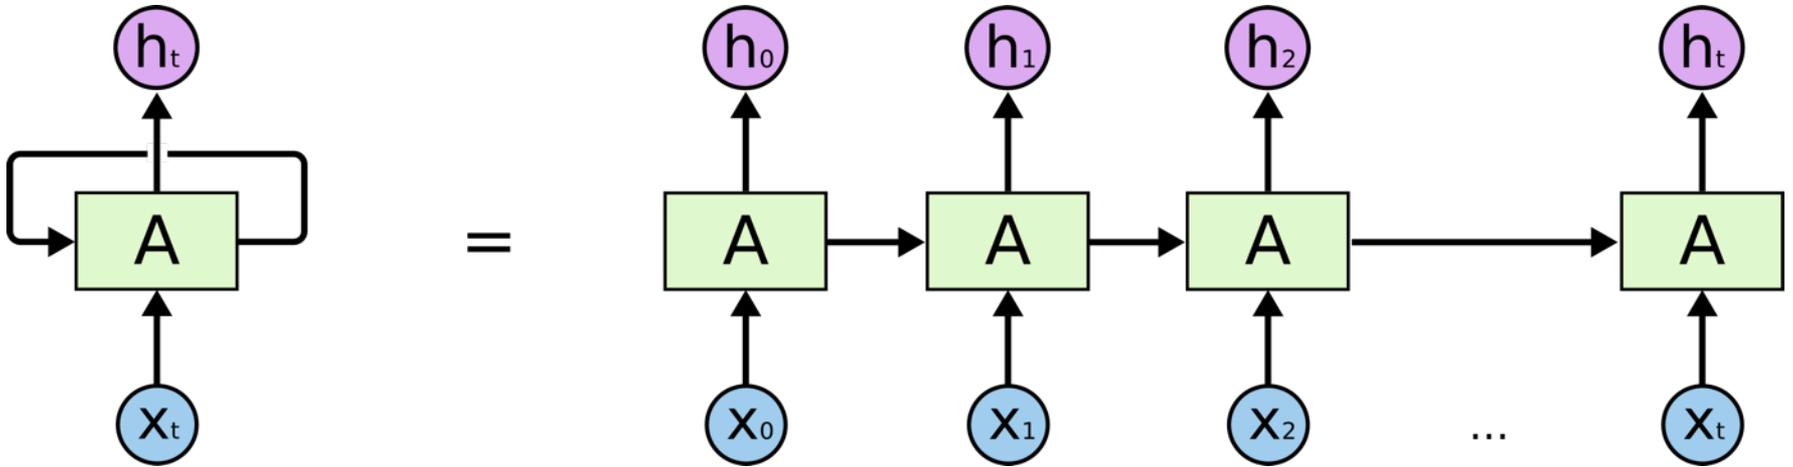
Natural language processing (NLP) is one of the most important technologies of the information age. Understanding complex language utterances is also a crucial part of artificial intelligence. Applications of NLP are everywhere because people communicate most everything in language: web search, advertisement, 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,

<http://cs224d.stanford.edu/>

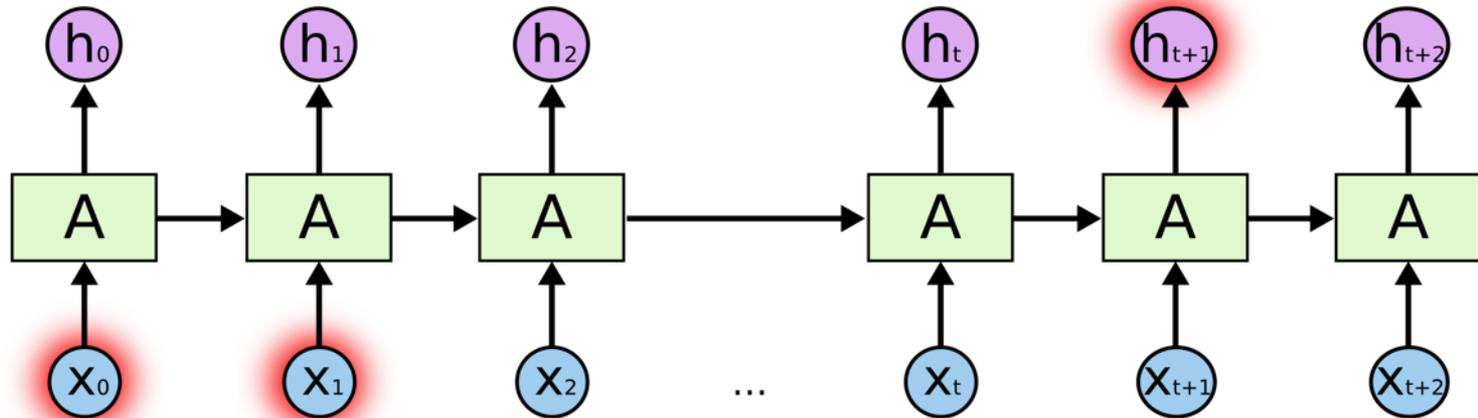
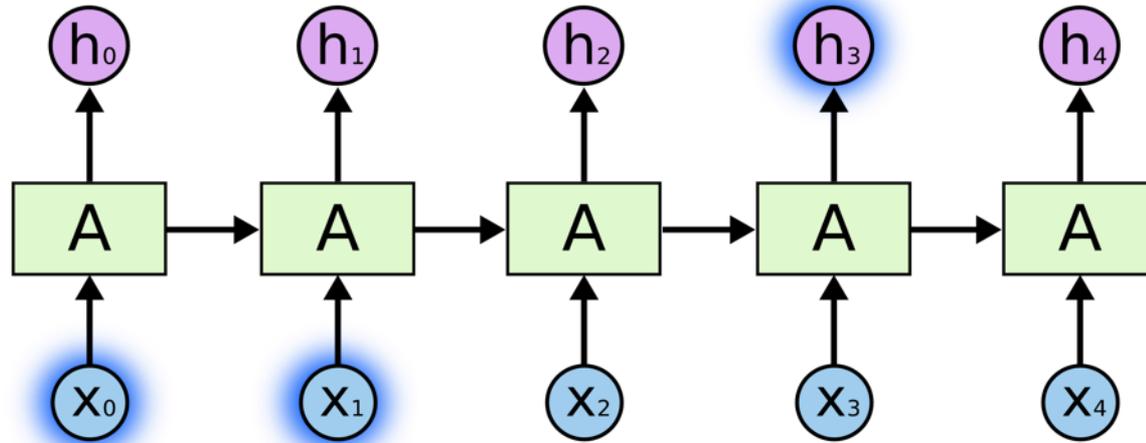
# Recurrent Neural Networks (RNNs)



# RNN

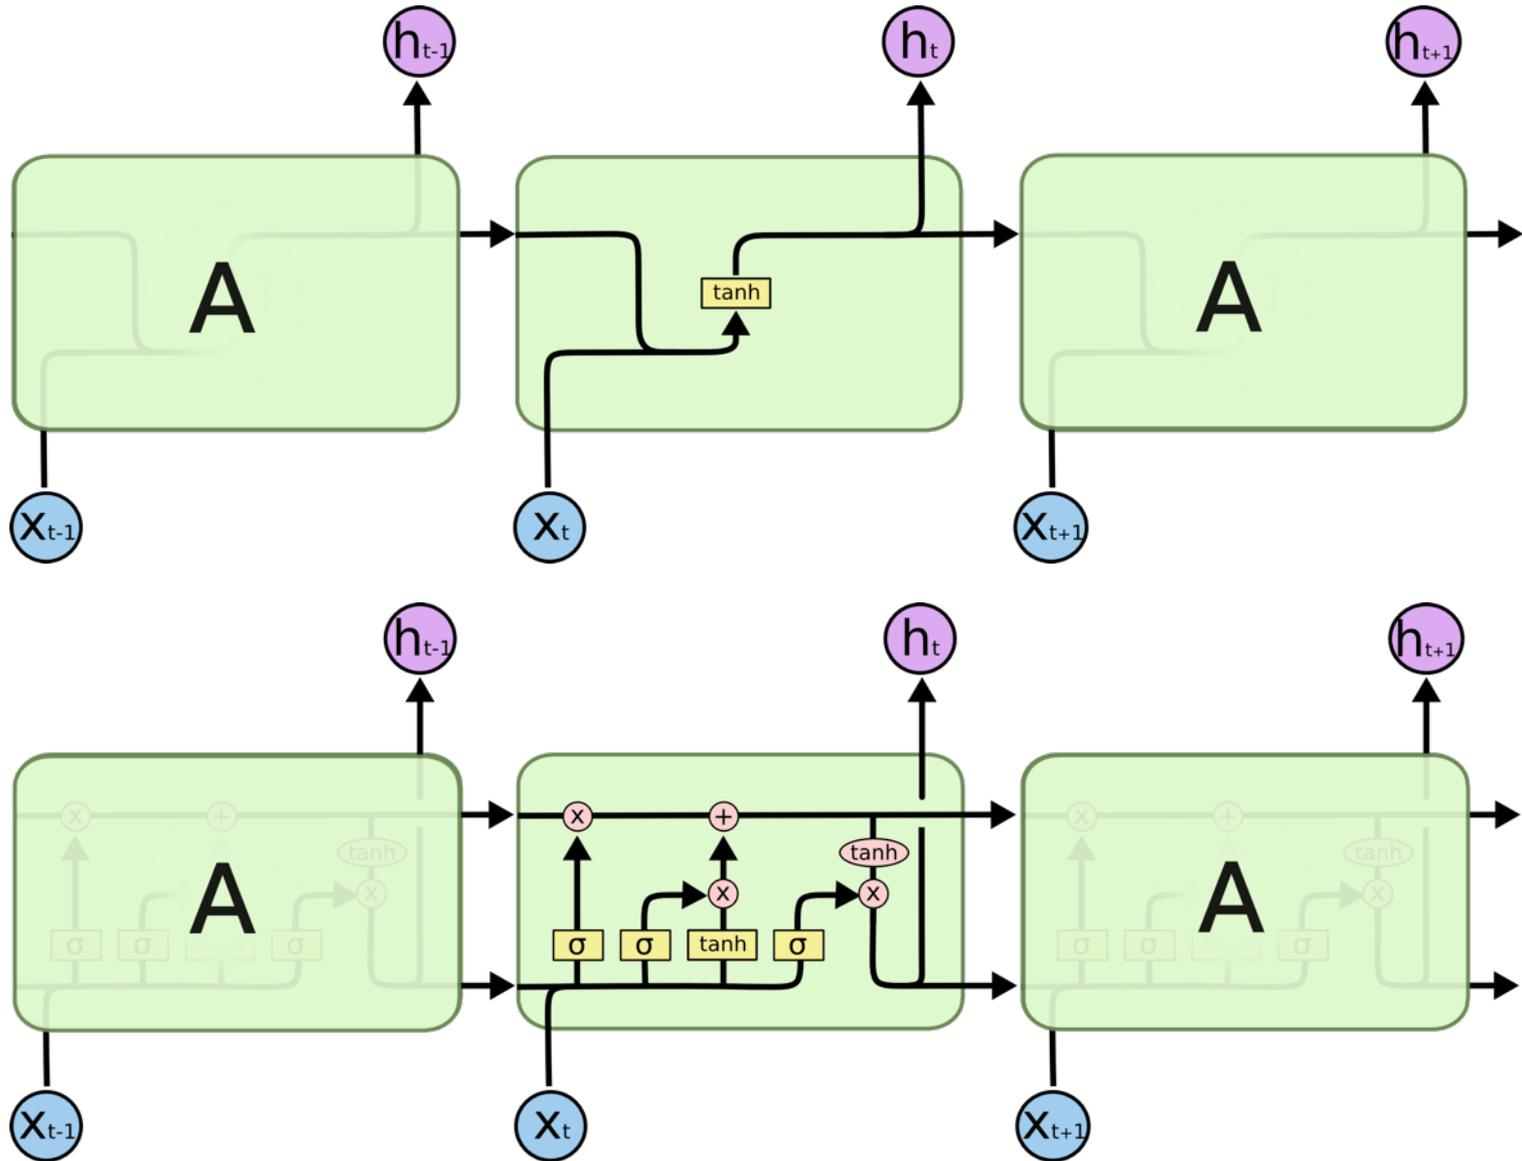


# RNN long-term dependencies

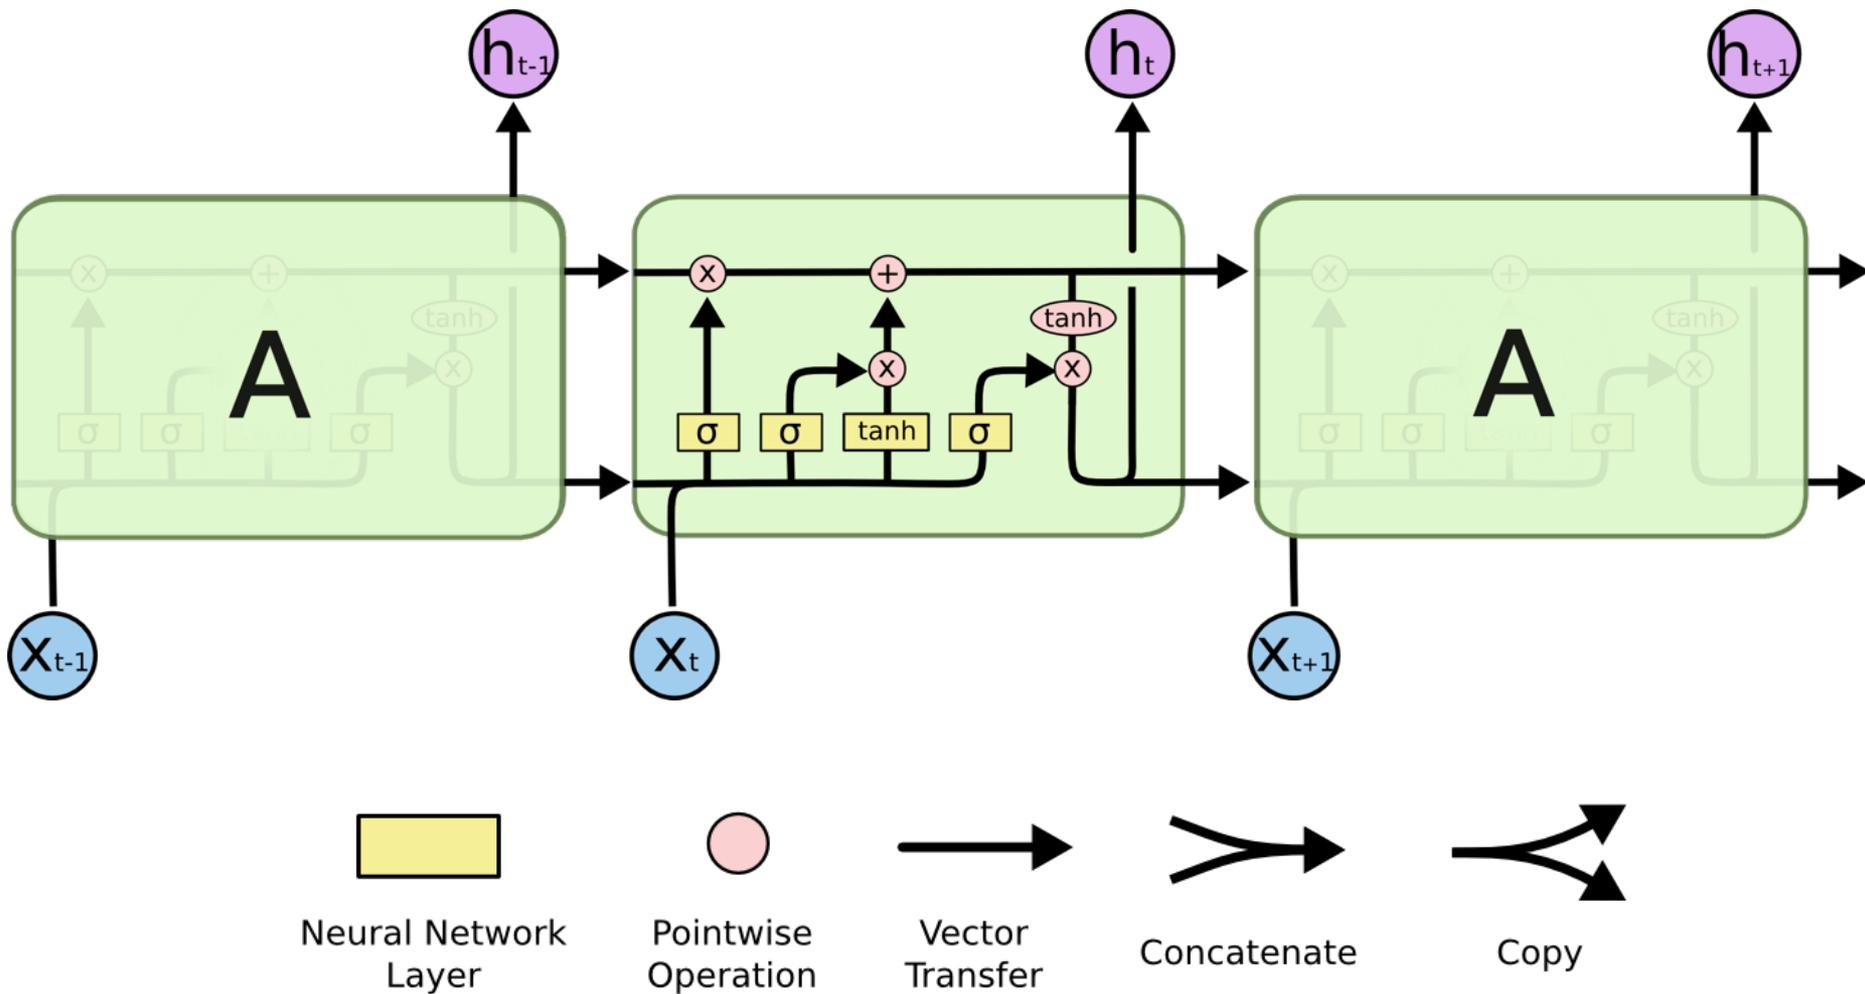


I grew up in France... I speak fluent French.

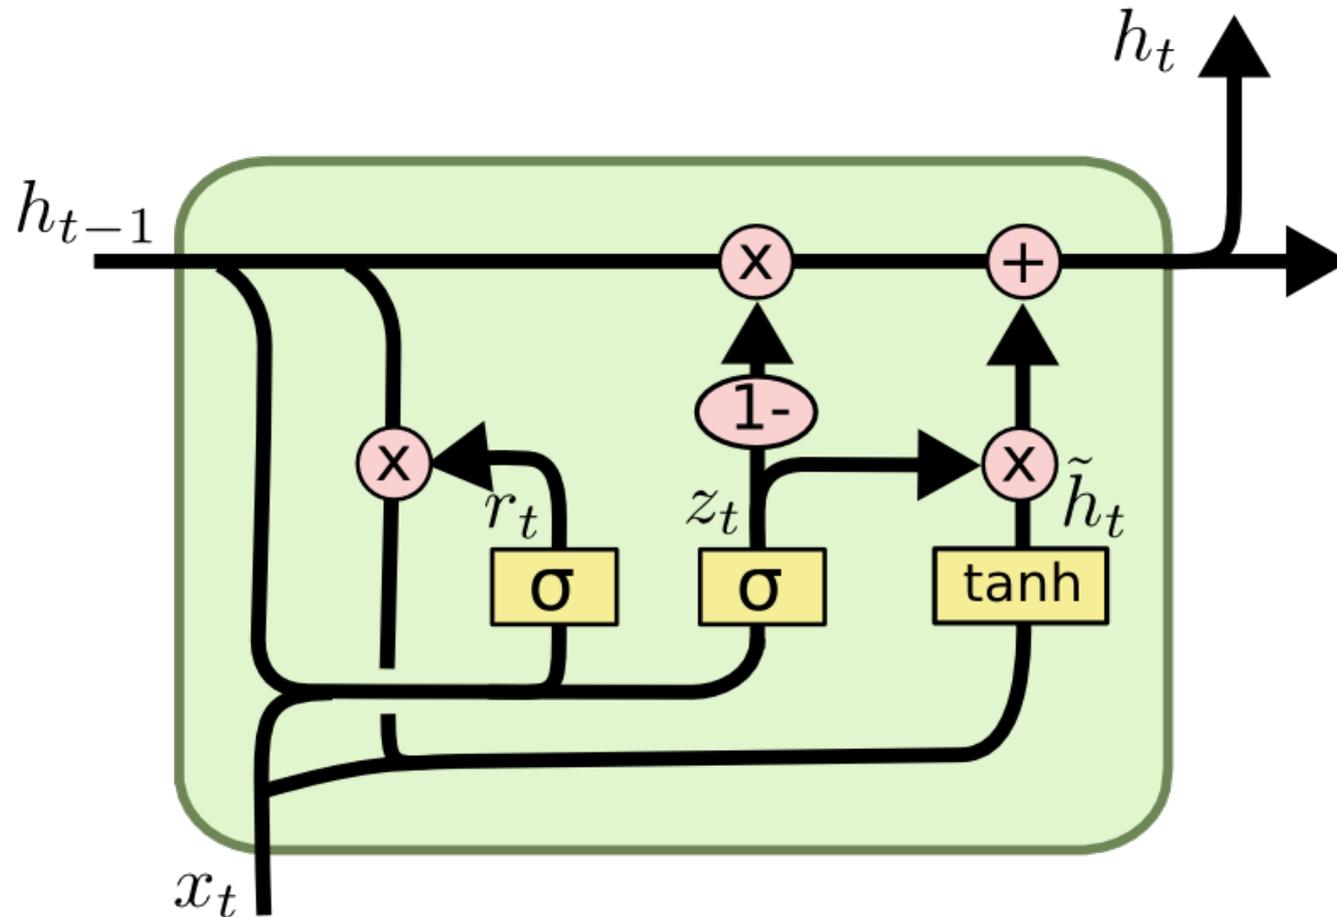
# RNN LSTM



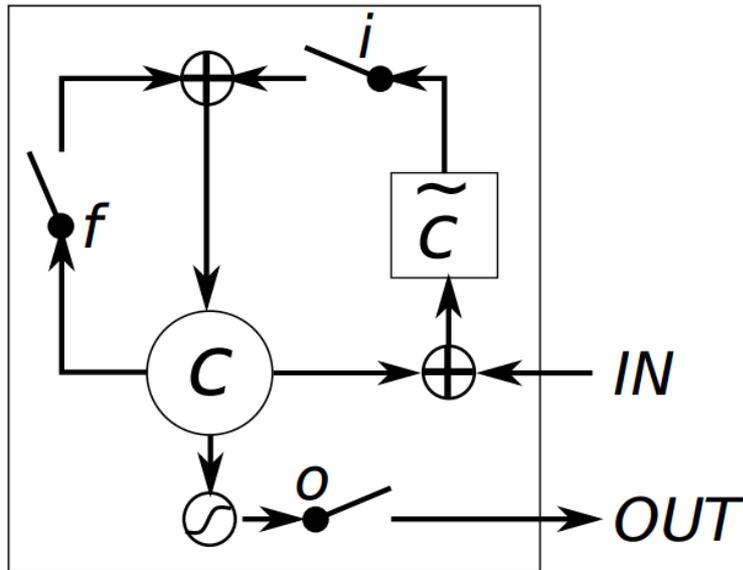
# Long Short Term Memory (LSTM)



# Gated Recurrent Unit (GRU)

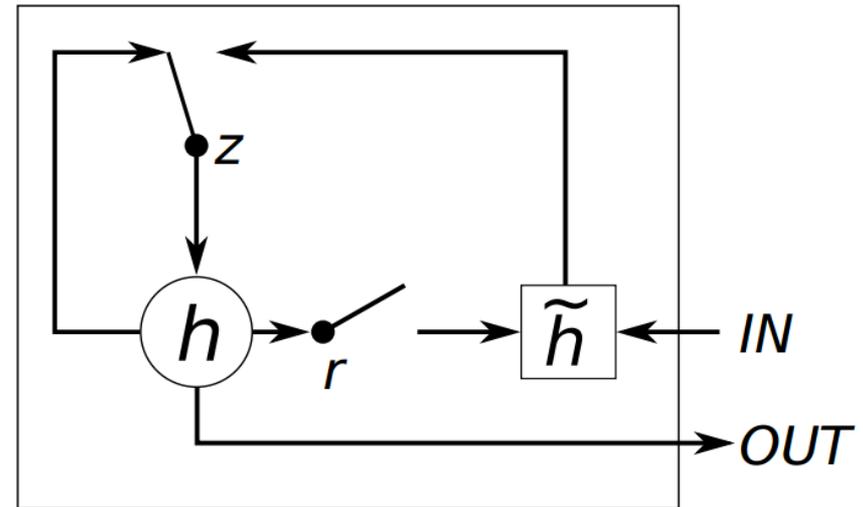


# LSTM vs GRU



## 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.

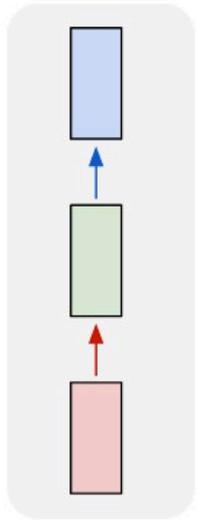


## GRU

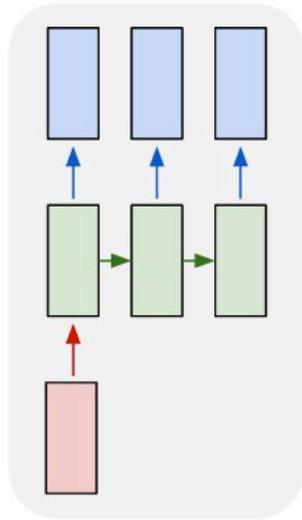
$r$  and  $z$  are the **reset** and **update** gates, and  $h$  and  $\tilde{h}$  are the activation and the candidate activation.

# LSTM Recurrent Neural Network

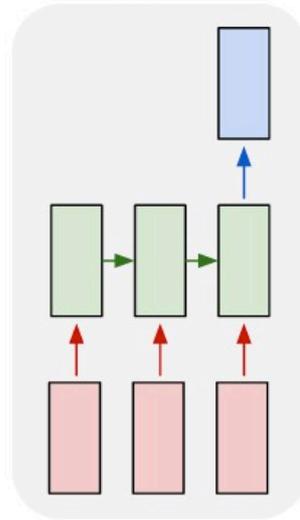
one to one



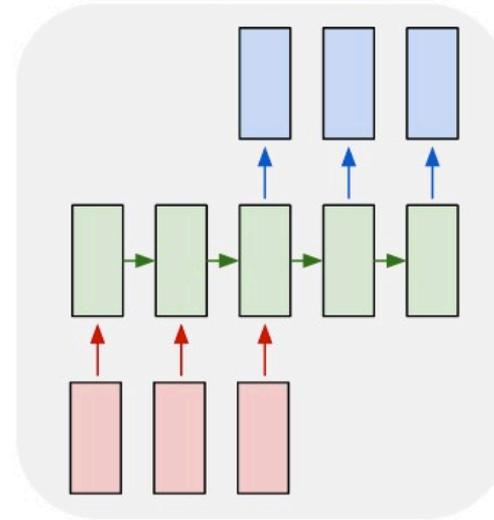
one to many



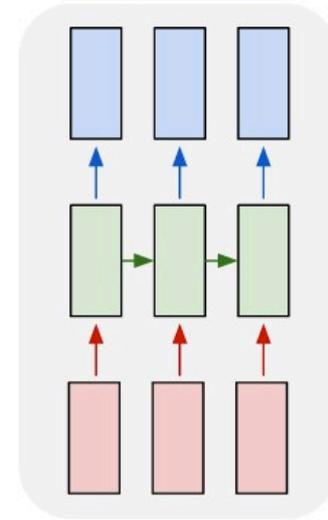
many to one



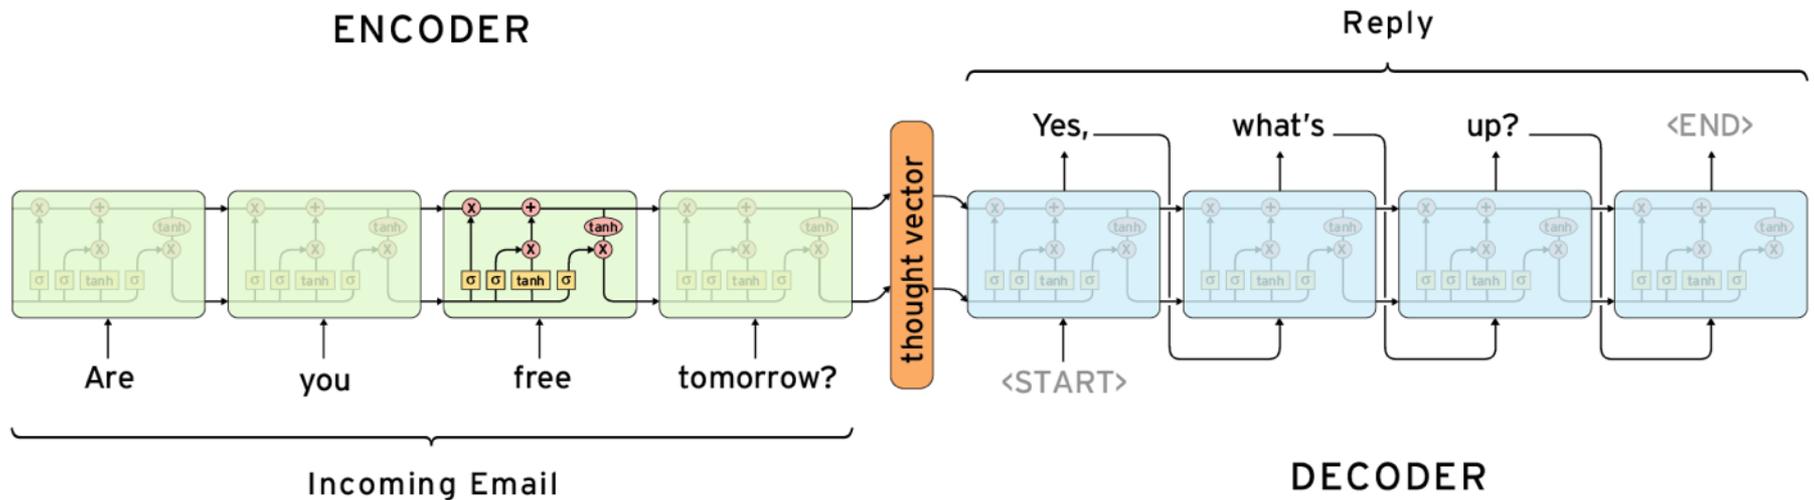
many to many



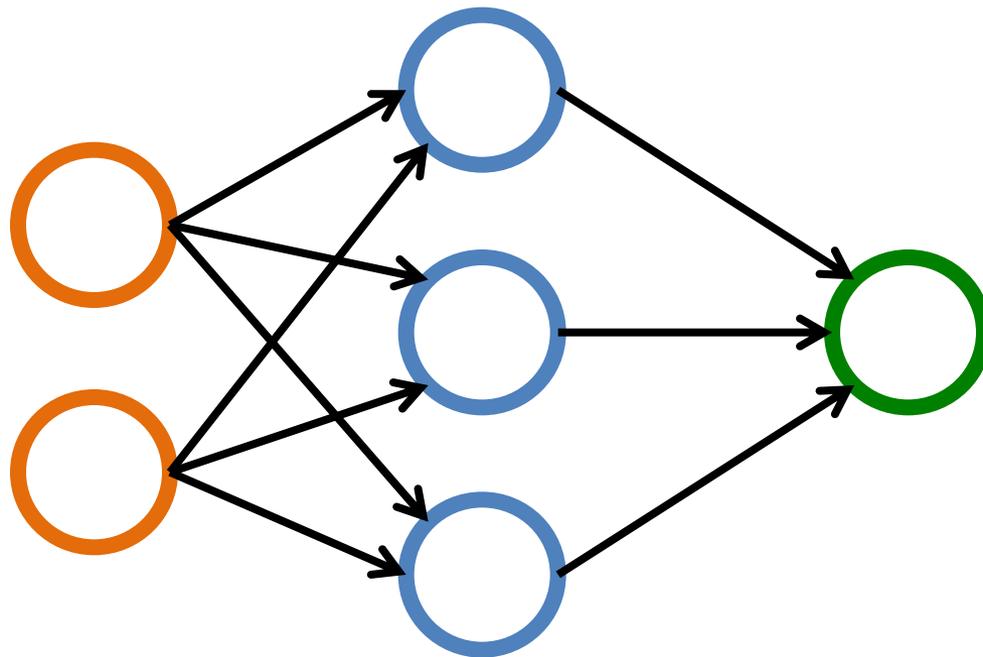
many to many



# The Sequence to Sequence model (seq2seq)



# Neural Networks

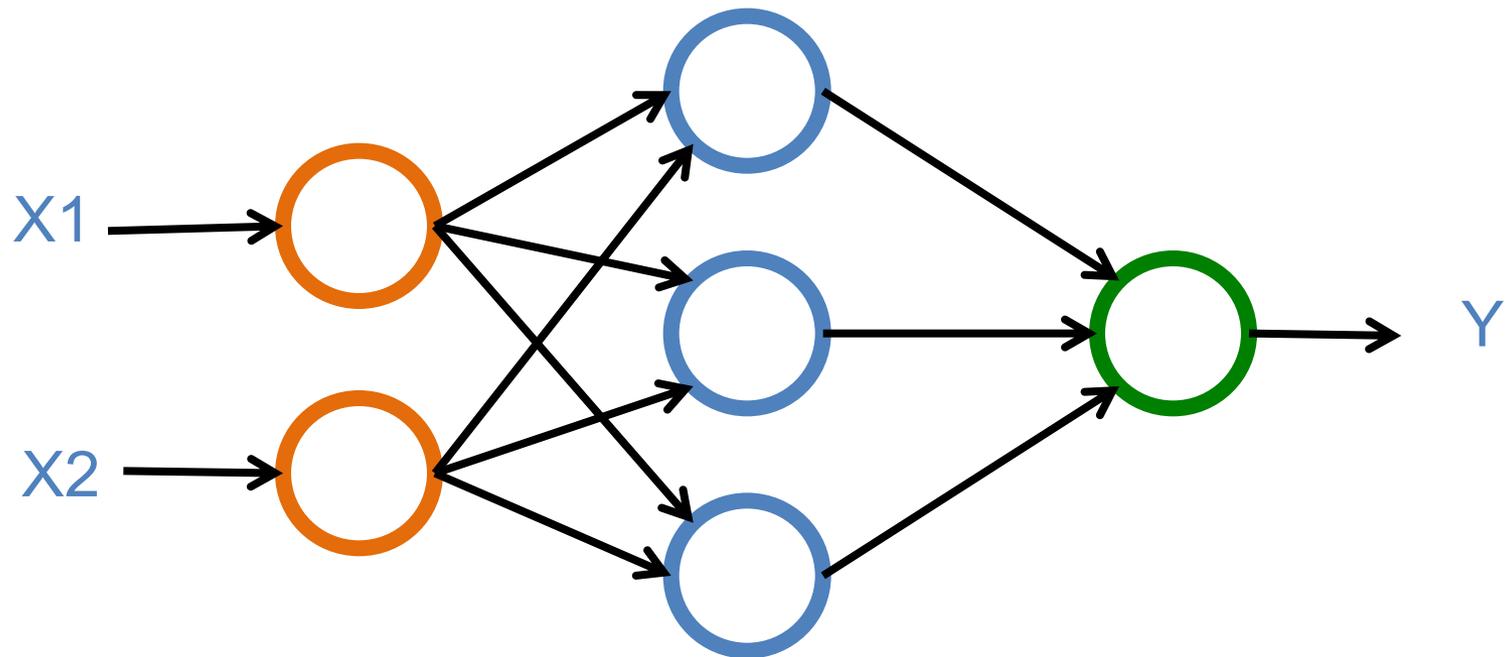


# Neural Networks

**Input Layer**  
(X)

**Hidden Layer**  
(H)

**Output Layer**  
(Y)



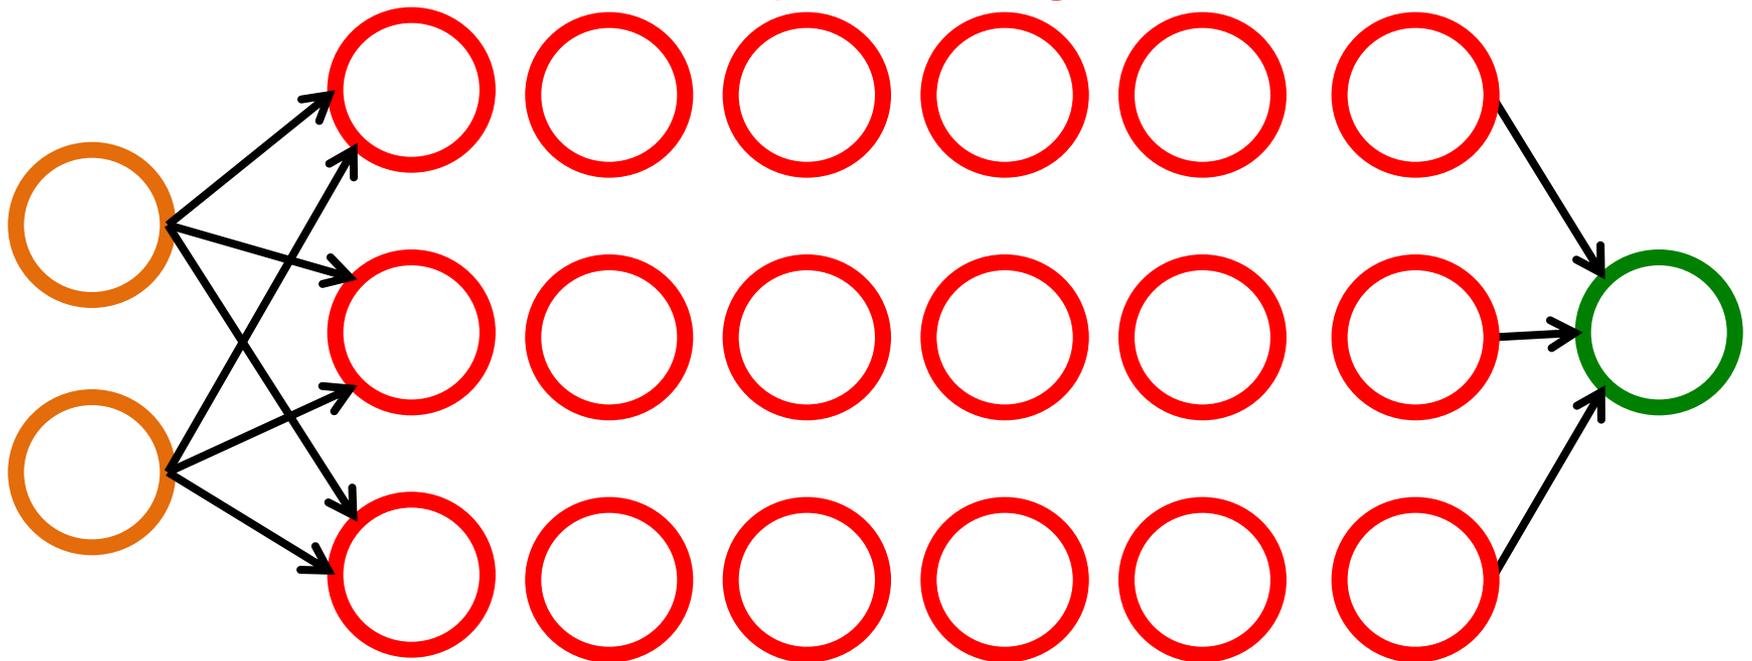
# Neural Networks

Input Layer  
(X)

Hidden Layers  
(H)

Output Layer  
(Y)

Deep Neural Networks  
Deep Learning

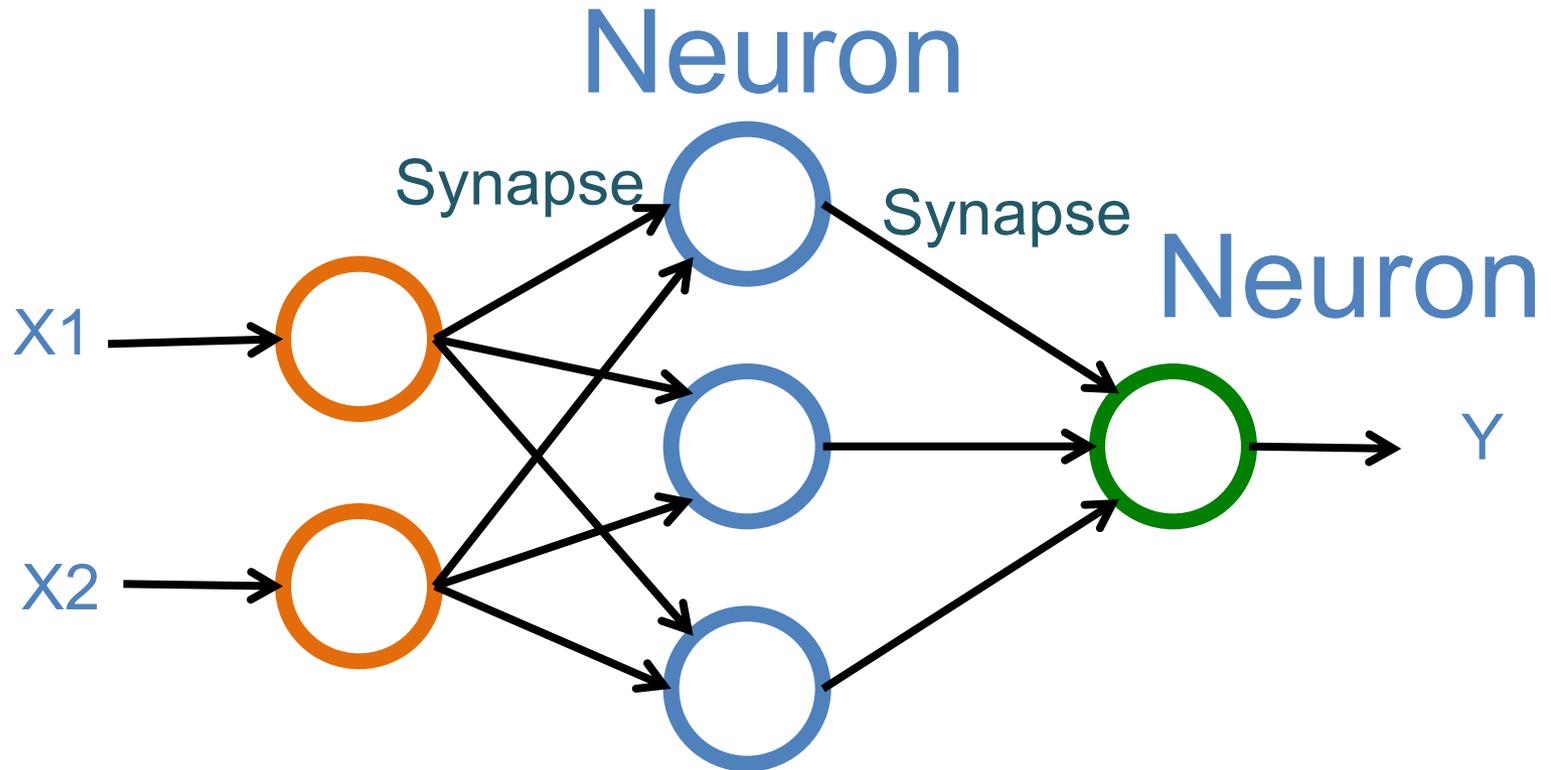


# Neural Networks

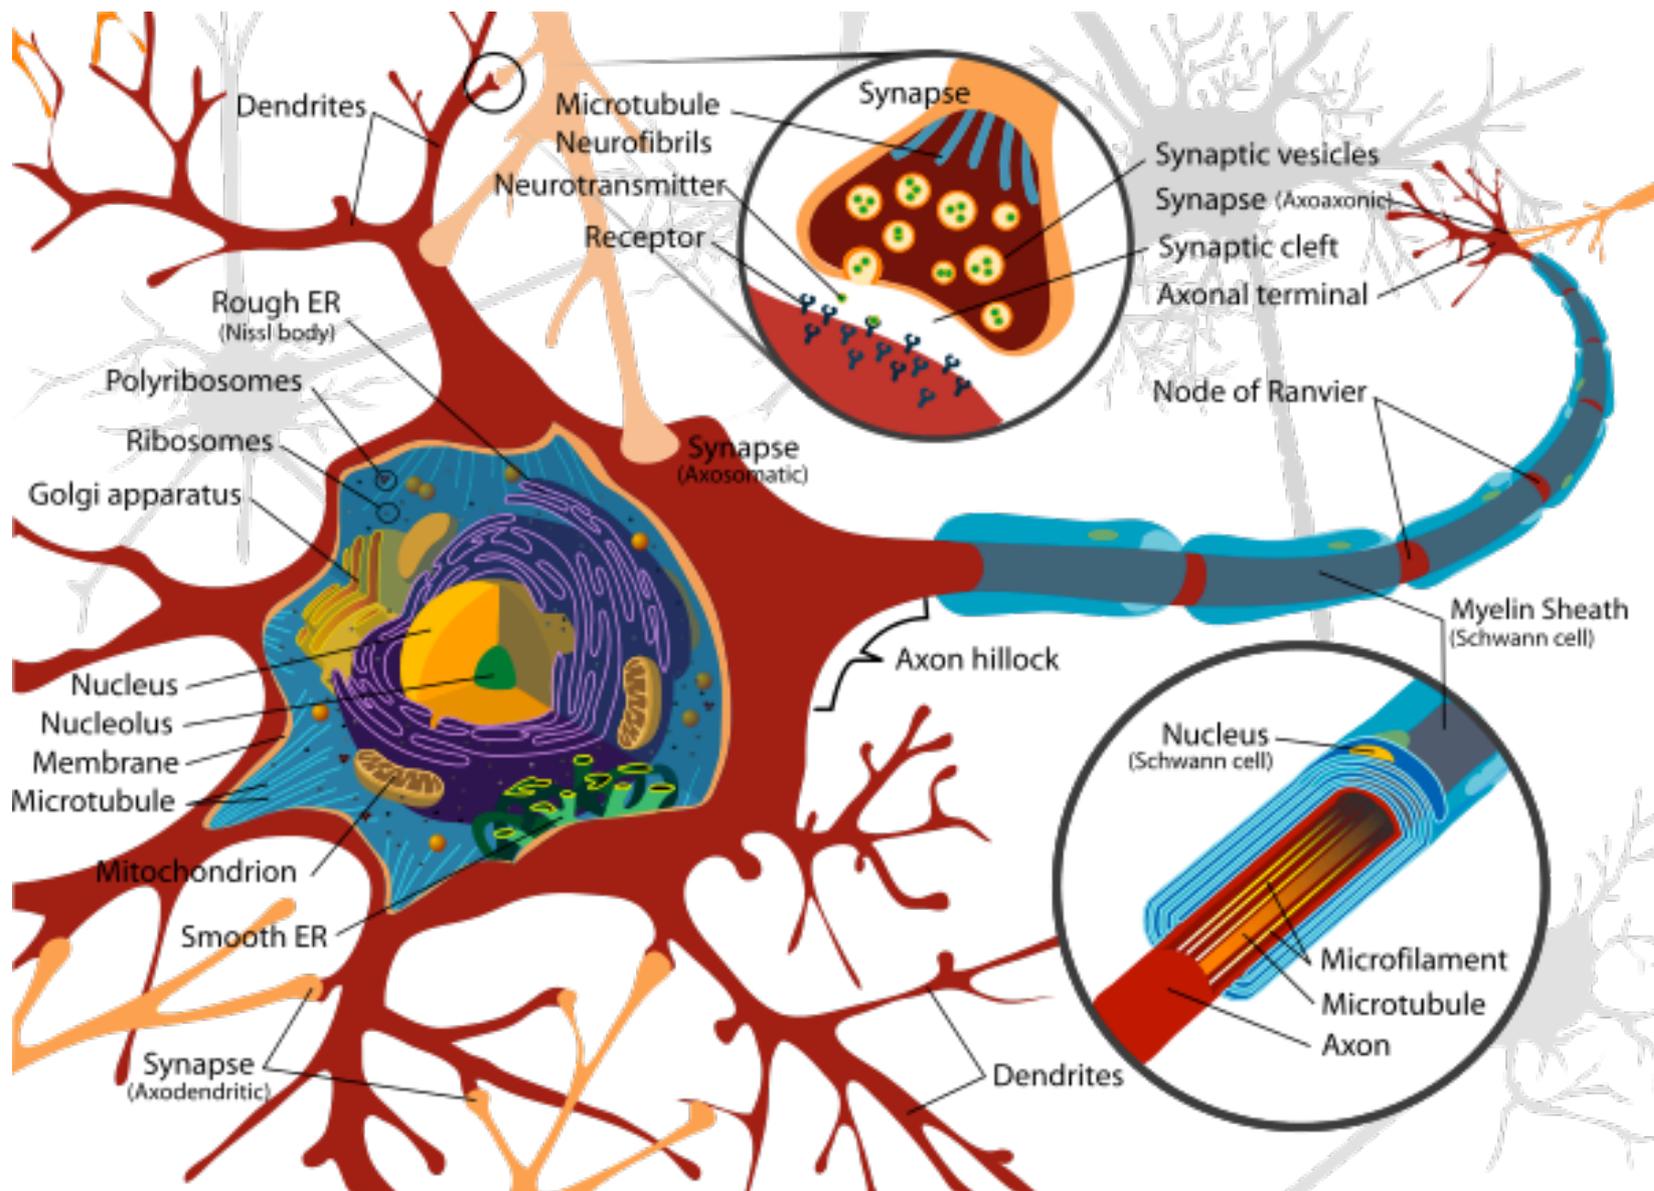
Input Layer  
(X)

Hidden Layer  
(H)

Output Layer  
(Y)

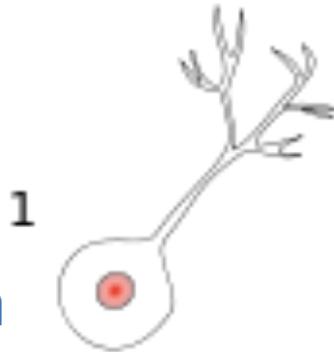


# Neuron and Synapse

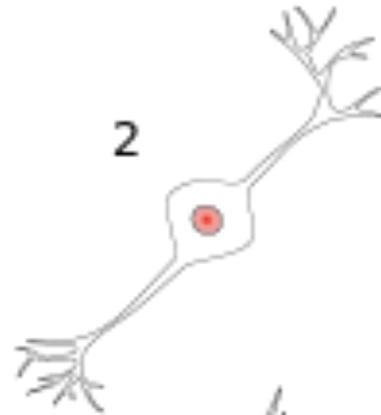


# Neurons

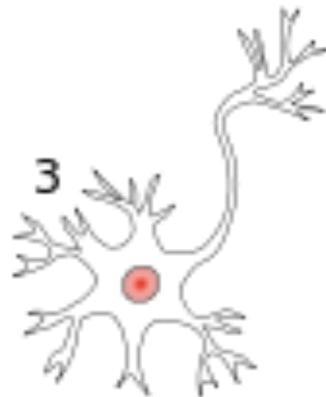
1 Unipolar neuron



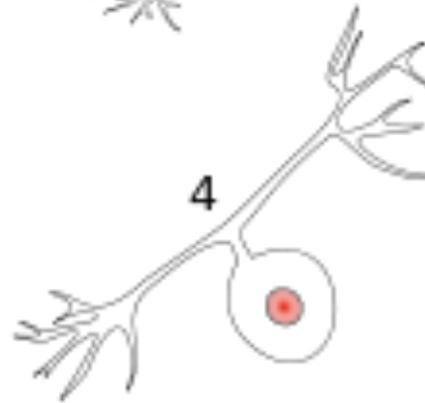
2 Bipolar neuron



3 Multipolar neuron



4 Pseudounipolar neuron

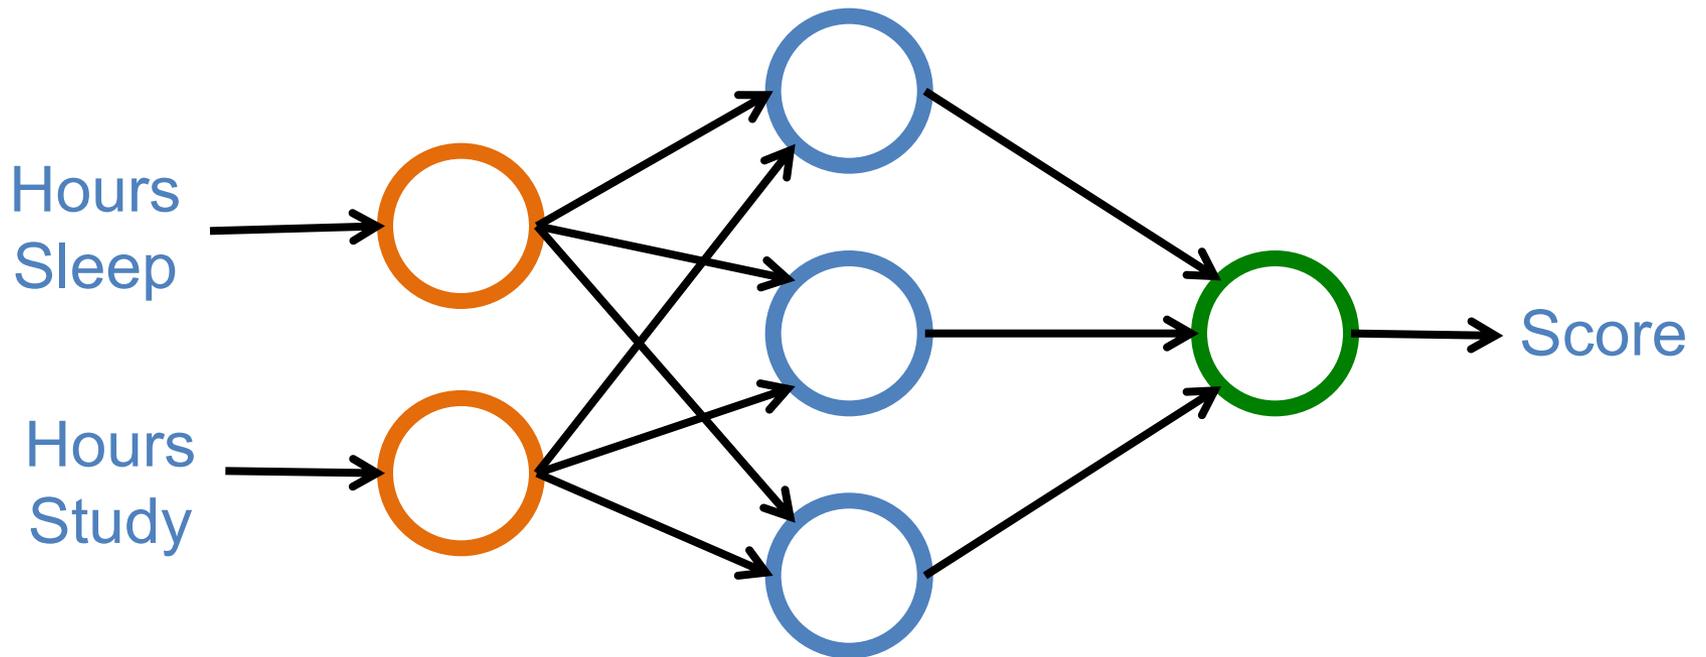


# Neural Networks

**Input Layer**  
(X)

**Hidden Layer**  
(H)

**Output Layer**  
(Y)

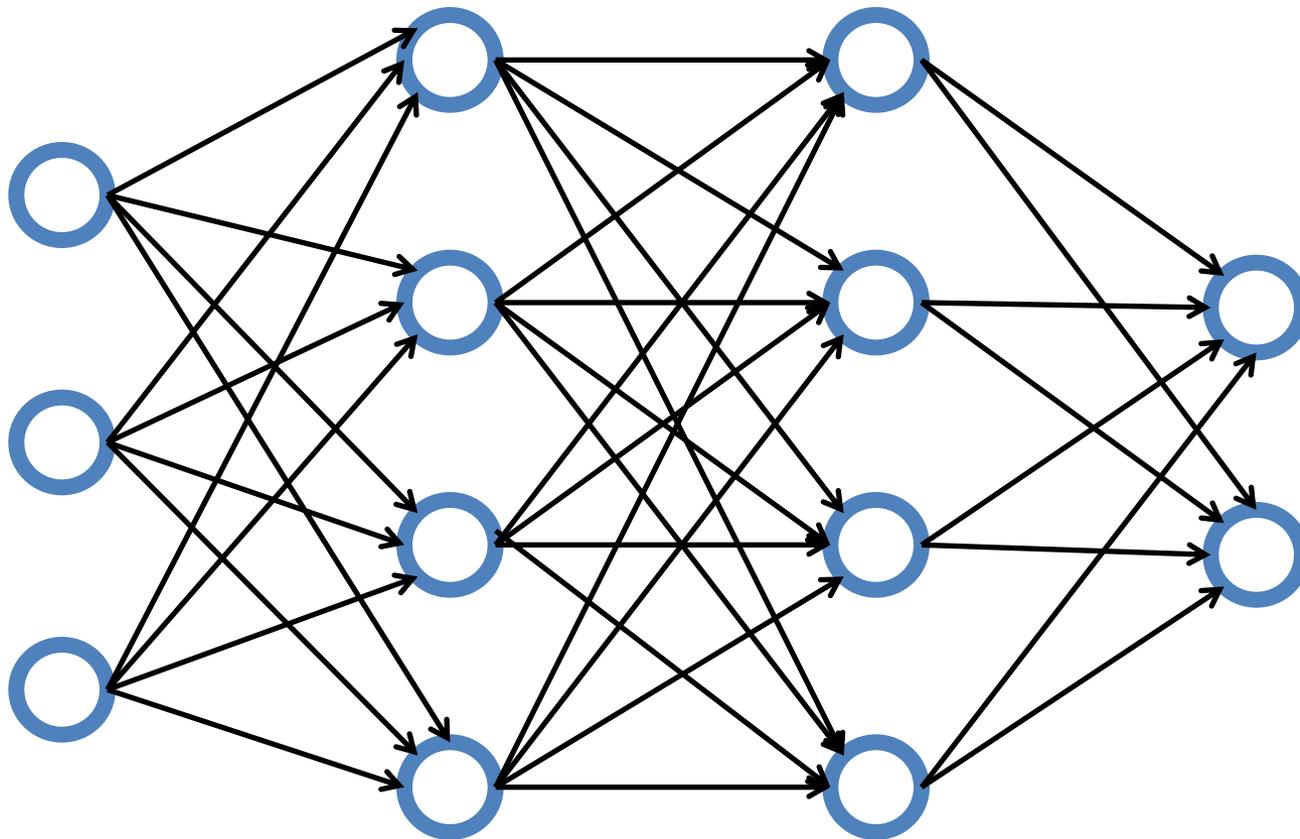


# Neural Networks

Input Layer  
(X)

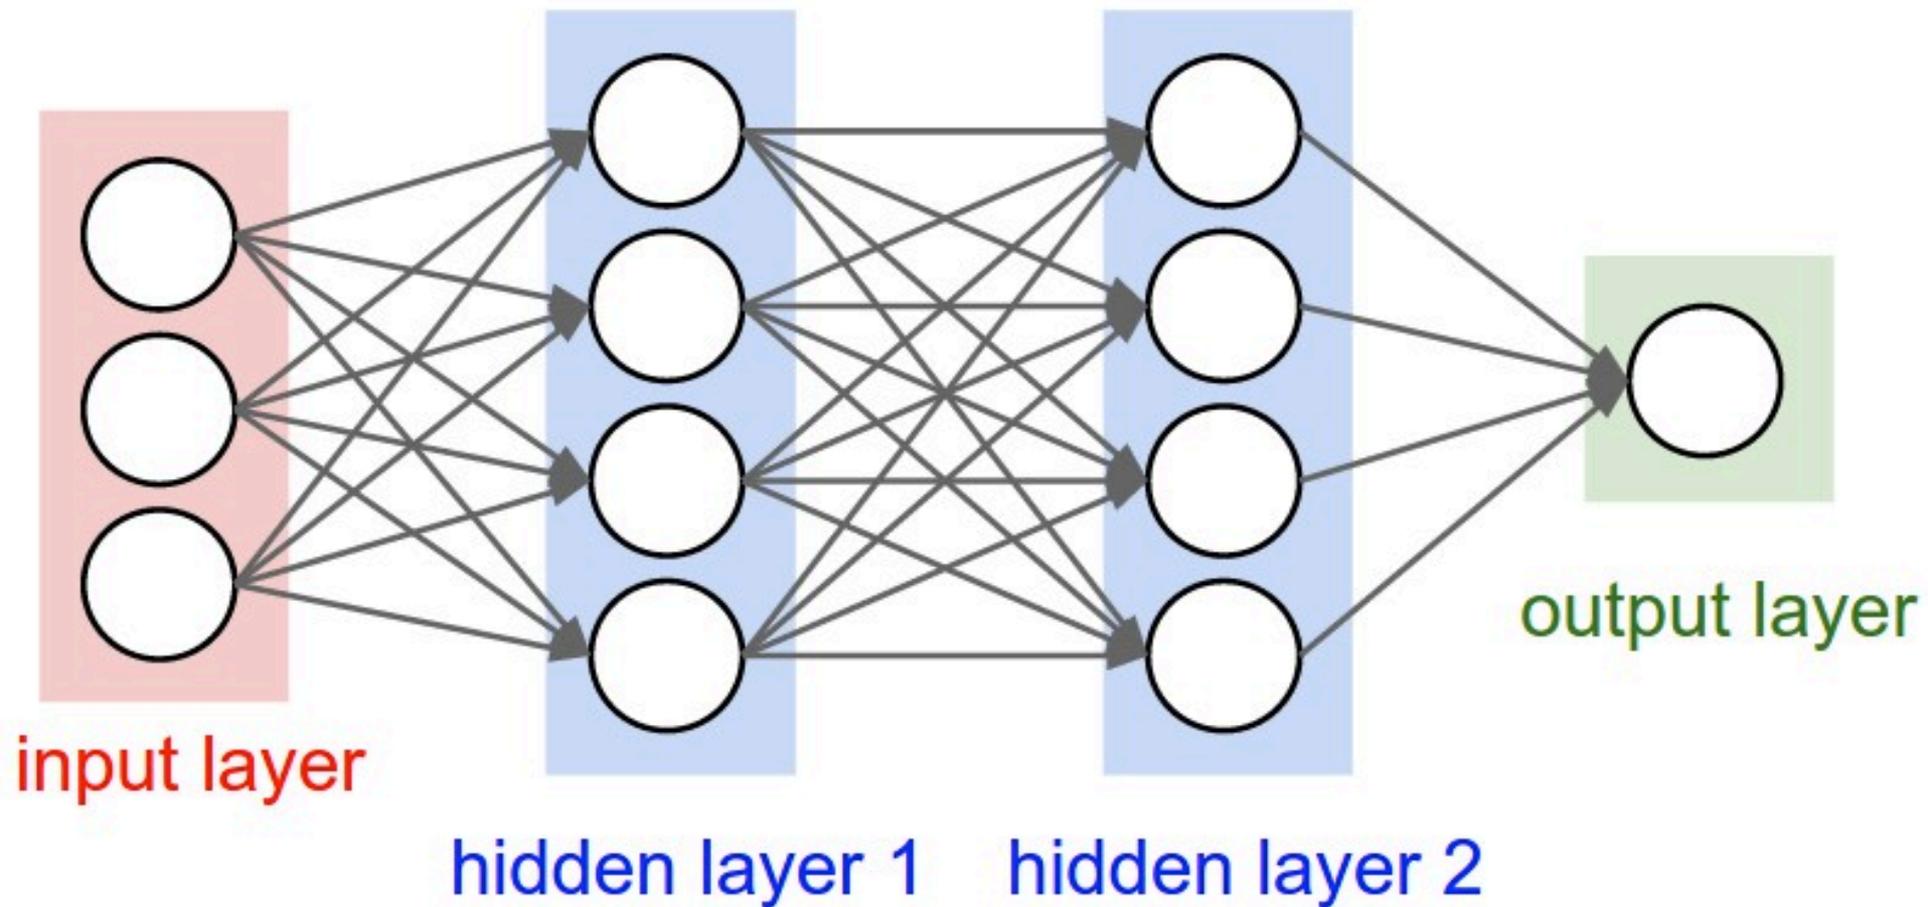
Hidden Layer  
(H)

Output Layer  
(Y)

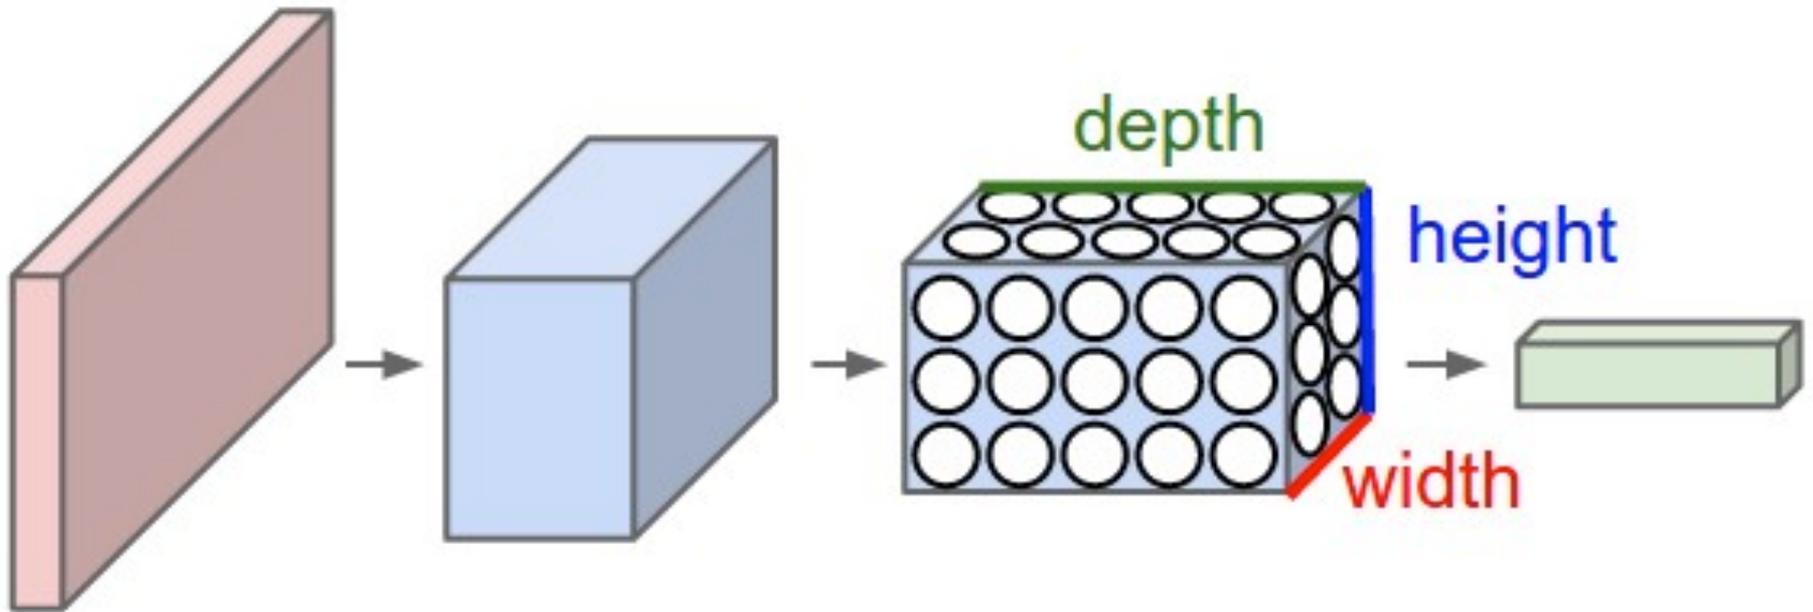


# Convolutional Neural Networks (CNNs / ConvNets)

# A regular 3-layer Neural Network

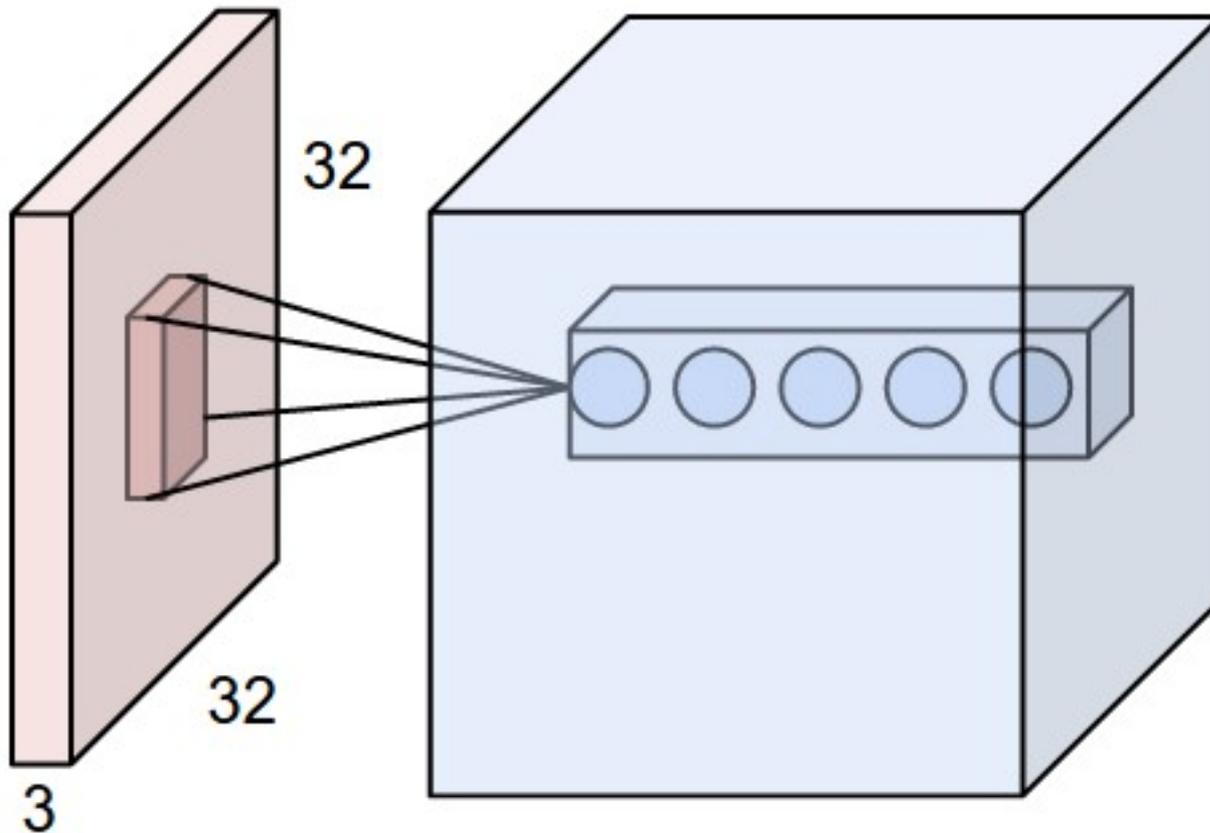


# A ConvNet arranges its neurons in three dimensions (width, height, depth)

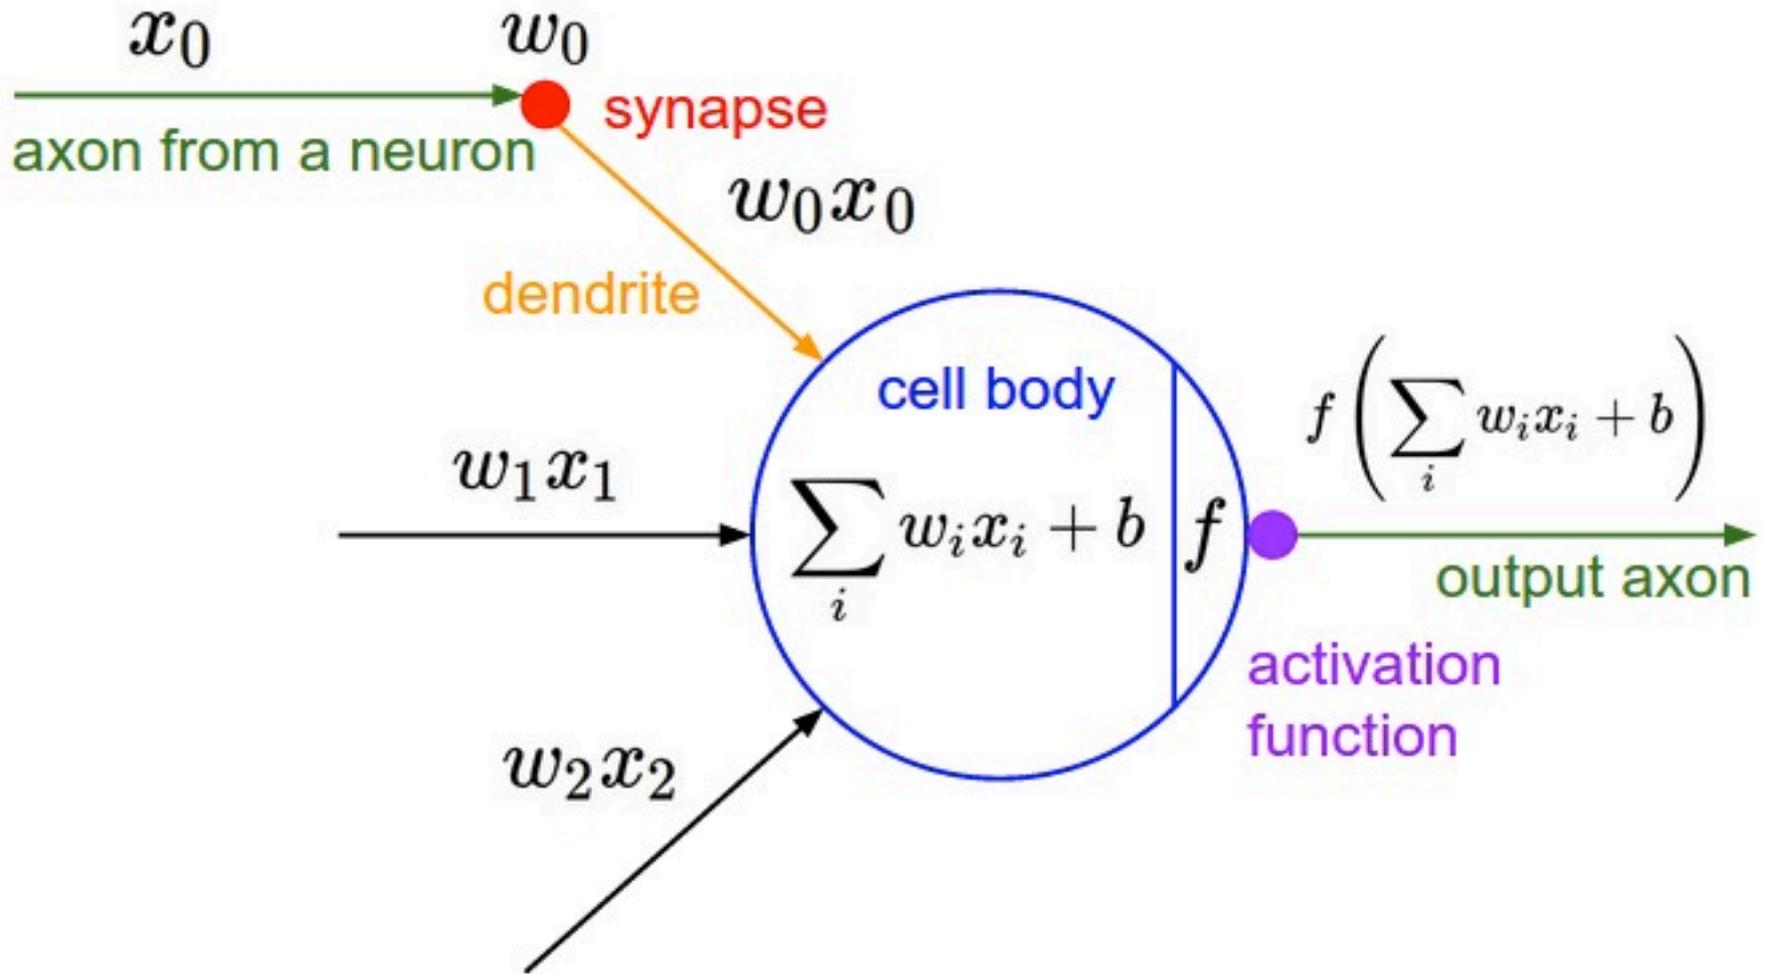




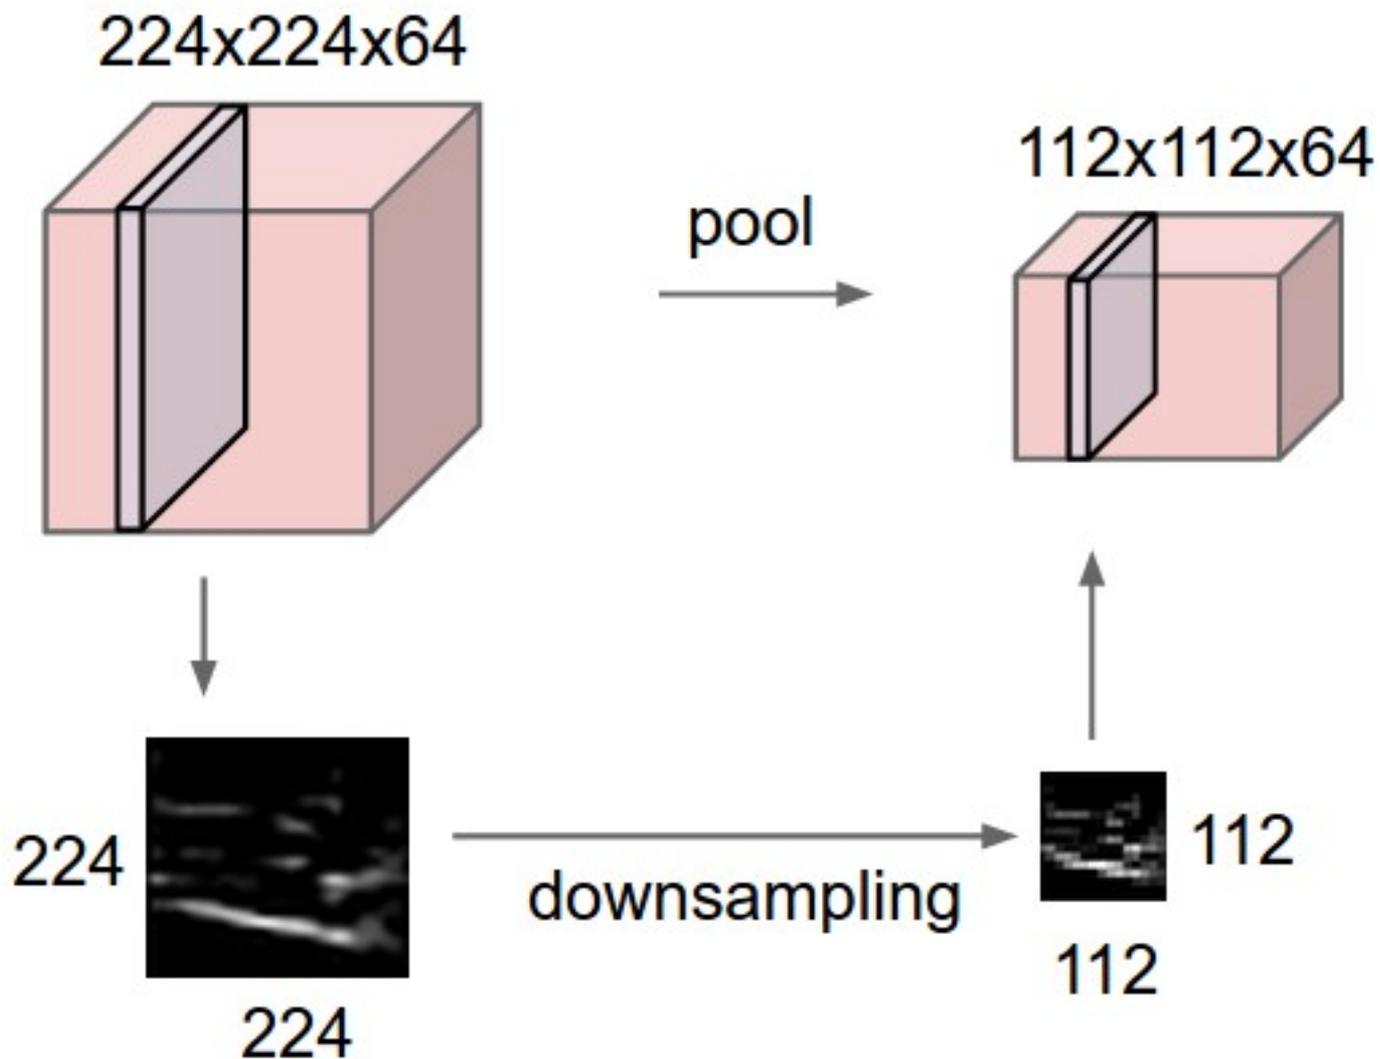
# ConvNets



# ConvNets



# ConvNets



# ConvNets

## max pooling

Single depth slice

x

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

y

max pool with 2x2 filters  
and stride 2



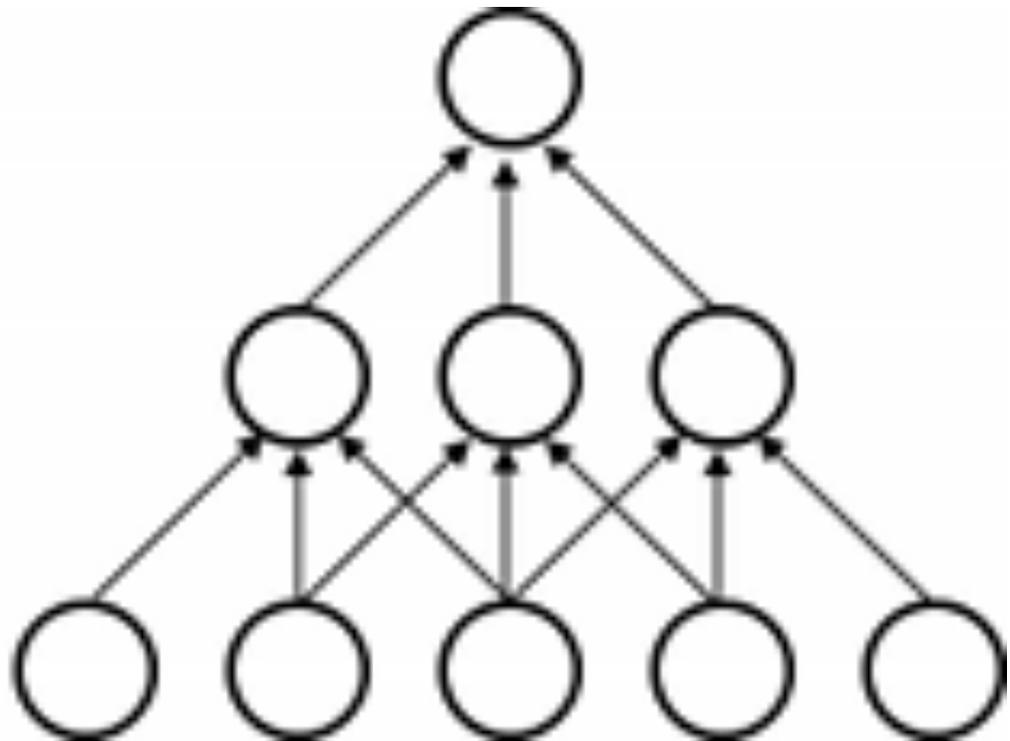
6	8
3	4

# Convolutional Neural Networks (CNN) (LeNet) Sparse Connectivity

layer  $m+1$

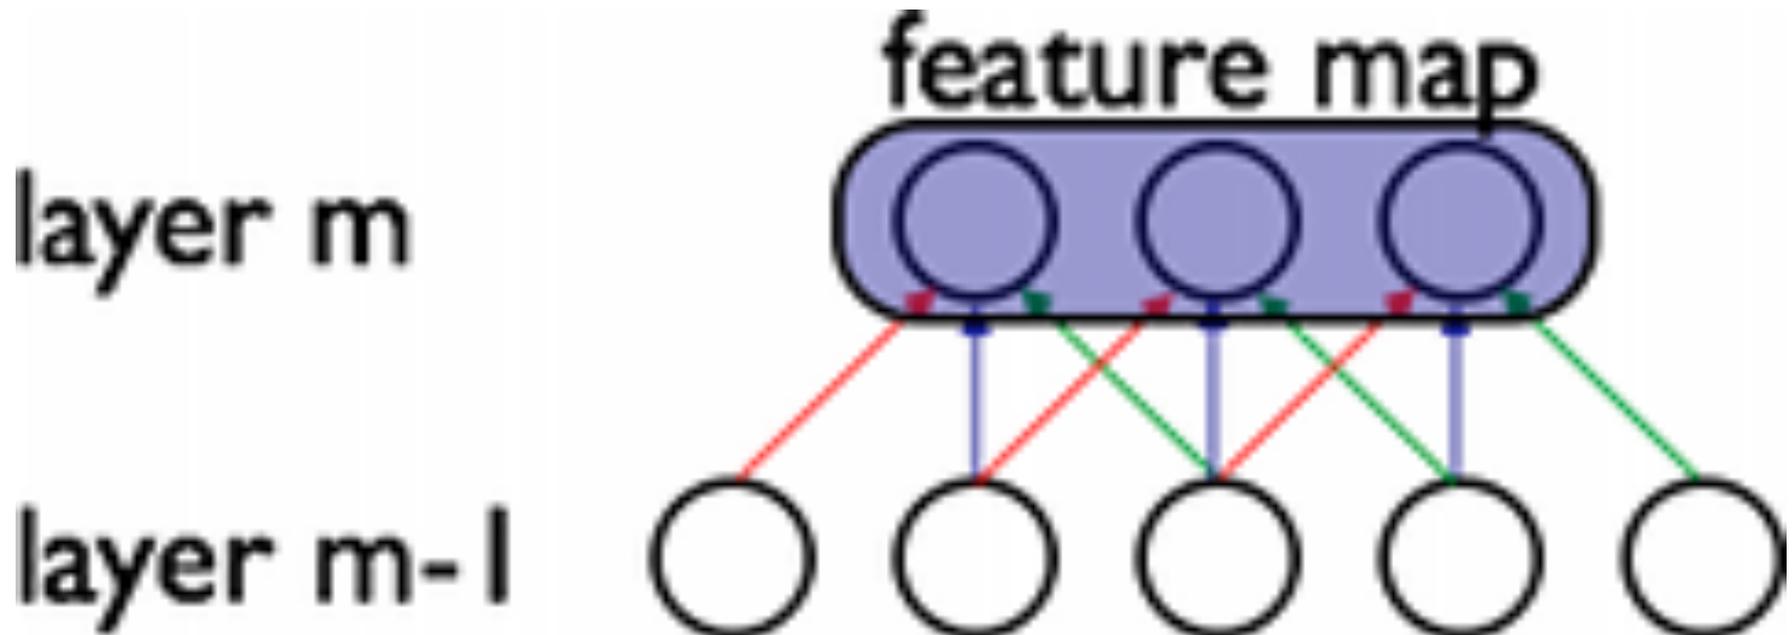
layer  $m$

layer  $m-1$



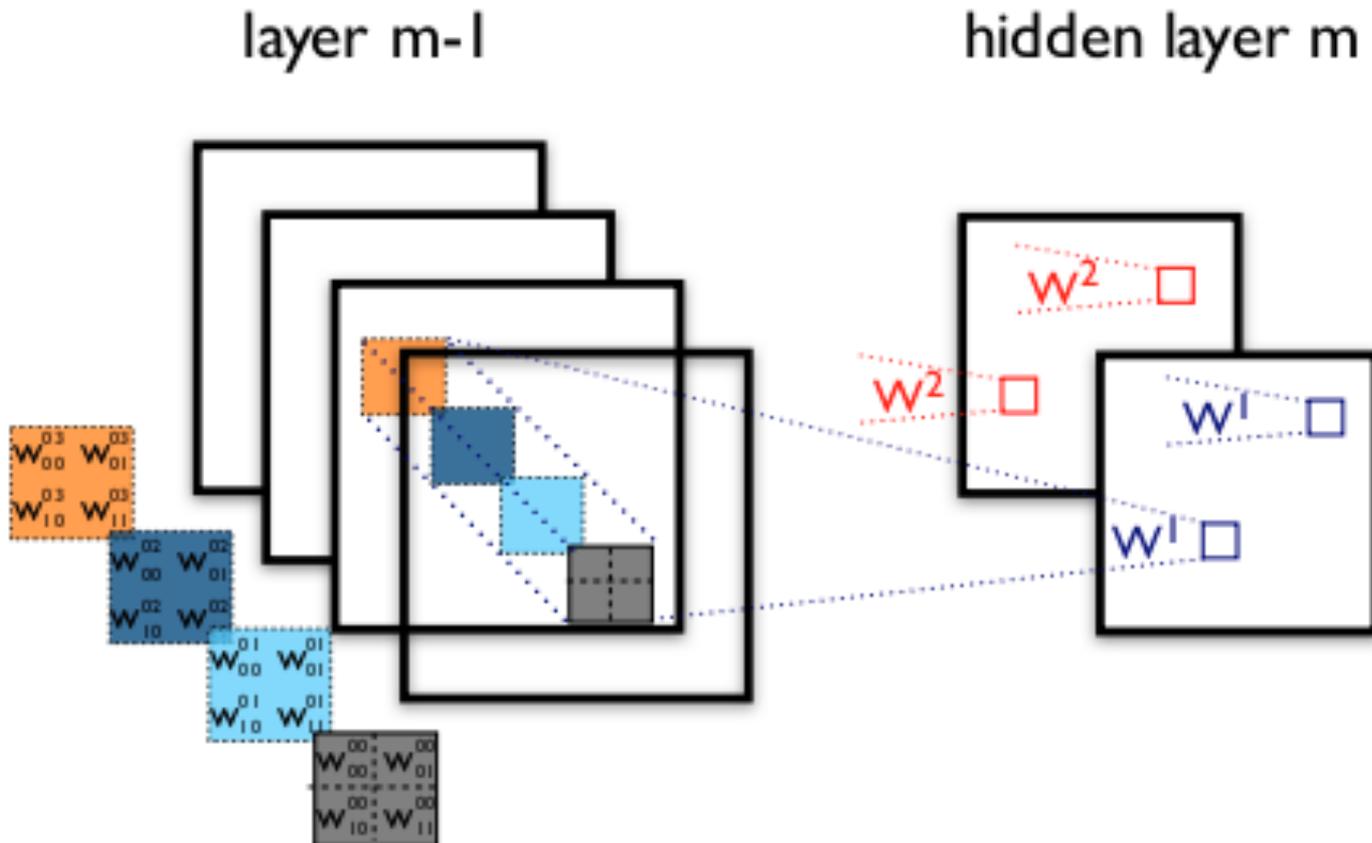
# Convolutional Neural Networks (CNN) (LeNet)

## Shared Weights



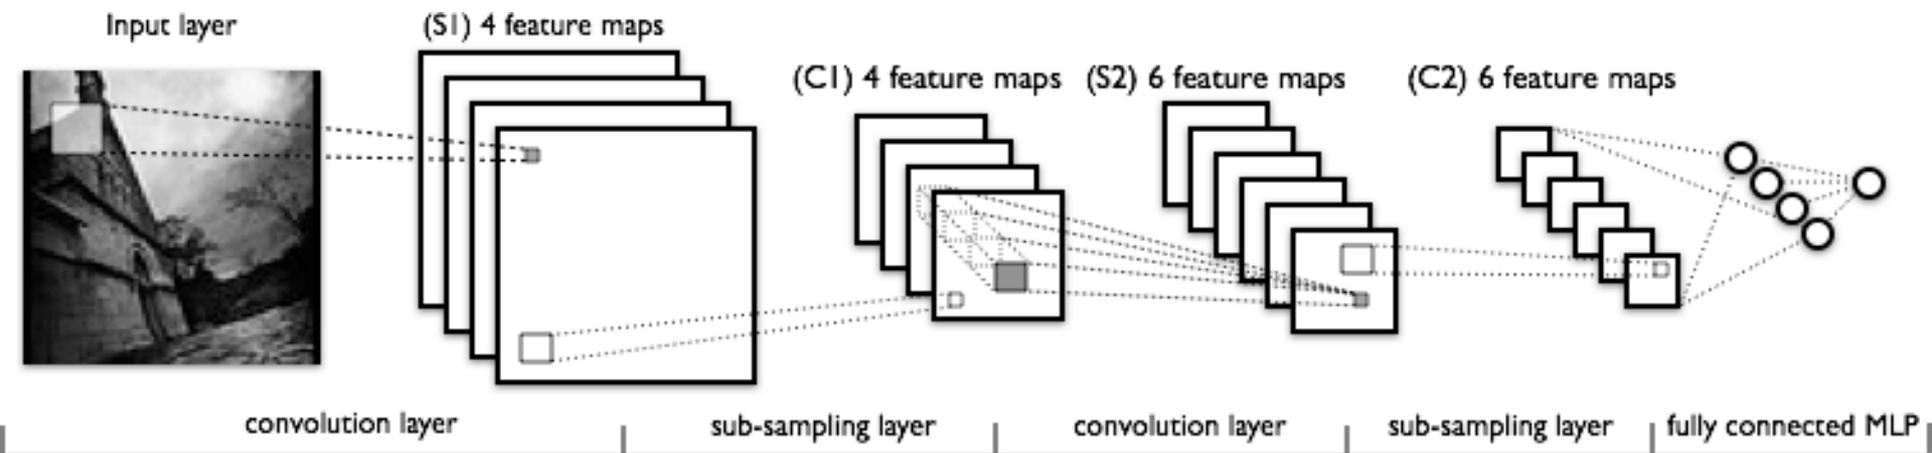
# Convolutional Neural Networks (CNN) (LeNet)

example of a convolutional layer



Source: <http://deeplearning.net/tutorial/lenet.html>

# Convolutional Neural Networks (CNN) (LeNet)



**show flights from Boston to New York today**

# Recurrent Neural Networks with Word Embeddings

## Semantic Parsing / Slot-Filling (Spoken Language Understanding)

Input (words)	show	flights	from	Boston	to	New	York	today
Output (labels)	O	O	O	B-dept	O	B-arr	I-arr	B-date

**show flights from Boston to New York today**

**show flights from Boston to New York today**

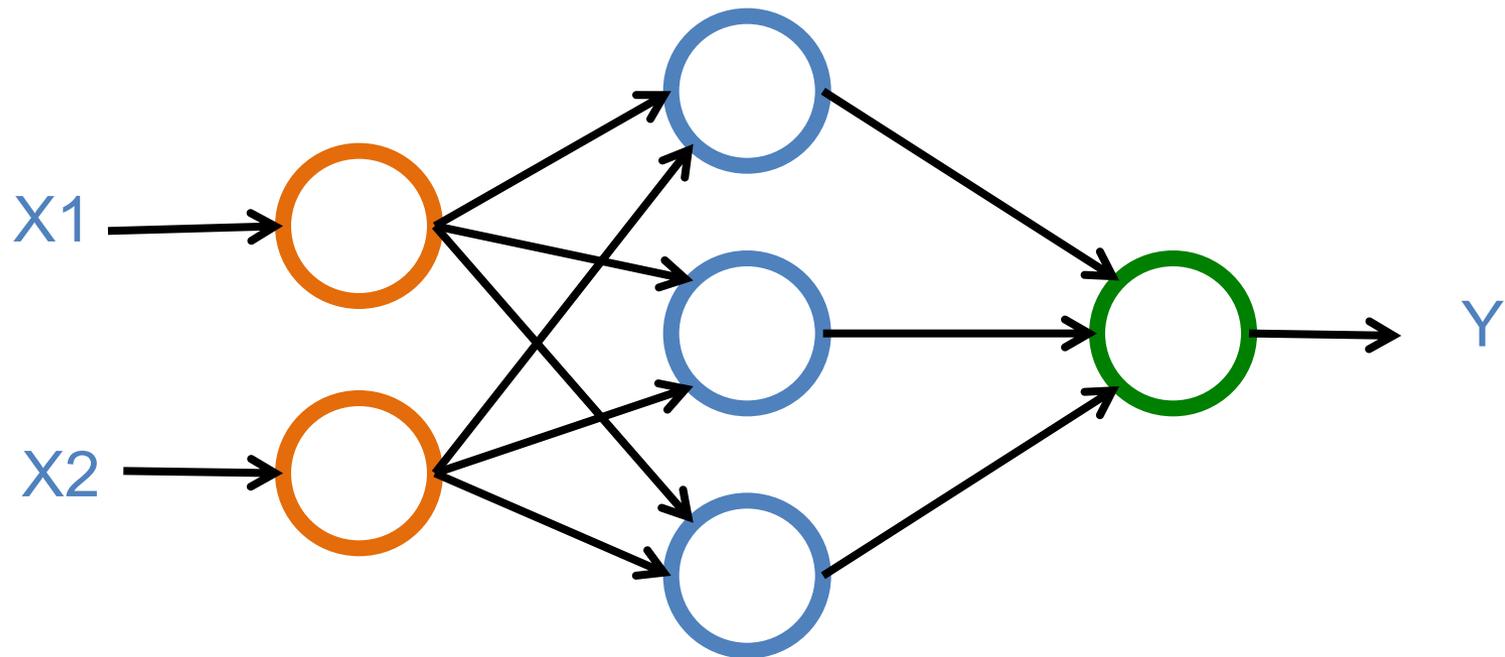
Input (words)	show	flights	from	Boston	to	New	York	today
Output (labels)	0	0	0	B-dept	0	B-arr	I-arr	B-date

# Neural Networks

**Input Layer**  
(X)

**Hidden Layer**  
(H)

**Output Layer**  
(Y)



<b>X</b>		<b>Y</b>
<b>Hours Sleep</b>	<b>Hours Study</b>	<b>Score</b>
<b>3</b>	<b>5</b>	<b>75</b>
<b>5</b>	<b>1</b>	<b>82</b>
<b>10</b>	<b>2</b>	<b>93</b>
<b>8</b>	<b>3</b>	<b>?</b>

	<b>X</b>		<b>Y</b>
	<b>Hours Sleep</b>	<b>Hours Study</b>	<b>Score</b>
<b>Training</b>	<b>3</b>	<b>5</b>	<b>75</b>
	<b>5</b>	<b>1</b>	<b>82</b>
	<b>10</b>	<b>2</b>	<b>93</b>
<b>Testing</b>	<b>8</b>	<b>3</b>	<b>?</b>

**Training a Network**  
**=**  
**Minimize the Cost Function**

# Training a Network

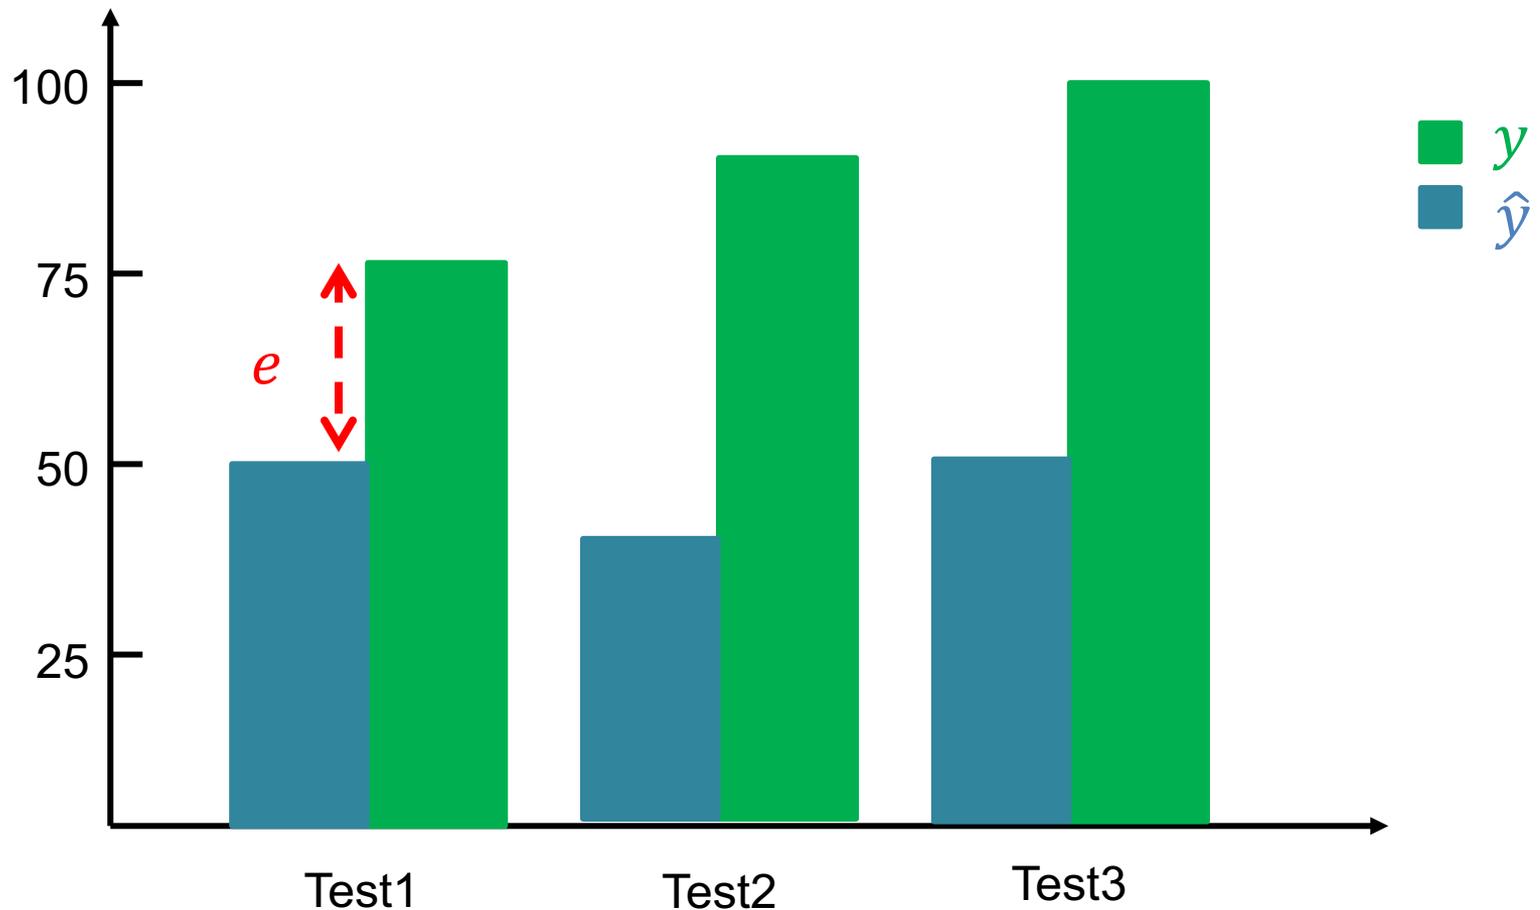
=

Minimize the **Cost** Function

Minimize the **Loss** Function

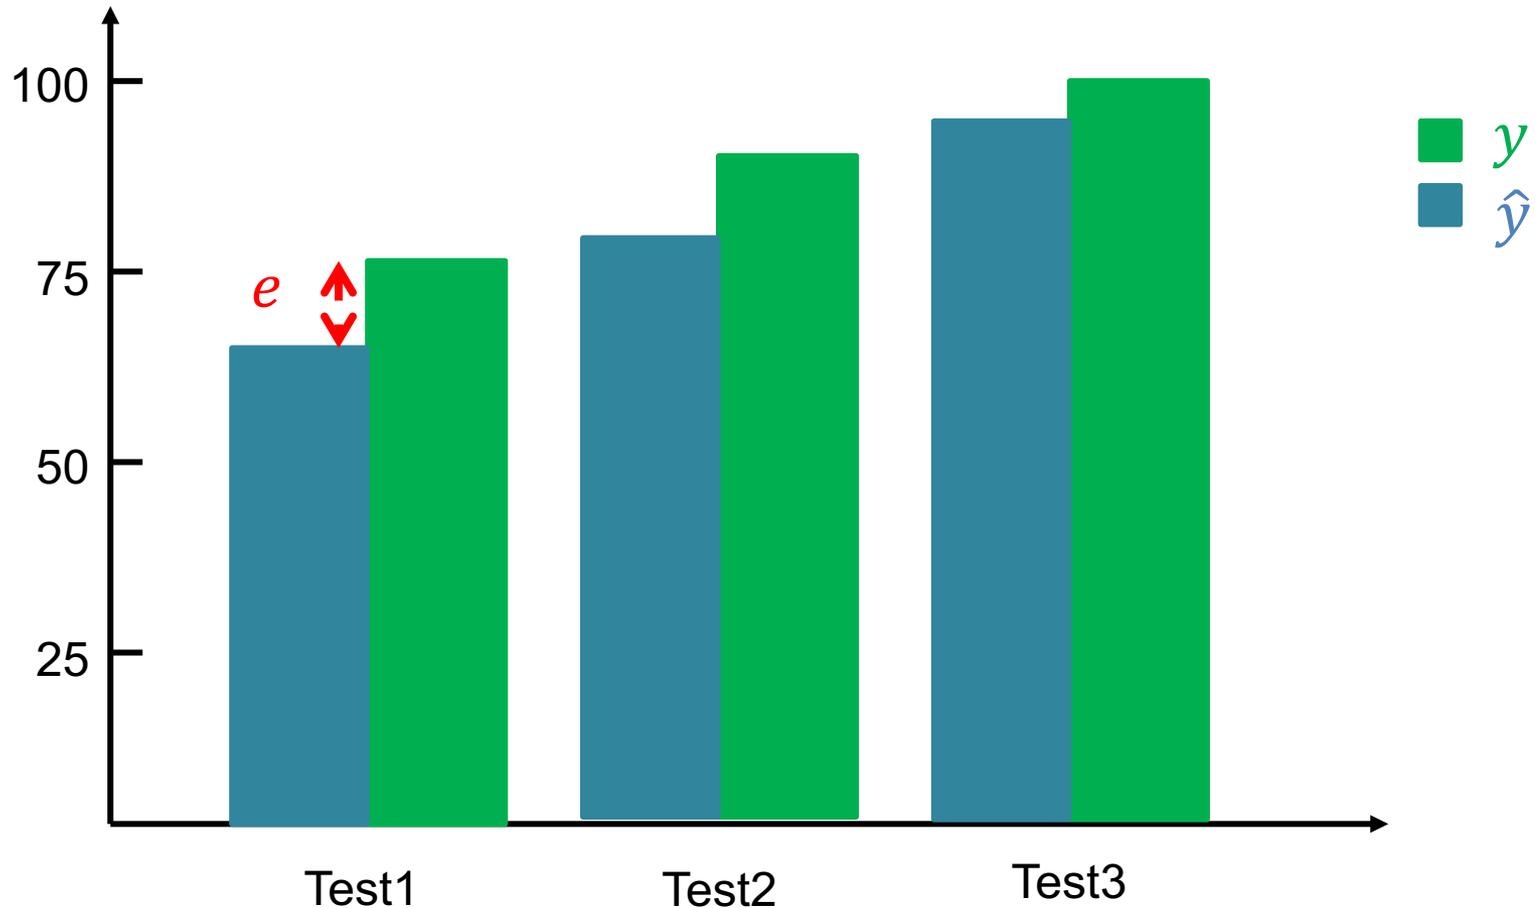
**Error = Predict Y - Actual Y**

**Error : Cost : Loss**



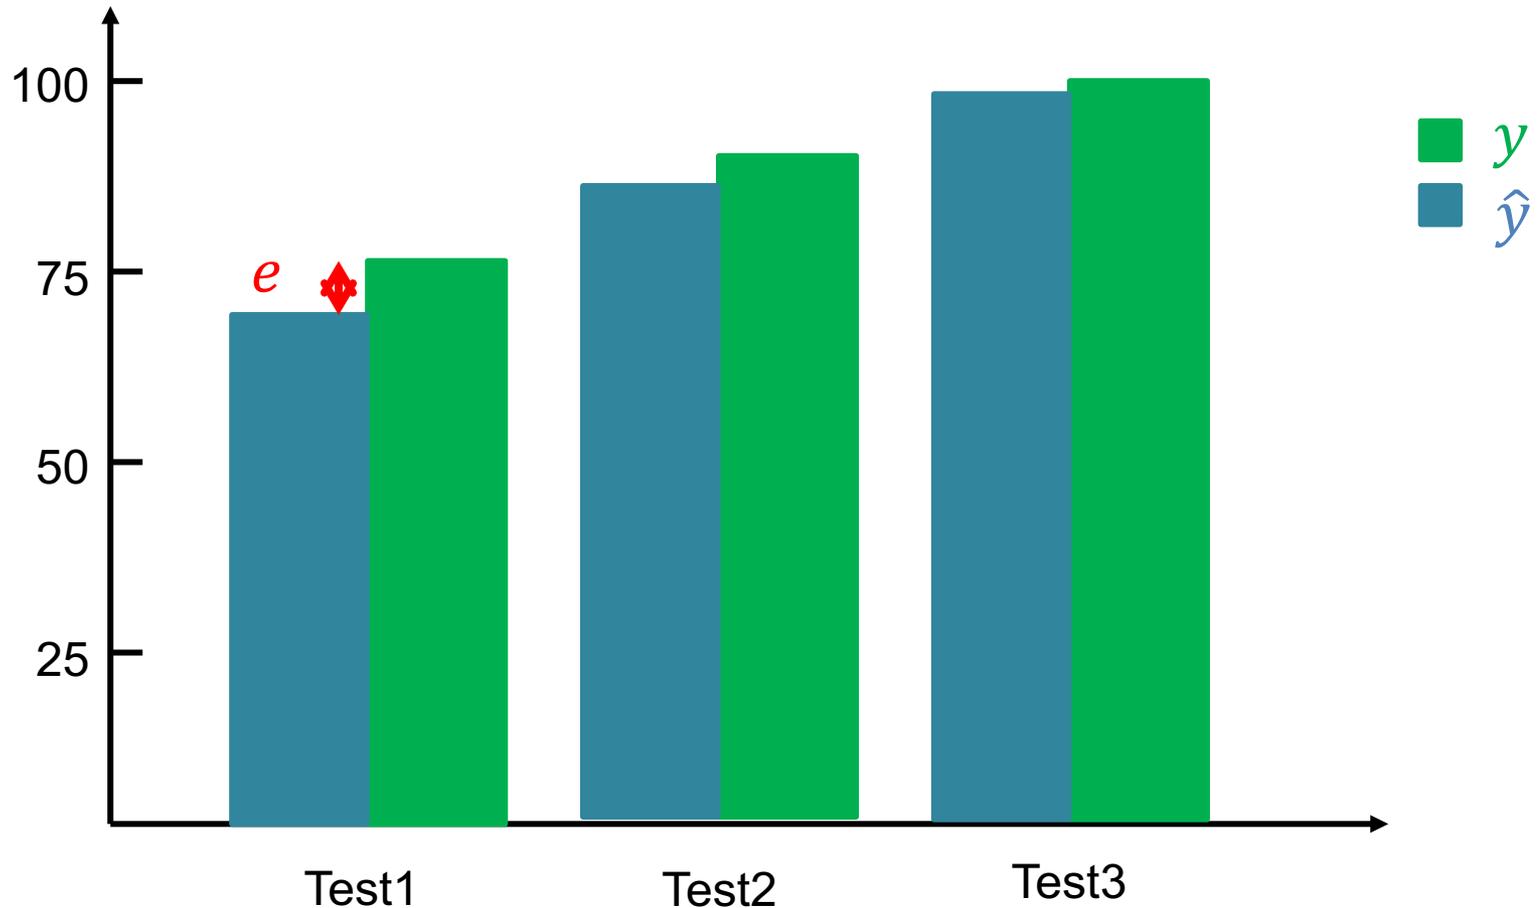
**Error = Predict Y - Actual Y**

**Error : Cost : Loss**



**Error = Predict Y - Actual Y**

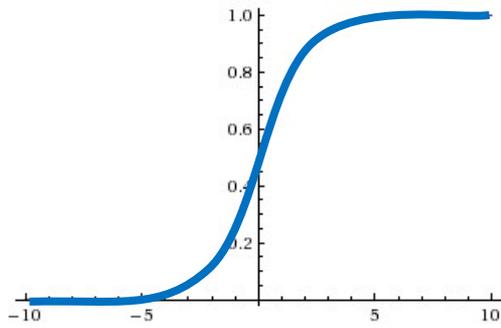
**Error : Cost : Loss**



# Activation Functions

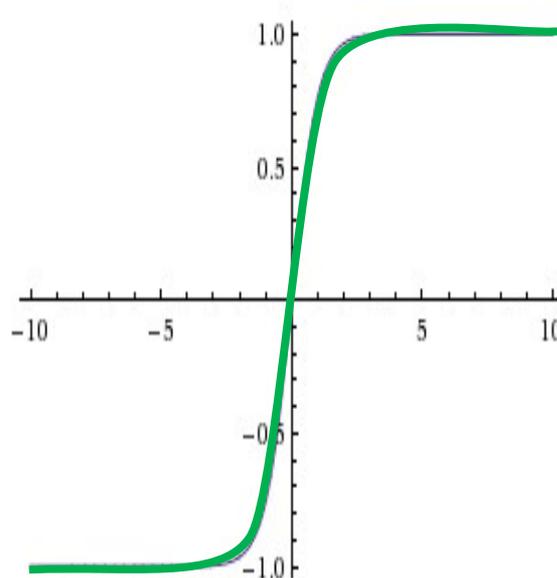
# Activation Functions

**Sigmoid**



**[0, 1]**

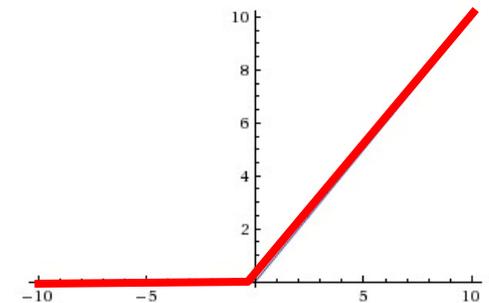
**TanH**



**[-1, 1]**

**ReLU**

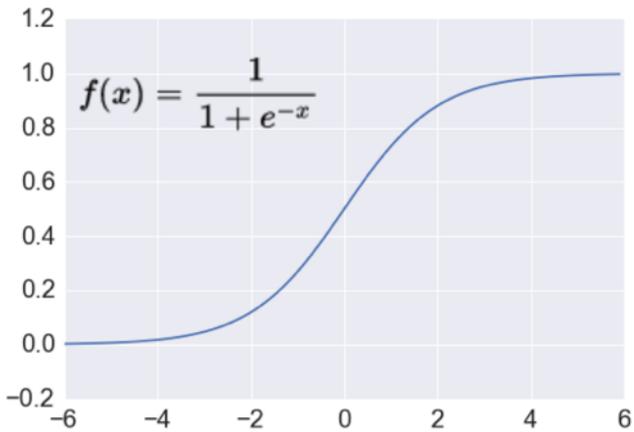
(Rectified Linear Unit)



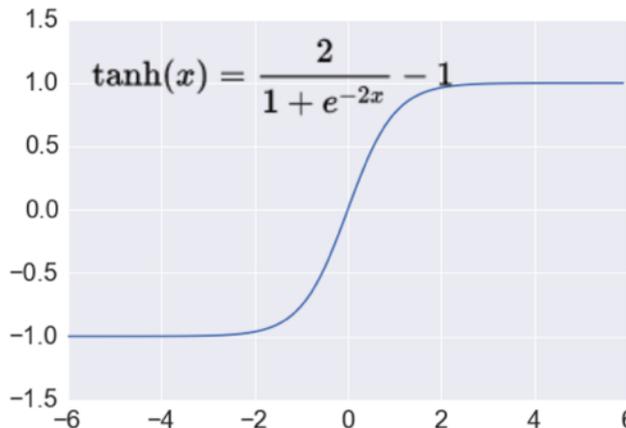
**$f(x) = \max(0, x)$**

# Activation Functions

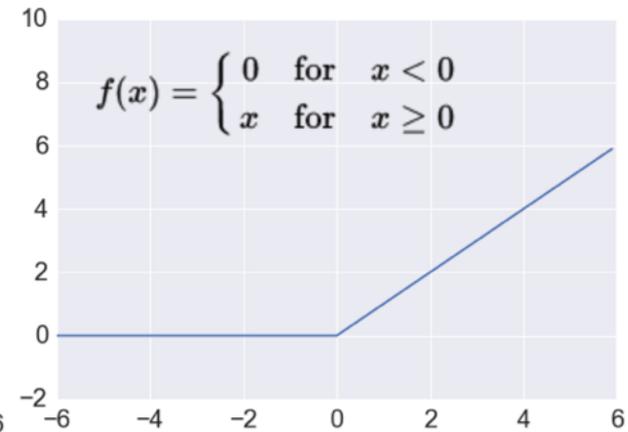
Sigmoid



TanH



ReLU



# Loss Function

# Binary Classification: 2 Class

**Activation Function:  
Sigmoid**

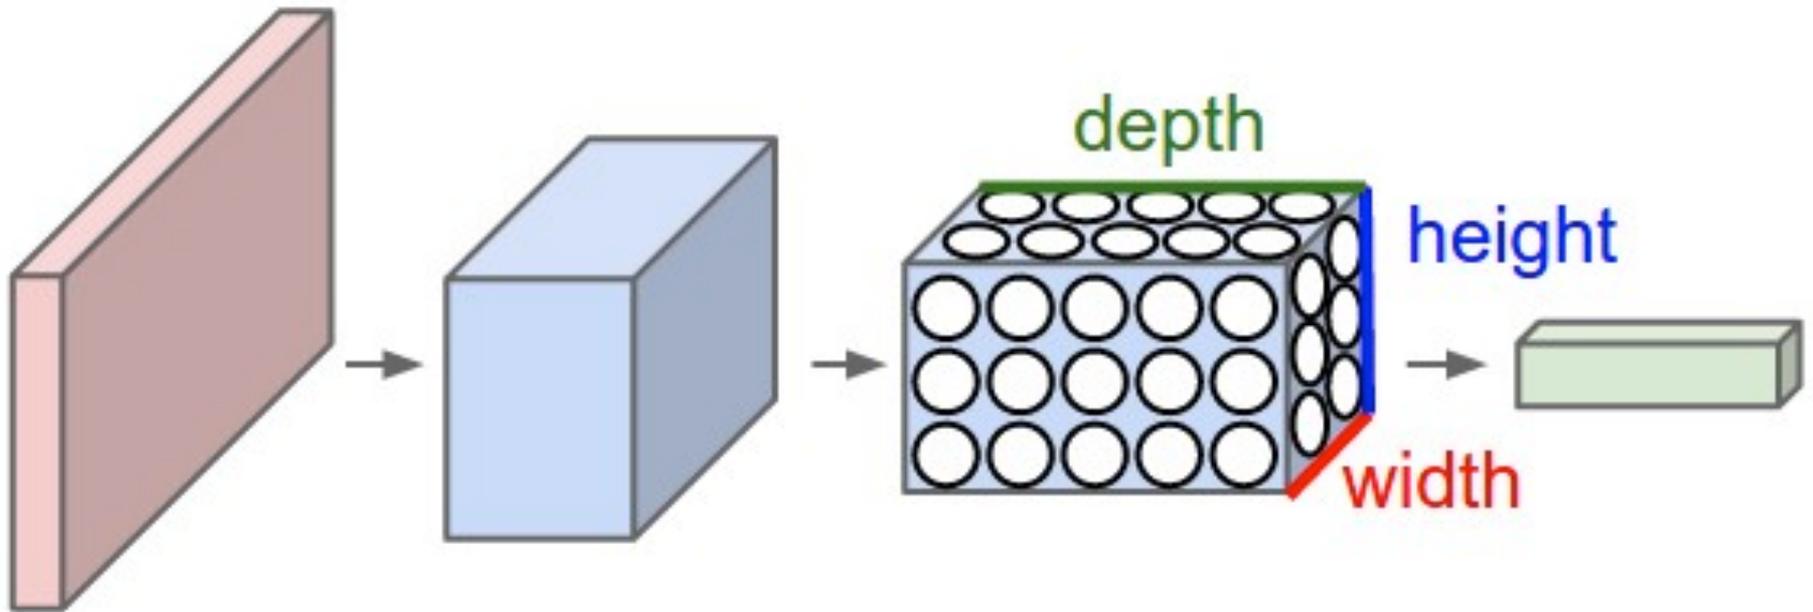
**Loss Function:  
Binary Cross-Entropy**

**Multiple Classification: 10 Class**

**Activation Function:  
SoftMAX**

**Loss Function:  
Categorical Cross-Entropy**

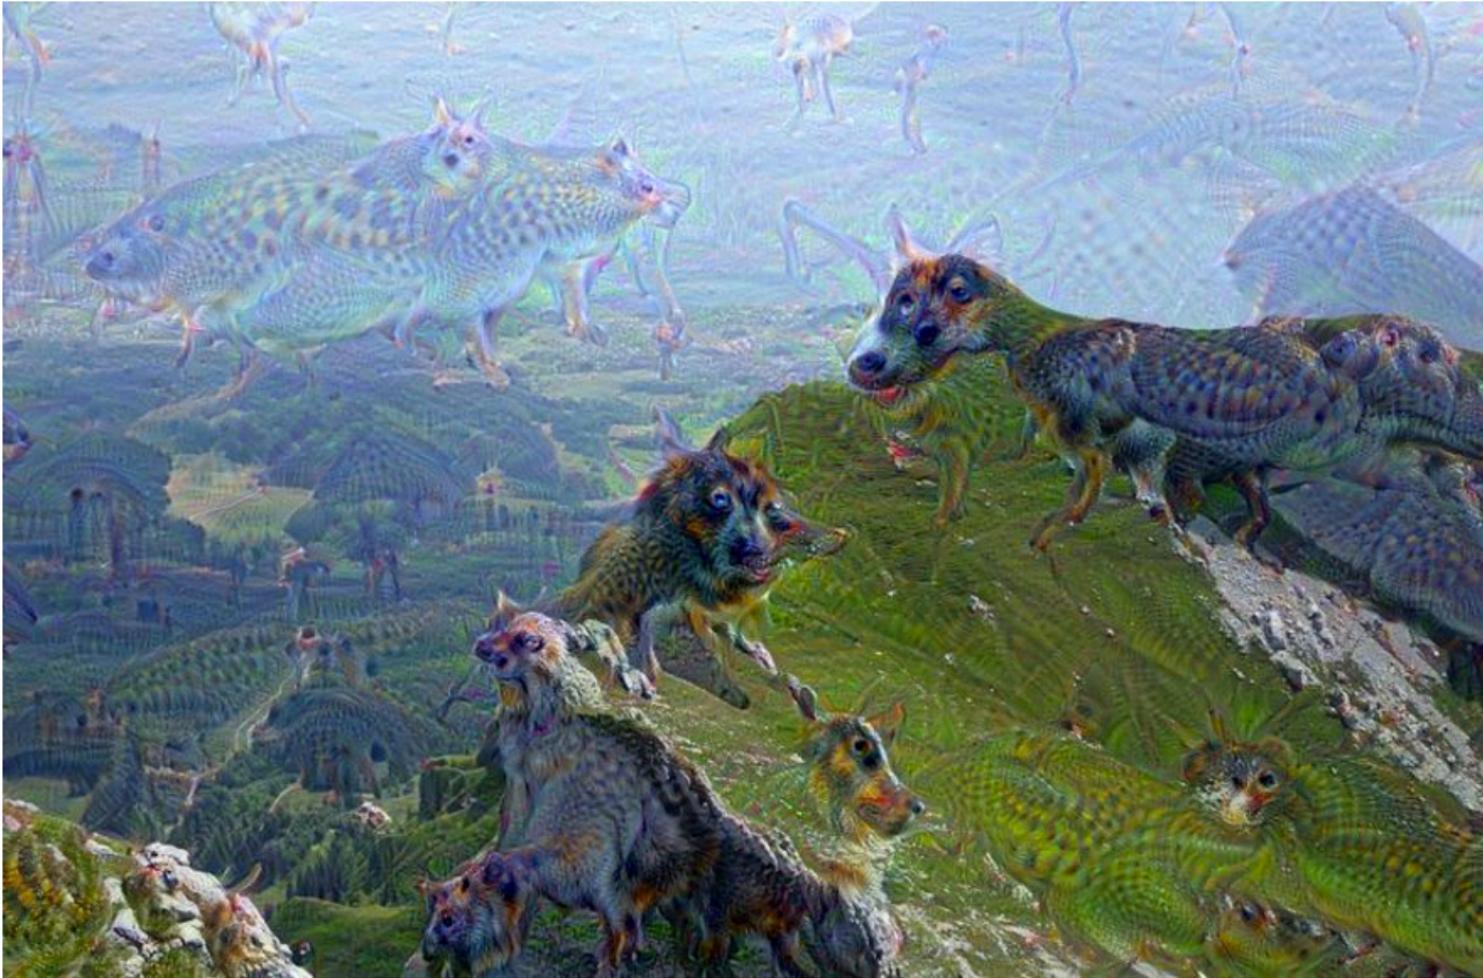
# A ConvNet arranges its neurons in three dimensions (width, height, depth)



# DeepDream

GitHub, Inc. [US] <https://github.com/tensorflow/tensorflow/blob/master/tensorflow/examples/tutorials/deepdream/deepdream.ipynb>

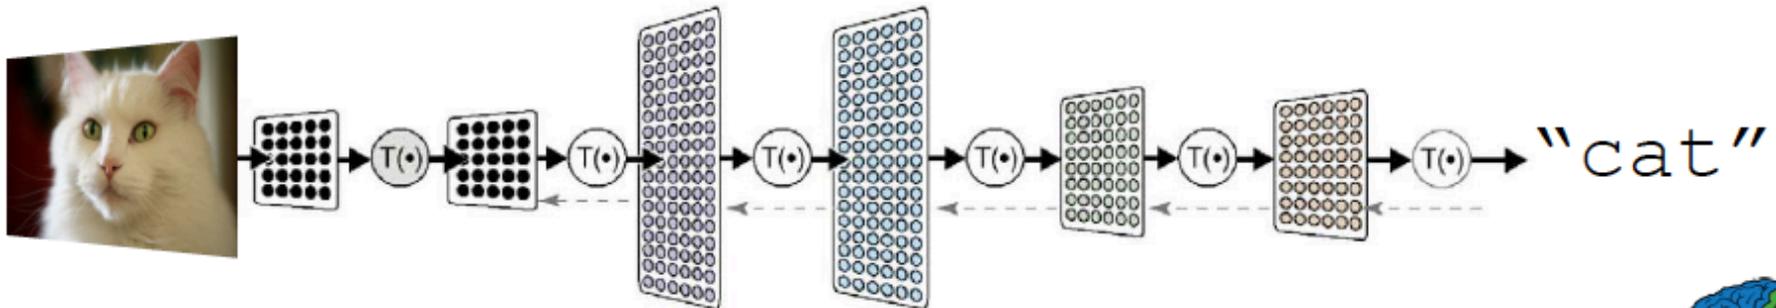
```
In [15]: render_deepdream(tf.square(T('mixed4c')), img0)
```



Note that results can differ from the [Caffe's](#) implementation, as we are using an independently trained network. Still, the network seems to like dogs and animal-like features due to the nature of the ImageNet dataset.

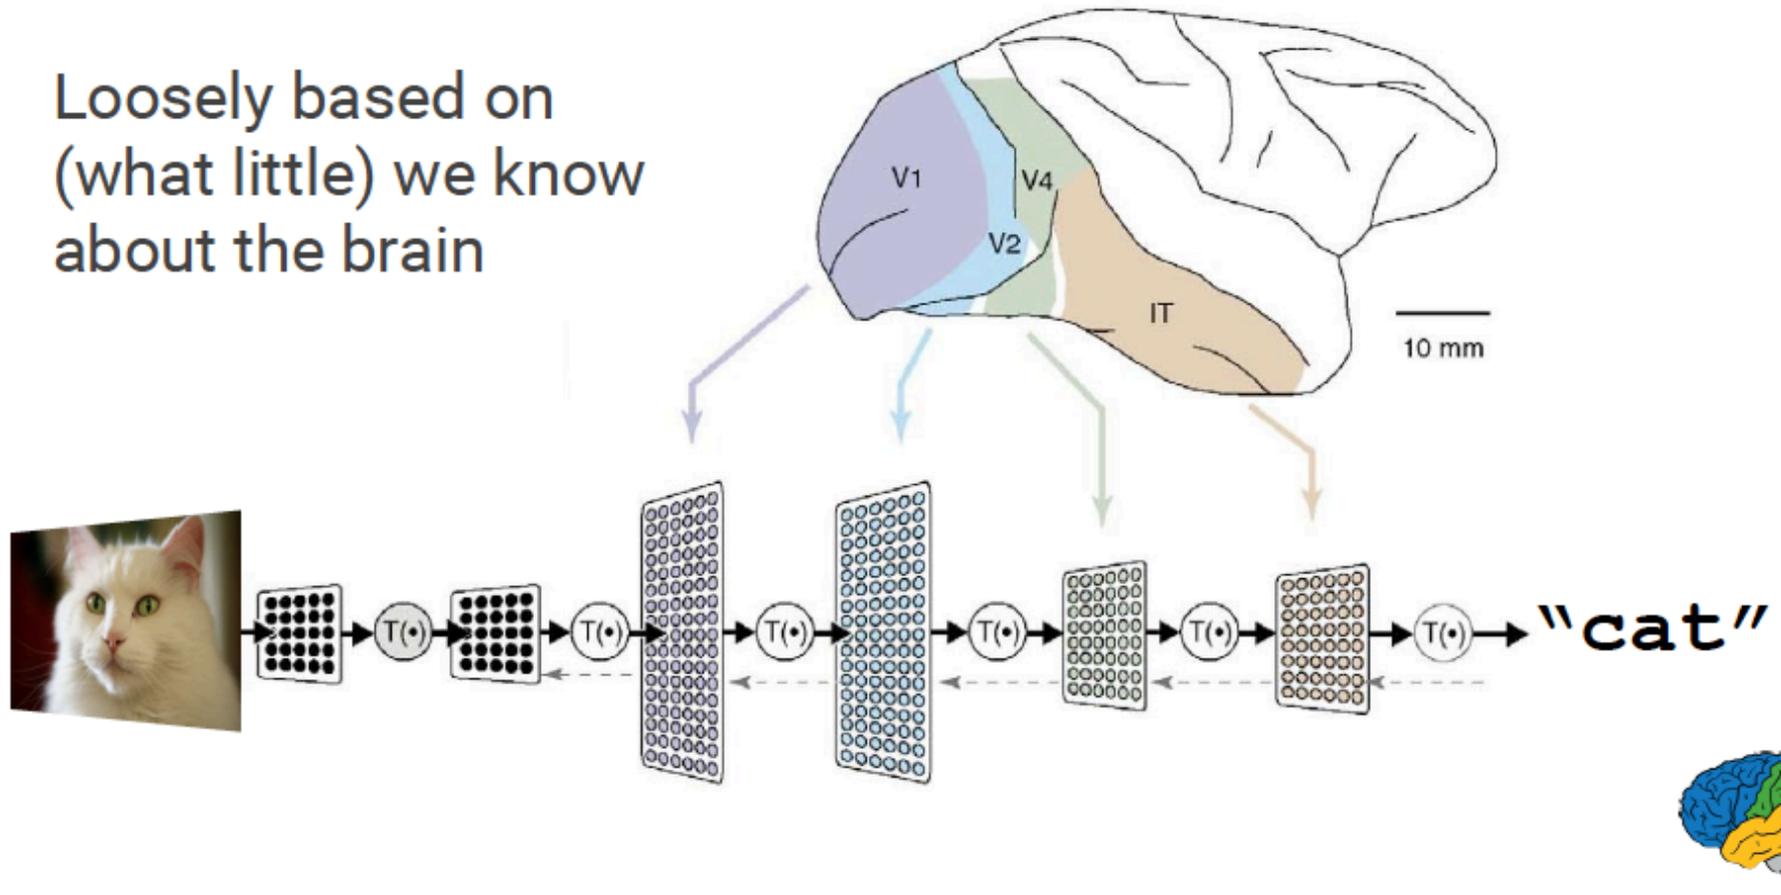
# Deep Learning

- 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

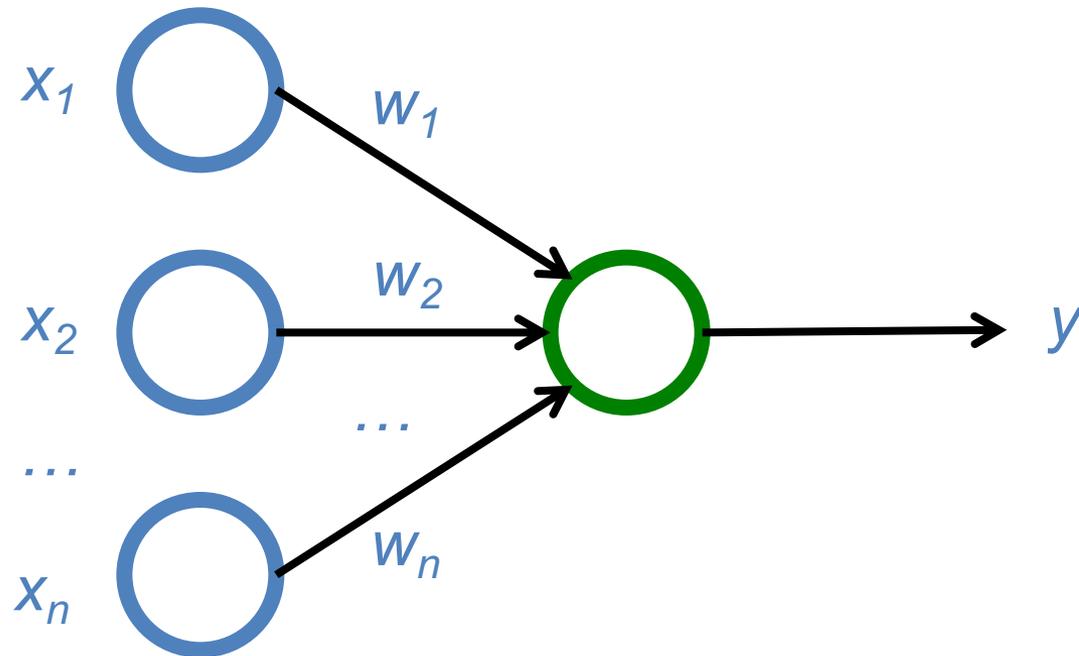


# What is Deep Learning?

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

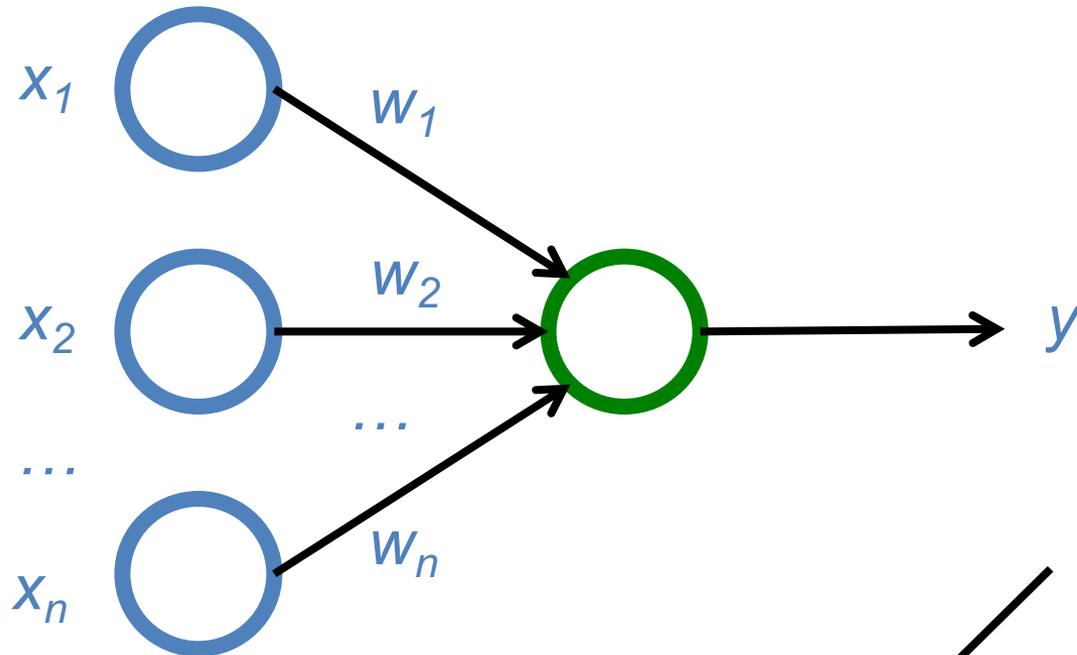


# The Neuron



# The Neuron

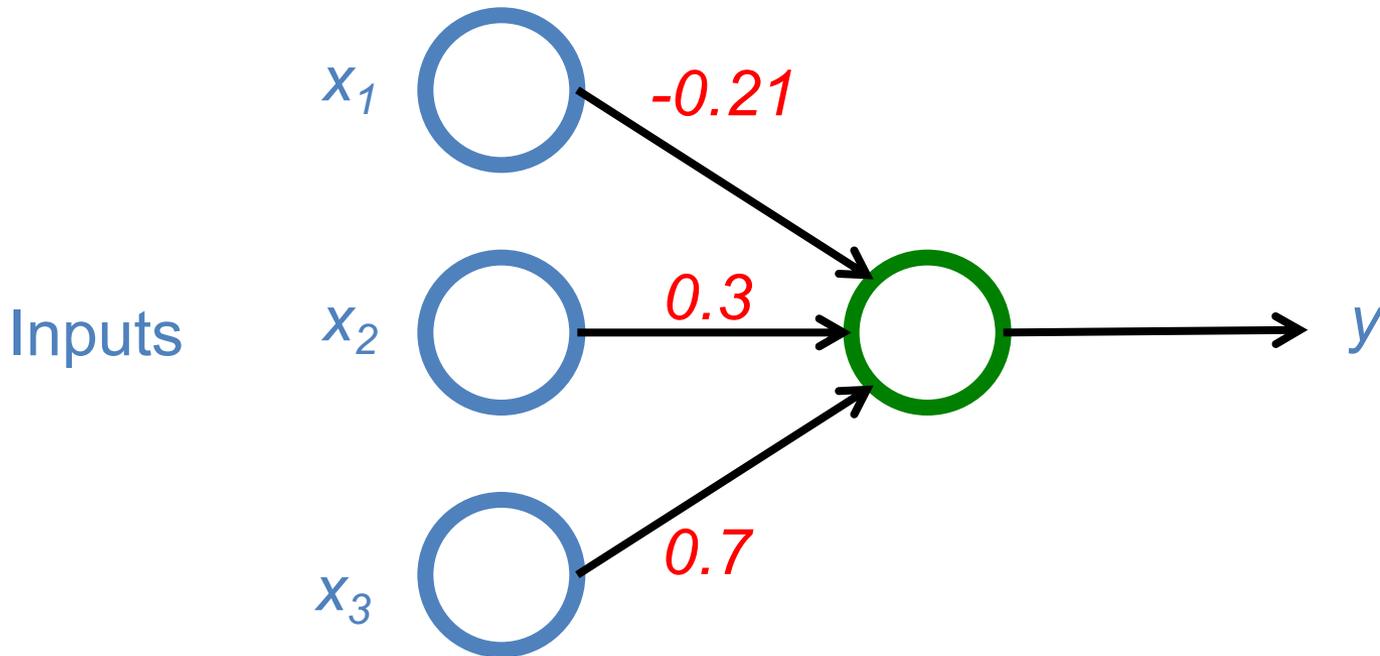
$$y = F\left(\sum_i w_i x_i\right)$$



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

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

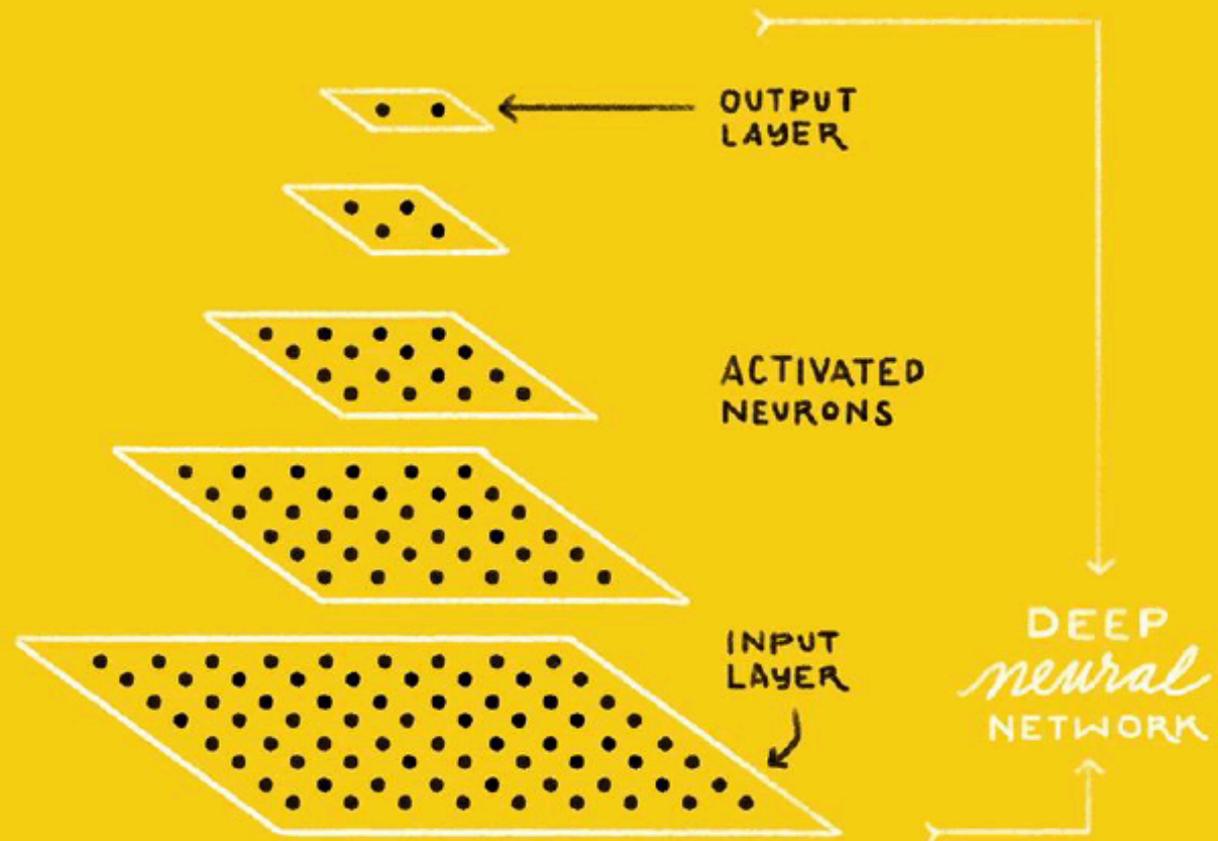
Weights



IS THIS A  
**CAT** or **DOG**?



**CAT**   **DOG**



# Learning Algorithm

While not done:

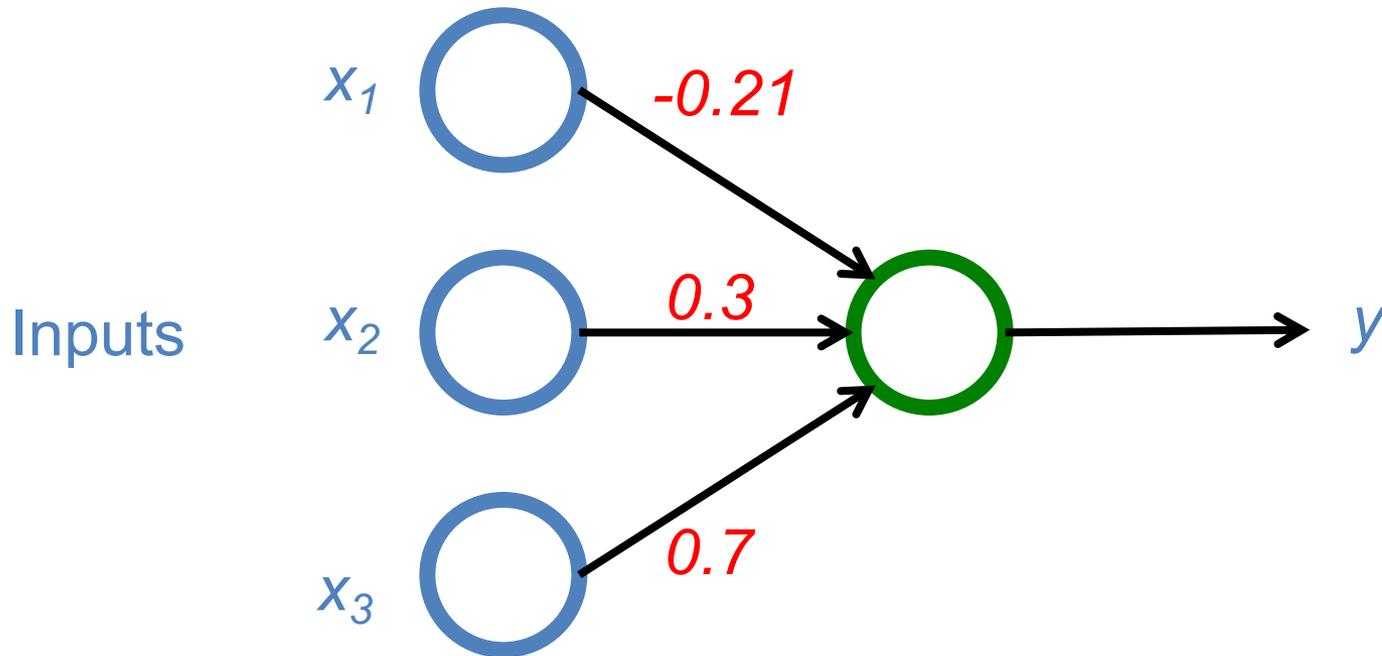
Pick a random training example “(input, label)”

Run neural network on “input”

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

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

Weights

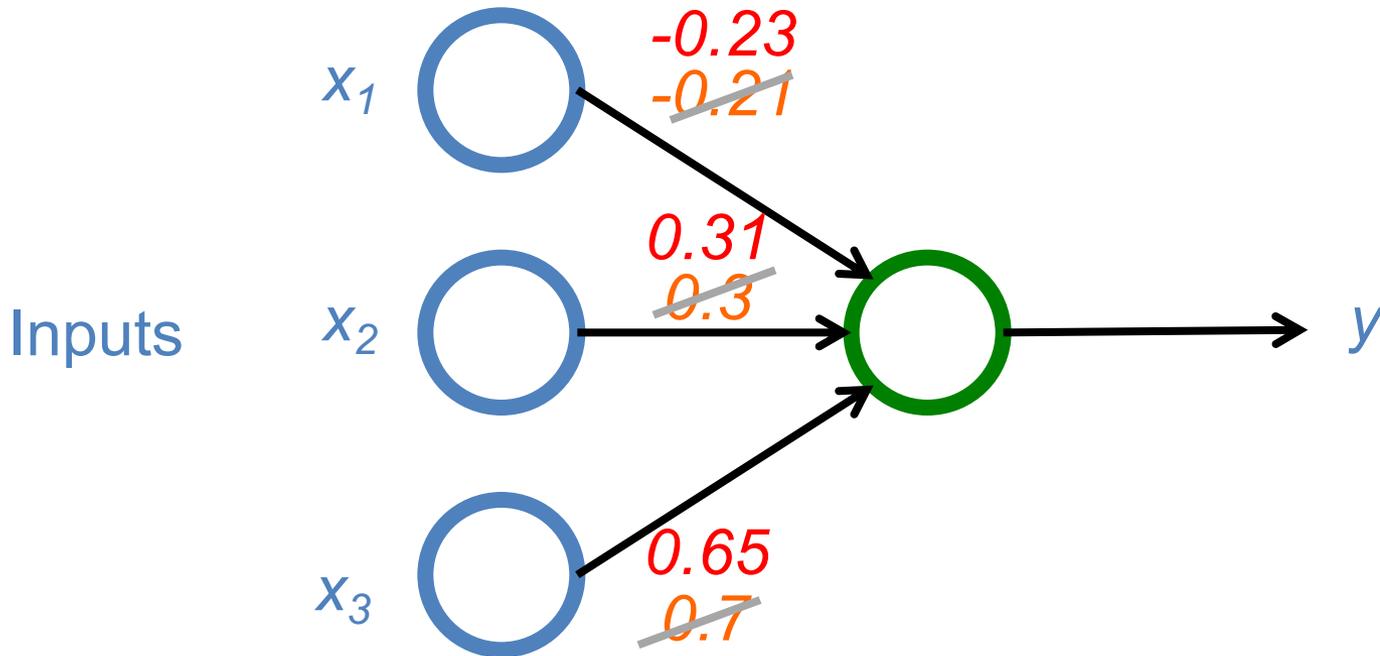


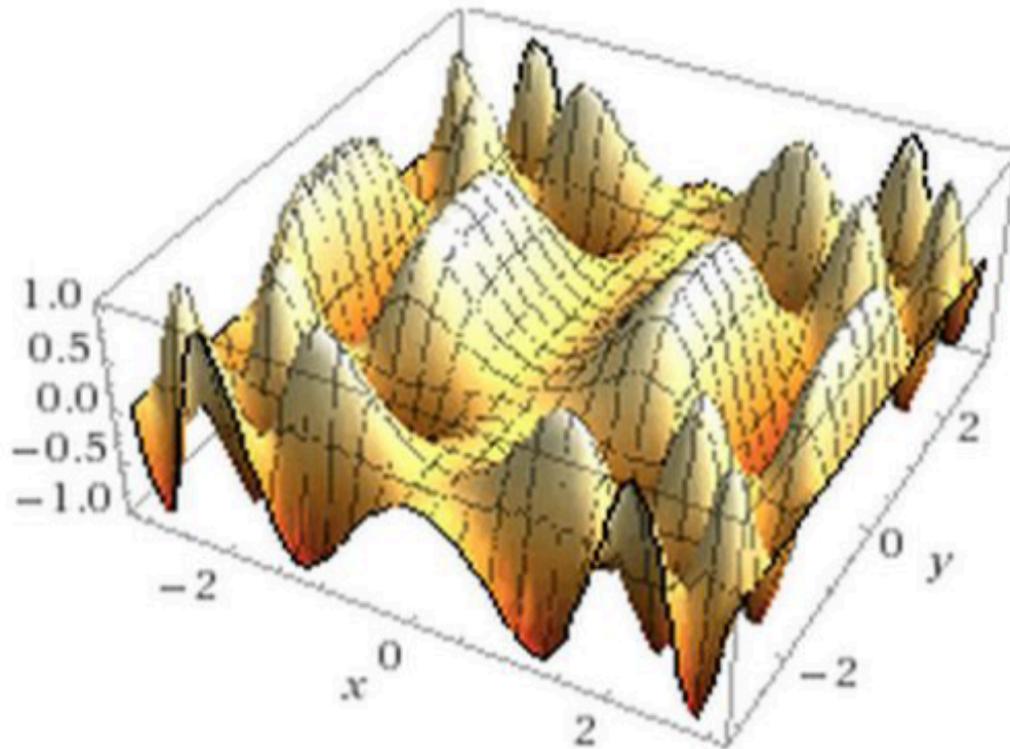
Next time:

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

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

Weights





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

# Important Property of Neural Networks

Results get better with

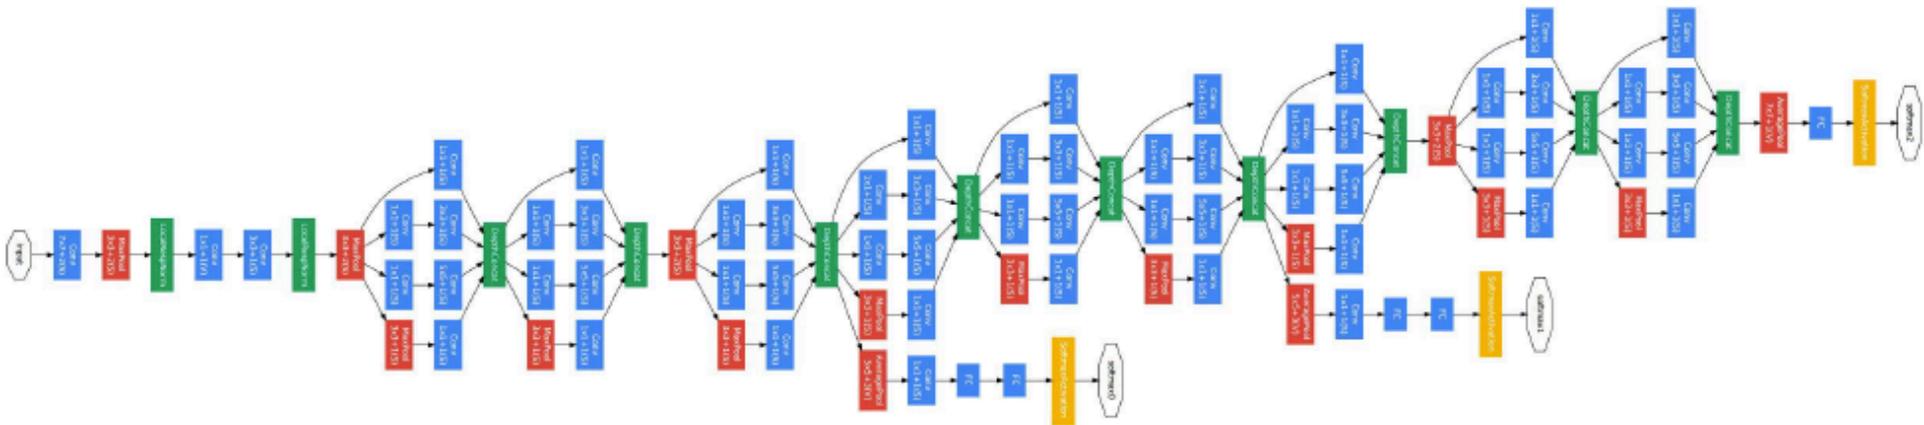
**More data +**

**Bigger models +**

**More computation**

(Better algorithms, new insights  
and improved techniques always help, too!)

# The Inception Architecture (GoogLeNet, 2014)



## Going Deeper with Convolutions

Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, Andrew Rabinovich

ArXiv 2014, CVPR 2015



# Deep Learning Software

- Keras
  - Deep Learning library for **Theano** and **TensorFlow**
- Tensorflow
  - TensorFlow™ is an open source software library for numerical computation using data flow graphs.
- Theano
  - CPU/GPU symbolic expression compiler in python (from MILA lab at University of Montreal)

**Deep Learning**

**with**

**Google**

**TensorFlow**

# Deep Learning Libraries: Tensorflow and Keras

## Deep learning libraries: GitHub activity from February 11 to April 12, 2017

new contributors from 2017-02-11 to 2017-04-12			new forks from 2017-02-11 to 2017-04-12		
#1: 131		tensorflow/tensorflow	#1: 4192		tensorflow/tensorflow
#2: 63		fchollet/keras	#2: 991		fchollet/keras
#3: 51		pytorch/pytorch	#3: 810		BVLC/caffe
#4: 49		dmlc/mxnet	#4: 517		deeplearning4j/deeplearning4j
#5: 18		Theano/Theano	#5: 414		dmlc/mxnet
#6: 11		BVLC/caffe	#6: 307		pytorch/pytorch
#7: 11		Microsoft/CNTK	#7: 244		Microsoft/CNTK
#8: 9		tflearn/tflearn	#8: 211		tflearn/tflearn
#9: 9		pfnet/chainer	#9: 134		torch/torch7
#10: 8		deeplearning4j/deeplearning4j	#10: 131		Theano/Theano
#11: 5		NVIDIA/DIGITS	#11: 116		baidu/paddle
#12: 4		baidu/paddle	#12: 88		NVIDIA/DIGITS
#13: 3			#13: 55		pfnet/chainer

new issues from 2017-02-11 to 2017-04-12			aggregate activity from 2017-02-11 to 2017-04-12		
#1: 1175		tensorflow/tensorflow	#1: 36.64		tensorflow/tensorflow
#2: 568		fchollet/keras	#2: 12.52		fchollet/keras
#3: 499		dmlc/mxnet	#3: 8.53		dmlc/mxnet
#4: 286		pytorch/pytorch	#4: 6.09		BVLC/caffe
#5: 257		Microsoft/CNTK	#5: 5.92		pytorch/pytorch
#6: 239		deeplearning4j/deeplearning4j	#6: 5.12		deeplearning4j/deeplearning4j
#7: 219		baidu/paddle	#7: 4.12		Microsoft/CNTK
#8: 173		Theano/Theano	#8: 2.93		Theano/Theano
#9: 171		BVLC/caffe	#9: 2.86		baidu/paddle
#10: 112		NVIDIA/DIGITS	#10: 2.17		tflearn/tflearn
#11: 84		tflearn/tflearn	#11: 1.68		NVIDIA/DIGITS
#12: 57		pfnet/chainer	#12: 1.38		torch/torch7
#13: 47		torch/torch7	#13: 1.12		pfnet/chainer



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**  
is an  
**Open Source**  
**Software Library**  
for  
**Machine Intelligence**

# numerical computation using data flow graphs

# Tensor

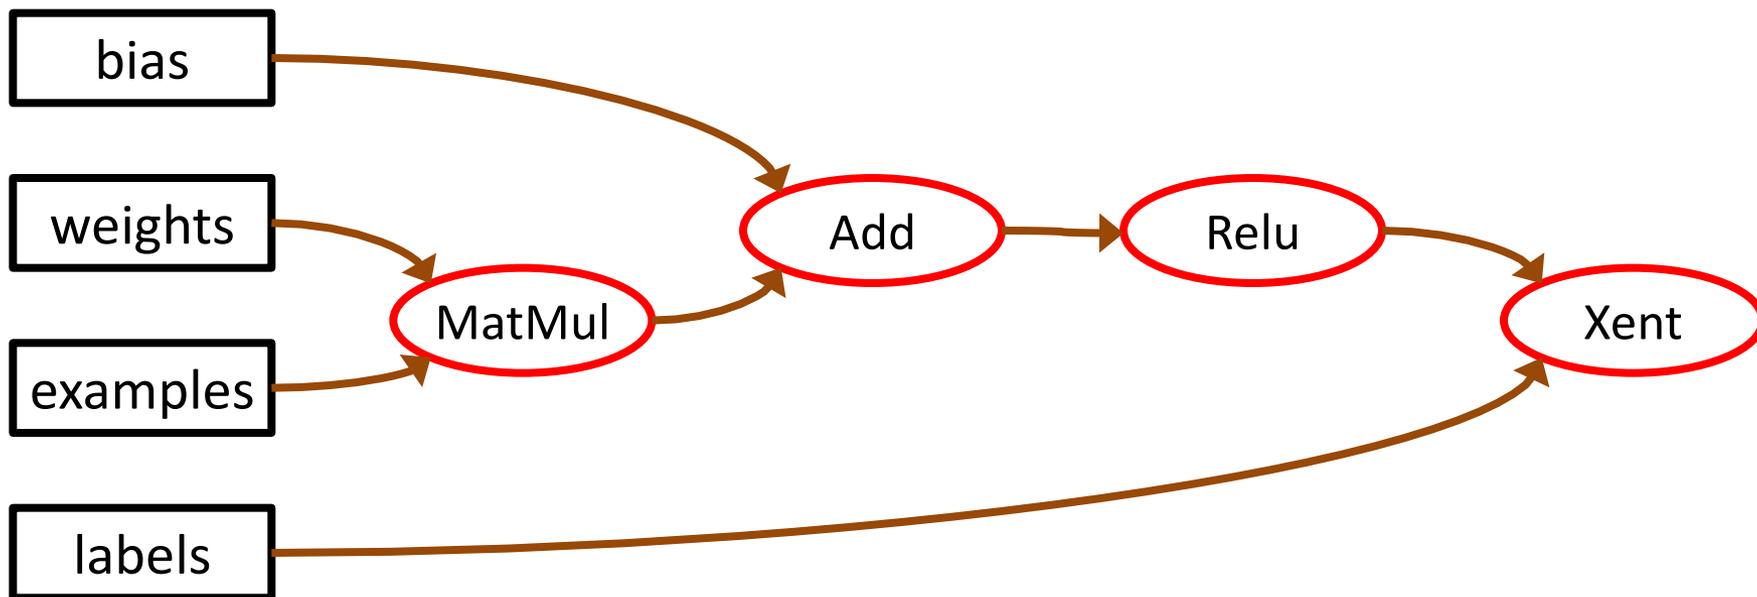
- **3**
  - # a rank 0 tensor; this is a **scalar** with shape []
- **[1., 2., 3.]**
  - # a rank 1 tensor; this is a **vector** with shape [3]
- **[[1., 2., 3.], [4., 5., 6.]]**
  - # a rank 2 tensor; a **matrix** with shape [2, 3]
- **[[[1., 2., 3.]], [[7., 8., 9.]]]**
  - # a rank 3 **tensor** with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

**Nodes:**  
**mathematical operations**

**edges:**  
**multidimensional data arrays**  
**(tensors)**  
**communicated between nodes**

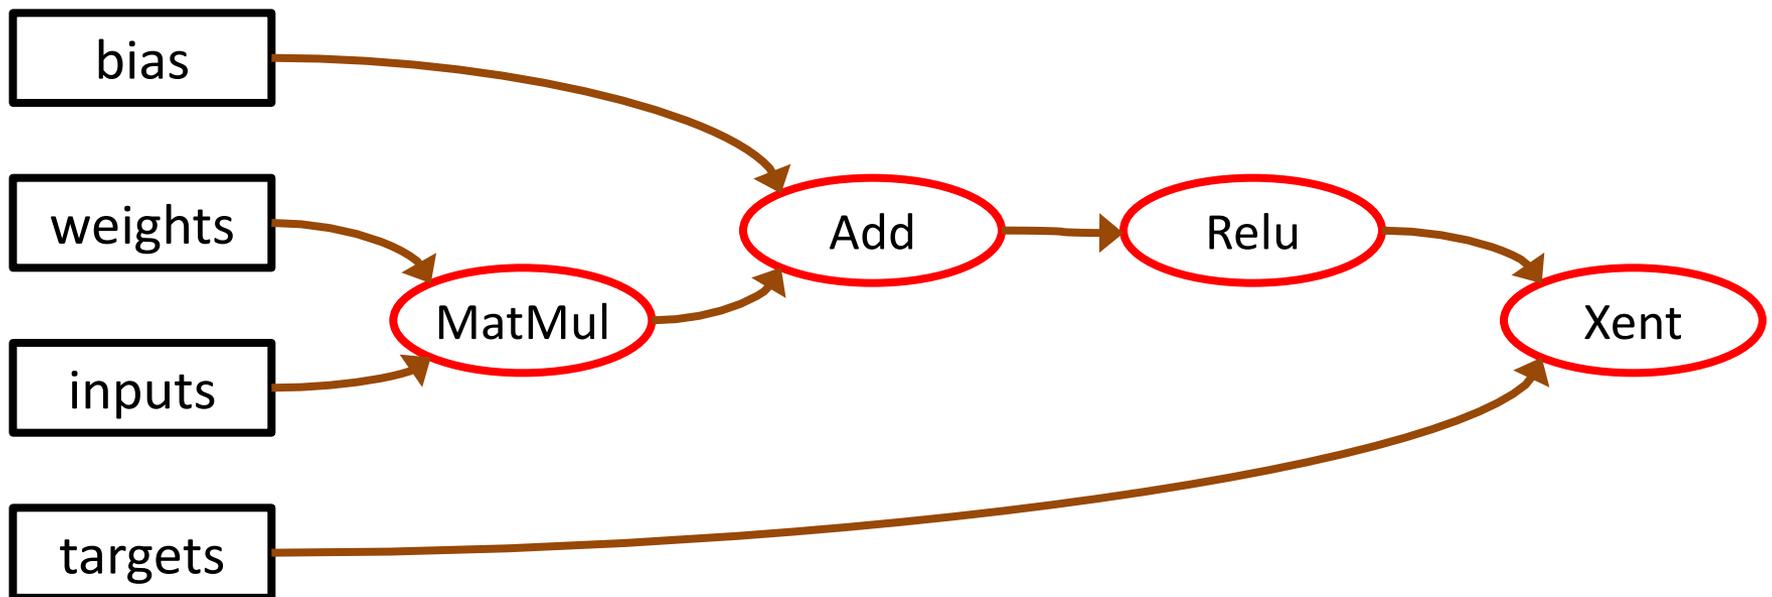
# Computation is a Dataflow Graph

Graph of **Nodes**,  
also called **Operations** or **ops**.

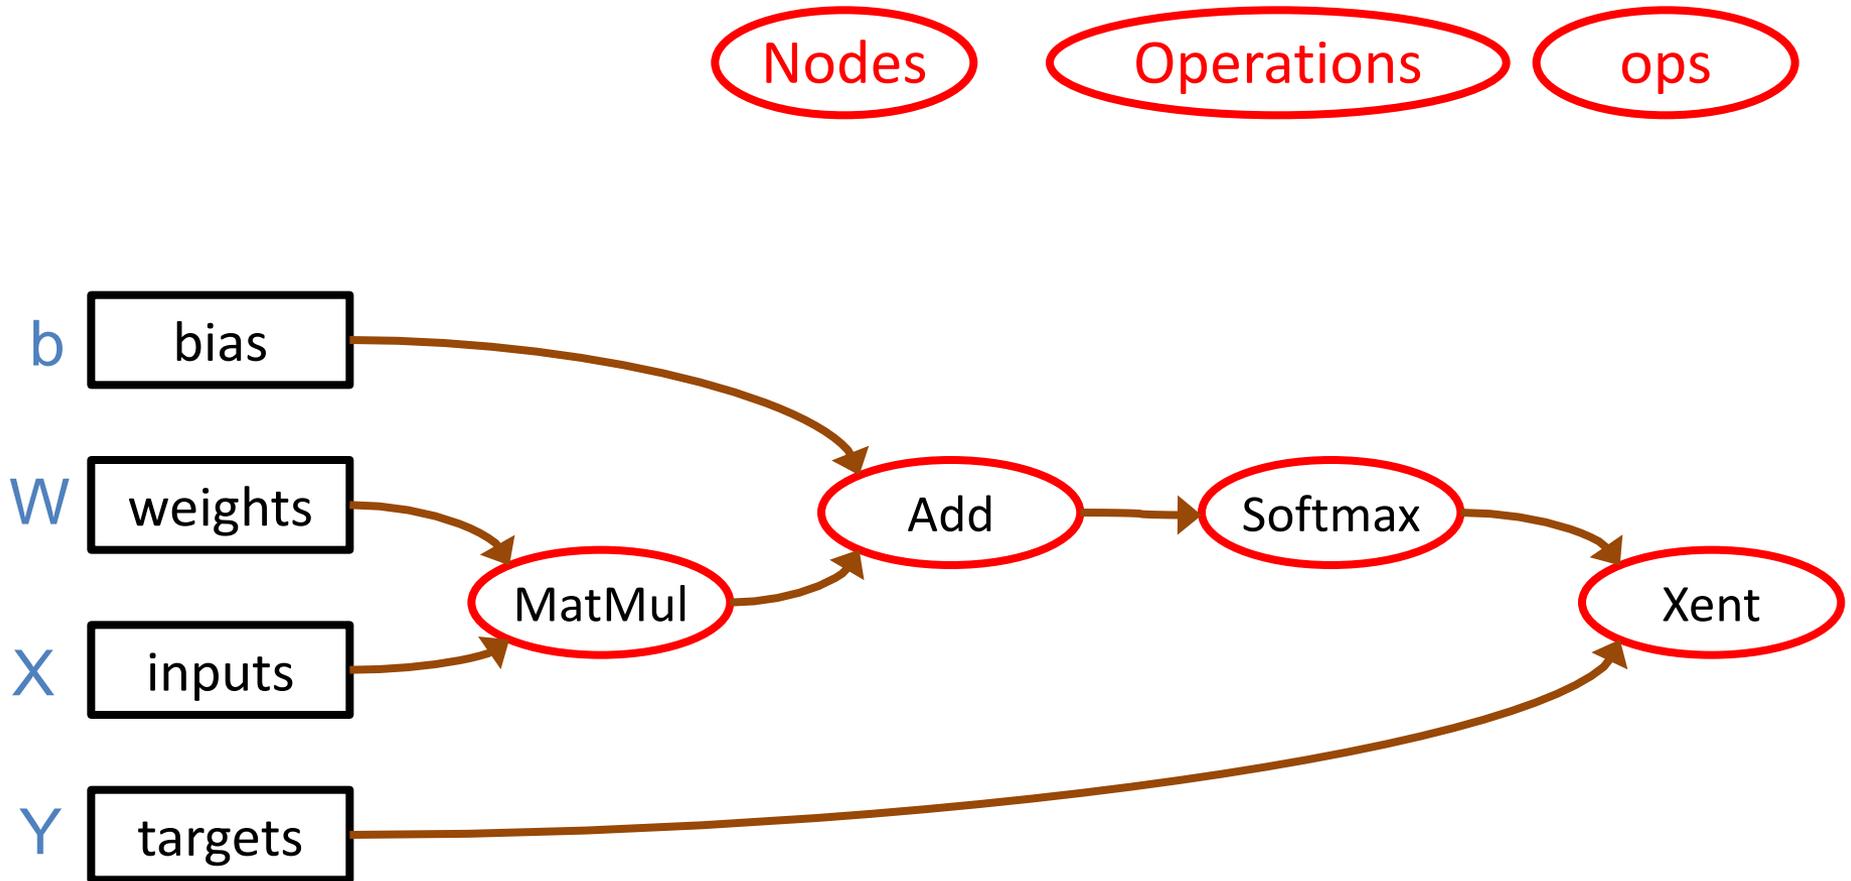


# Computation is a Dataflow Graph

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



# Logistic Regression as Dataflow Graph

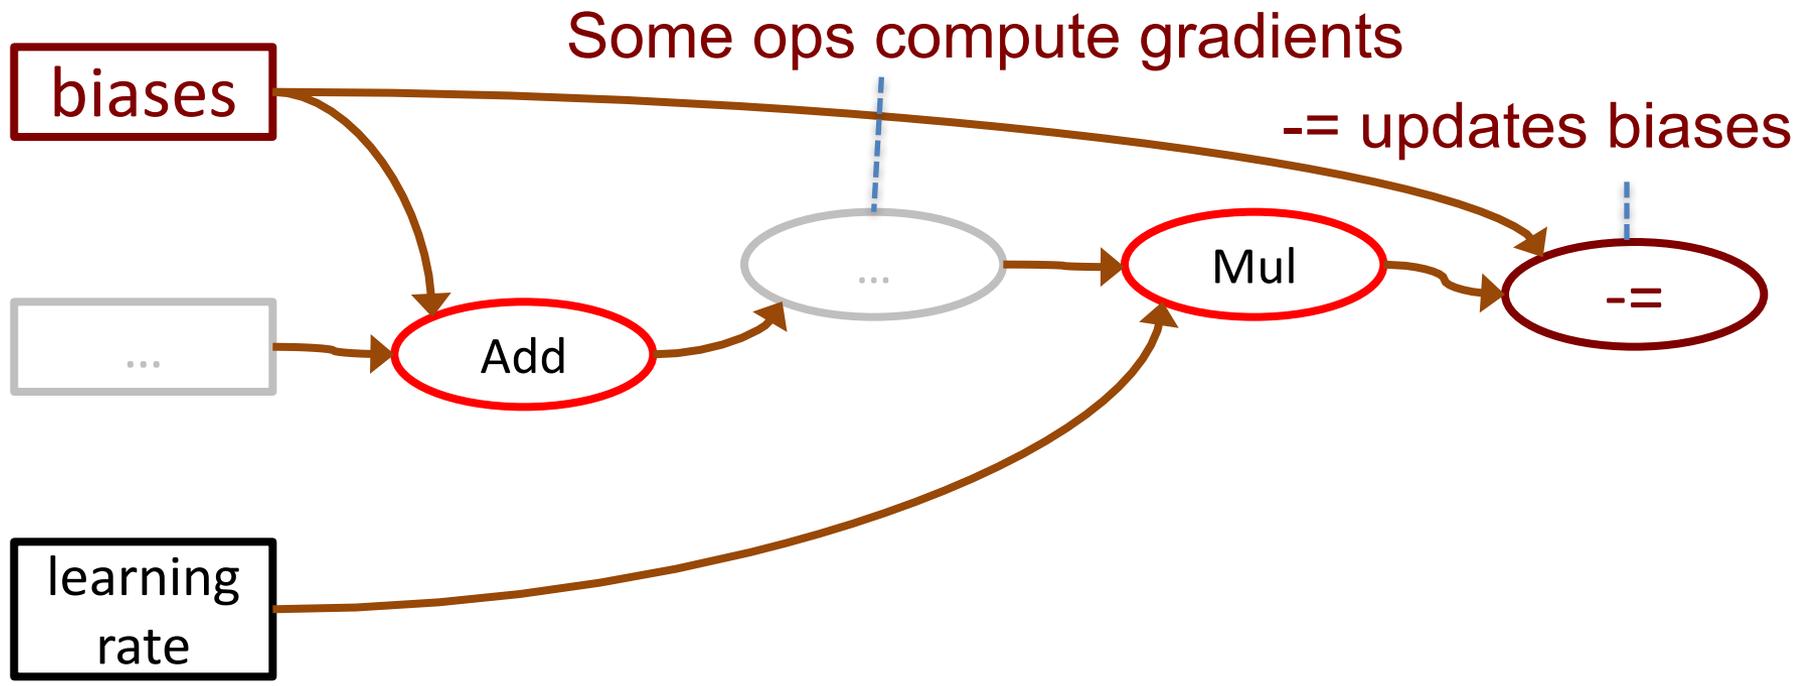


Edges are N-dimensional arrays: **Tensors**

# Computation is a Dataflow Graph

**with state**

'Biases' is a variable

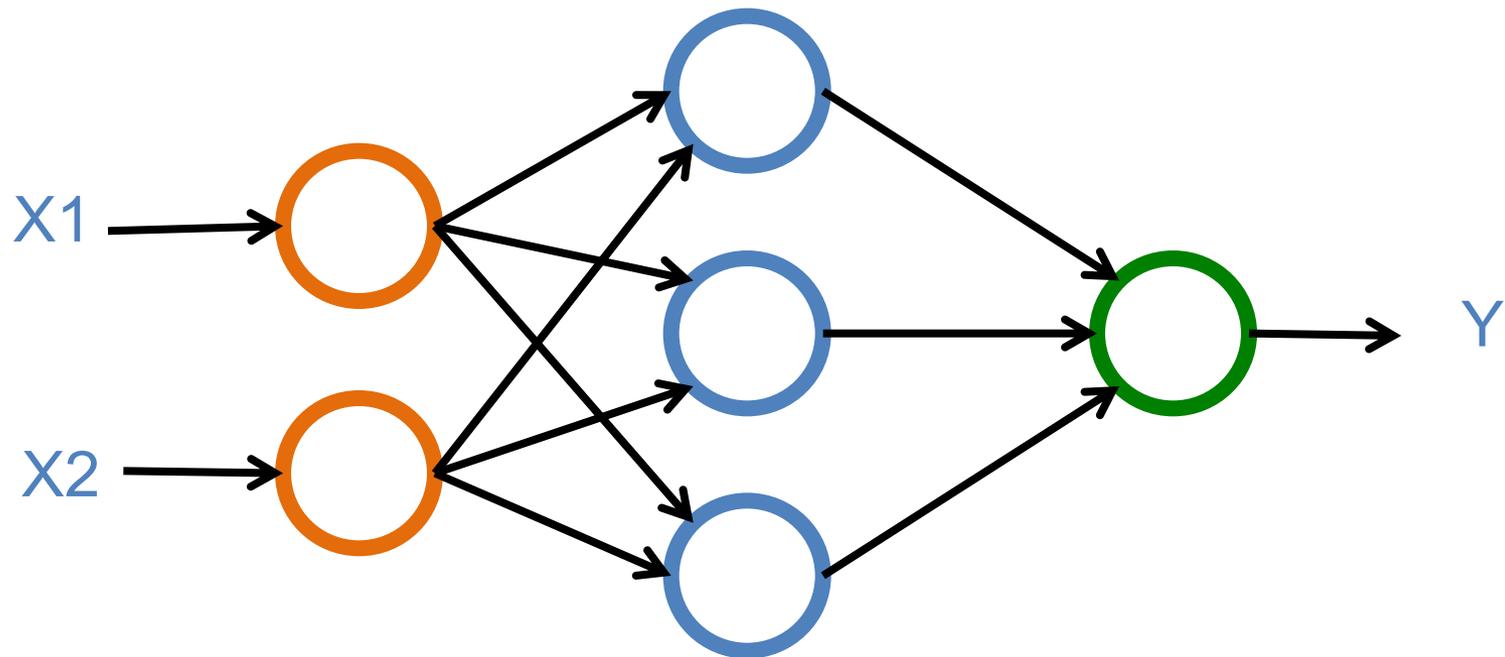


# Neural Networks

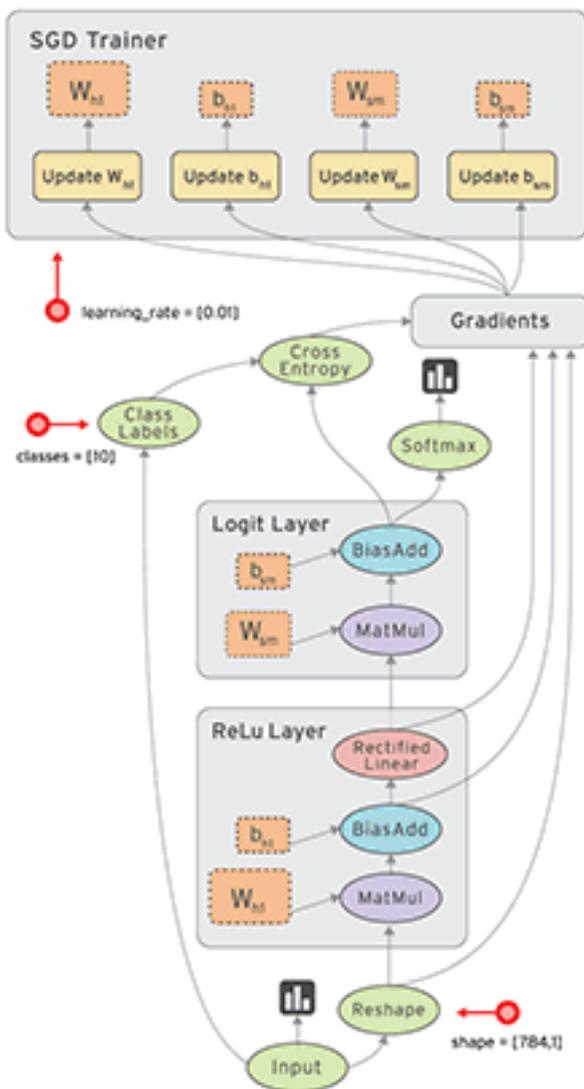
**Input Layer**  
(X)

**Hidden Layer**  
(H)

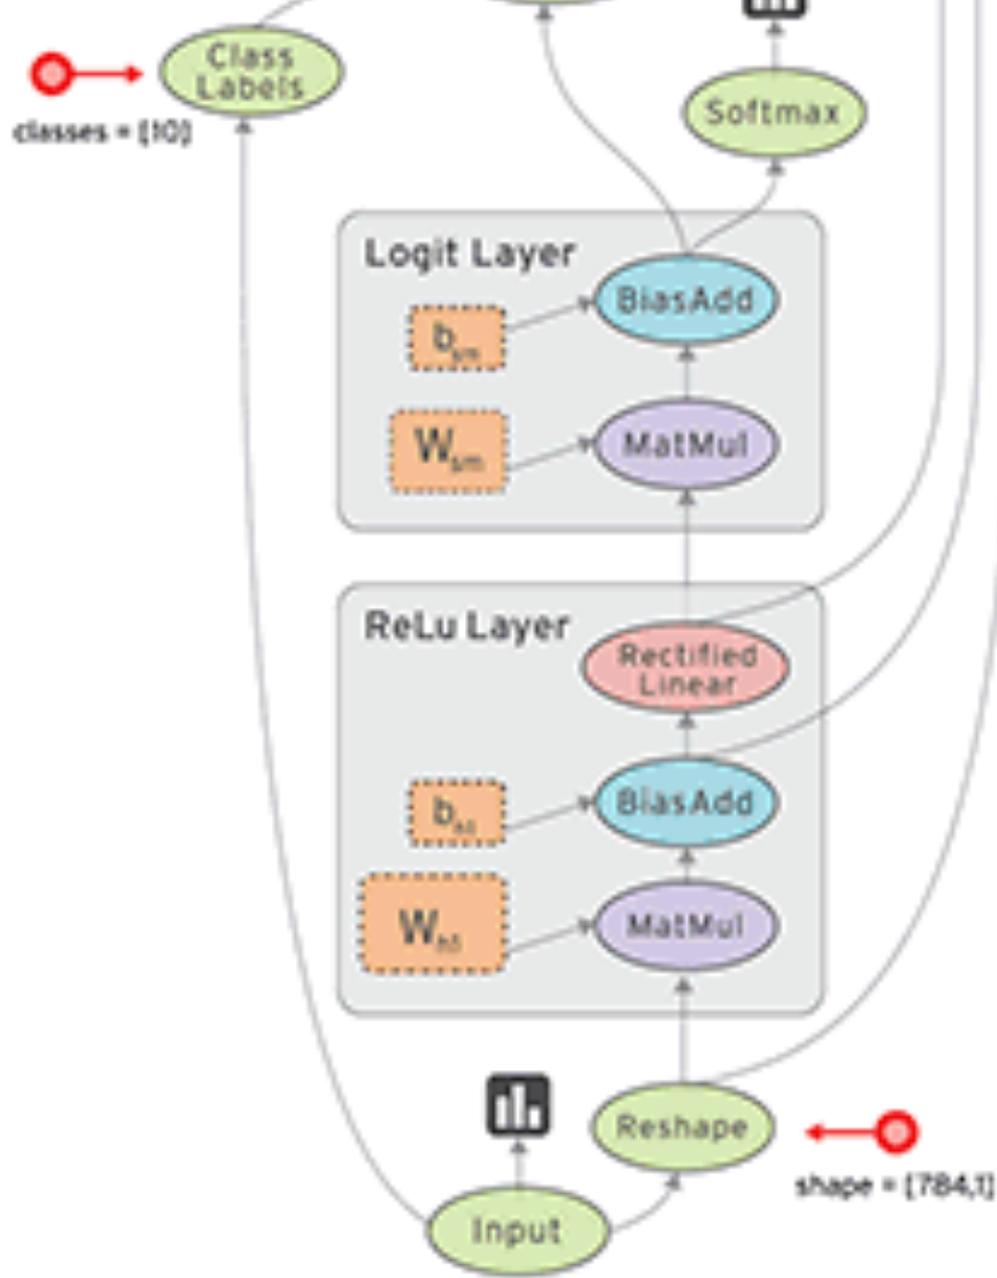
**Output Layer**  
(Y)



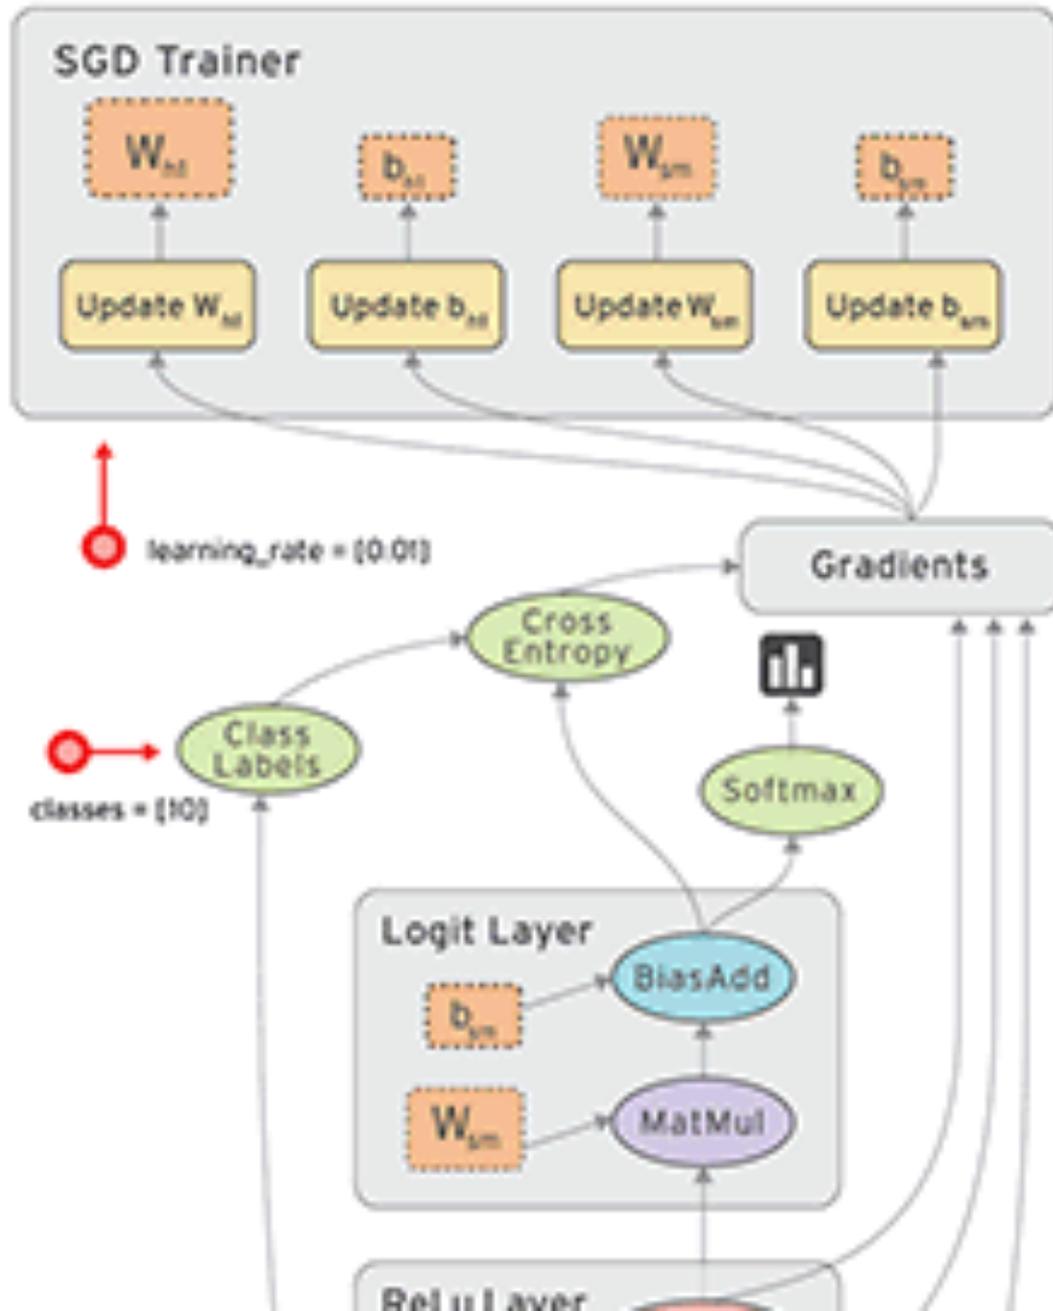
# Data Flow Graph



# Data Flow Graph



# Data Flow Graph



# TensorFlow

```
conda info --envs
```

```
conda --version
```

```
python --version
```

```
conda list
```

```
conda create -n tensorflow python=3.5
```

```
source activate tensorflow
```

```
activate tensorflow
```

```
sudo pip install keras
```

```
pip install keras
```

```
pip install tensorflow
```

```
pip install ipython[all]
```



# TensorFlow Playground

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



Iterations  
000,582

Learning rate  
0.03

Activation  
Tanh

Regularization  
None

Regularization rate  
0

Problem type  
Classification

## DATA

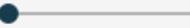
Which dataset do you want to use?



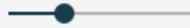
Ratio of training to test data: 50%



Noise: 0



Batch size: 10

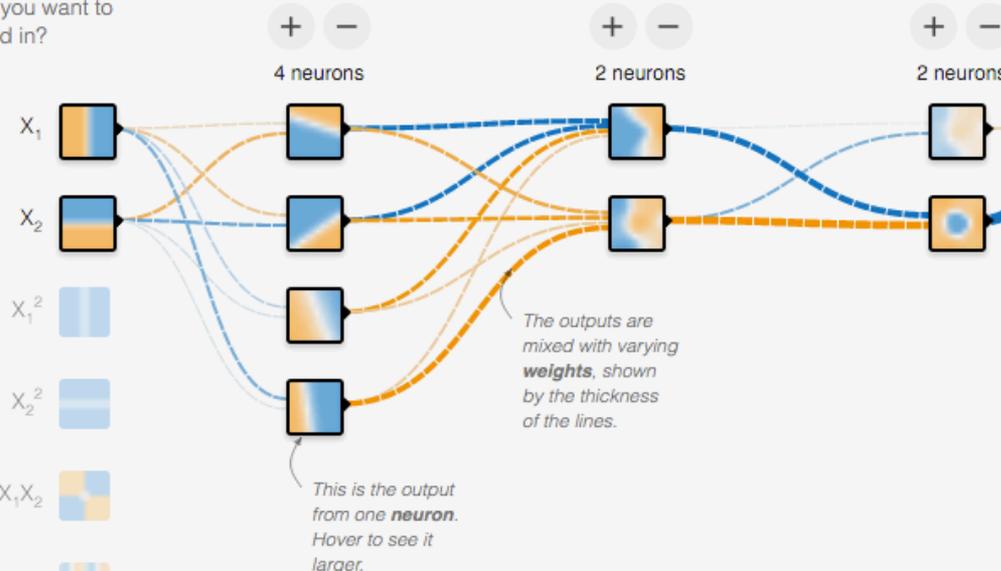


## INPUT

Which properties do you want to feed in?

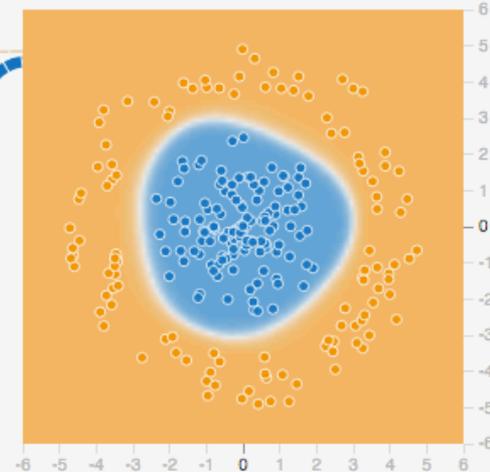


## 3 HIDDEN LAYERS



## OUTPUT

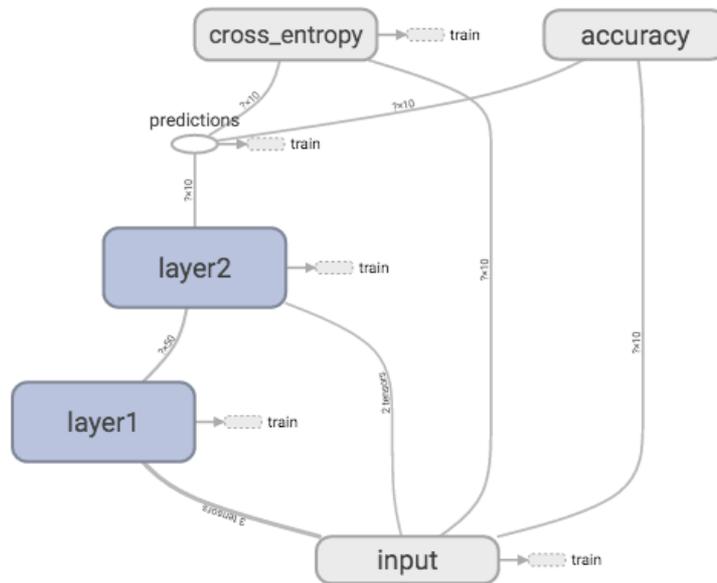
Test loss 0.000  
Training loss 0.000



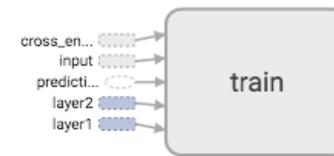
# TensorBoard

Fit to screen  
 Download PNG  
 Run train (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



# Try your first TensorFlow

```
$ python
```

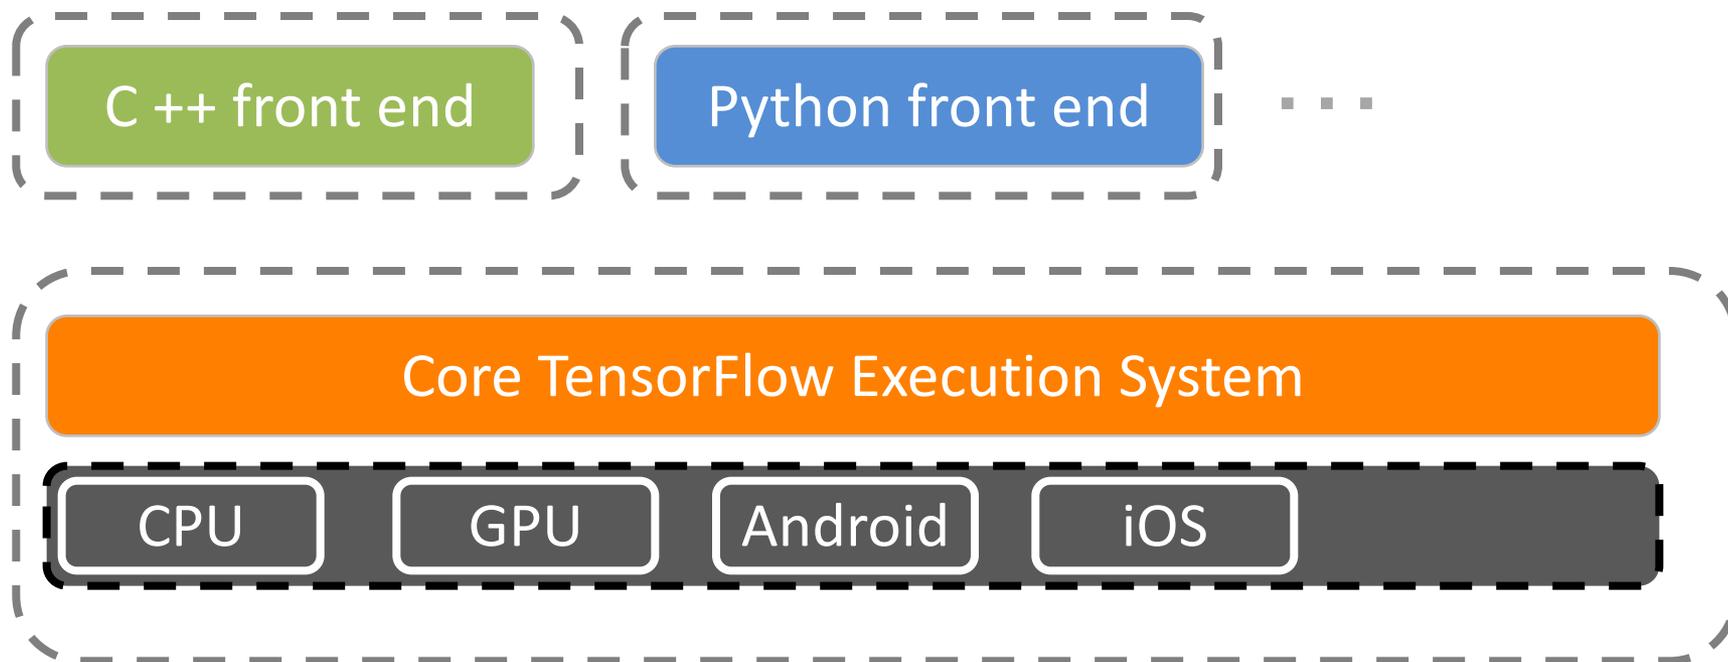
```
>>> import tensorflow as tf
>>> hello = tf.constant('Hello, TensorFlow!')
>>> sess = tf.Session()
>>> sess.run(hello)
Hello, TensorFlow!
>>> a = tf.constant(10)
>>> b = tf.constant(32)
>>> sess.run(a+b)
42
>>>
```

# 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  
>>>
```

# Architecture of TensorFlow



# TensorFlow and Deep Learning

# TensorFlow and Deep Learning

## 1 Overview

Preparation: Install

2 TensorFlow, get the sample code

3 Theory: train a neural network

4 Theory: a 1-layer neural network

5 Theory: gradient descent

6 Lab: let's jump into the code

7 Lab: adding layers

8 Lab: special care for deep networks

9 Lab: learning rate decay

10 Lab: dropout, overfitting

11 Theory: convolutional networks

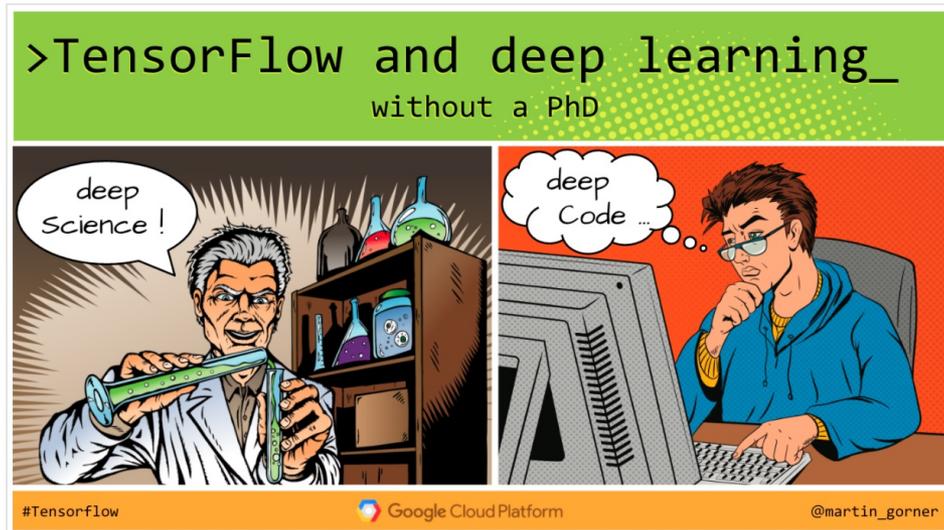
Did you find a mistake? [Please file a bug.](#)

Lab: a convolutional

← TensorFlow and deep learning, without a PhD

🕒 149 min remaining

## 1. Overview



In this codelab, you will learn how to build and train a neural network that recognises handwritten digits. Along the way, as you enhance your neural network to achieve 99% accuracy, you will also discover the tools of the trade that deep learning professionals use to train their models efficiently.

This codelab uses the [MNIST](#) dataset, a collection of 60,000 labeled digits that has kept generations of PhDs busy for almost two decades. You will solve the problem with less than 100 lines of Python / TensorFlow code.

### What you'll learn

# TensorFlow MNIST Tutorial



Features Business Explore Pricing

This repository

Search

Sign in or Sign up

martin-gorner / tensorflow-mnist-tutorial

Watch

50

★ Star

489

Fork

204

Code

Issues 6

Pull requests 2

Projects 0

Pulse

Graphs

Sample code for "Tensorflow and deep learning, without a PhD" presentation and code lab.

102 commits

1 branch

0 releases

4 contributors

Apache-2.0

Branch: master

New pull request

Find file

Clone or download

	<b>martin-gorner</b> committed on <b>GitHub</b> Update INSTALL.txt ...	Latest commit ed331aa 25 days ago
<a href="#">mlengine</a>	added example using the Tensorflow high level layers API	26 days ago
<a href="#">.gitignore</a>	small bug fix in batch norm	6 months ago
<a href="#">CONTRIBUTING.md</a>	initial commit 2	4 months ago
<a href="#">INSTALL.txt</a>	Update INSTALL.txt	25 days ago
<a href="#">LICENSE</a>	Initial commit	a year ago
<a href="#">README.md</a>	better image URL	3 months ago
<a href="#">mnist_1.0_softmax.py</a>	global_variables_initializer used everywhere instead of initalize_al...	2 months ago
<a href="#">mnist_2.0_five_layers_sigmoid.py</a>	Fix spacing in the network structure comment	a month ago
<a href="#">mnist_2.1_five_layers_relu_lrdecay...</a>	Fix spacing in the network structure comment	a month ago

# TensorFlow and Deep Learning

- What is a neural network and how to train it
- How to build a basic 1-layer neural network using TensorFlow
- How to add more layers
- Training tips and tricks: overfitting, dropout, learning rate decay ...
- How to troubleshoot deep neural networks
- How to build convolutional networks

# TensorFlow MNIST Tutorial

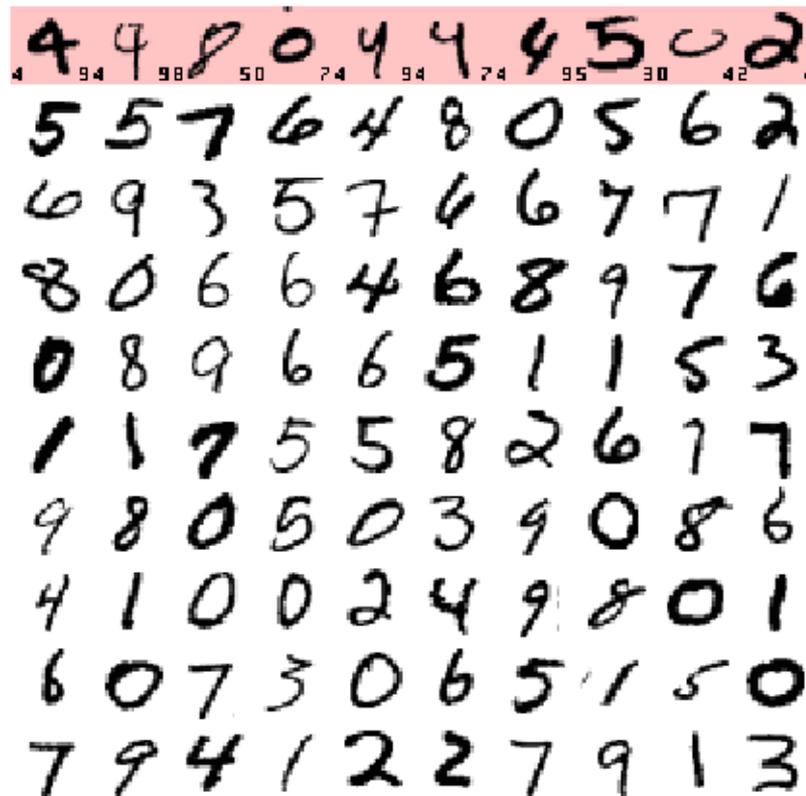
```
git clone https://github.com/martin-gorner/tensorflow-mnist-tutorial.git
```

```
cd tensorflow-mnist-tutorial
```

```
python3 mnist_1.0_softmax.py
```

# MNIST dataset: 60,000 labeled digits

Training digits



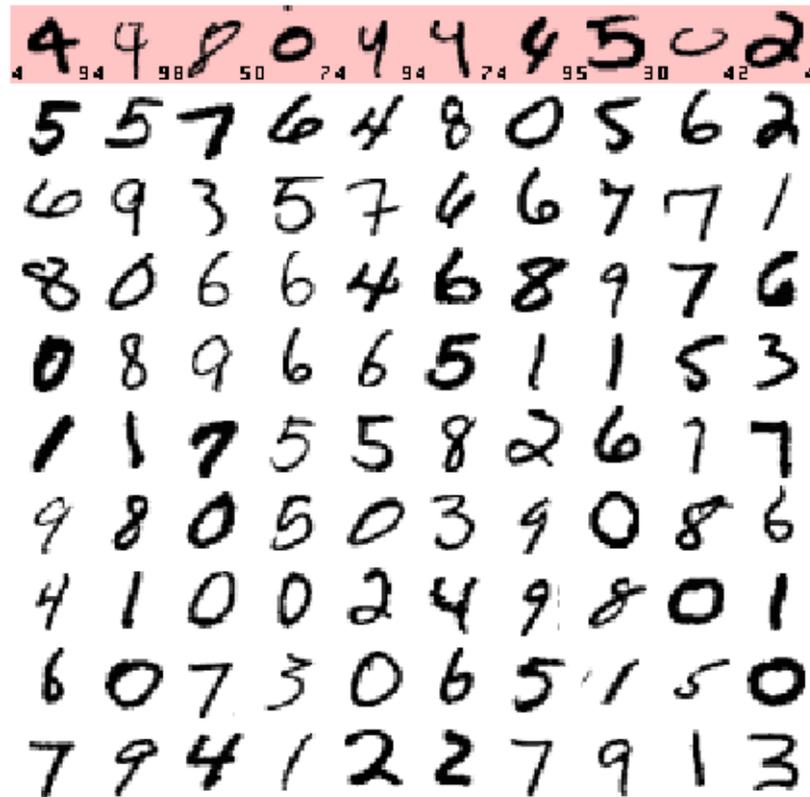


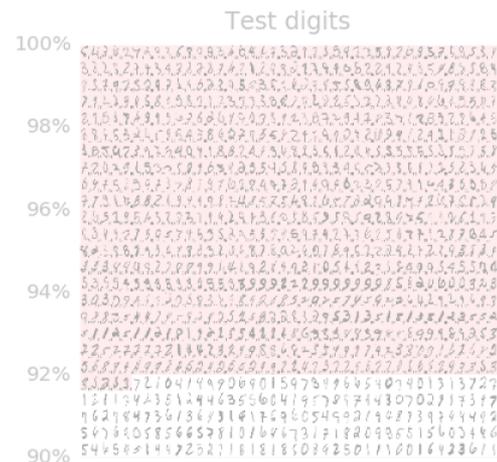
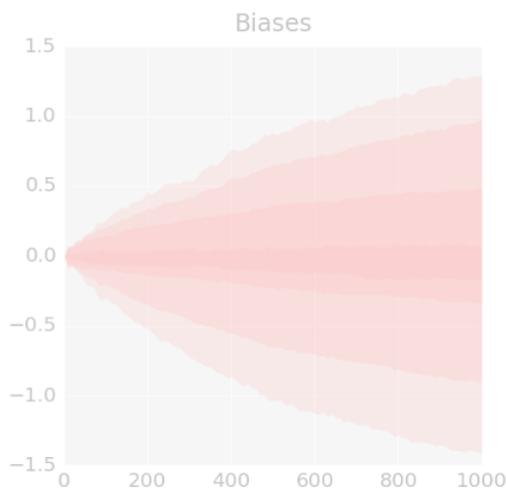
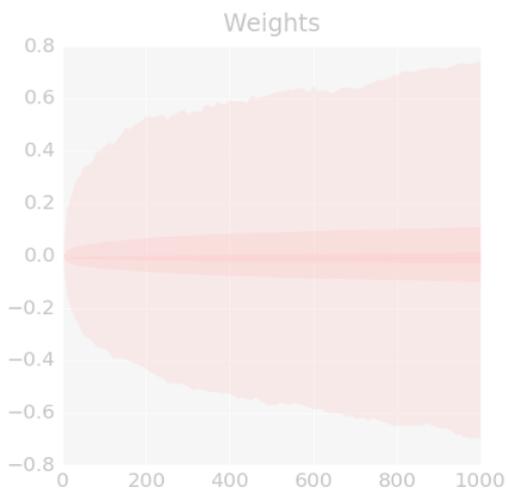
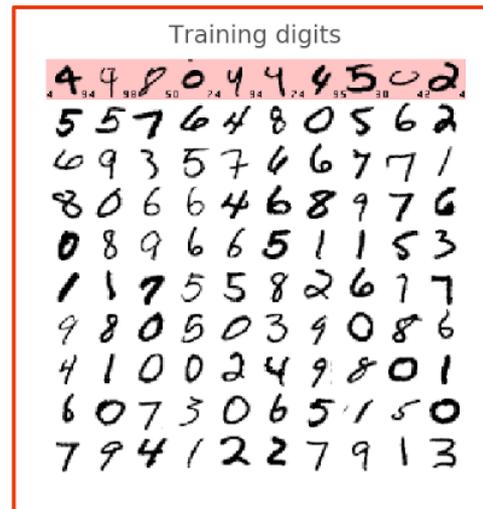
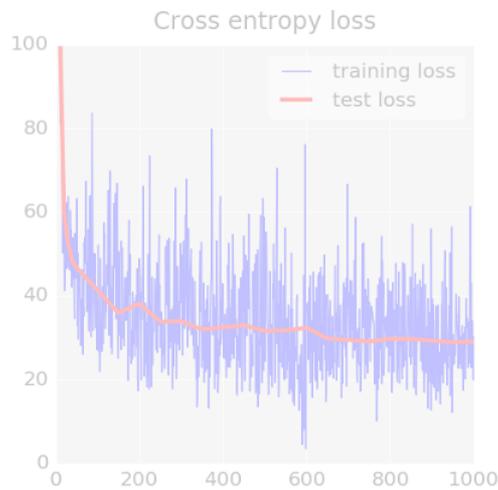
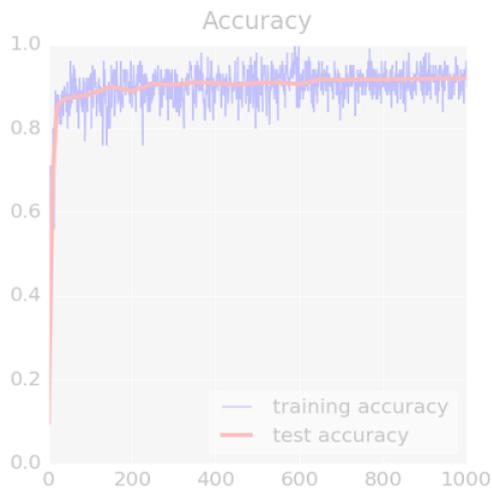
# Train a Neural Network

Training digits

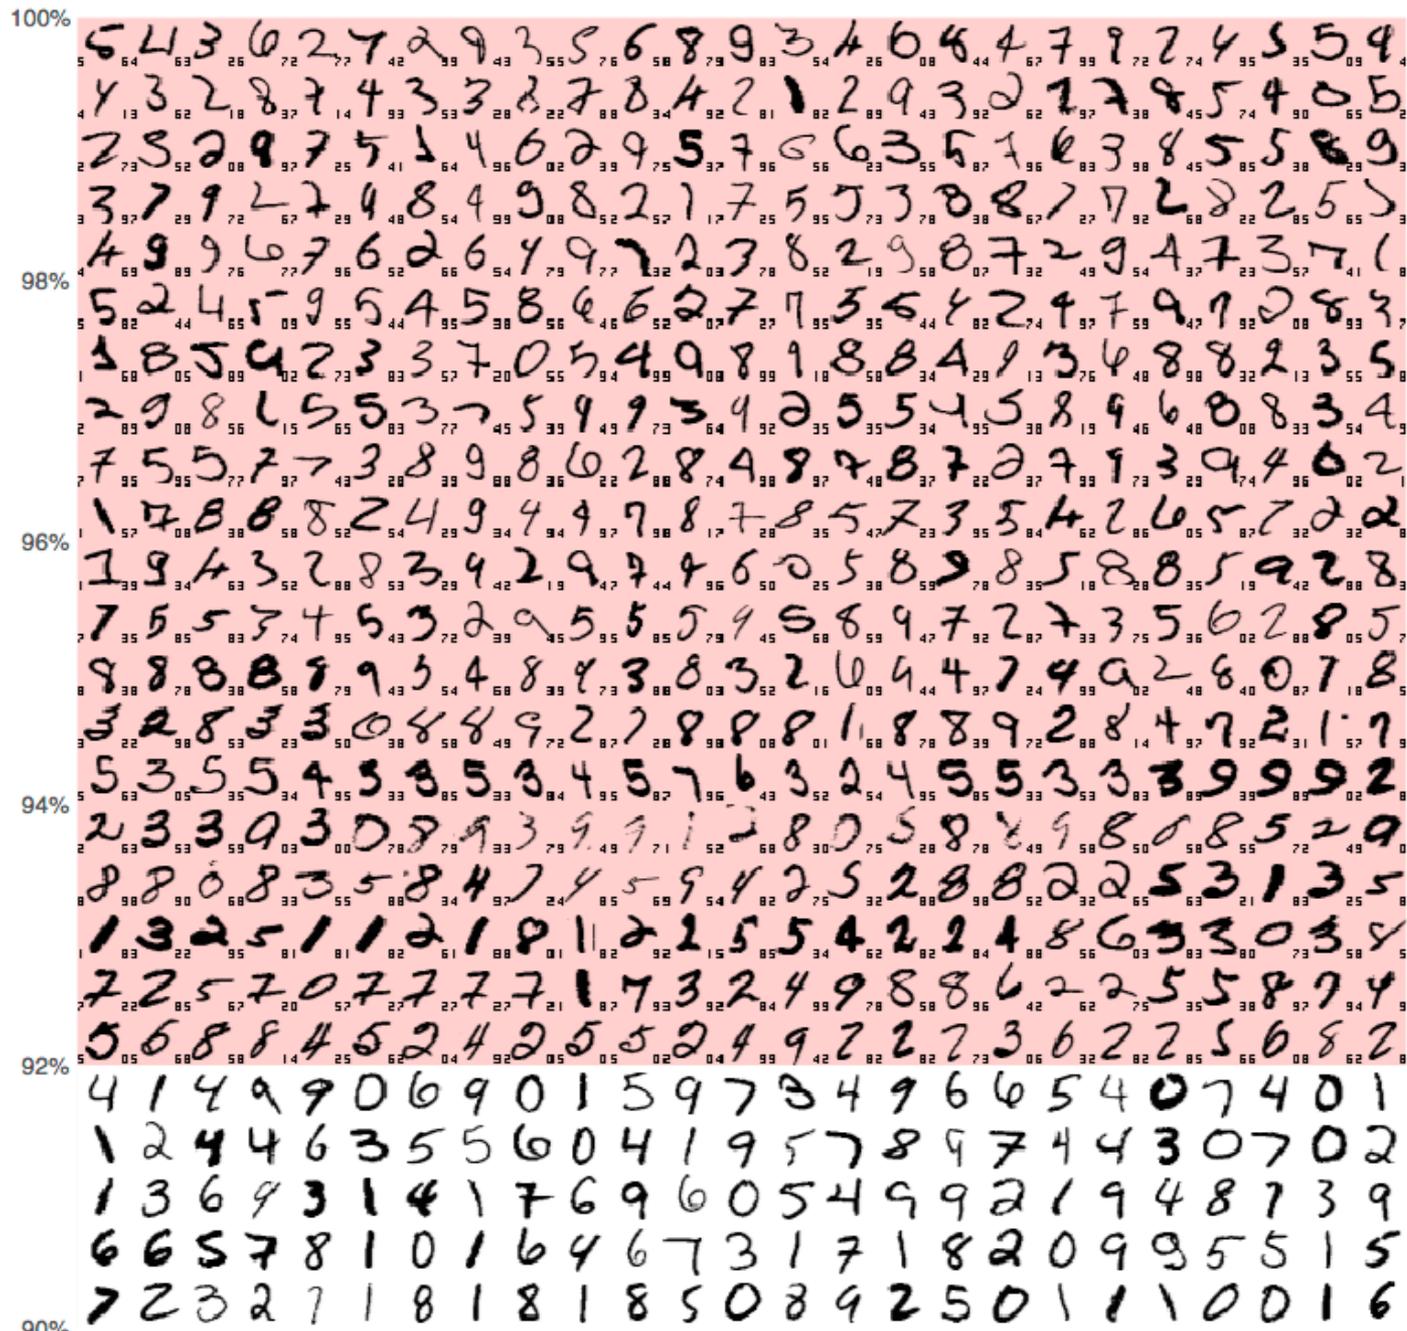
updates to **weights** and **biases** =>  
better recognition (loop)

## Training digits

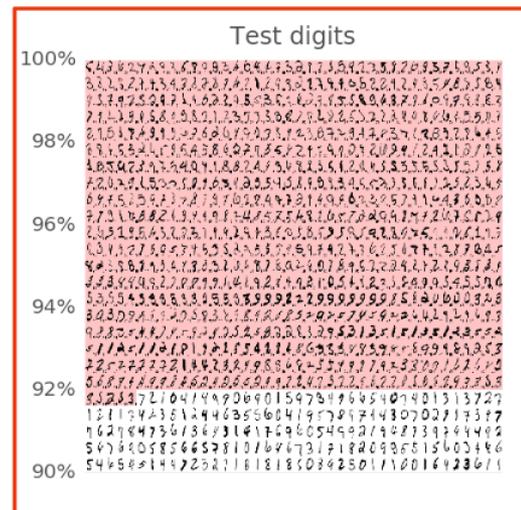
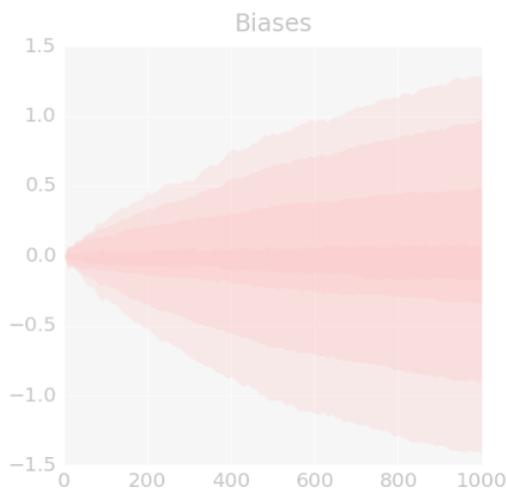
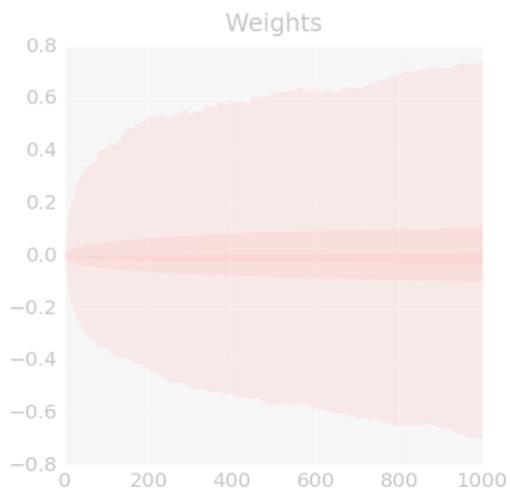
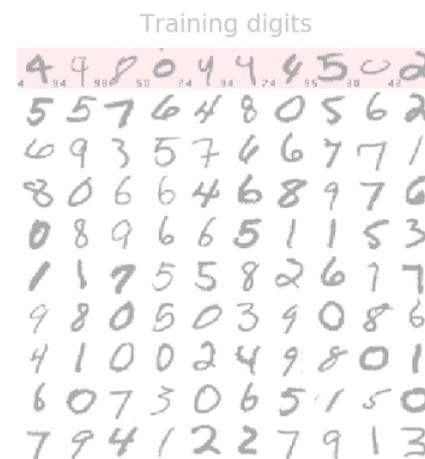
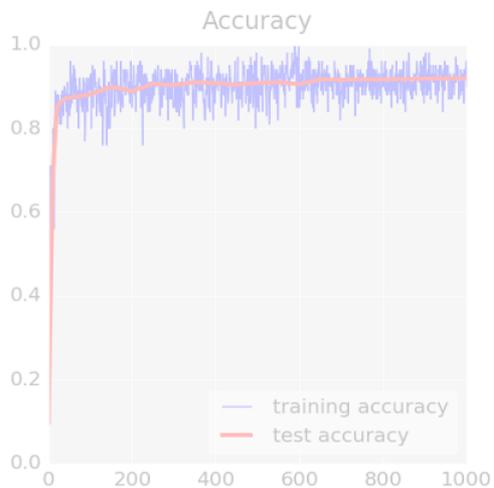


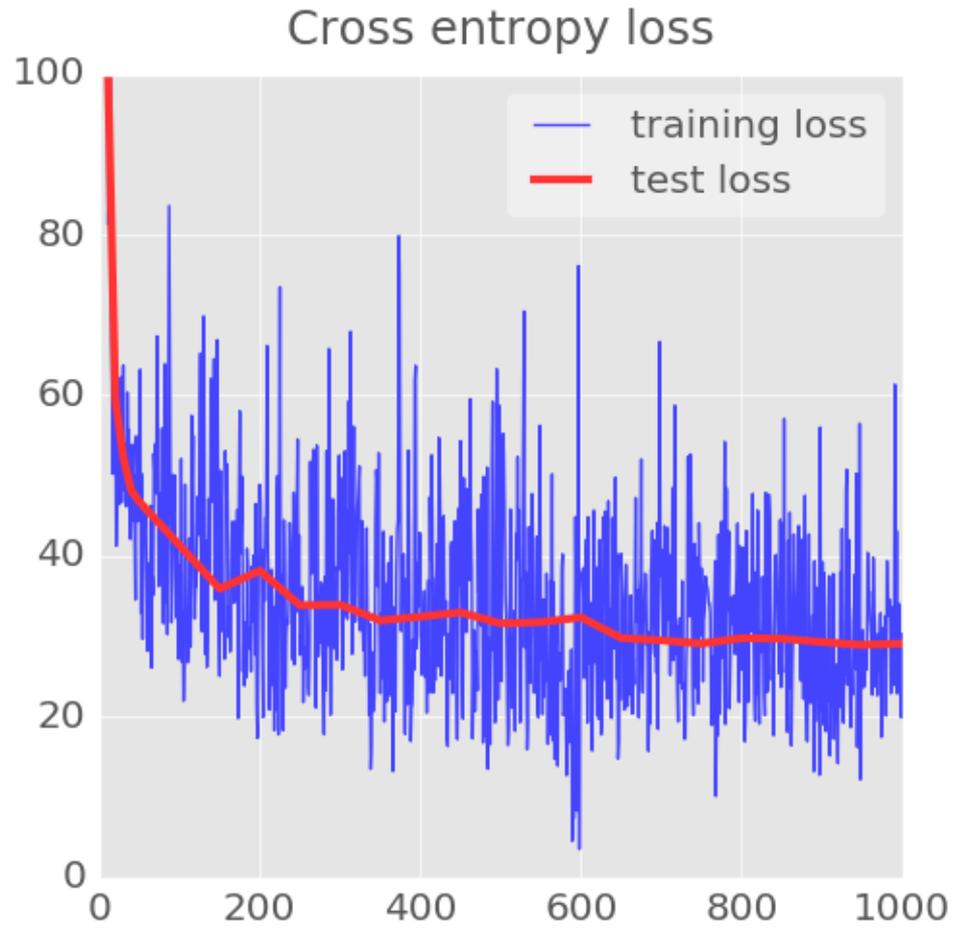


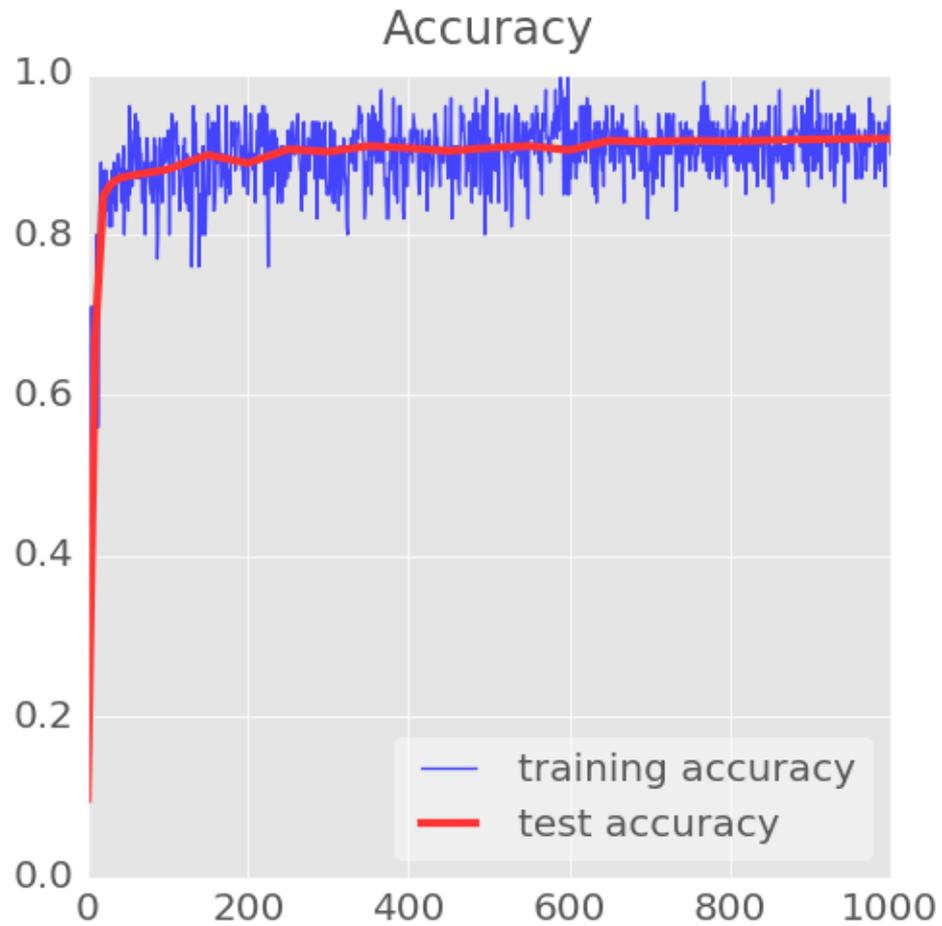
# Test digits

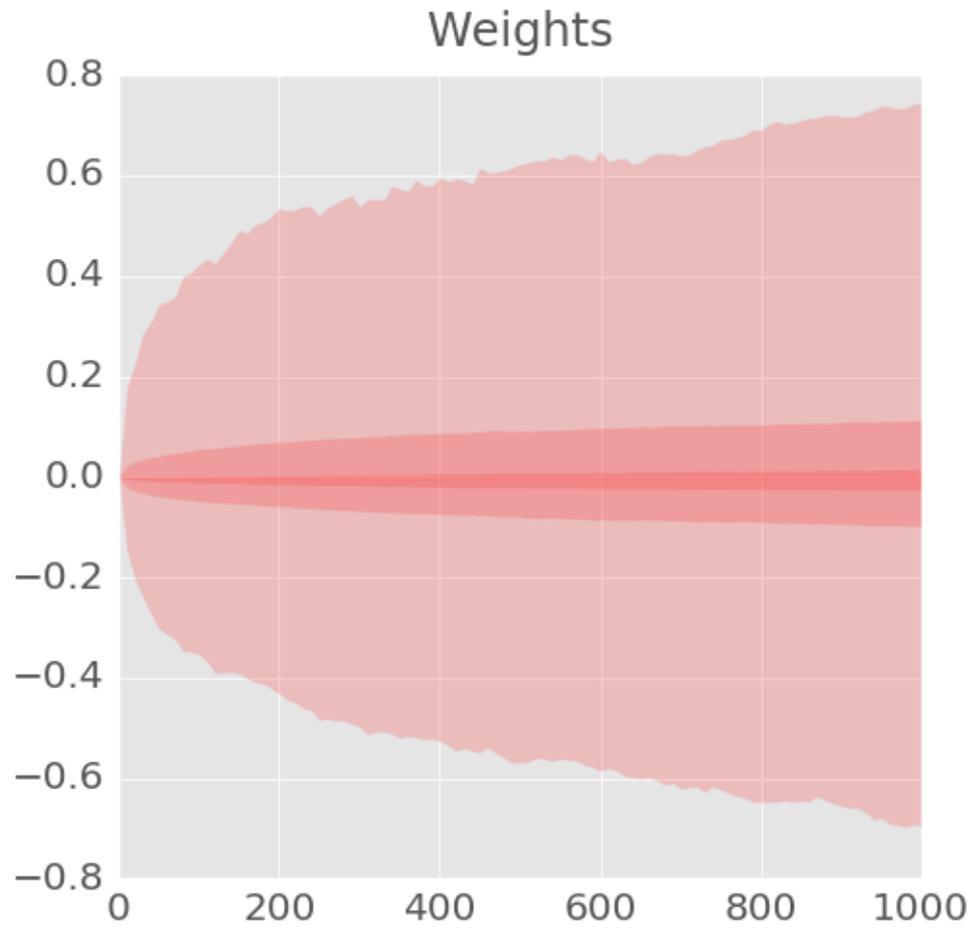


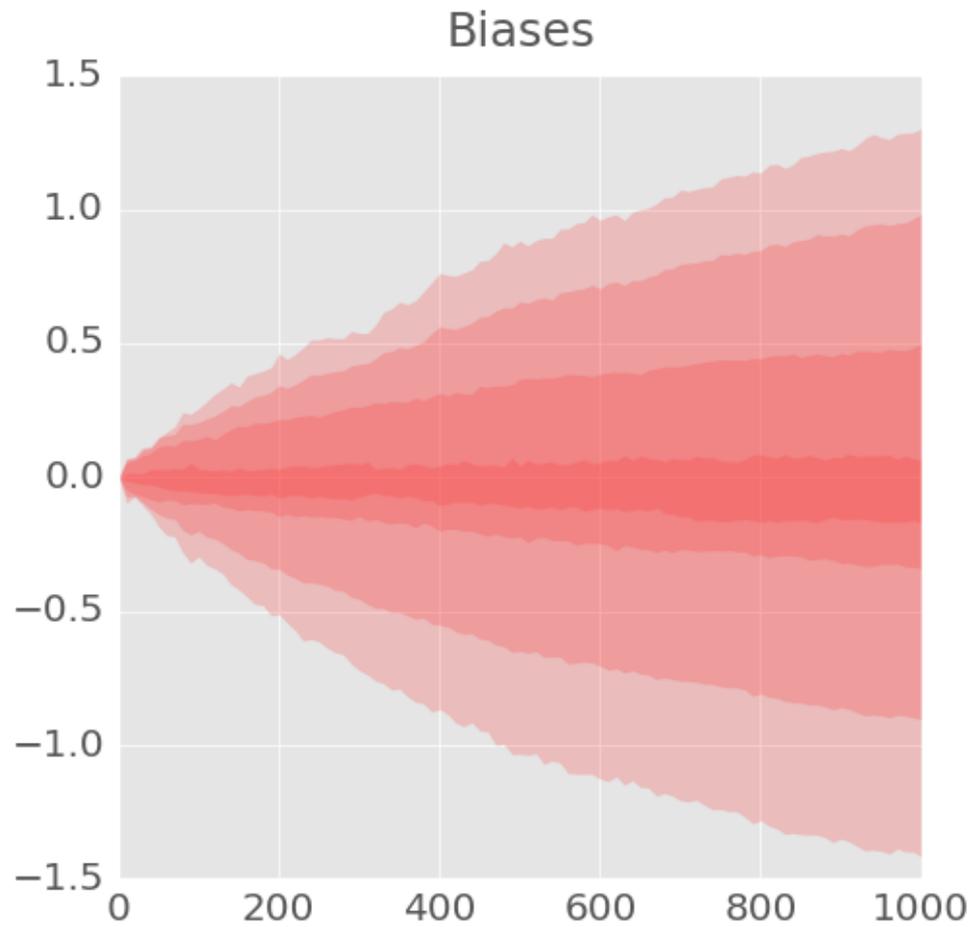
Source: <https://codelabs.developers.google.com/codelabs/cloud-tensorflow-mnist/#2>











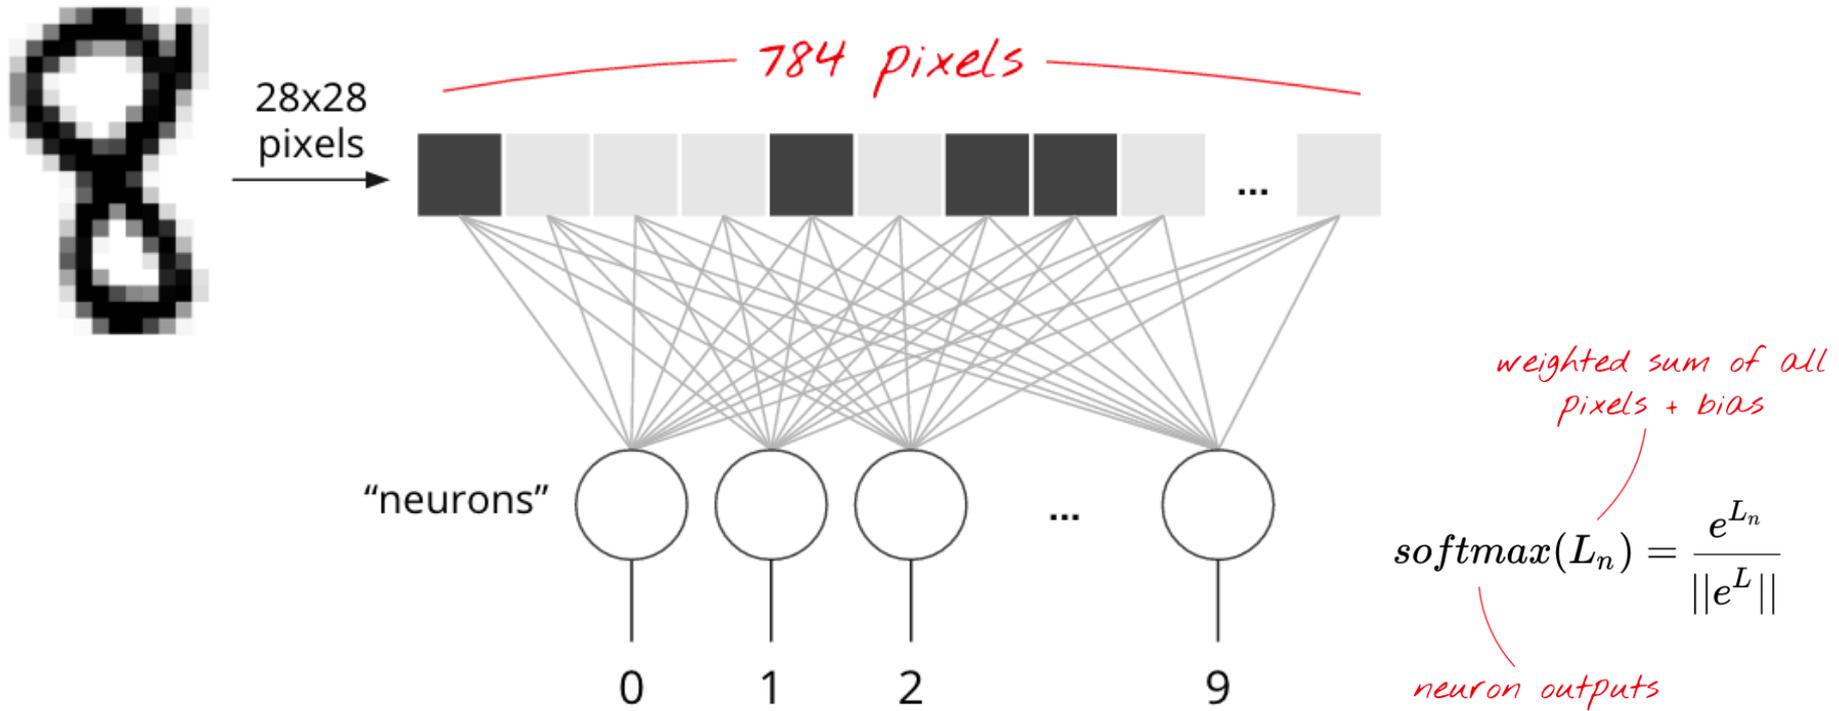


# Cookbook

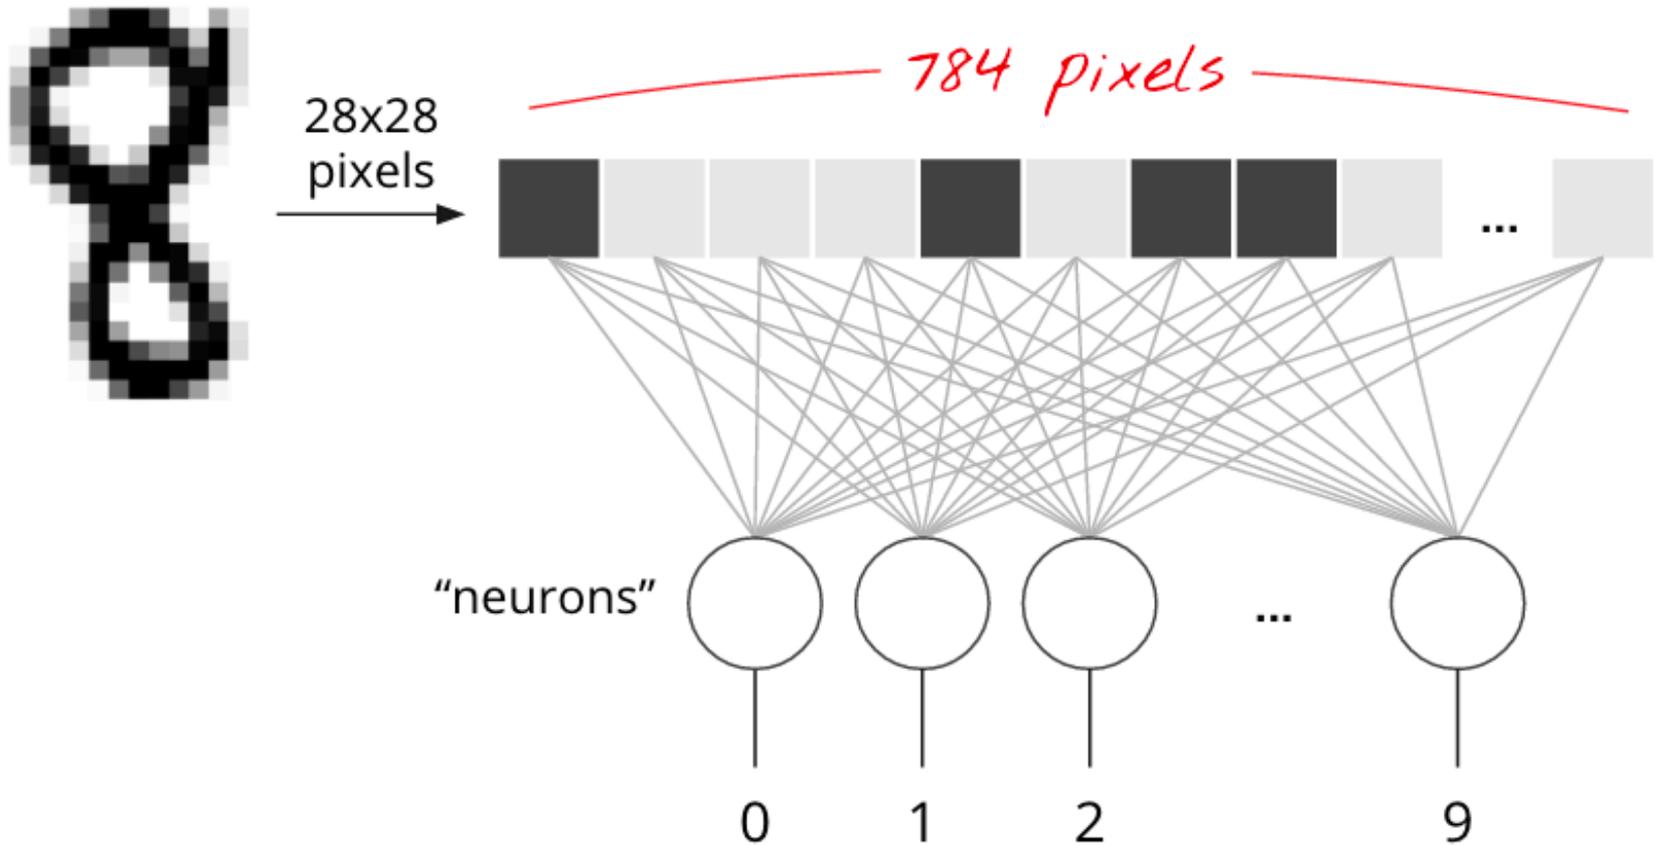
Softmax  
Cross-entropy  
Mini-batch



# Very Simple Model: Softmax Classification



# Very Simple Model: Softmax Classification



# Very Simple Model: Softmax Classification

*weighted sum of all  
pixels + bias*

$$\text{softmax}(L_n) = \frac{e^{L_n}}{\|e^L\|}$$

*neuron outputs*

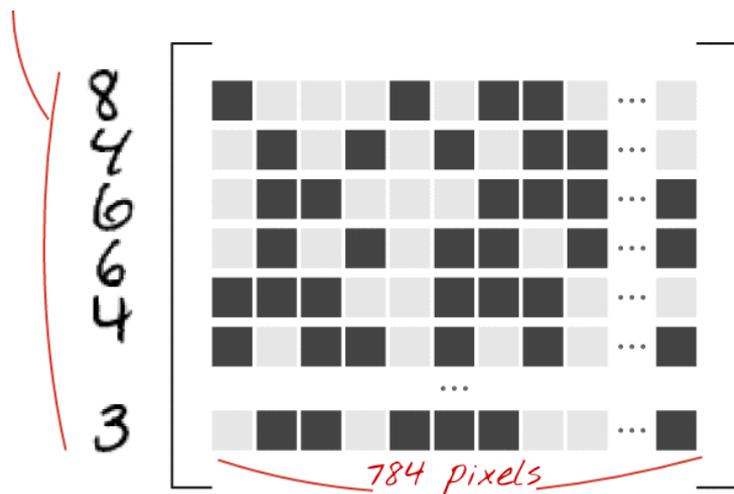
# In Matrix notation, 100 images at a time

*X: 100 images,  
one per line,  
flattened*

*10 columns*

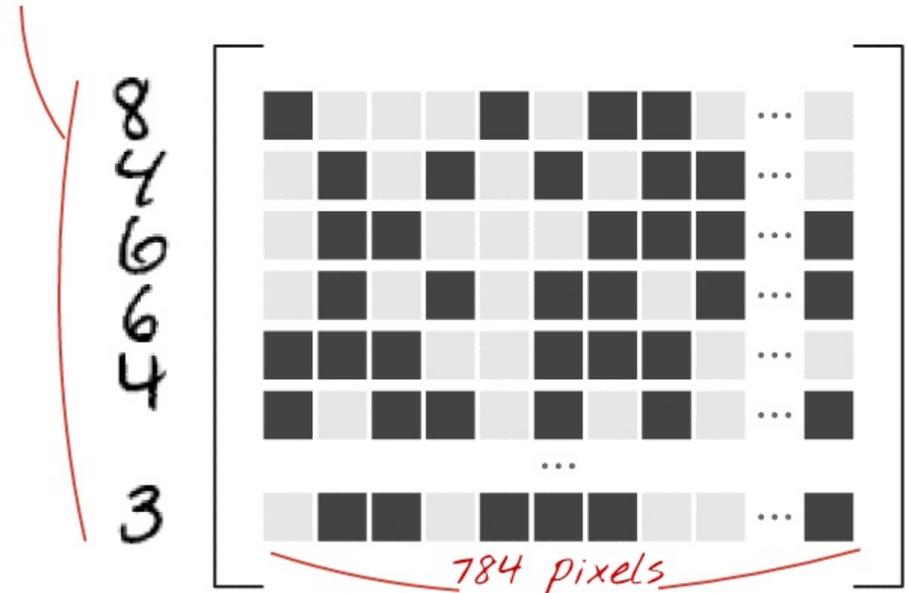
$W_{0,0}$	$W_{0,1}$	$W_{0,2}$	$W_{0,3}$	...	$W_{0,9}$
$W_{1,0}$	$W_{1,1}$	$W_{1,2}$	$W_{1,3}$	...	$W_{1,9}$
$W_{2,0}$	$W_{2,1}$	$W_{2,2}$	$W_{2,3}$	...	$W_{2,9}$
$W_{3,0}$	$W_{3,1}$	$W_{3,2}$	$W_{3,3}$	...	$W_{3,9}$
$W_{4,0}$	$W_{4,1}$	$W_{4,2}$	$W_{4,3}$	...	$W_{4,9}$
$W_{5,0}$	$W_{5,1}$	$W_{5,2}$	$W_{5,3}$	...	$W_{5,9}$
$W_{6,0}$	$W_{6,1}$	$W_{6,2}$	$W_{6,3}$	...	$W_{6,9}$
$W_{7,0}$	$W_{7,1}$	$W_{7,2}$	$W_{7,3}$	...	$W_{7,9}$
$W_{8,0}$	$W_{8,1}$	$W_{8,2}$	$W_{8,3}$	...	$W_{8,9}$
...					
$W_{783,0}$	$W_{783,1}$	$W_{783,2}$	...		$W_{783,9}$

*784 lines*

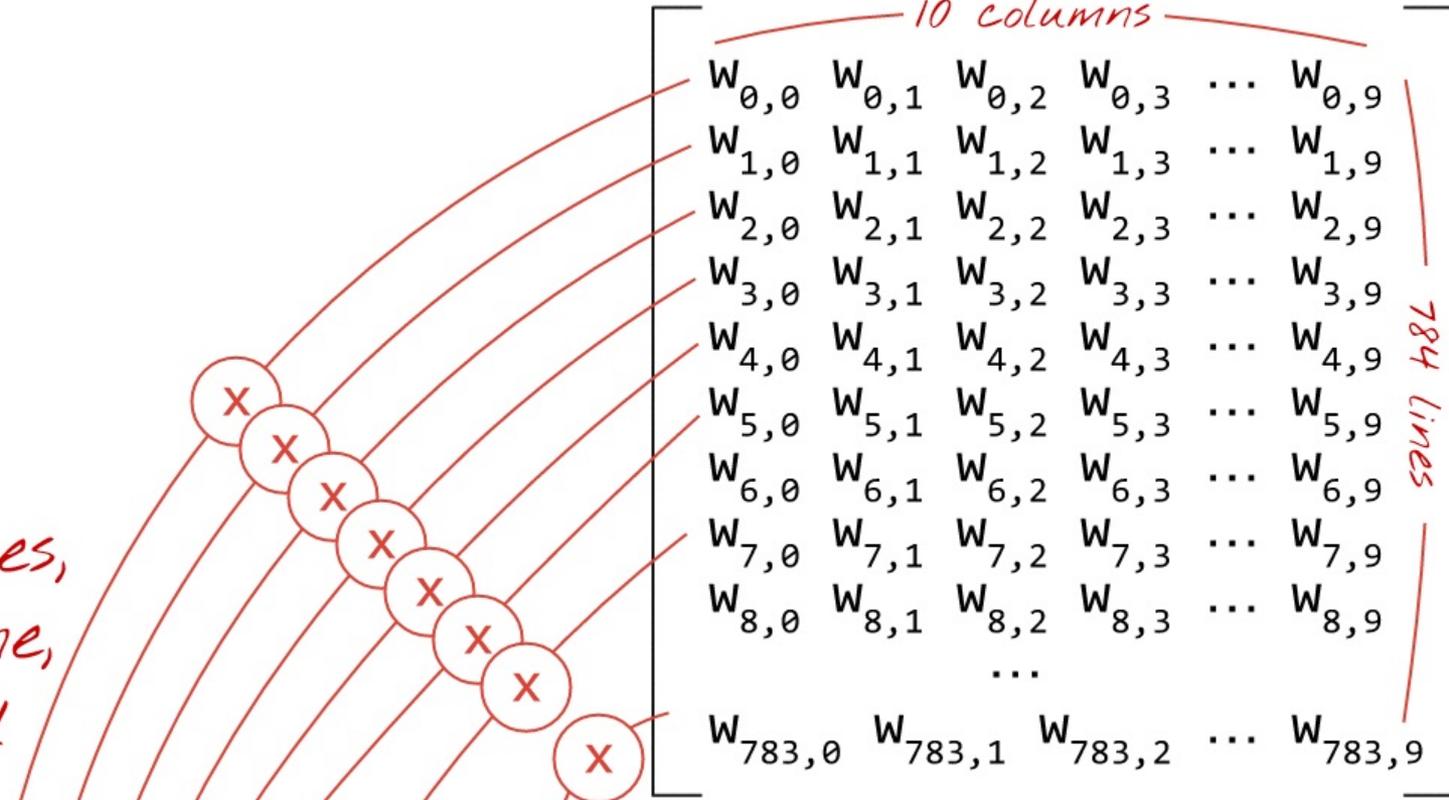
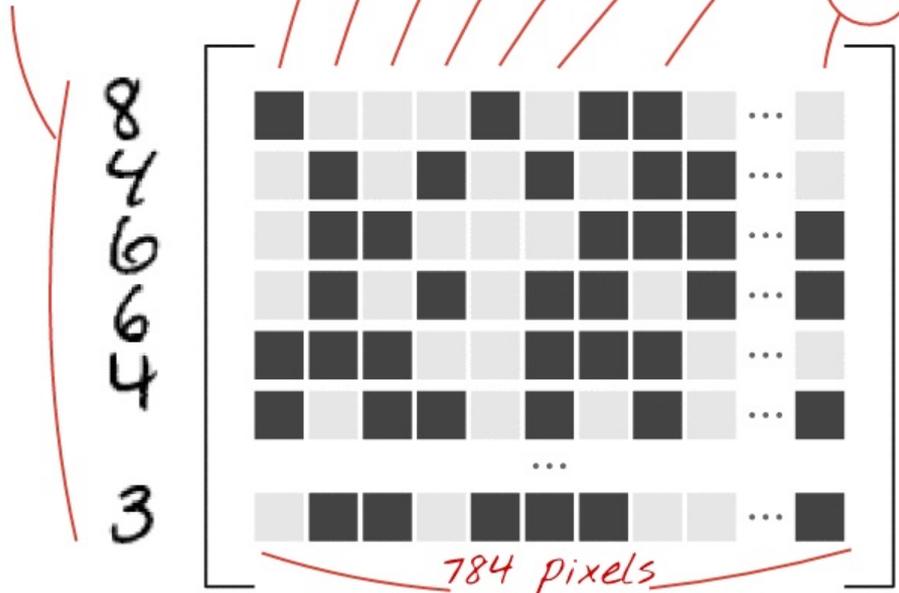


*X: 100 images,  
one per line,  
flattened*

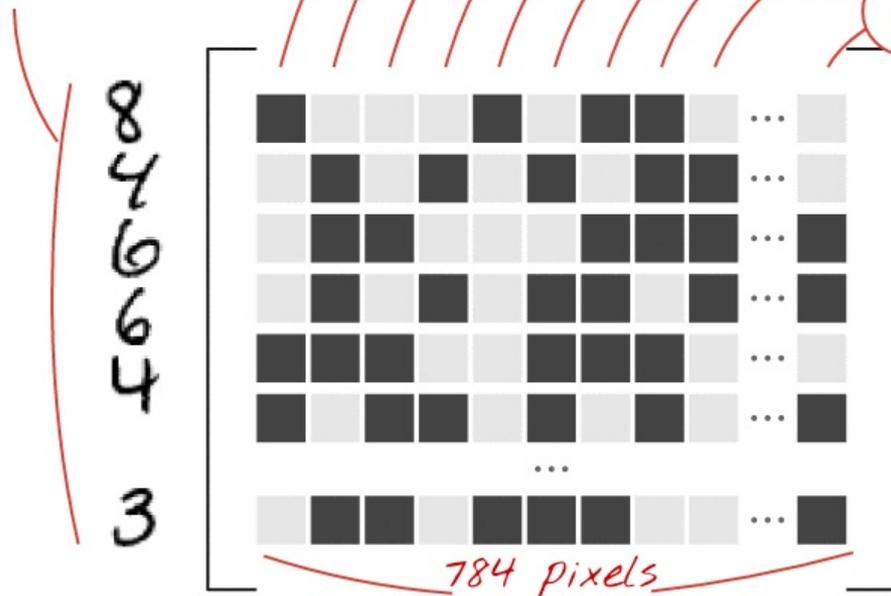
<i>10 columns</i>									
$W_{0,0}$	$W_{0,1}$	$W_{0,2}$	$W_{0,3}$	...	$W_{0,9}$	<i>784 lines</i>			
$W_{1,0}$	$W_{1,1}$	$W_{1,2}$	$W_{1,3}$	...	$W_{1,9}$				
$W_{2,0}$	$W_{2,1}$	$W_{2,2}$	$W_{2,3}$	...	$W_{2,9}$				
$W_{3,0}$	$W_{3,1}$	$W_{3,2}$	$W_{3,3}$	...	$W_{3,9}$				
$W_{4,0}$	$W_{4,1}$	$W_{4,2}$	$W_{4,3}$	...	$W_{4,9}$				
$W_{5,0}$	$W_{5,1}$	$W_{5,2}$	$W_{5,3}$	...	$W_{5,9}$				
$W_{6,0}$	$W_{6,1}$	$W_{6,2}$	$W_{6,3}$	...	$W_{6,9}$				
$W_{7,0}$	$W_{7,1}$	$W_{7,2}$	$W_{7,3}$	...	$W_{7,9}$				
$W_{8,0}$	$W_{8,1}$	$W_{8,2}$	$W_{8,3}$	...	$W_{8,9}$				
...									
$W_{783,0}$	$W_{783,1}$	$W_{783,2}$	...	$W_{783,9}$					



X: 100 images,  
one per line,  
flattened



X: 100 images,  
one per line,  
flattened



10 columns

$W_{0,0}$	$W_{0,1}$	$W_{0,2}$	$W_{0,3}$	...	$W_{0,9}$
$W_{1,0}$	$W_{1,1}$	$W_{1,2}$	$W_{1,3}$	...	$W_{1,9}$
$W_{2,0}$	$W_{2,1}$	$W_{2,2}$	$W_{2,3}$	...	$W_{2,9}$
$W_{3,0}$	$W_{3,1}$	$W_{3,2}$	$W_{3,3}$	...	$W_{3,9}$
$W_{4,0}$	$W_{4,1}$	$W_{4,2}$	$W_{4,3}$	...	$W_{4,9}$
$W_{5,0}$	$W_{5,1}$	$W_{5,2}$	$W_{5,3}$	...	$W_{5,9}$
$W_{6,0}$	$W_{6,1}$	$W_{6,2}$	$W_{6,3}$	...	$W_{6,9}$
$W_{7,0}$	$W_{7,1}$	$W_{7,2}$	$W_{7,3}$	...	$W_{7,9}$
$W_{8,0}$	$W_{8,1}$	$W_{8,2}$	$W_{8,3}$	...	$W_{8,9}$
...	...	...	...	...	...
$W_{783,0}$	$W_{783,1}$	$W_{783,2}$	...	...	$W_{783,9}$

784 lines



**What are "weights" and "biases" ?**  
**How is the "cross-entropy"**  
**computed ?**  
**How exactly does the training**  
**algorithm work ?**

$$Y = f(X)$$

Predictions

$Y[100, 10]$

Images

$X[100, 784]$

Weights

$W[784, 10]$

Biases

$b[10]$

$$Y = \text{softmax}(X \cdot W + b)$$

applied line  
by line

matrix multiply

broadcast  
on all lines

tensor shapes in [ ]

# Cross Entropy

0	1	2	3	4	5	6	7	8	9
0	0	0	0	0	0	1	0	0	0

actual probabilities, "one-hot" encoded

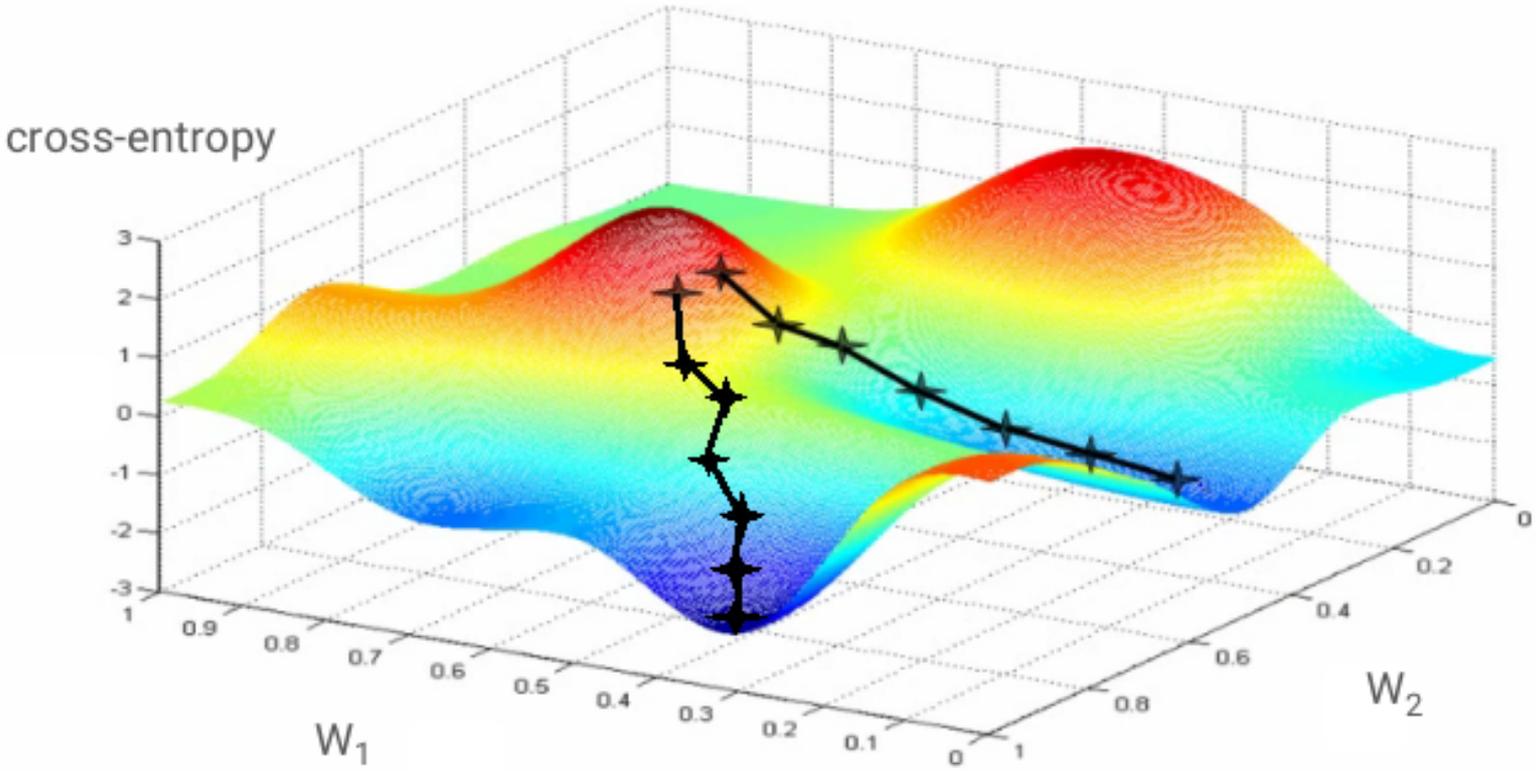
Cross entropy:  $-\sum Y_i' \cdot \log(Y_i)$

computed probabilities

0.1	0.2	0.1	0.3	0.2	0.1	0.9	0.2	0.1	0.1
0	1	2	3	4	5	6	7	8	9

this is a "6"

# Cross Entropy



# Training Loop

**Training digits and labels**

**=> loss function**

**=> gradient (partial derivatives)**

**=> steepest descent**

**=> update weights and biases**

**=> repeat with next mini-batch of training images and labels**

**"mini-batches":**  
**100 images and labels**

```
import tensorflow as tf
```

# mnist\_1.0\_softmax.py

```
import tensorflow as tf
X = tf.placeholder(tf.float32, [None, 28, 28, 1])
W = tf.Variable(tf.zeros([784, 10]))
b = tf.Variable(tf.zeros([10]))

init = tf.initialize_all_variables()
```

# mnist\_1.0\_softmax.py

```
# model
Y = tf.nn.softmax(tf.matmul(tf.reshape(X, [-1, 784]), W) + b)
# placeholder for correct labels
Y_ = tf.placeholder(tf.float32, [None, 10])

# loss function
cross_entropy = -tf.reduce_sum(Y_ * tf.log(Y))
# % of correct answers found in batch
is_correct = tf.equal(tf.argmax(Y, 1), tf.argmax(Y_, 1))
accuracy = tf.reduce_mean(tf.cast(is_correct, tf.float32))
```

# mnist\_1.0\_softmax.py

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

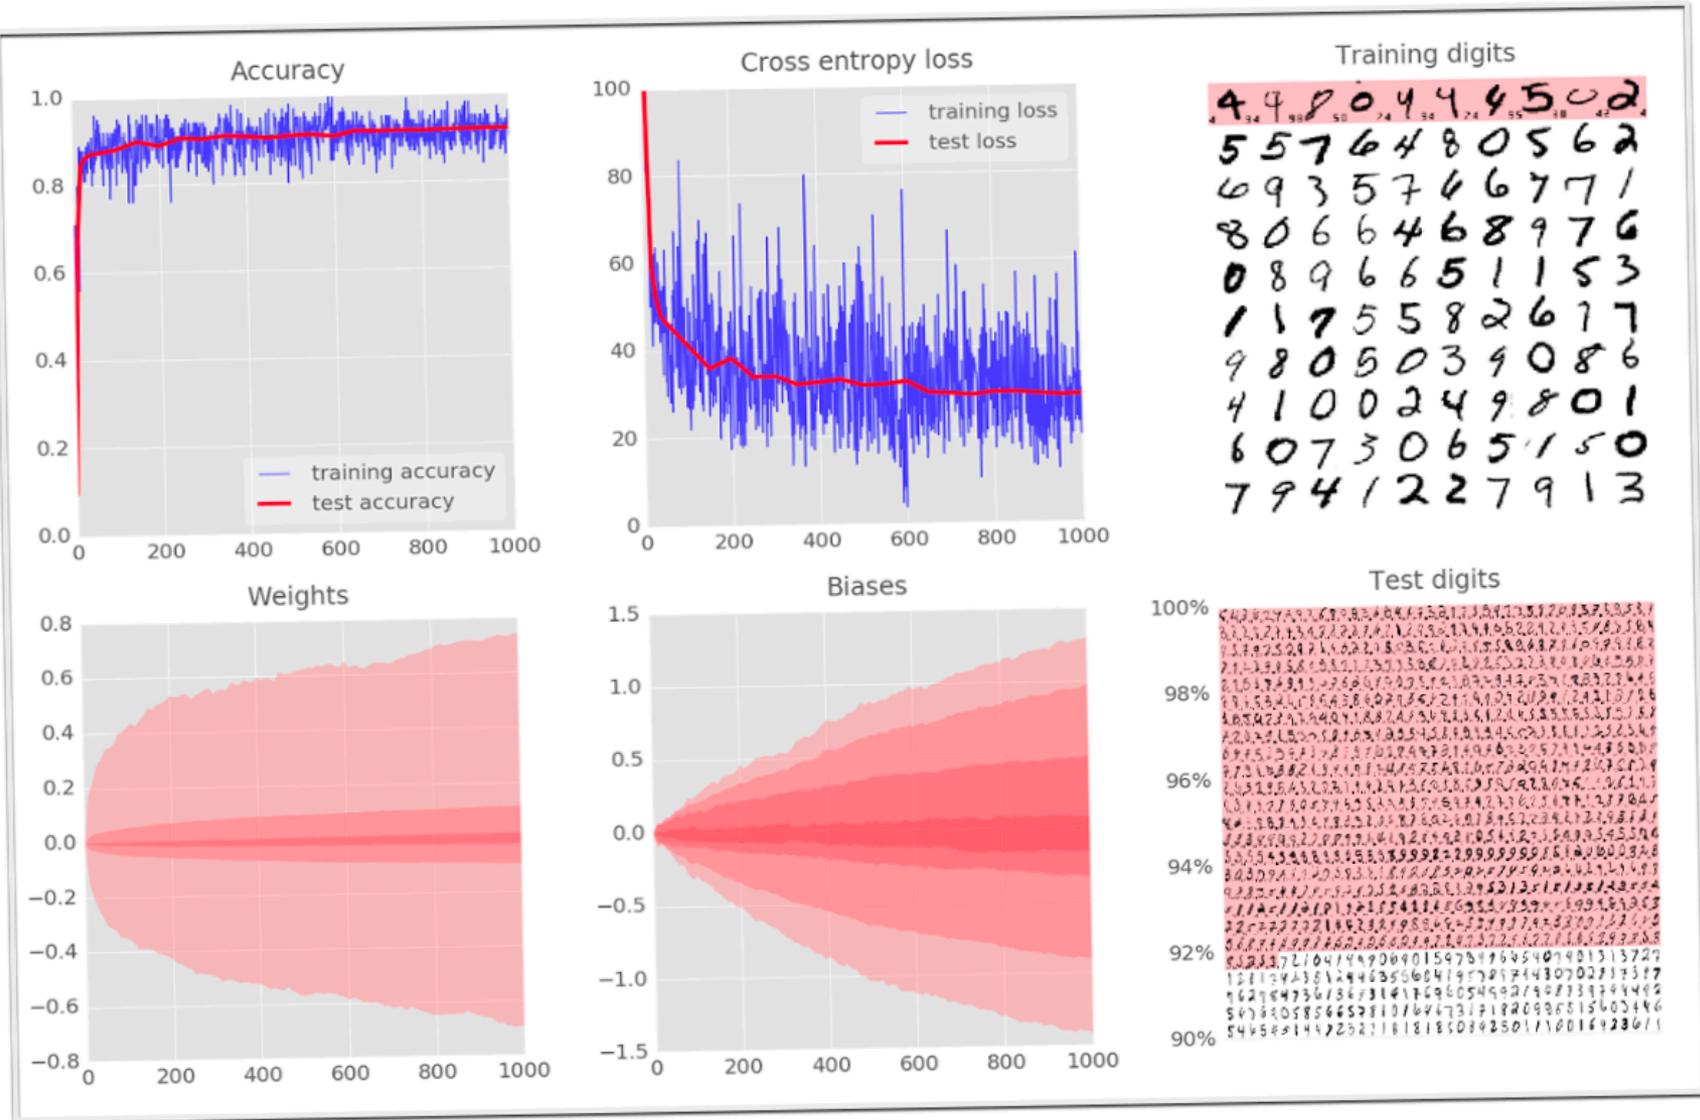
for i in range(1000):
    # load batch of images and correct answers
    batch_X, batch_Y = mnist.train.next_batch(100)
    train_data={X: batch_X, Y_: batch_Y}

    # train
    sess.run(train_step, feed_dict=train_data)
```

# mnist\_1.0\_softmax.py

```
# success ?  
a,c = sess.run([accuracy, cross_entropy],  
feed_dict=train_data)  
  
# success on test data ?  
test_data={X: mnist.test.images, Y_: mnist.test.labels}  
a,c = sess.run([accuracy, cross_entropy], feed=test_data)
```

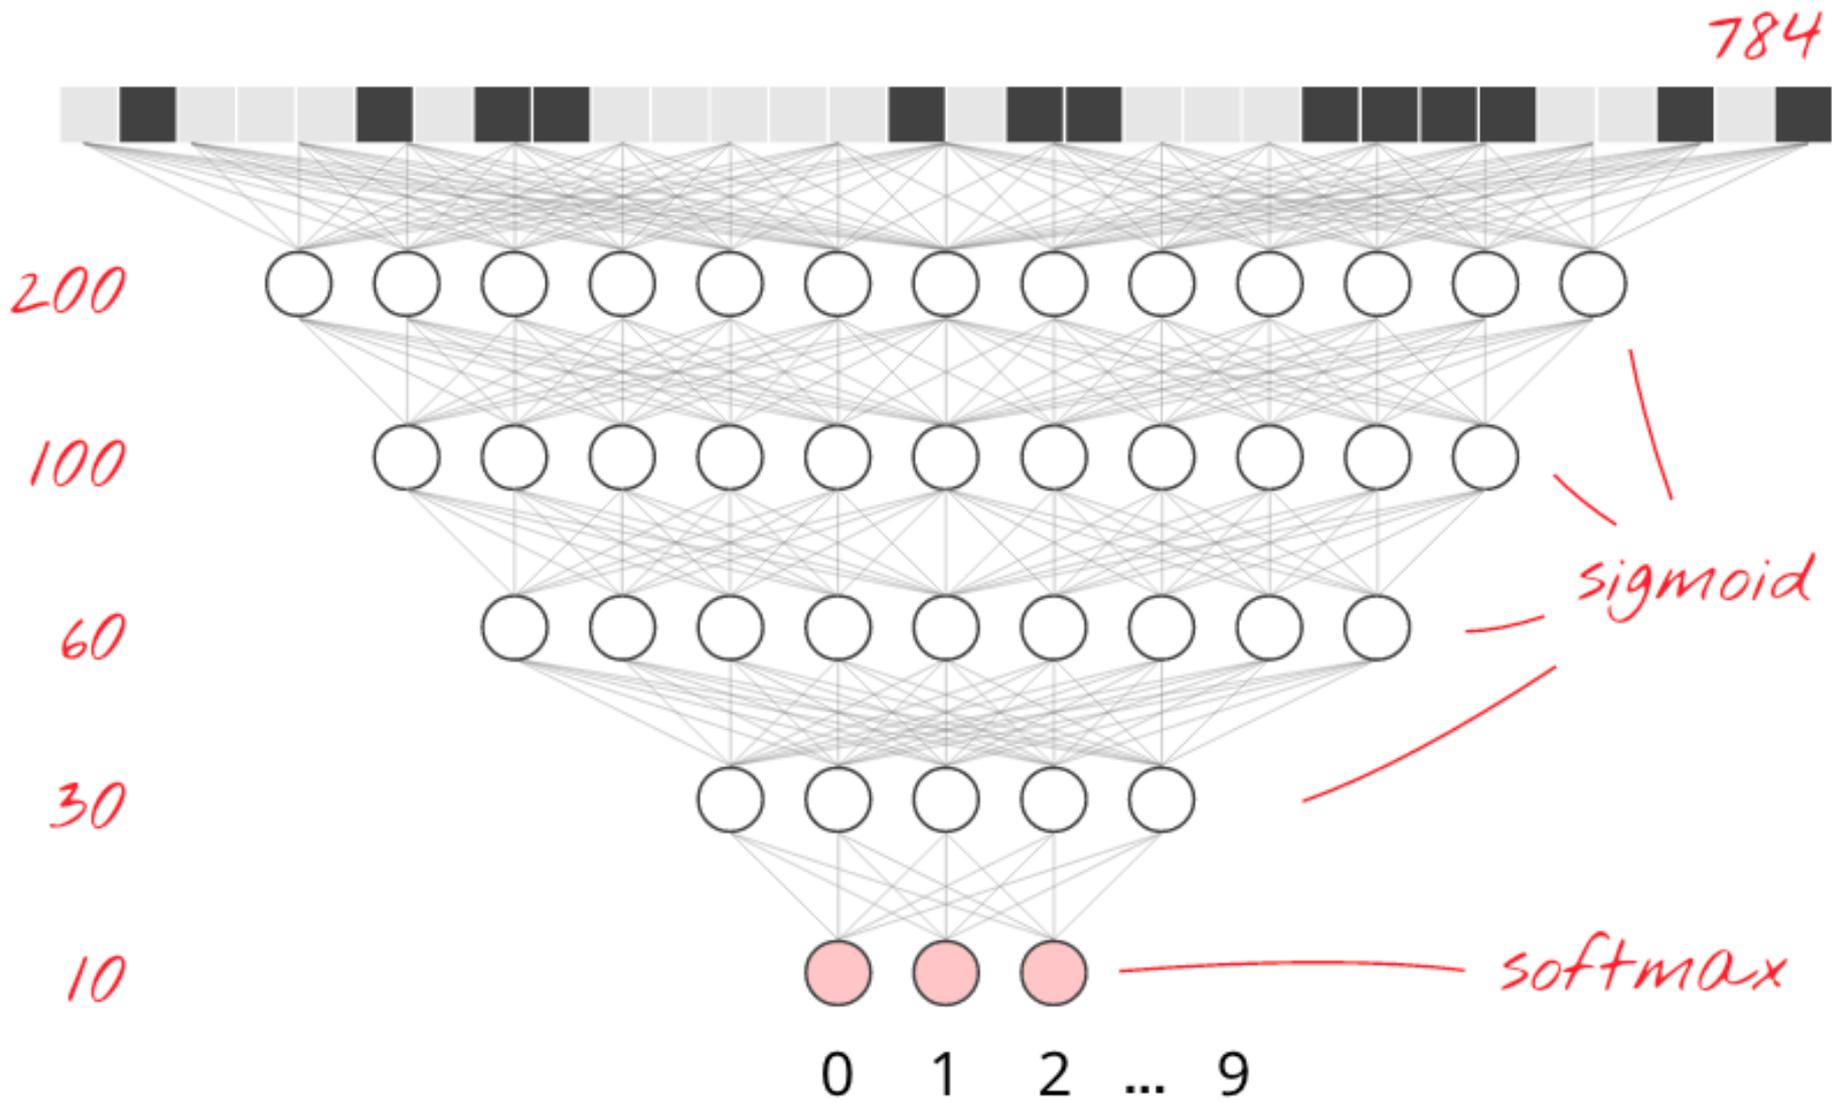
# mnist\_1.0\_softmax.py



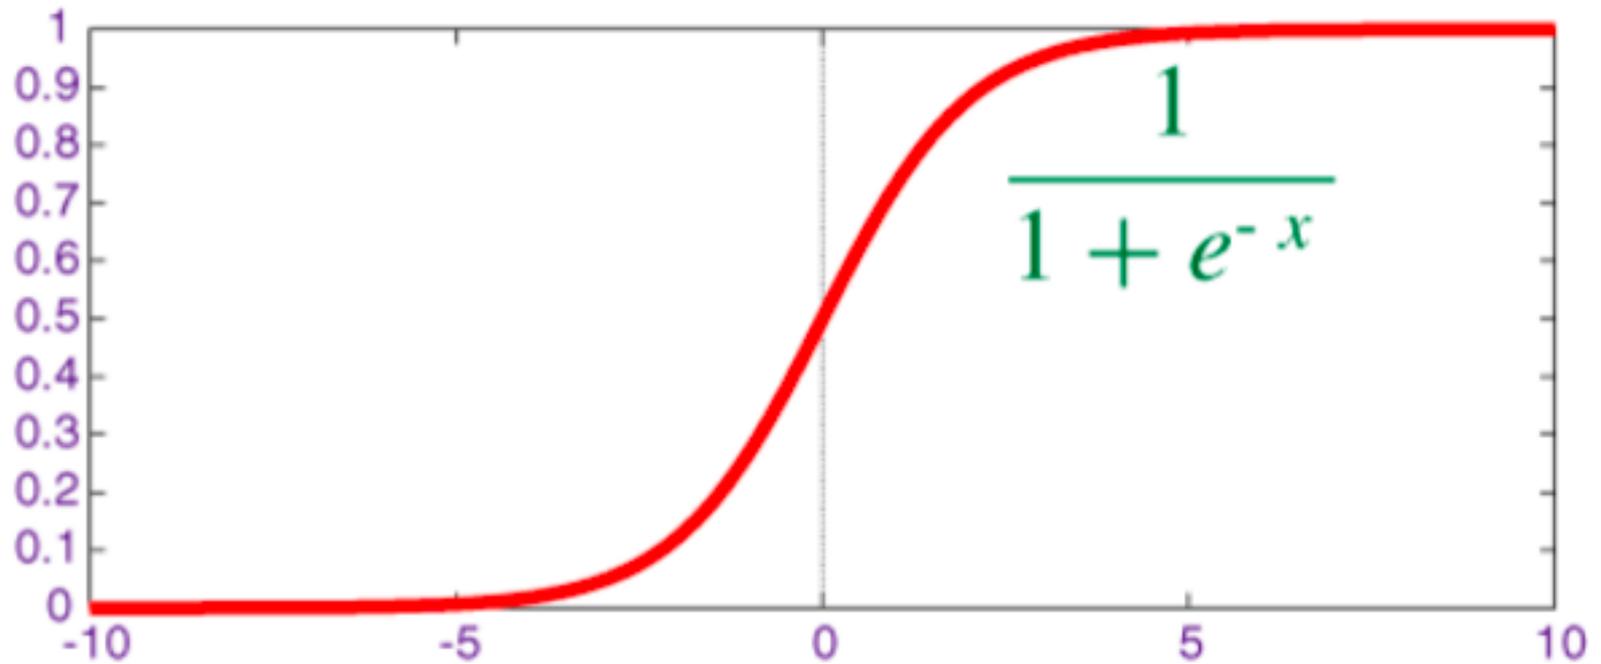
92!  
😊

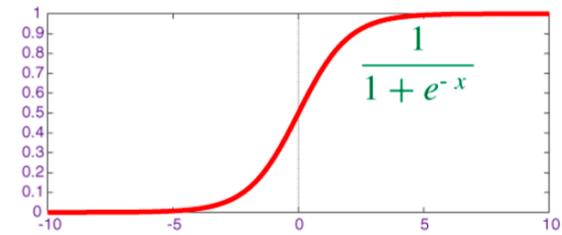
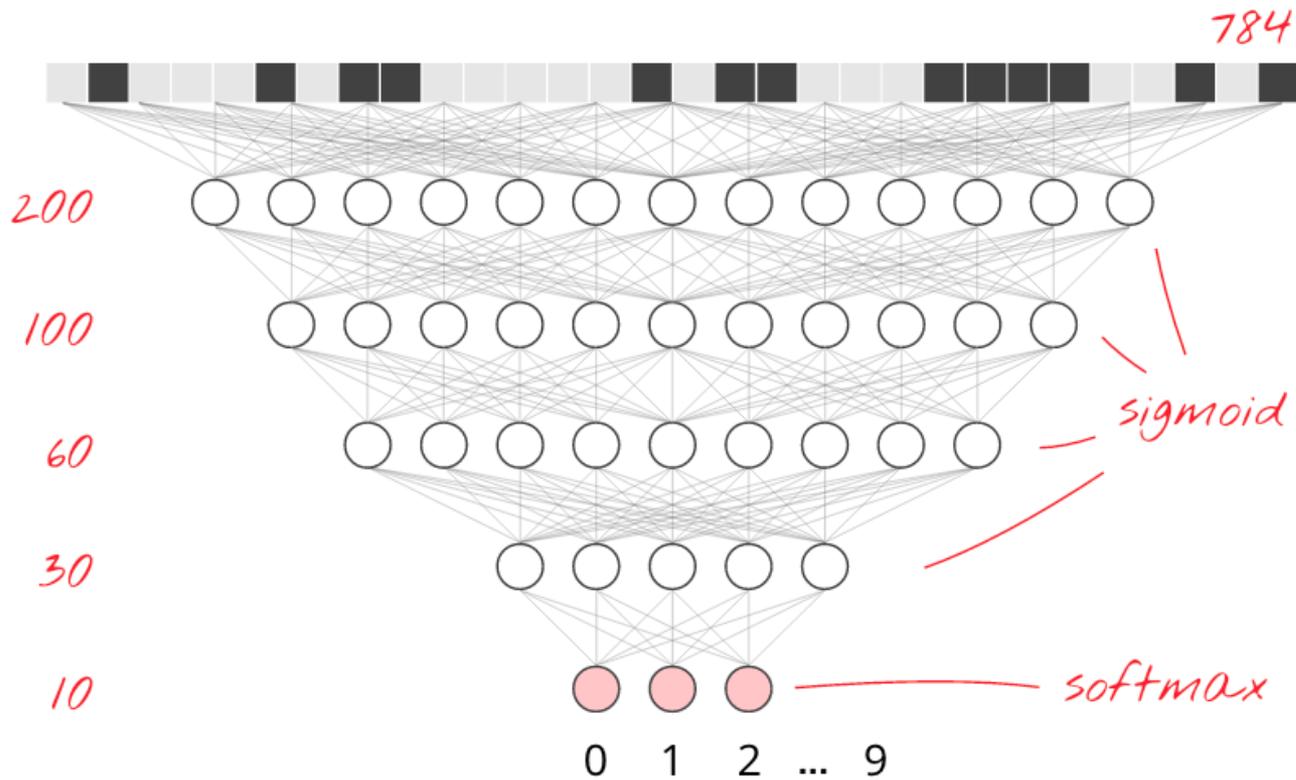
# Deep Learning

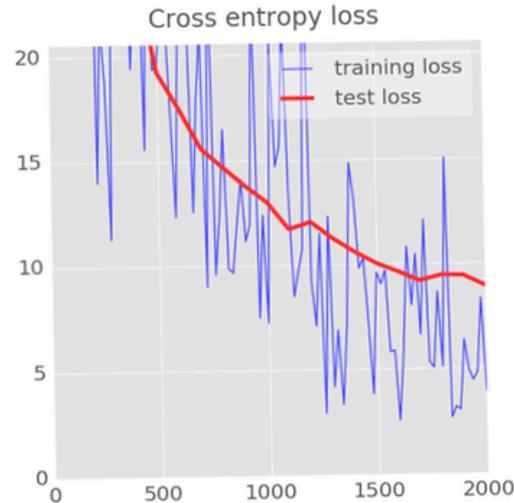
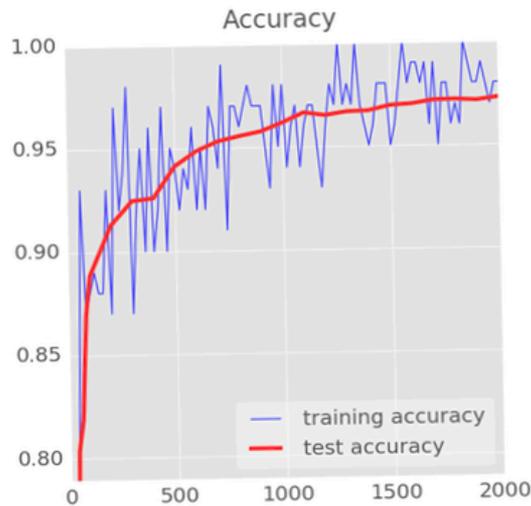




# Sigmoid



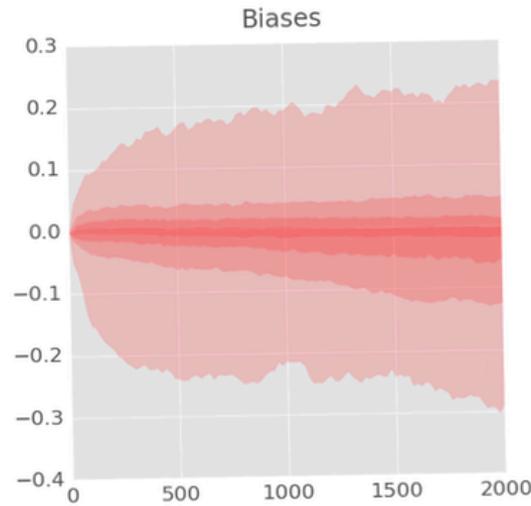
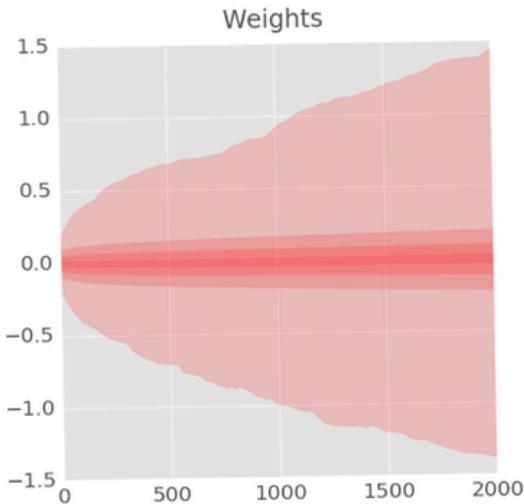




### Training digits

```

8 3 8 7 3 9 4 4 8 7
0 0 0 3 1 9 0 4 1 8
8 5 4 6 6 0 1 6 9 7
8 7 5 0 2 7 6 9 9 6
8 8 2 6 8 2 1 4 1 0
2 8 0 2 0 4 0 3 4 4
0 2 4 7 7 5 8 2 3 4
7 5 4 3 1 6 5 8 2 3
6 8 2 6 2 6 4 8 2 7
7 4 5 2 5 8 4 8 4 5
  
```



### Test digits

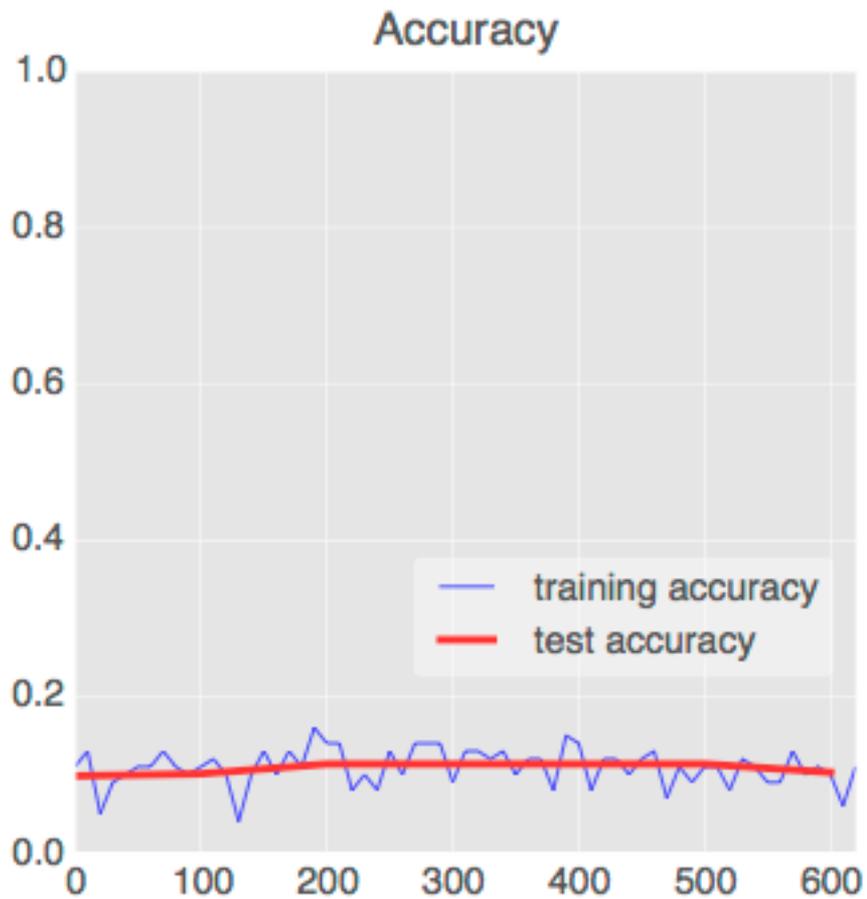
```

100% 6468733082219468555966116976311017100025
51117930276495141371671135116839020181
225375131901326511255222140805488123
77227247683514520325529110207051554313
471658493112830834311470111048999701583
8833599000291121001177133911118096899
81822222283200222100118888011128872104144
906901597347645407401313472712113435124
46356641977897184307027173217162744736
13671417694054990719873974442507644058
5665710164773122009865154034665445851
4472321181818502425011105011692361115952
4459390365939272811338879224158933042
41407712237711818030199419212915264152
920400287124027433005194935193591071121
5339785138105131594751194421506537208
854114031618274281982544238524603171574
7172140292049148154949332100352647533391
2680566622227583618412381975408491052379
101395213136572226226189773038314394642
12254534002327368744963098042063493313
327102170653301639099420657860202711755
101462474939224278591802051131671162039
4641067143741931734769137336427881144931
90% 01707944855408211404061132692693114025120
  
```

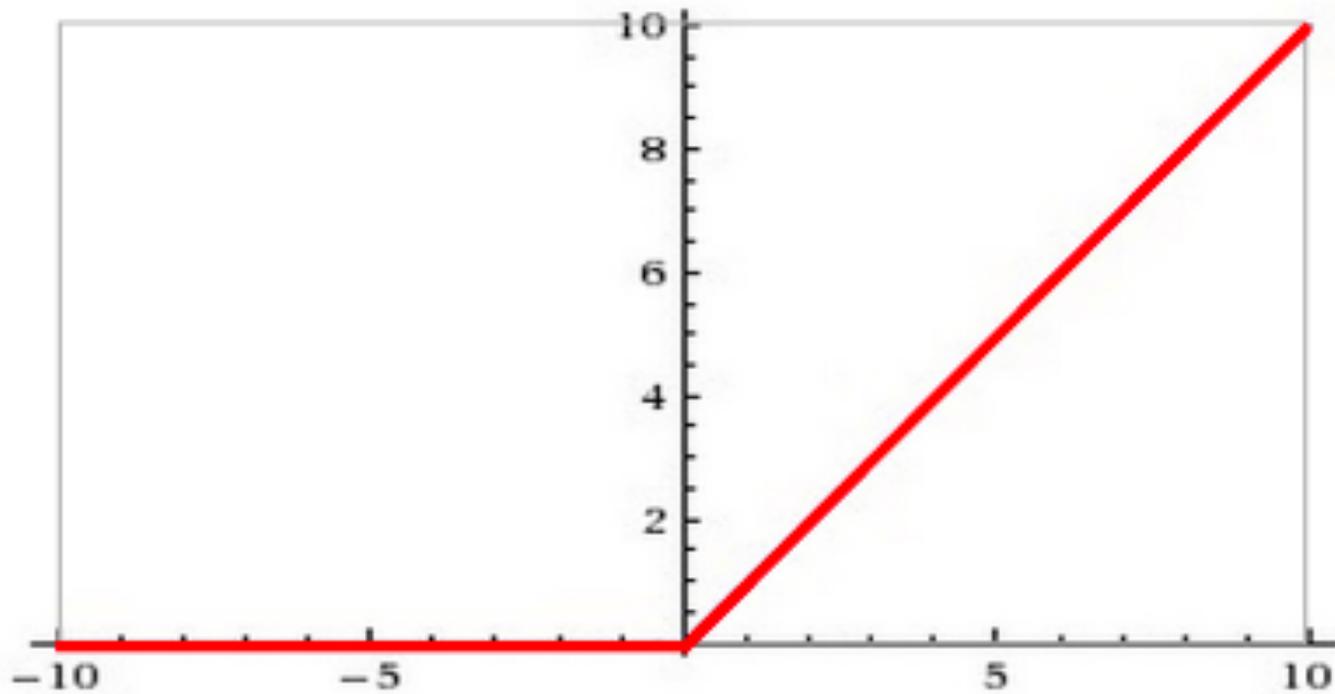
97!  
😊

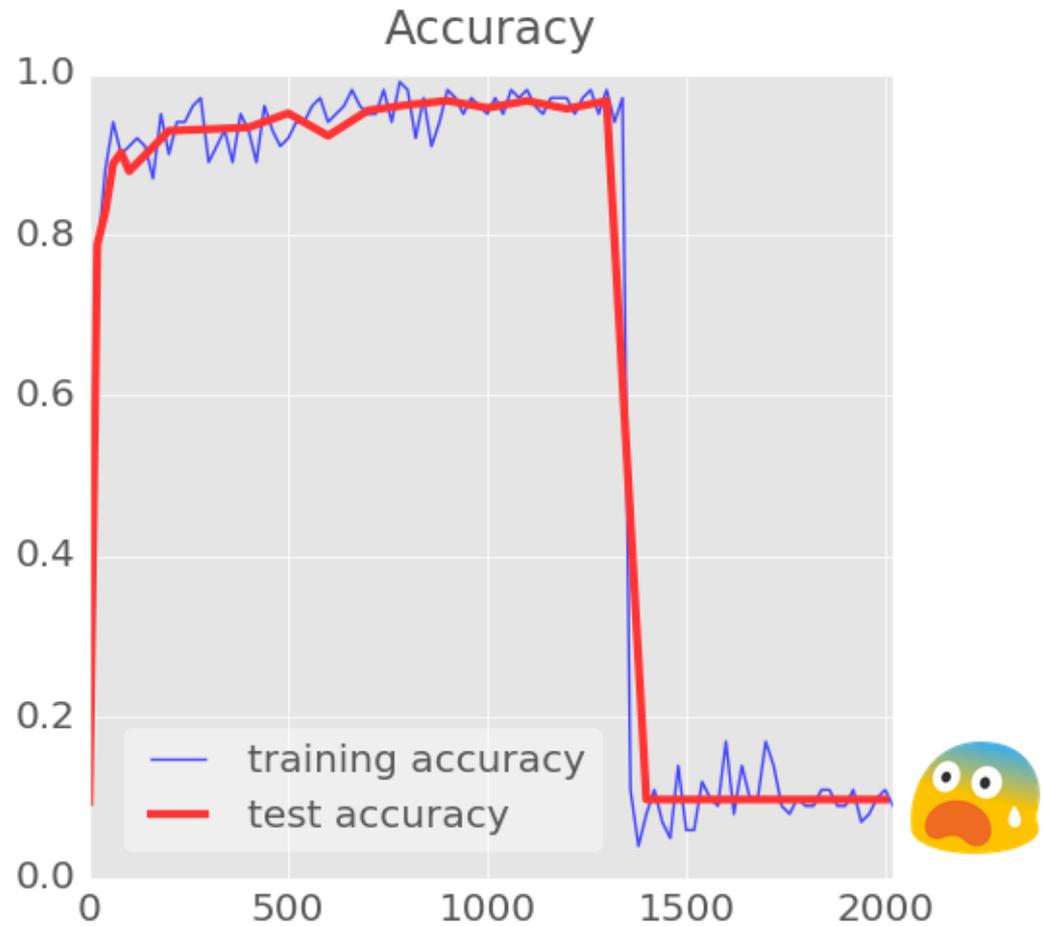
# ReLU





# ReLU





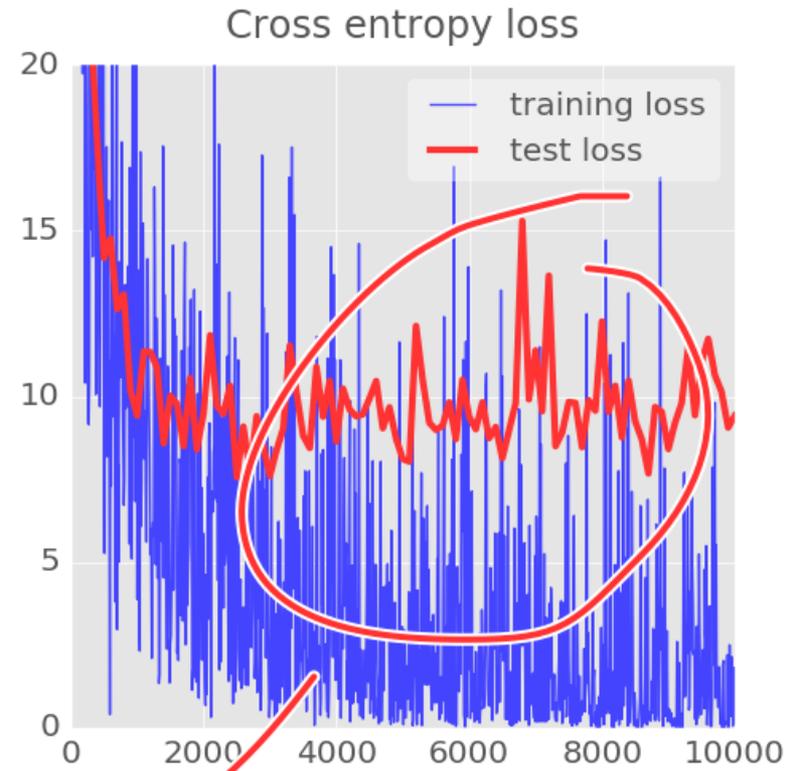
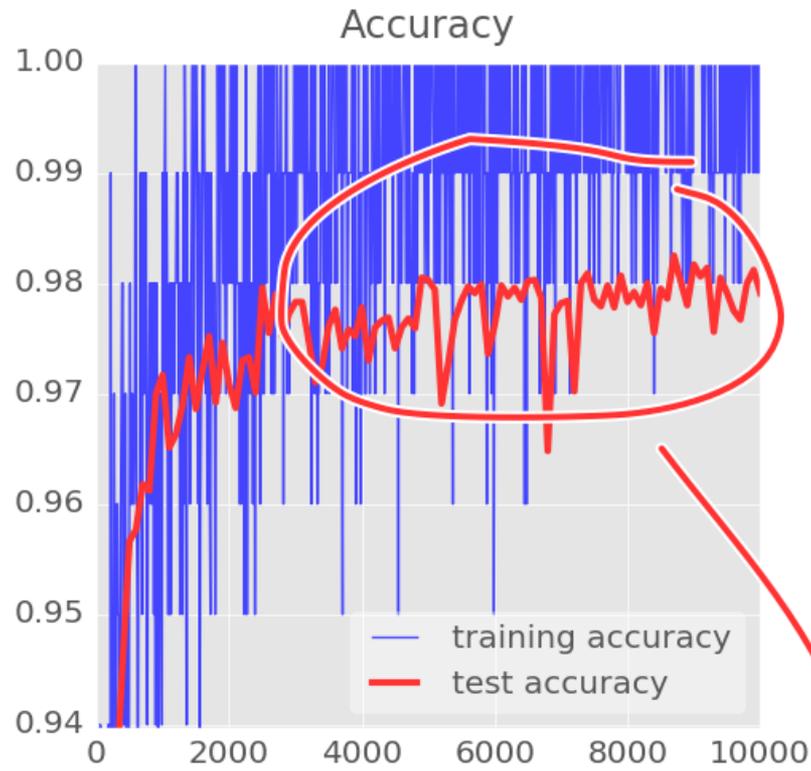
# Learning Rate

Slow down...

Learning  
rate decay

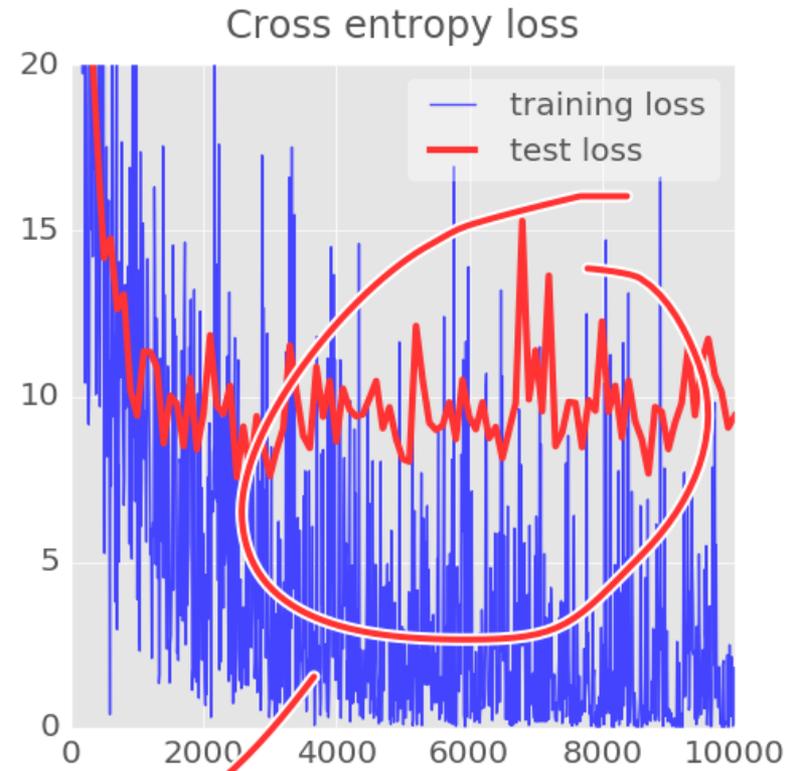
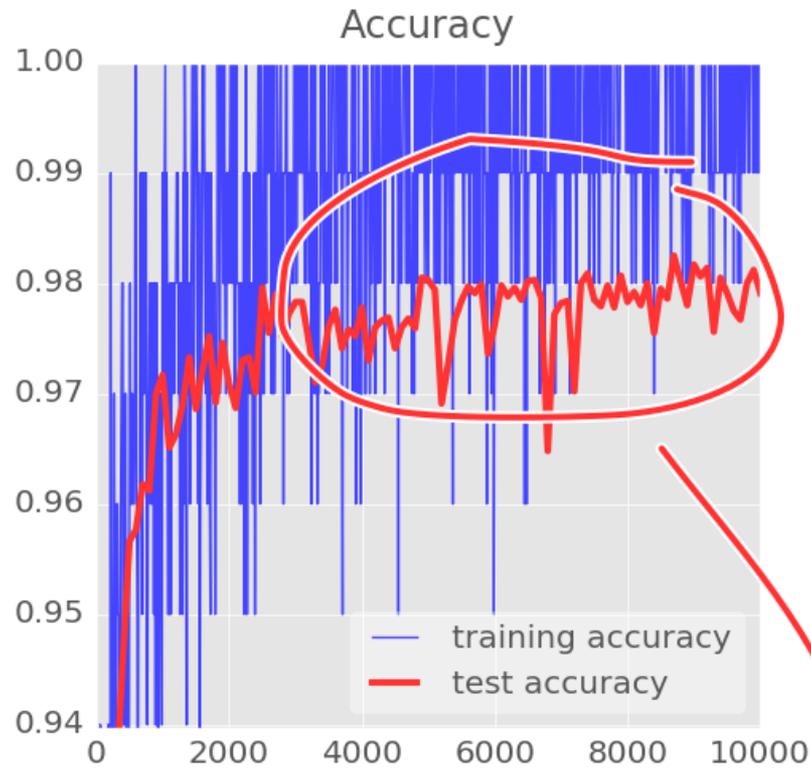


LR =  
0.003

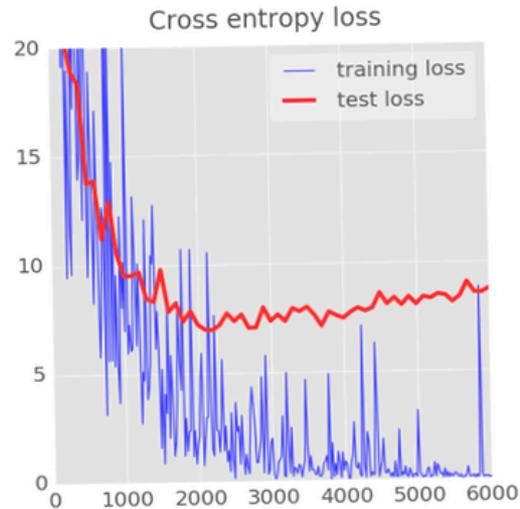
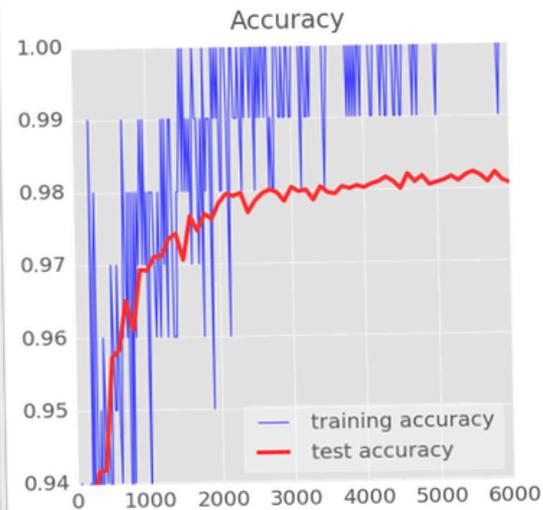


yuck!

LR =  
0.003



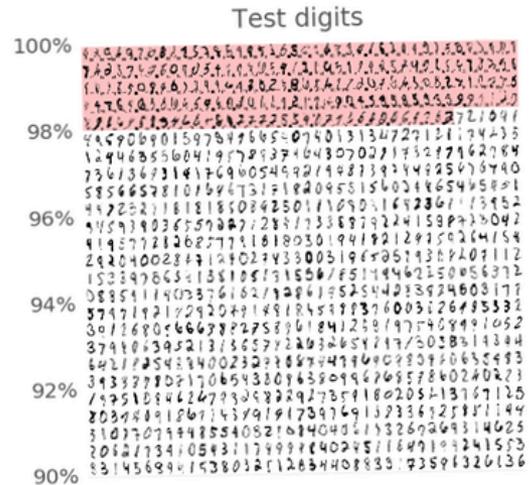
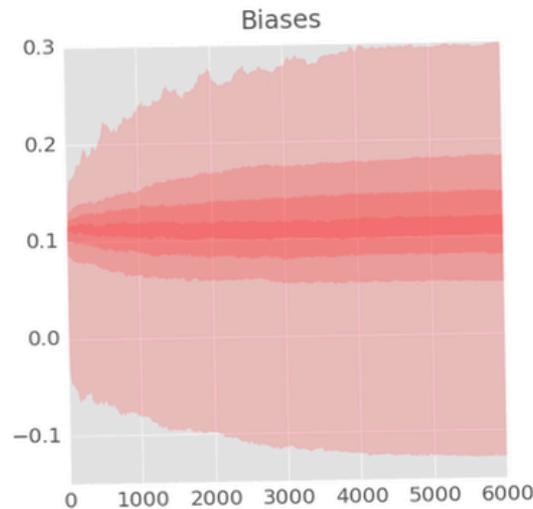
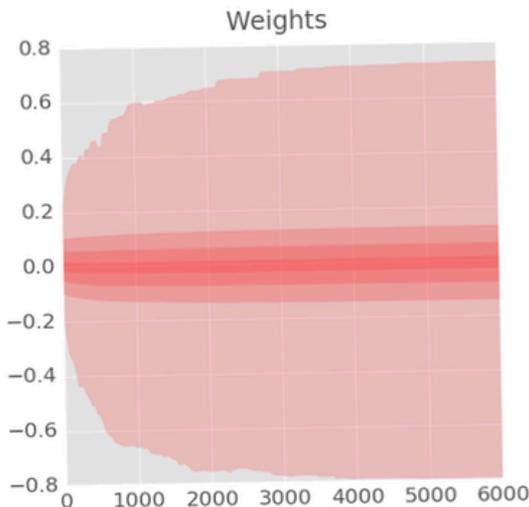
yuck!



### Training digits

```

5 8 0 6 2 8 2 8 6 5
8 4 4 / 6 7 9 9 4 5
0 7 8 7 8 0 7 1 0 7
5 5 7 5 5 8 3 8 5 9
6 6 2 4 0 8 6 7 5 1
3 9 1 0 5 4 0 5 0 1
6 4 8 3 7 3 7 8 1 1
1 6 8 8 0 2 6 5 7 0
0 6 6 7 5 4 8 6 4 5
1 8 7 2 5 0 1 5 2 9
  
```

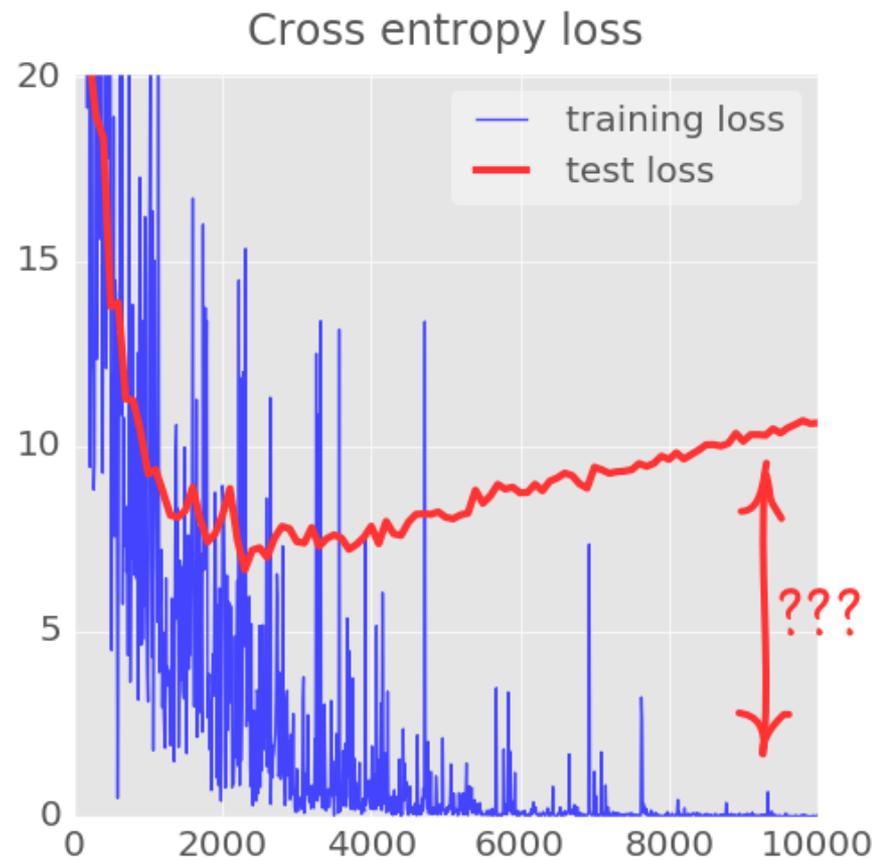
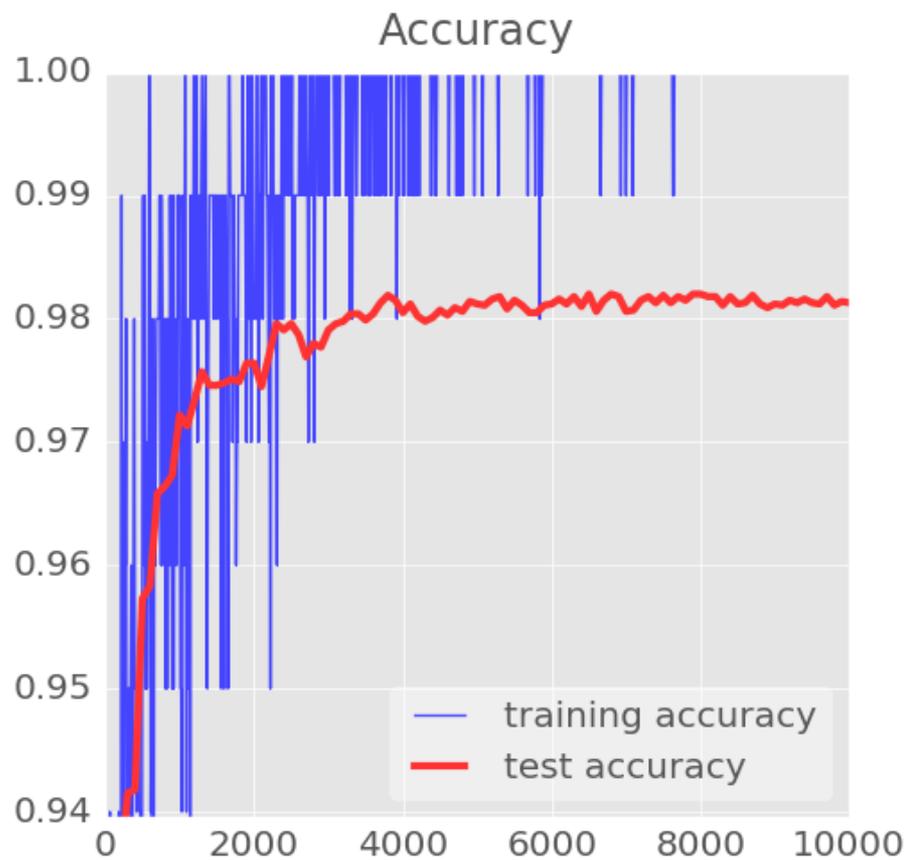


98!  
😊

# Dropout

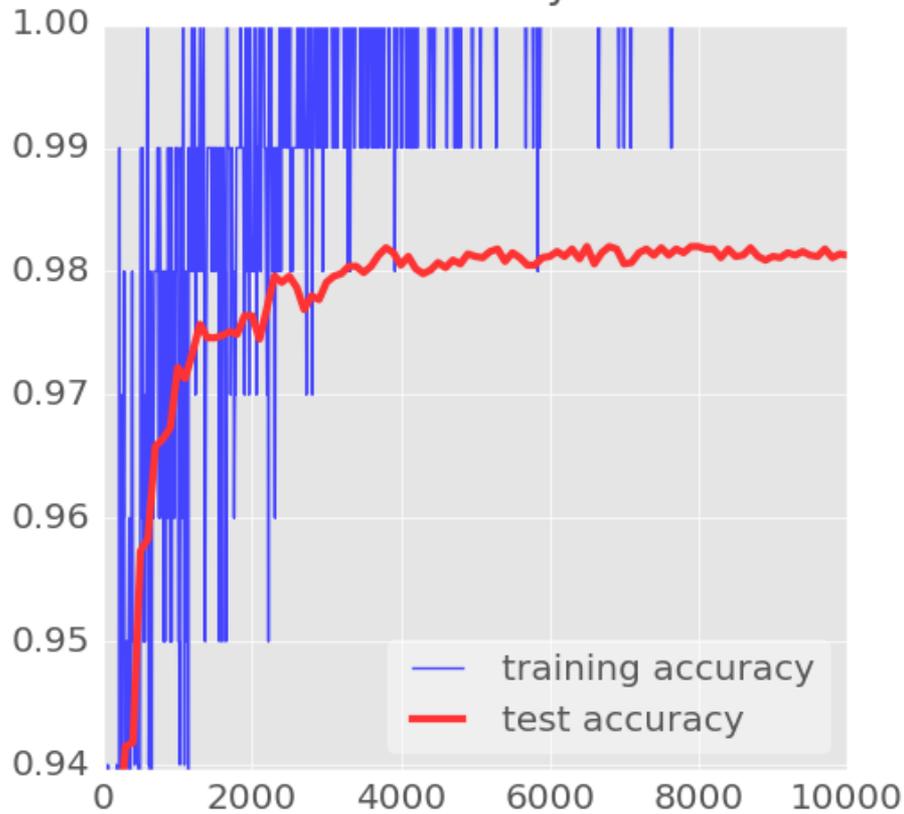
Dropout



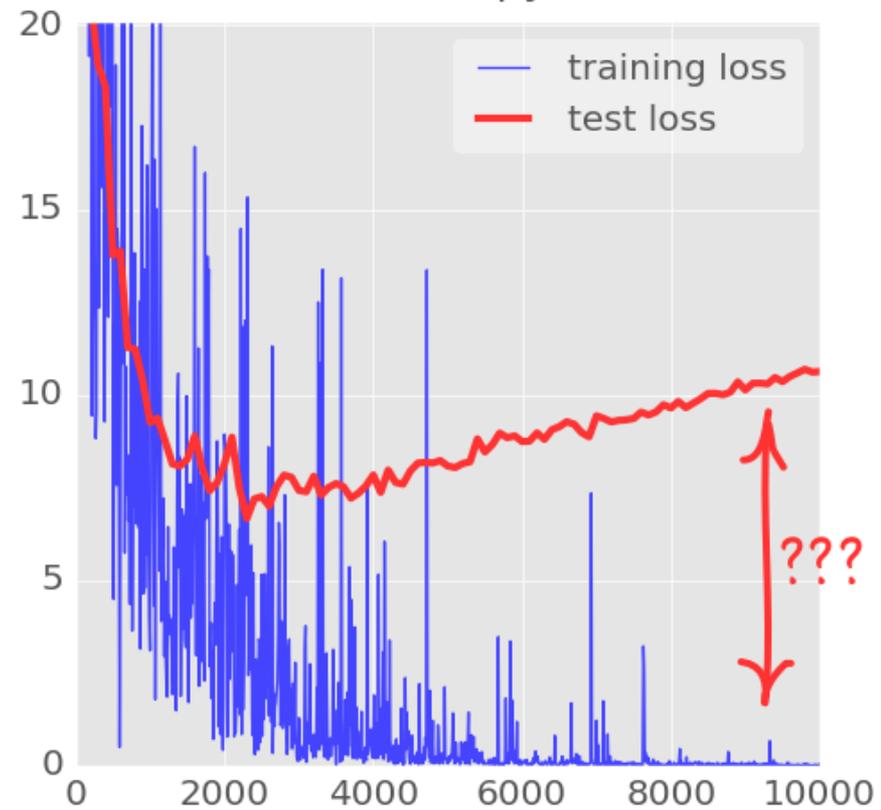


# Overfitting

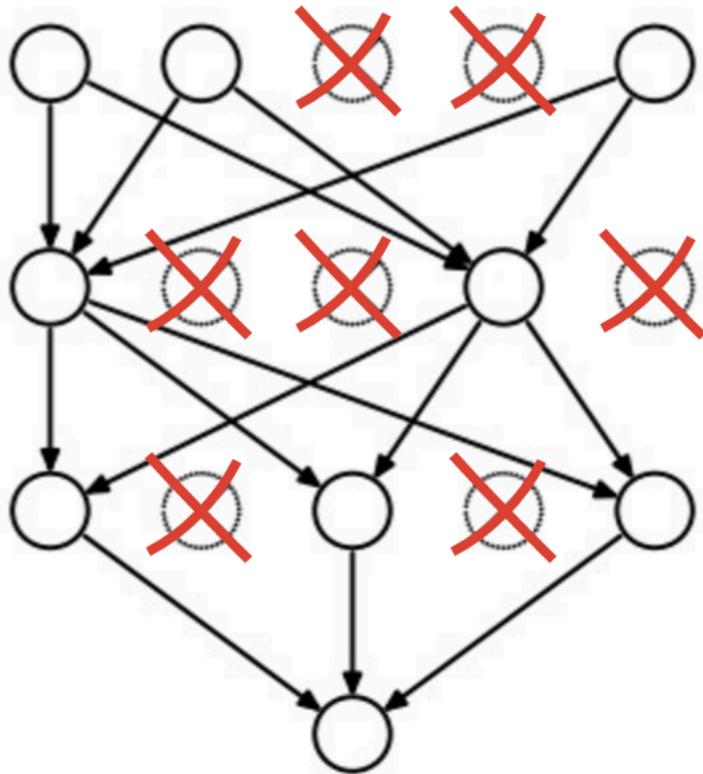
Accuracy



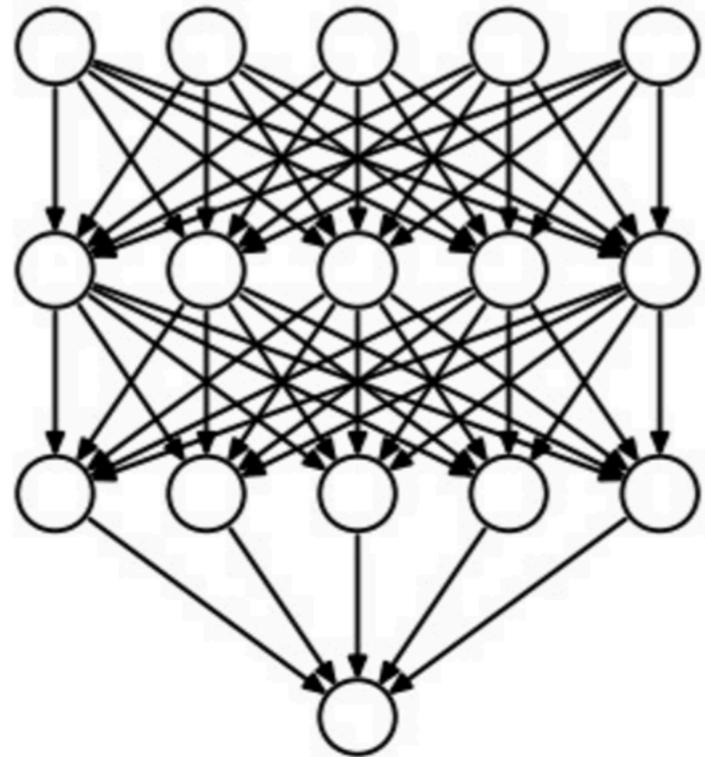
Cross entropy loss



# Dropout

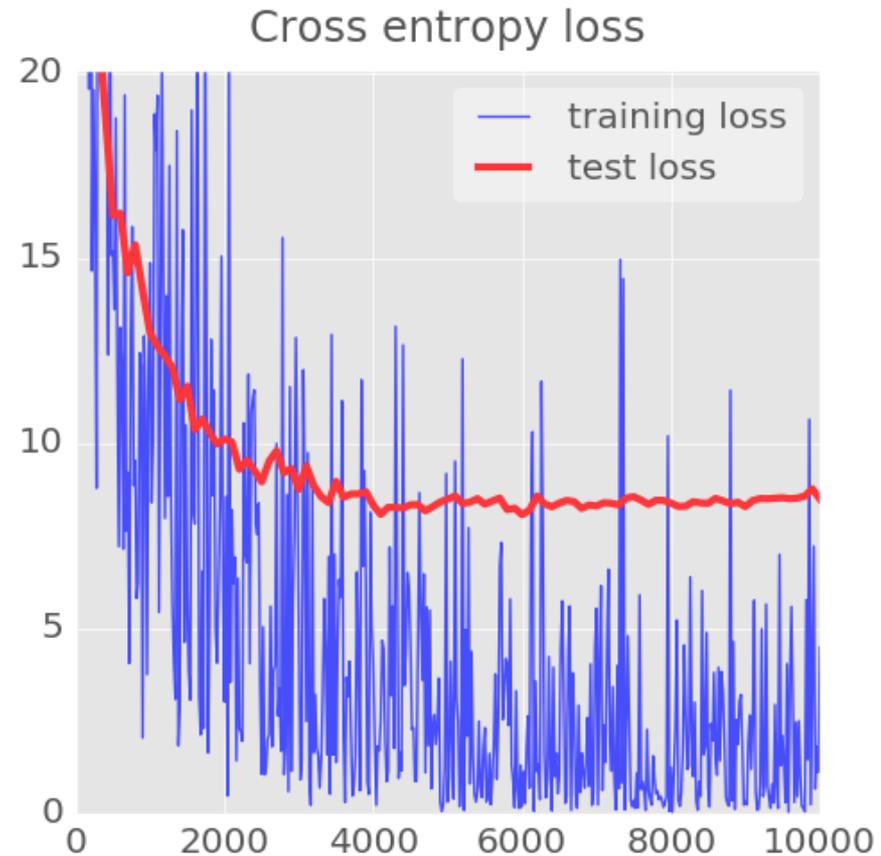
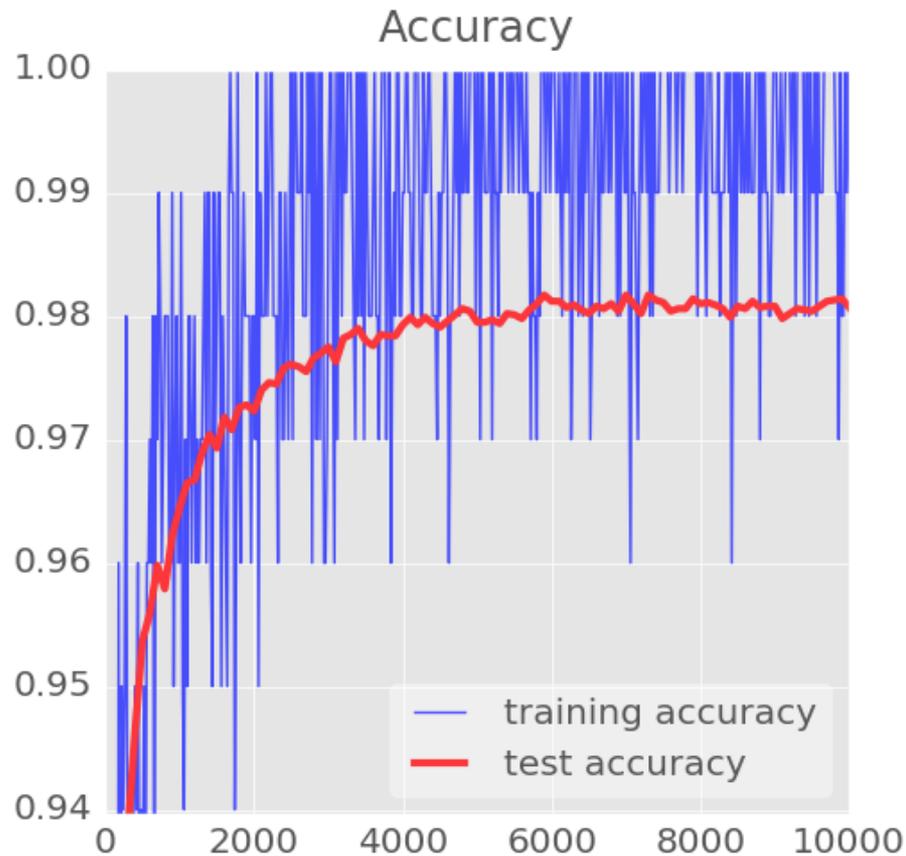


Training  
 $p_{\text{keep}} = 0.75$

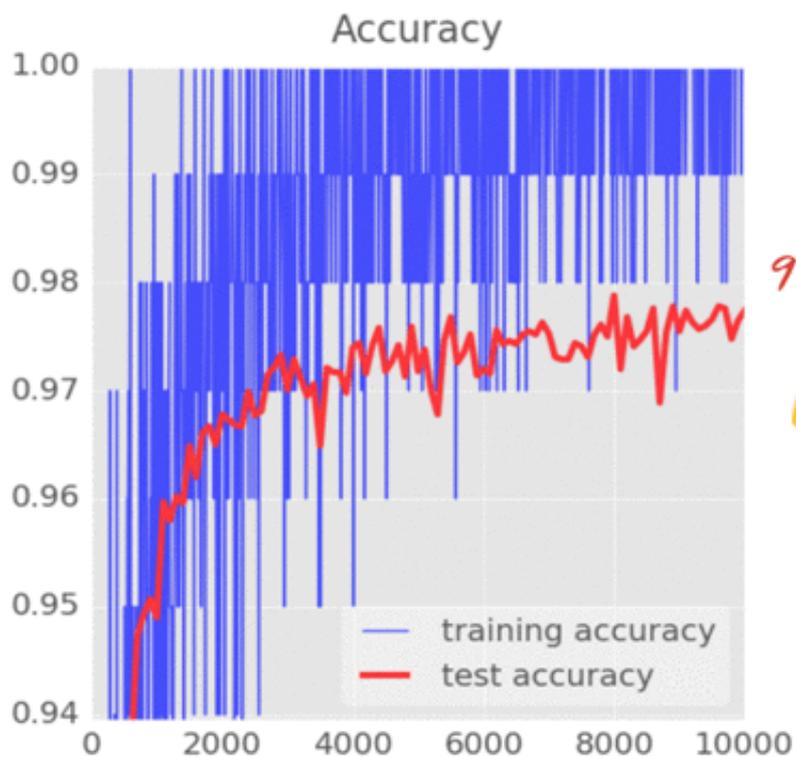


Test  
 $p_{\text{keep}} = 1.0$

# TensorFlow MNIST Tutorial

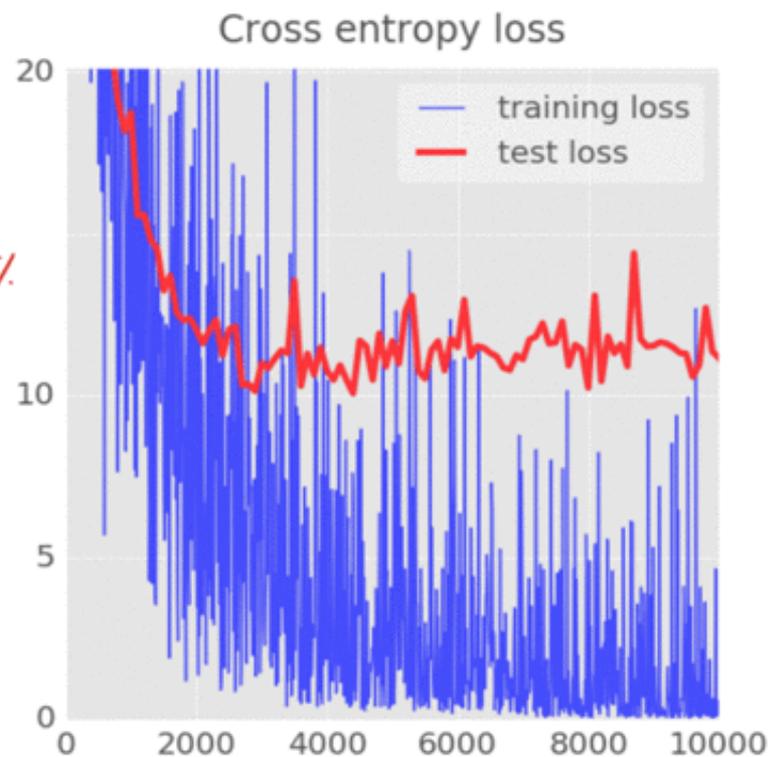


# TensorFlow MNIST Tutorial

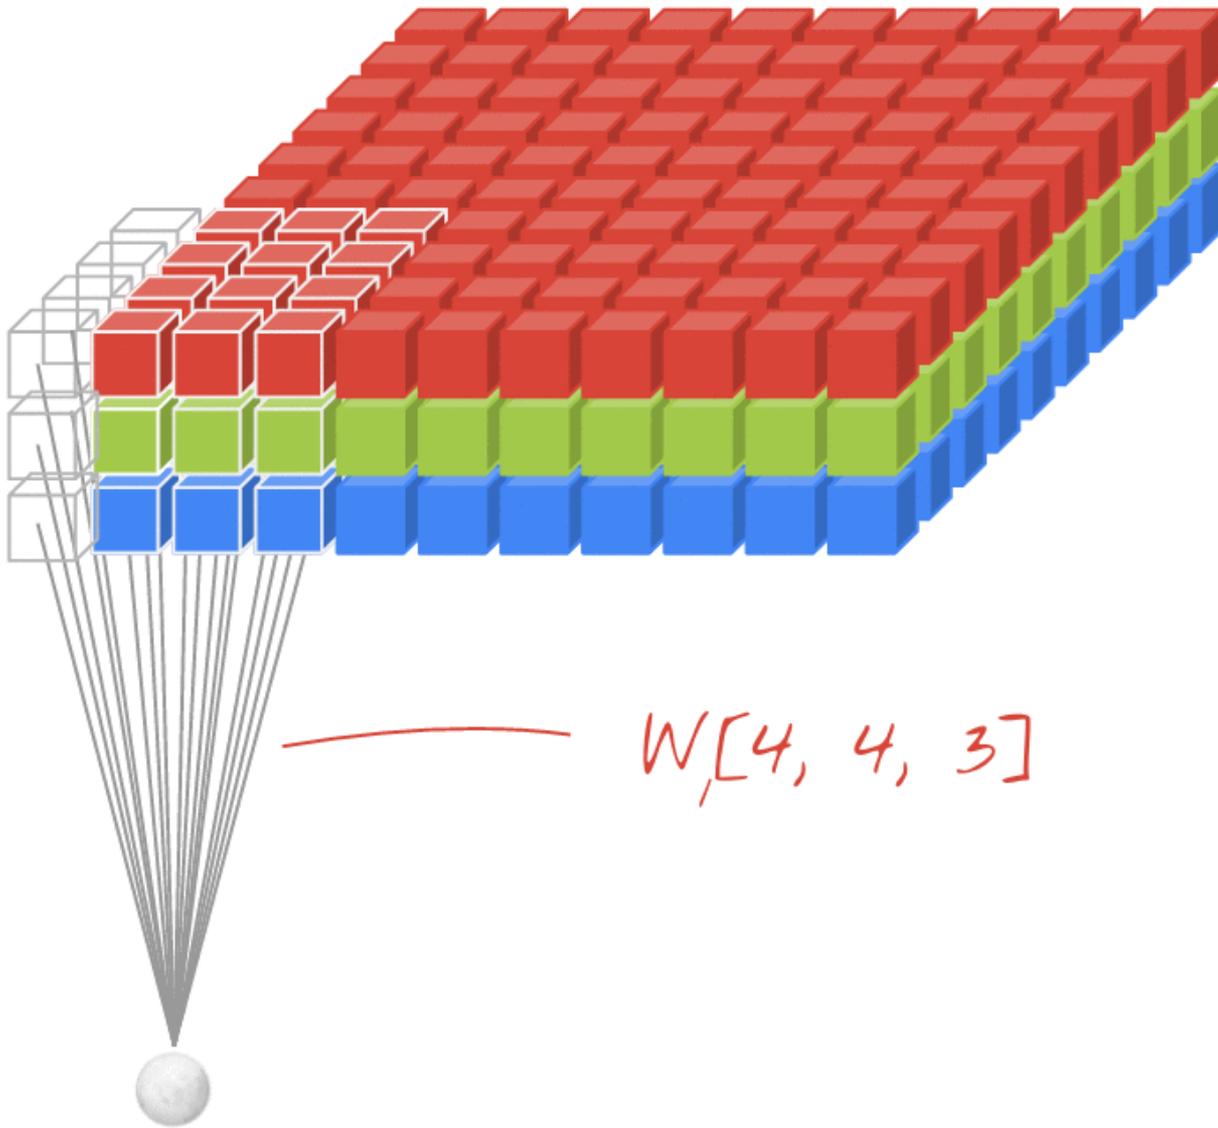


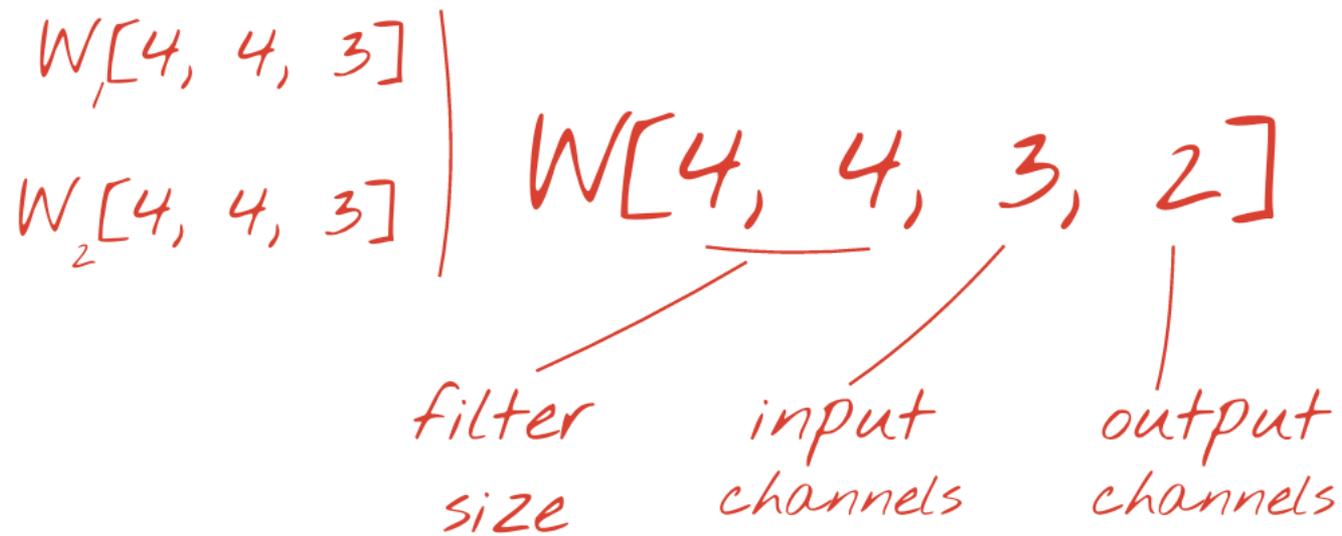
5 layers  
sigmoid

97.9%

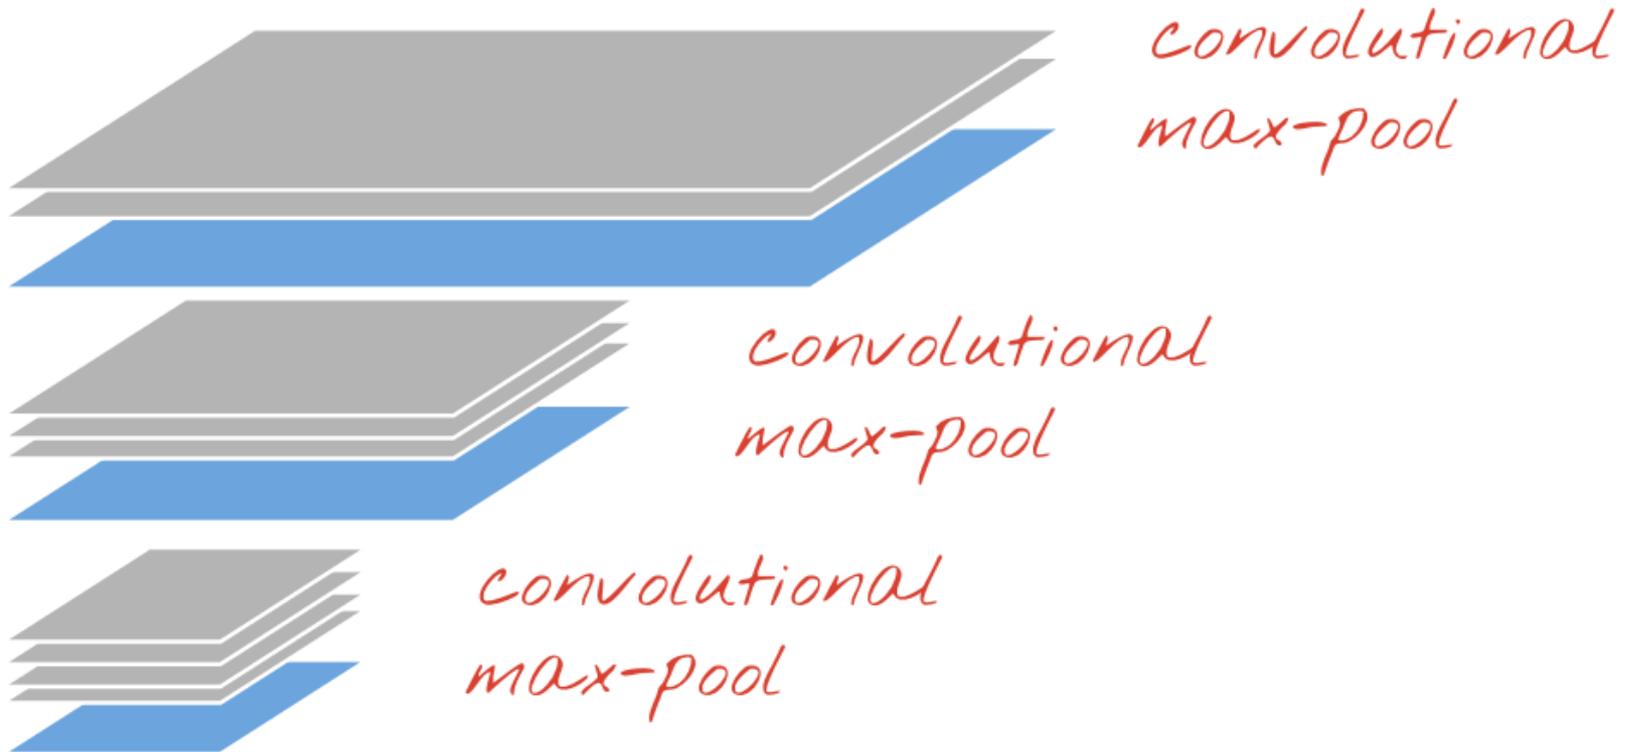








# Convolutional Max-Pool



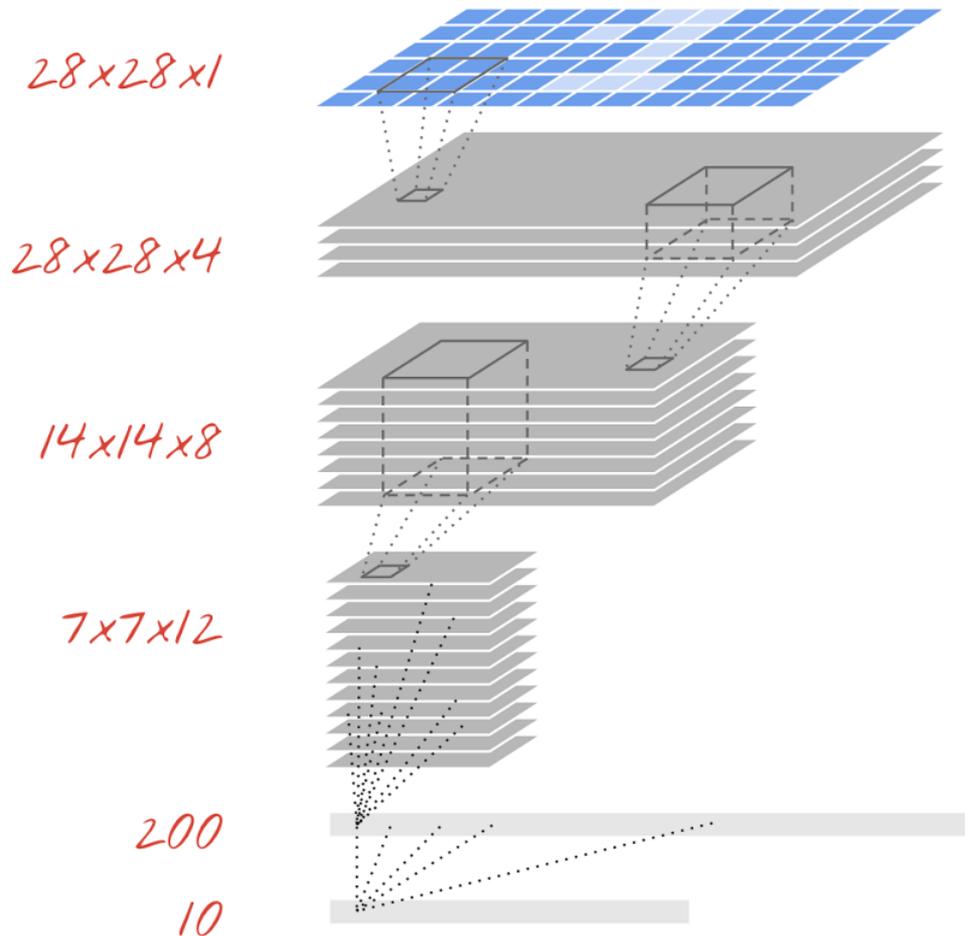
# All Convolutional

Hacker's tip



ALL  
Convolutional

+ biases on  
all layers



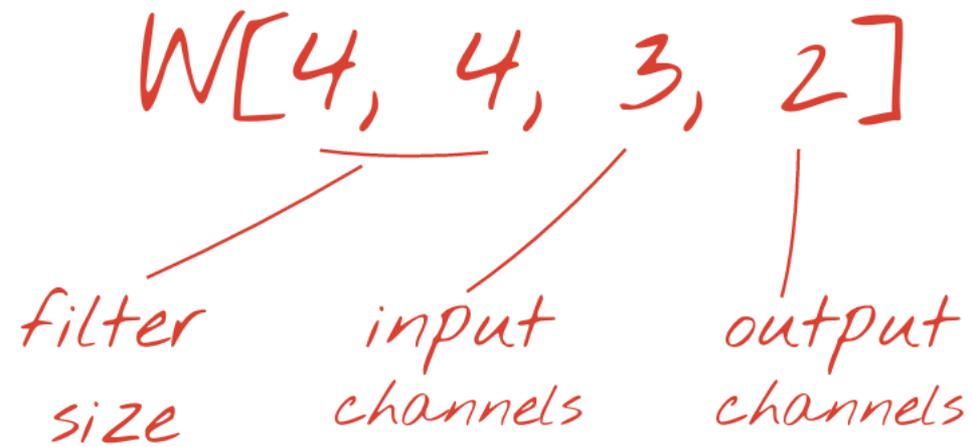
convolutional layer, 4 channels  
 $W1[5, 5, 1, 4]$  stride 1

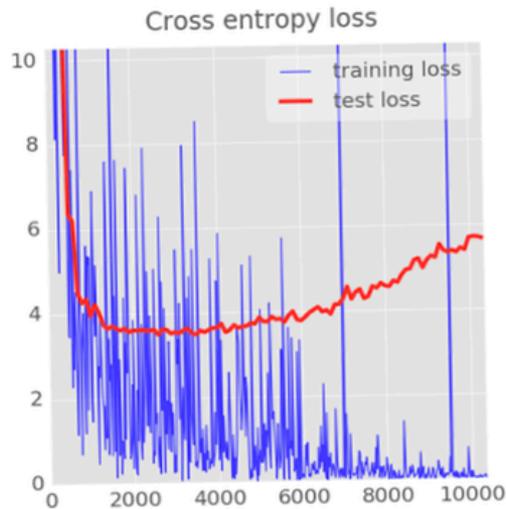
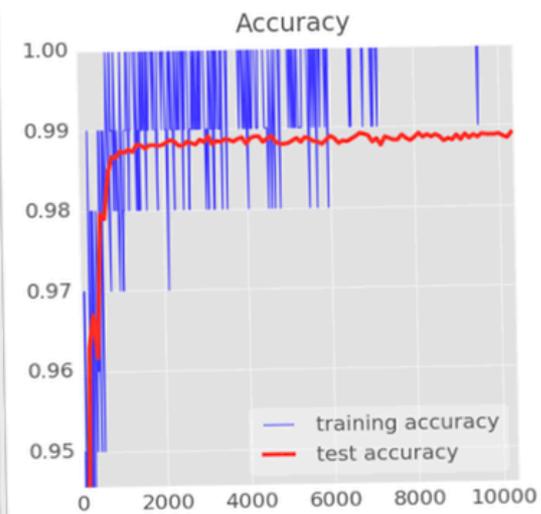
convolutional layer, 8 channels  
 $W2[4, 4, 4, 8]$  stride 2

convolutional layer, 12 channels  
 $W3[4, 4, 8, 12]$  stride 2

fully connected layer  
 $W4[7 \times 7 \times 12, 200]$

softmax readout layer  
 $W5[200, 10]$

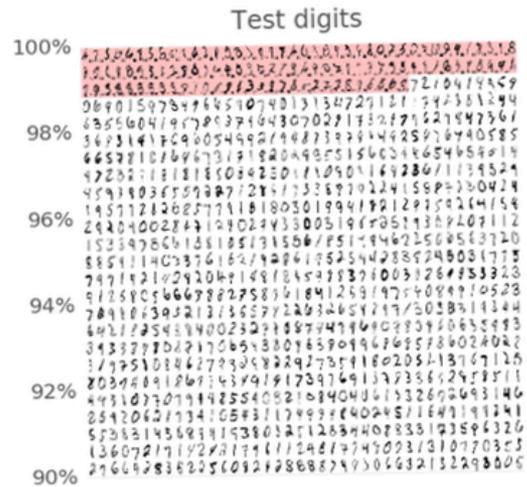
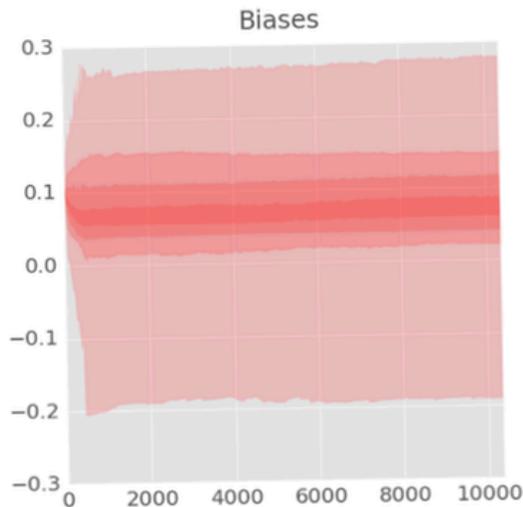
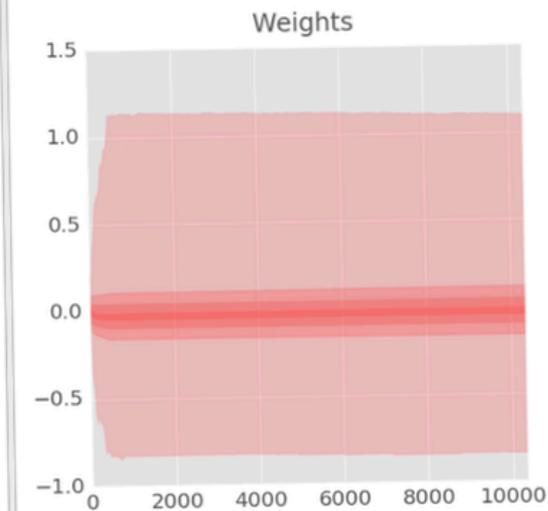




### Training digits

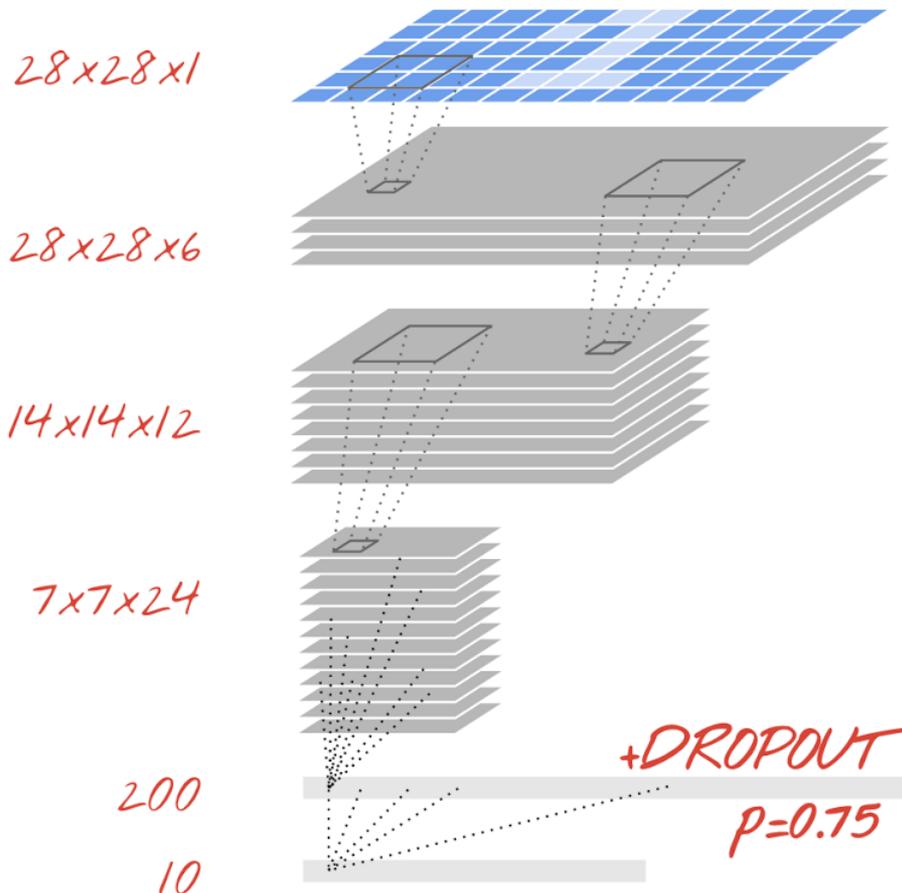
```

7 4 9 6 2 9 1 6 1 1
1 0 6 6 6 6 1 3 3 3
7 0 7 2 6 1 9 5 4 6
4 4 7 8 9 5 2 7 5 2
7 0 9 0 4 5 0 6 5 4
3 1 5 2 6 2 5 2 7 7
1 2 1 7 0 6 9 0 3 9
5 9 9 9 7 2 7 7 6 5
9 9 2 8 7 1 3 7 4 7
1 0 1 6 5 7 8 7 8 6
  
```



98.9%  
😊

+ biases on  
all layers



convolutional layer, 6 channels  
 $W1[6, 6, 1, 6]$  stride 1

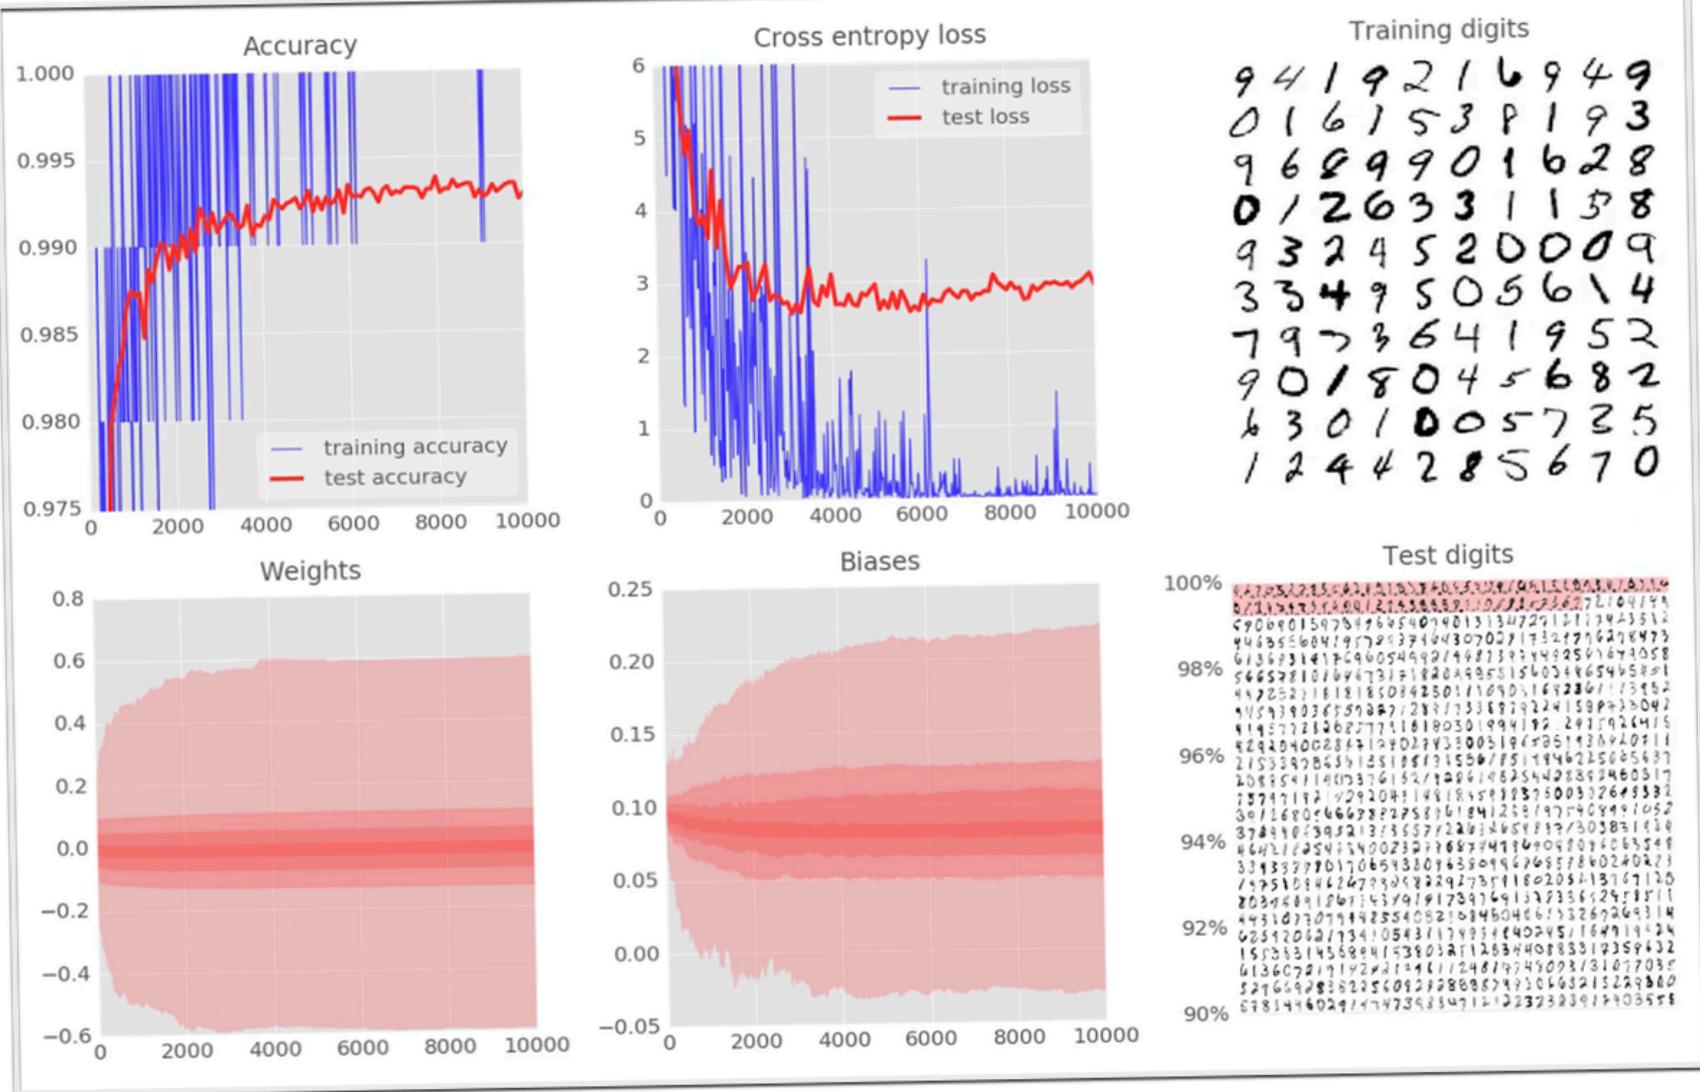
convolutional layer, 12 channels  
 $W2[5, 5, 6, 12]$  stride 2

convolutional layer, 24 channels  
 $W3[4, 4, 12, 24]$  stride 2

fully connected layer  
 $W4[7 \times 7 \times 24, 200]$

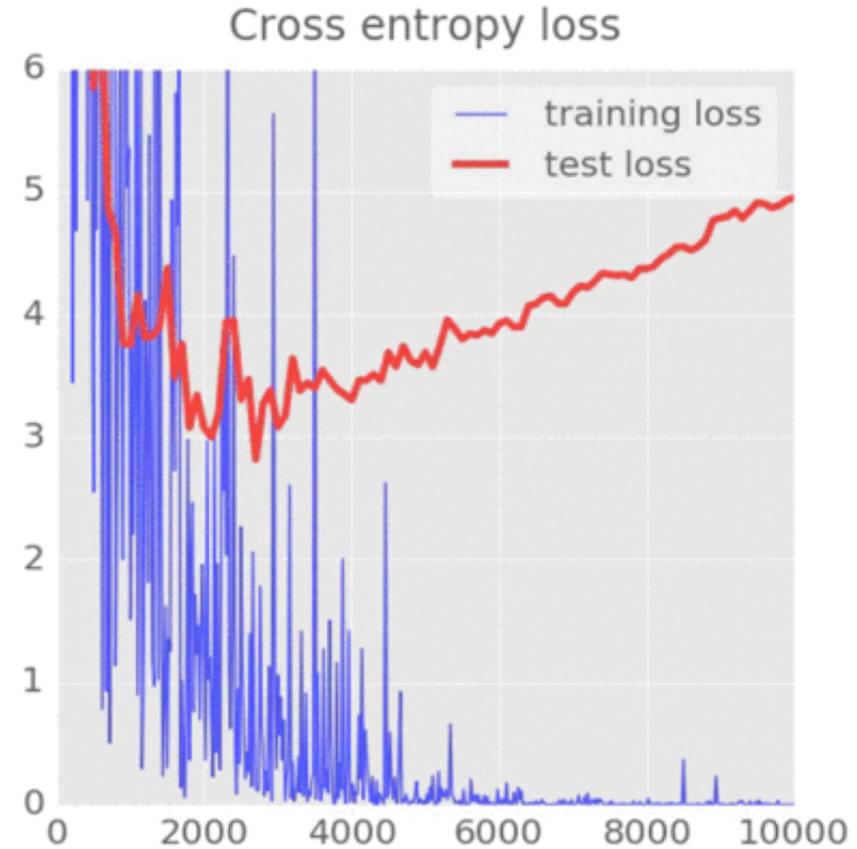
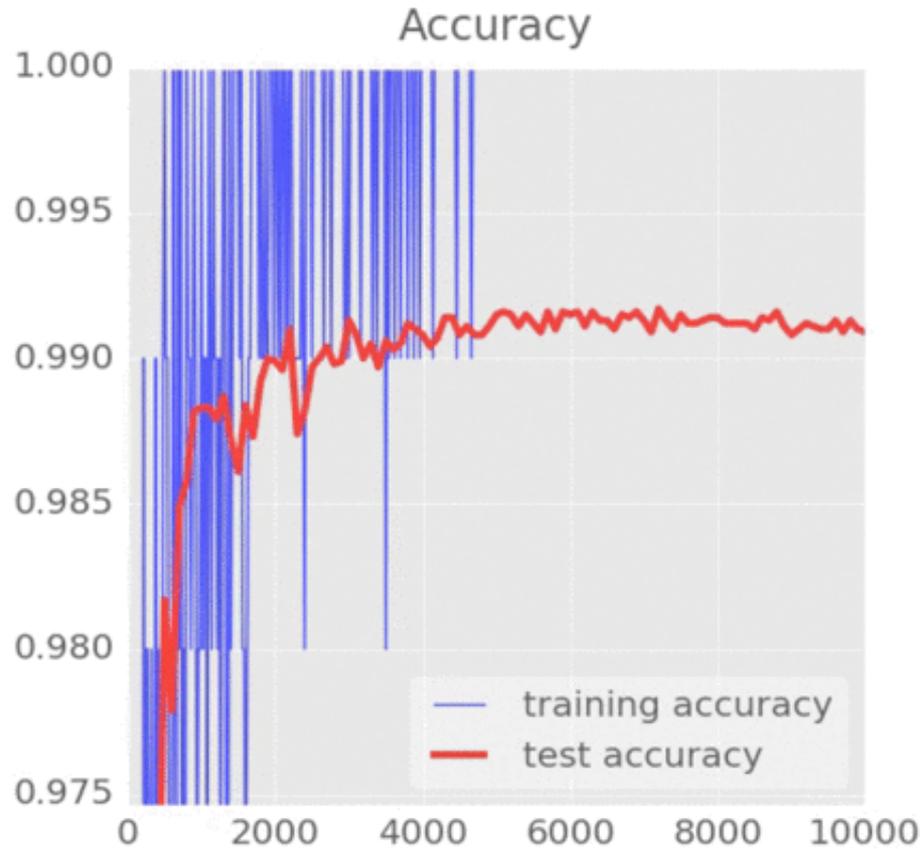
softmax readout layer  
 $W5[200, 10]$

# TensorFlow MNIST Tutorial



99.3!

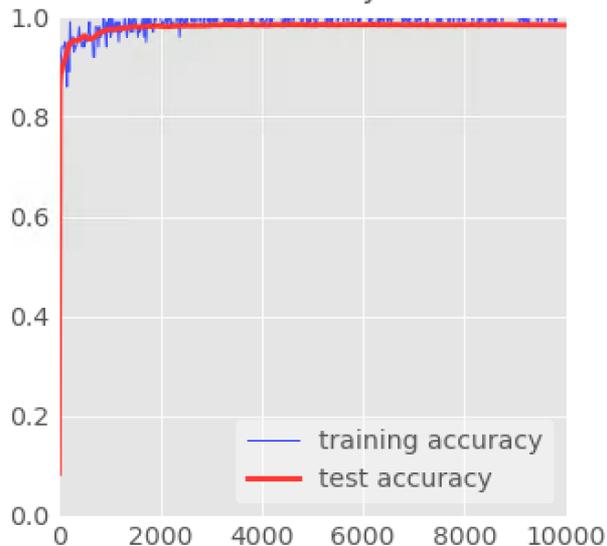
# TensorFlow MNIST Tutorial



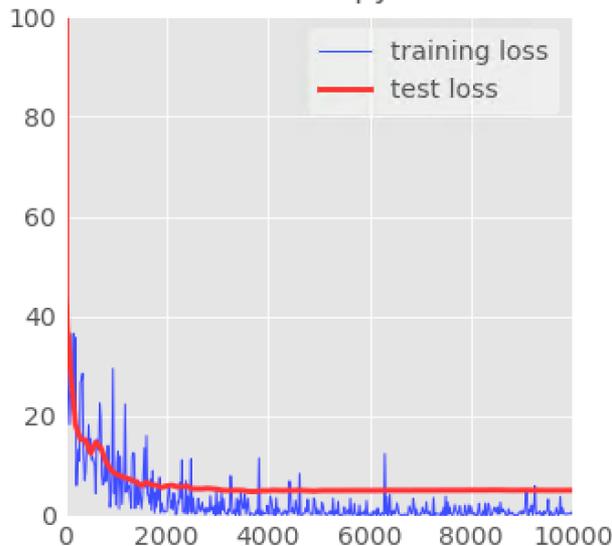
larger convolutional network

# TensorFlow MNIST Tutorial

Accuracy



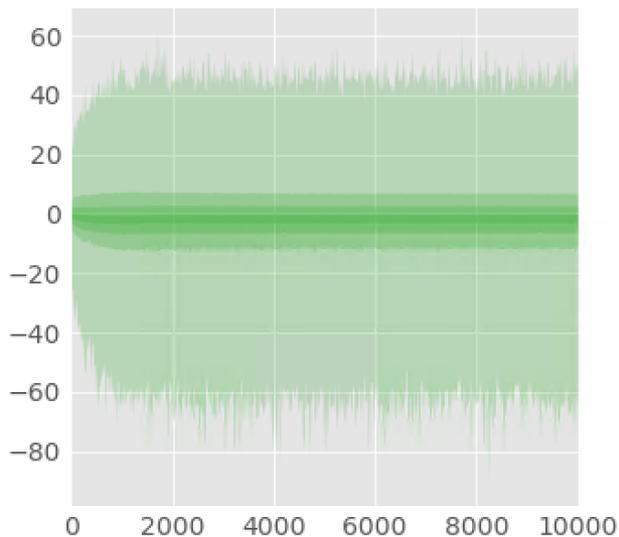
Cross entropy loss



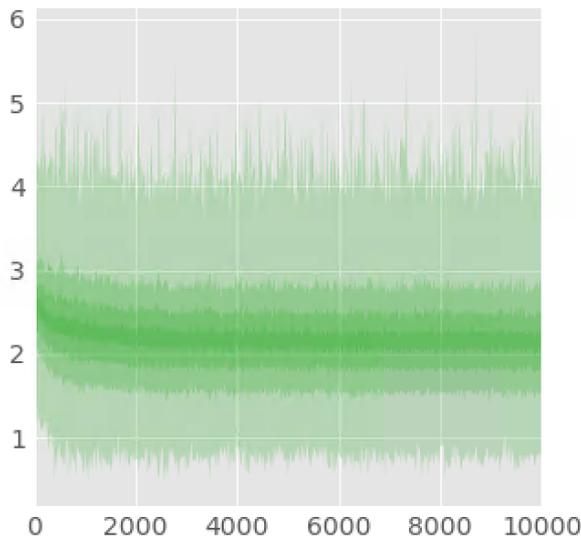
Training digits

6 6 4 3 5 5 0 7 1 0  
 4 5 4 4 1 2 7 4 3 4  
 7 4 2 2 4 0 9 2 3 0  
 9 5 8 0 4 1 4 0 1 0  
 2 7 5 4 5 7 1 7 5 9  
 0 9 1 6 4 7 5 5 1 3  
 7 5 2 9 5 9 2 9 9 4  
 6 1 8 0 0 7 8 0 2 3  
 7 1 7 6 0 2 7 1 1 0  
 5 4 0 6 3 4 4 9 9 3

Logits

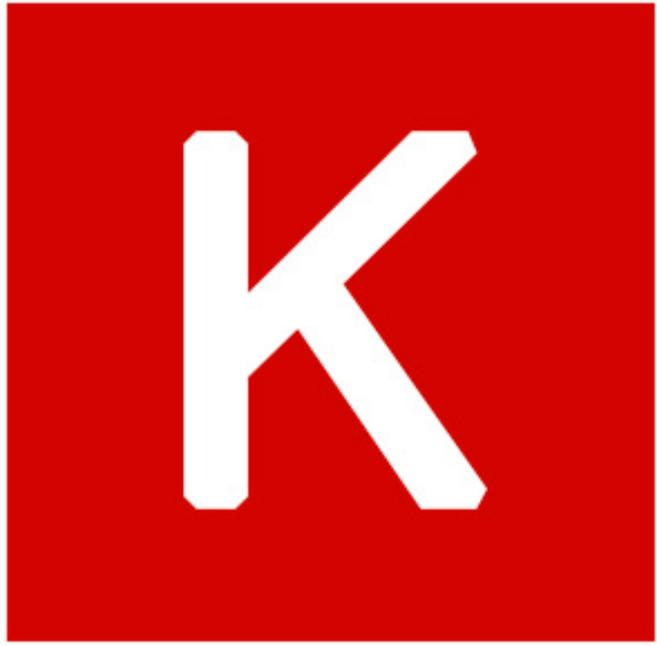


Max activations across batch



Test digits





**K**



**Keras**



# Keras

- Keras is a **high-level neural networks API**
- Written in Python and capable of running on top of either **TensorFlow** 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.

# Keras

**K** Keras Documentation

Search docs

Home

- Keras: Deep Learning library for Theano and TensorFlow
- You have just found Keras.
- Guiding principles
- Getting started: 30 seconds to Keras
- Installation
- Switching from TensorFlow to Theano
- Support
- Why this name, Keras?

Getting started

- Guide to the Sequential model
- Guide to the Functional API
- FAQ

Models

- About Keras models
- Sequential
- Model (functional API)

Layers

- About Keras layers

GitHub

Next »

Docs » Home

[Edit on GitHub](#)

## Keras: Deep Learning library for Theano and TensorFlow

### You have just found Keras.

Keras is a high-level neural networks API, written in Python and capable of running on top of either [TensorFlow](#) 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 2.7-3.5**.

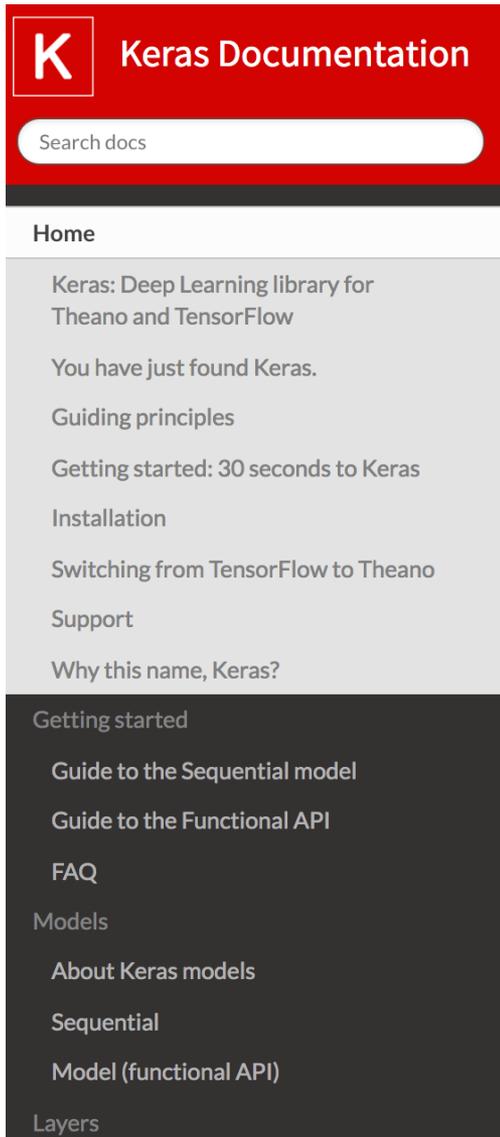
### Guiding principles

- **User friendliness.** Keras is an API designed for human beings, not machines. It puts user experience front and center. Keras follows best practices for reducing cognitive load: it offers consistent & simple APIs, it minimizes the number of user actions required for common use cases, and it provides clear and actionable feedback upon user error.
- **Modularity.** A model is understood as a sequence or a graph of standalone, fully-configurable modules that can

<http://keras.io/>

# Deep Learning with Keras

# Keras Installation



The image shows a vertical navigation menu for the Keras Documentation website. At the top is a red header with a white 'K' logo and the text 'Keras Documentation'. Below this is a search bar labeled 'Search docs'. The menu is divided into two sections: a light gray section for 'Home' and a dark gray section for 'Getting started'. The 'Home' section includes links for 'Keras: Deep Learning library for Theano and TensorFlow', 'You have just found Keras.', 'Guiding principles', 'Getting started: 30 seconds to Keras', 'Installation', 'Switching from TensorFlow to Theano', 'Support', and 'Why this name, Keras?'. The 'Getting started' section includes links for 'Guide to the Sequential model', 'Guide to the Functional API', 'FAQ', 'Models', 'About Keras models', 'Sequential', 'Model (functional API)', and 'Layers'.

## Installation

Keras uses the following dependencies:

- numpy, scipy
- yaml
- HDF5 and h5py (optional, required if you use model saving/loading functions)
- Optional but recommended if you use CNNs: cuDNN.

*When using the TensorFlow backend:*

- TensorFlow
  - See installation instructions.

*When using the Theano backend:*

- Theano
  - See installation instructions.

To install Keras, `cd` to the Keras folder and run the install command:

```
sudo python setup.py install
```

You can also install Keras from PyPI:

```
sudo pip install keras
```

```
conda info --envs
```

```
conda --version
```

```
python --version
```

```
conda list
```

```
conda create -n tensorflow python=3.5
```

```
source activate tensorflow
```

```
activate tensorflow
```

```
pip install Theano
```

```
conda install pydot
```

```
sudo pip install keras
```

```
pip install keras
```

```
pip install tensorflow
```

```
pip install ipython[all]
```

# Gensim

# pip install -U gensim

```
bash-3.2$ pip install -U gensim
Collecting gensim
  Downloading gensim-2.0.0-cp36-cp36m-macosx_10_6_intel.macosx_10_9_intel.macosx_10_9_x86_64.macosx_10_10_intel.macosx_10_10_x86_64.whl (5.6MB)
    100% |#####| 5.6MB 126kB/s
Requirement already up-to-date: six>=1.5.0 in ./anaconda/lib/python3.6/site-packages (from gensim)
Collecting scipy>=0.7.0 (from gensim)
  Downloading scipy-0.19.0-cp36-cp36m-macosx_10_6_intel.macosx_10_9_intel.macosx_10_9_x86_64.macosx_10_10_intel.macosx_10_10_x86_64.whl (16.2MB)
    100% |#####| 16.2MB 43kB/s
Collecting smart-open>=1.2.1 (from gensim)
  Downloading smart_open-1.5.2.tar.gz
Collecting numpy>=1.3 (from gensim)
  Downloading numpy-1.12.1-cp36-cp36m-macosx_10_6_intel.macosx_10_9_intel.macosx_10_9_x86_64.macosx_10_10_intel.macosx_10_10_x86_64.whl (4.4MB)
    100% |#####| 4.4MB 148kB/s
Collecting boto>=2.32 (from smart-open>=1.2.1->gensim)
  Downloading boto-2.46.1-py2.py3-none-any.whl (1.4MB)
    100% |#####| 1.4MB 372kB/s
Requirement already up-to-date: bz2file in ./anaconda/lib/python3.6/site-packages (from smart-open>=1.2.1->gensim)
Collecting requests (from smart-open>=1.2.1->gensim)
  Downloading requests-2.13.0-py2.py3-none-any.whl (584kB)
    100% |#####| 593kB 632kB/s
Building wheels for collected packages: smart-open
Running setup.py bdist_wheel for smart-open ... done
Stored in directory: /Users/imyday/Library/Caches/pip/wheels/02/44/43/68e963ce2b45baefa913a4e558bcd787403458afddffcf45ca
Successfully built smart-open
Installing collected packages: numpy, scipy, boto, requests, smart-open, gensim
Found existing installation: numpy 1.11.3
  Uninstalling numpy-1.11.3:
    Successfully uninstalled numpy-1.11.3
Found existing installation: scipy 0.18.1
  Uninstalling scipy-0.18.1:
    Successfully uninstalled scipy-0.18.1
Found existing installation: boto 2.45.0
  DEPRECATION: Uninstalling a distutils installed project (boto) has been deprecated and will be removed in a future version. This is due to the fact that uninstalling a distutils project will only partially uninstall the project.
  Uninstalling boto-2.45.0:
    Successfully uninstalled boto-2.45.0
Found existing installation: requests 2.12.4
  Uninstalling requests-2.12.4:
    Successfully uninstalled requests-2.12.4
Found existing installation: smart-open 1.4.0
  Uninstalling smart-open-1.4.0:
    Successfully uninstalled smart-open-1.4.0
Found existing installation: gensim 1.0.1
  Uninstalling gensim-1.0.1:
    Successfully uninstalled gensim-1.0.1
Successfully installed boto-2.46.1 gensim-2.0.0 numpy-1.12.1 requests-2.13.0 scipy-0.19.0 smart-open-1.5.2
bash-3.2$
```

# Keras

# sudo pip install keras

```
bash-3.2$ sudo pip install keras
Password:
The directory '/Users/imyday/Library/Caches/pip/http' or its parent directory is not owned by the current user and the cache has been disabled. Please check the permissions and owner of that directory. If executing pip with sudo, you may want sudo's -H flag.
The directory '/Users/imyday/Library/Caches/pip' or its parent directory is not owned by the current user and caching wheels has been disabled. check the permissions and owner of that directory. If executing pip with sudo, you may want sudo's -H flag.
Collecting keras
  Downloading Keras-2.0.3.tar.gz (196kB)
    100% |#####| 204kB 365kB/s
Collecting theano (from keras)
  Downloading Theano-0.9.0.tar.gz (3.1MB)
    100% |#####| 3.1MB 148kB/s
Requirement already satisfied: pyyaml in ./anaconda/lib/python3.6/site-packages (from keras)
Requirement already satisfied: six in ./anaconda/lib/python3.6/site-packages (from keras)
Requirement already satisfied: numpy>=1.9.1 in ./anaconda/lib/python3.6/site-packages (from theano->keras)
Requirement already satisfied: scipy>=0.14 in ./anaconda/lib/python3.6/site-packages (from theano->keras)
Installing collected packages: theano, keras
  Running setup.py install for theano ... done
  Running setup.py install for keras ... done
Successfully installed keras-2.0.3 theano-0.9.0
bash-3.2$
```

# TensorFlow



# Keras Github



Features Business Explore Pricing

This repository

Search

Sign in or Sign up

fchollet / keras

Watch

1,018

Star

14,893

Fork

5,180

Code

Issues 2,486

Pull requests 27

Projects 1

Wiki

Pulse

Graphs

Deep Learning library for Python. Convnets, recurrent neural networks, and more. Runs on TensorFlow or Theano.

<http://keras.io/>

deep-learning

tensorflow

theano

neural-networks

machine-learning

data-science

3,503 commits

4 branches

28 releases

424 contributors

Branch: master

New pull request

Find file

Clone or download

 <b>phiple</b> committed with <b>fchollet</b> Added logsumexp to backend. (#6346)	Latest commit 7d52af6 a day ago
 <a href="#">docker</a>	Update docker files to TensorFlow 1, Theano 0.9 (#6116) 20 days ago
 <a href="#">docs</a>	fix stateful RNNs FAQ link (#6336) 3 days ago
 <a href="#">examples</a>	Spelling errors (#6232) 11 days ago
 <a href="#">keras</a>	Added logsumexp to backend. (#6346) a day ago
 <a href="#">tests</a>	Added logsumexp to backend. (#6346) a day ago
 <a href="#">.gitignore</a>	Fix FAQ question a month ago
 <a href="#">.travis.yml</a>	Update Travis config 9 days ago
 <a href="#">CONTRIBUTING.md</a>	Mention requests for contribution in CONTRIBUTING.md a month ago

<https://github.com/fchollet/keras>

# Keras Examples



Features Business Explore Pricing

This repository

Search

Sign in or Sign up

fchollet / keras

Watch 1,018

Star 14,893

Fork 5,181

Code

Issues 2,486

Pull requests 27

Projects 1

Wiki

Pulse

Graphs

Branch: master

keras / examples /

Create new file

Find file

History



Mohanson committed with fchollet Spelling errors (#6232)

Latest commit 5bd3976 11 days ago

..

<a href="#">README.md</a>	Adding mnist_acgan.py example link in README (#4876)	4 months ago
<a href="#">addition_rnn.py</a>	Spelling errors (#6232)	11 days ago
<a href="#">antirectifier.py</a>	Style fix for examples. (#5980)	28 days ago
<a href="#">babi_memnn.py</a>	Style fixes in example scripts	a month ago
<a href="#">babi_rnn.py</a>	Style fixes in example scripts	a month ago
<a href="#">cifar10_cnn.py</a>	fix rmsprop learning rate for convergence (#6182)	17 days ago
<a href="#">conv_filter_visualization.py</a>	Finish updating examples.	a month ago
<a href="#">conv_lstm.py</a>	Update a number of example scripts.	2 months ago
<a href="#">deep_dream.py</a>	Finish updating examples.	a month ago
<a href="#">image_ocr.py</a>	Fixed URL for wordlist.tgz in image_ocr.py (#6136)	20 days ago
<a href="#">imdb_bidirectional_lstm.py</a>	Finish updating examples.	a month ago
<a href="#">imdb_cnn.py</a>	Finish updating examples.	a month ago
<a href="#">imdb_cnn_lstm.py</a>	Style fix for examples. (#5980)	28 days ago

# Keras MNIST CNN

localhost:8888/notebooks/Documents/SCDBA/DL/Keras\_mnist\_cnn.ipynb

Jupyter Keras\_mnist\_cnn Last Checkpoint: an hour ago (autosaved)

Python 3 Logout

File Edit View Insert Cell Kernel Widgets Help

Python 3

Code CellToolbar

```
from __future__ import print_function
import keras
from keras.datasets import mnist
from keras.models import Sequential
from keras.layers import Dense, Dropout, Flatten
from keras.layers import Conv2D, MaxPooling2D
from keras import backend as K

batch_size = 128
num_classes = 10
epochs = 12

# input image dimensions
img_rows, img_cols = 28, 28

# the data, shuffled and split between train and test sets
(x_train, y_train), (x_test, y_test) = mnist.load_data()

if K.image_data_format() == 'channels_first':
    x_train = x_train.reshape(x_train.shape[0], 1, img_rows, img_cols)
    x_test = x_test.reshape(x_test.shape[0], 1, img_rows, img_cols)
    input_shape = (1, img_rows, img_cols)
else:
    x_train = x_train.reshape(x_train.shape[0], img_rows, img_cols, 1)
    x_test = x_test.reshape(x_test.shape[0], img_rows, img_cols, 1)
    input_shape = (img_rows, img_cols, 1)

x_train = x_train.astype('float32')
x_test = x_test.astype('float32')
x_train /= 255
x_test /= 255
print('x_train shape:', x_train.shape)
print(x_train.shape[0], 'train samples')
print(x_test.shape[0], 'test samples')
```

# Keras MNIST CNN

localhost:8888/notebooks/Documents/SCDBA/DL/Keras\_mnist\_cnn.ipynb

Jupyter Keras\_mnist\_cnn Last Checkpoint: an hour ago (autosaved)

Logout

File Edit View Insert Cell Kernel Widgets Help

Python 3

```
# convert class vectors to binary class matrices
y_train = keras.utils.to_categorical(y_train, num_classes)
y_test = keras.utils.to_categorical(y_test, num_classes)

model = Sequential()
model.add(Conv2D(32, kernel_size=(3, 3),
                 activation='relu',
                 input_shape=input_shape))
model.add(Conv2D(64, (3, 3), activation='relu'))
model.add(MaxPooling2D(pool_size=(2, 2)))
model.add(Dropout(0.25))
model.add(Flatten())
model.add(Dense(128, activation='relu'))
model.add(Dropout(0.5))
model.add(Dense(num_classes, activation='softmax'))

model.compile(loss=keras.losses.categorical_crossentropy,
              optimizer=keras.optimizers.Adadelta(),
              metrics=['accuracy'])

model.fit(x_train, y_train,
          batch_size=batch_size,
          epochs=epochs,
          verbose=1,
          validation_data=(x_test, y_test))
score = model.evaluate(x_test, y_test, verbose=0)
print('Test loss:', score[0])
print('Test accuracy:', score[1])
```

Using TensorFlow backend.

```
Downloading data from https://s3.amazonaws.com/img-datasets/mnist.npz
x_train shape: (60000, 28, 28, 1)
60000 train samples
10000 test samples
Train on 60000 samples, validate on 10000 samples
Epoch 1/12
```

# Keras MNIST CNN

```
localhost:8888/notebooks/Documents/SCDBA/DL/Keras_mnist_cnn.ipynb

jupyter Keras_mnist_cnn Last Checkpoint: an hour ago (autosaved) Python 3

File Edit View Insert Cell Kernel Widgets Help

Using TensorFlow backend.

Downloading data from https://s3.amazonaws.com/img-datasets/mnist.npz
x_train shape: (60000, 28, 28, 1)
60000 train samples
10000 test samples
Train on 60000 samples, validate on 10000 samples
Epoch 1/12
60000/60000 [=====] - 200s - loss: 0.3155 - acc: 0.9028 - val_loss: 0.0756 - val_acc: 0.9761
Epoch 2/12
60000/60000 [=====] - 209s - loss: 0.1106 - acc: 0.9681 - val_loss: 0.0523 - val_acc: 0.9837
Epoch 3/12
60000/60000 [=====] - 220s - loss: 0.0834 - acc: 0.9749 - val_loss: 0.0416 - val_acc: 0.9852
Epoch 4/12
60000/60000 [=====] - 224s - loss: 0.0700 - acc: 0.9795 - val_loss: 0.0392 - val_acc: 0.9879
Epoch 5/12
60000/60000 [=====] - 229s - loss: 0.0614 - acc: 0.9818 - val_loss: 0.0358 - val_acc: 0.9871
Epoch 6/12
60000/60000 [=====] - 227s - loss: 0.0558 - acc: 0.9828 - val_loss: 0.0345 - val_acc: 0.9880
Epoch 7/12
60000/60000 [=====] - 217s - loss: 0.0498 - acc: 0.9850 - val_loss: 0.0337 - val_acc: 0.9883
Epoch 8/12
60000/60000 [=====] - 217s - loss: 0.0473 - acc: 0.9865 - val_loss: 0.0294 - val_acc: 0.9899
Epoch 9/12
60000/60000 [=====] - 217s - loss: 0.0439 - acc: 0.9872 - val_loss: 0.0316 - val_acc: 0.9889
Epoch 10/12
60000/60000 [=====] - 217s - loss: 0.0415 - acc: 0.9871 - val_loss: 0.0319 - val_acc: 0.9897
Epoch 11/12
60000/60000 [=====] - 217s - loss: 0.0380 - acc: 0.9889 - val_loss: 0.0275 - val_acc: 0.9904
Epoch 12/12
60000/60000 [=====] - 215s - loss: 0.0376 - acc: 0.9889 - val_loss: 0.0285 - val_acc: 0.9905
Test loss: 0.0285460013417
Test accuracy: 0.9905
```

# Keras MINST CNN

```
from __future__ import print_function
import keras
from keras.datasets import mnist
from keras.models import Sequential
from keras.layers import Dense, Dropout, Flatten
from keras.layers import Conv2D, MaxPooling2D
from keras import backend as K

batch_size = 128
num_classes = 10
epochs = 12

# input image dimensions
img_rows, img_cols = 28, 28

# the data, shuffled and split between train and test sets
(x_train, y_train), (x_test, y_test) = mnist.load_data()

if K.image_data_format() == 'channels_first':
    x_train = x_train.reshape(x_train.shape[0], 1, img_rows, img_cols)
    x_test = x_test.reshape(x_test.shape[0], 1, img_rows, img_cols)
    input_shape = (1, img_rows, img_cols)
else:
    x_train = x_train.reshape(x_train.shape[0], img_rows, img_cols, 1)
    x_test = x_test.reshape(x_test.shape[0], img_rows, img_cols, 1)
    input_shape = (img_rows, img_cols, 1)

x_train = x_train.astype('float32')
x_test = x_test.astype('float32')
x_train /= 255
x_test /= 255
print('x_train shape:', x_train.shape)
print(x_train.shape[0], 'train samples')
print(x_test.shape[0], 'test samples')

# convert class vectors to binary class matrices
y_train = keras.utils.to_categorical(y_train, num_classes)
y_test = keras.utils.to_categorical(y_test, num_classes)

model = Sequential()
model.add(Conv2D(32, kernel_size=(3, 3),
                 activation='relu',
                 input_shape=input_shape))
model.add(Conv2D(64, (3, 3), activation='relu'))
model.add(MaxPooling2D(pool_size=(2, 2)))
model.add(Dropout(0.25))
model.add(Flatten())
model.add(Dense(128, activation='relu'))
model.add(Dropout(0.5))
model.add(Dense(num_classes, activation='softmax'))

model.compile(loss=keras.losses.categorical_crossentropy,
              optimizer=keras.optimizers.Adadelta(),
              metrics=['accuracy'])

model.fit(x_train, y_train,
          batch_size=batch_size,
          epochs=epochs,
          verbose=1,
          validation_data=(x_test, y_test))
score = model.evaluate(x_test, y_test, verbose=0)
print('Test loss:', score[0])
print('Test accuracy:', score[1])
```

# Keras MNIST CNN

```
from __future__ import print_function
import keras
from keras.datasets import mnist
from keras.models import Sequential
from keras.layers import Dense, Dropout, Flatten
from keras.layers import Conv2D, MaxPooling2D
from keras import backend as K
```

# Keras MNIST CNN

```
batch_size = 128
num_classes = 10
epochs = 12

# input image dimensions
img_rows, img_cols = 28, 28

# the data, shuffled and split between train and test sets
(x_train, y_train), (x_test, y_test) = mnist.load_data()

if K.image_data_format() == 'channels_first':
    x_train = x_train.reshape(x_train.shape[0], 1, img_rows, img_cols)
    x_test = x_test.reshape(x_test.shape[0], 1, img_rows, img_cols)
    input_shape = (1, img_rows, img_cols)
else:
    x_train = x_train.reshape(x_train.shape[0], img_rows, img_cols, 1)
    x_test = x_test.reshape(x_test.shape[0], img_rows, img_cols, 1)
    input_shape = (img_rows, img_cols, 1)
```

# Keras MNIST CNN

```
x_train = x_train.astype('float32')
x_test = x_test.astype('float32')
x_train /= 255
x_test /= 255
print('x_train shape:', x_train.shape)
print(x_train.shape[0], 'train samples')
print(x_test.shape[0], 'test samples')

# convert class vectors to binary class matrices
y_train = keras.utils.to_categorical(y_train, num_classes)
y_test = keras.utils.to_categorical(y_test, num_classes)
```

# Keras MNIST CNN

```
model = Sequential()
model.add(Conv2D(32, kernel_size=(3, 3),
                activation='relu',
                input_shape=input_shape))
model.add(Conv2D(64, (3, 3), activation='relu'))
model.add(MaxPooling2D(pool_size=(2, 2)))
model.add(Dropout(0.25))
model.add(Flatten())
model.add(Dense(128, activation='relu'))
model.add(Dropout(0.5))
model.add(Dense(num_classes, activation='softmax'))

model.compile(loss=keras.losses.categorical_crossentropy,
              optimizer=keras.optimizers.Adadelta(),
              metrics=['accuracy'])
```

# Keras MNIST CNN

```
model.fit(x_train, y_train,  
          batch_size=batch_size,  
          epochs=epochs,  
          verbose=1,  
          validation_data=(x_test, y_test))  
score = model.evaluate(x_test, y_test, verbose=0)  
print('Test loss:', score[0])  
print('Test accuracy:', score[1])
```

# Keras MNIST CNN

python mnist\_cnn.py

Using TensorFlow backend.

Downloading data from <https://s3.amazonaws.com/img-datasets/mnist.npz>

x\_train shape: (60000, 28, 28, 1)

60000 train samples

10000 test samples

Train on 60000 samples, validate on 10000 samples

Epoch 1/12

60000/60000 [=====] - 108s - loss: 0.3510 - acc: 0.8921 - val\_loss: 0.0880 - val\_acc: 0.9738

Epoch 2/12

60000/60000 [=====] - 106s - loss: 0.1200 - acc: 0.9649 - val\_loss: 0.0567 - val\_acc: 0.9820

Epoch 3/12

60000/60000 [=====] - 104s - loss: 0.0889 - acc: 0.9735 - val\_loss: 0.0438 - val\_acc: 0.9856

Epoch 4/12

60000/60000 [=====] - 106s - loss: 0.0744 - acc: 0.9783 - val\_loss: 0.0392 - val\_acc: 0.9862

Epoch 5/12

60000/60000 [=====] - 106s - loss: 0.0648 - acc: 0.9807 - val\_loss: 0.0363 - val\_acc: 0.9873

Epoch 6/12

60000/60000 [=====] - 109s - loss: 0.0574 - acc: 0.9840 - val\_loss: 0.0348 - val\_acc: 0.9884

Epoch 7/12

60000/60000 [=====] - 104s - loss: 0.0522 - acc: 0.9842 - val\_loss: 0.0324 - val\_acc: 0.9890

Epoch 8/12

60000/60000 [=====] - 104s - loss: 0.0484 - acc: 0.9856 - val\_loss: 0.0315 - val\_acc: 0.9894

Epoch 9/12

60000/60000 [=====] - 104s - loss: 0.0447 - acc: 0.9870 - val\_loss: 0.0296 - val\_acc: 0.9902

Epoch 10/12

60000/60000 [=====] - 109s - loss: 0.0419 - acc: 0.9877 - val\_loss: 0.0338 - val\_acc: 0.9894

Epoch 11/12

60000/60000 [=====] - 104s - loss: 0.0405 - acc: 0.9879 - val\_loss: 0.0301 - val\_acc: 0.9896

Epoch 12/12

60000/60000 [=====] - 127s - loss: 0.0391 - acc: 0.9883 - val\_loss: 0.0304 - val\_acc: 0.9899

Test loss: 0.030424870987

Test accuracy: 0.9899

# Keras IMDB CNN

localhost:8888/notebooks/Documents/SCDBA/DL/Keras\_imdb\_cnn.ipynb

Jupyter Keras\_imdb\_cnn Last Checkpoint: 15 minutes ago (unsaved changes)

Logout

File Edit View Insert Cell Kernel Widgets Help

Python 3

Code CellToolbar

```
from __future__ import print_function
from keras.preprocessing import sequence
from keras.models import Sequential
from keras.layers import Dense, Dropout, Activation
from keras.layers import Embedding
from keras.layers import Conv1D, GlobalMaxPooling1D
from keras.datasets import imdb

# set parameters:
max_features = 5000
maxlen = 400
batch_size = 32
embedding_dims = 50
filters = 250
kernel_size = 3
hidden_dims = 250
epochs = 2

print('Loading data...')
(x_train, y_train), (x_test, y_test) = imdb.load_data(num_words=max_features)
print(len(x_train), 'train sequences')
print(len(x_test), 'test sequences')

print('Pad sequences (samples x time)')
x_train = sequence.pad_sequences(x_train, maxlen=maxlen)
x_test = sequence.pad_sequences(x_test, maxlen=maxlen)
print('x_train shape:', x_train.shape)
print('x_test shape:', x_test.shape)

print('Build model...')
model = Sequential()

# we start off with an efficient embedding layer which maps
# our vocab indices into embedding_dims dimensions
model.add(Embedding(max_features,
                    embedding_dims,
```

# Keras IMDB CNN

localhost:8888/notebooks/Documents/SCDBA/DL/Keras\_imdb\_cnn.ipynb

Jupyter Keras\_imdb\_cnn Last Checkpoint: 19 minutes ago (autosaved)



File Edit View Insert Cell Kernel Widgets Help

Python 3

Code CellToolbar

```
model.add(Embedding(max_features,
                    embedding_dims,
                    input_length=maxlen))
model.add(Dropout(0.2))

# we add a Convolution1D, which will learn filters
# word group filters of size filter_length:
model.add(Conv1D(filters,
                 kernel_size,
                 padding='valid',
                 activation='relu',
                 strides=1))

# we use max pooling:
model.add(GlobalMaxPooling1D())

# We add a vanilla hidden layer:
model.add(Dense(hidden_dims))
model.add(Dropout(0.2))
model.add(Activation('relu'))

# We project onto a single unit output layer, and squash it with a sigmoid:
model.add(Dense(1))
model.add(Activation('sigmoid'))

model.compile(loss='binary_crossentropy',
              optimizer='adam',
              metrics=['accuracy'])
model.fit(x_train, y_train,
          batch_size=batch_size,
          epochs=epochs,
          validation_data=(x_test, y_test))
```

Using TensorFlow backend.

Loading data...

Downloading data from <https://s3.amazonaws.com/text-datasets/imdb.npz>

25000 train sequences

# Keras IMDB CNN

localhost:8888/notebooks/Documents/SCDBA/DL/Keras\_imdb\_cnn.ipynb

jupyter Keras\_imdb\_cnn Last Checkpoint: 13 minutes ago (autosaved)



File Edit View Insert Cell Kernel Widgets Help

Python 3

Code CellToolbar

```
# we use max pooling:
model.add(GlobalMaxPooling1D())

# We add a vanilla hidden layer:
model.add(Dense(hidden_dims))
model.add(Dropout(0.2))
model.add(Activation('relu'))

# We project onto a single unit output layer, and squash it with a sigmoid:
model.add(Dense(1))
model.add(Activation('sigmoid'))

model.compile(loss='binary_crossentropy',
              optimizer='adam',
              metrics=['accuracy'])
model.fit(x_train, y_train,
        batch_size=batch_size,
        epochs=epochs,
        validation_data=(x_test, y_test))
```

Using TensorFlow backend.

Loading data...

Downloading data from <https://s3.amazonaws.com/text-datasets/imdb.npz>

25000 train sequences

25000 test sequences

Pad sequences (samples x time)

x\_train shape: (25000, 400)

x\_test shape: (25000, 400)

Build model...

Train on 25000 samples, validate on 25000 samples

Epoch 1/2

25000/25000 [=====] - 266s - loss: 0.4110 - acc: 0.8012 - val\_loss: 0.2965 - val\_acc: 0.8739

Epoch 2/2

25000/25000 [=====] - 286s - loss: 0.2429 - acc: 0.9020 - val\_loss: 0.2726 - val\_acc: 0.8862

Out[1]: <keras.callbacks.History at 0x11dc37b00>

# Keras IMDB CNN

```
python imdb_cnn.py
Using TensorFlow backend.
Loading data...
Downloading data from https://s3.amazonaws.com/text-datasets/imdb.npz
25000 train sequences
25000 test sequences
Pad sequences (samples x time)
x_train shape: (25000, 400)
x_test shape: (25000, 400)
Build model...
Train on 25000 samples, validate on 25000 samples
Epoch 1/2
25000/25000 [=====] - 157s - loss: 0.4050 - acc: 0.8065 - val_loss: 0.2924 - val_acc: 0.8750
Epoch 2/2
25000/25000 [=====] - 128s - loss: 0.2433 - acc: 0.9040 - val_loss: 0.2701 - val_acc: 0.8865
Exception ignored in: <bound method BaseSession.__del__ of <tensorflow.python.client.session.Session object at 0x0000019F153C2A20>>
Traceback (most recent call last):
  File "C:\Program Files\Anaconda3\lib\site-packages\tensorflow\python\client\session.py", line 587, in __del__
AttributeError: 'NoneType' object has no attribute 'TF_NewStatus'
```

# Keras IMDB LSTM

```
from __future__ import print_function
from keras.preprocessing import sequence
from keras.models import Sequential
from keras.layers import Dense, Embedding
from keras.layers import LSTM
from keras.datasets import imdb

max_features = 20000
maxlen = 80 # cut texts after this number of words (among top max_features most common words)
batch_size = 32

print('Loading data...')
(x_train, y_train), (x_test, y_test) = imdb.load_data(num_words=max_features)
print(len(x_train), 'train sequences')
print(len(x_test), 'test sequences')

print('Pad sequences (samples x time)')
x_train = sequence.pad_sequences(x_train, maxlen=maxlen)
x_test = sequence.pad_sequences(x_test, maxlen=maxlen)
print('x_train shape:', x_train.shape)
print('x_test shape:', x_test.shape)

print('Build model...')
model = Sequential()
model.add(Embedding(max_features, 128))
model.add(LSTM(128, dropout=0.2, recurrent_dropout=0.2))
model.add(Dense(1, activation='sigmoid'))

# try using different optimizers and different optimizer configs
model.compile(loss='binary_crossentropy',
              optimizer='adam',
              metrics=['accuracy'])

print('Train...')
model.fit(x_train, y_train,
         batch_size=batch_size,
         epochs=15,
         validation_data=(x_test, y_test))
score, acc = model.evaluate(x_test, y_test,
                           batch_size=batch_size)
print('Test score:', score)
print('Test accuracy:', acc)
```

# Keras IMDB LSTM

```
from __future__ import print_function
from keras.preprocessing import sequence
from keras.models import Sequential
from keras.layers import Dense, Embedding
from keras.layers import LSTM
from keras.datasets import imdb
```

# Keras IMDB LSTM

```
max_features = 20000
maxlen = 80 # cut texts after this number of words (among top
max_features most common words)
batch_size = 32

print('Loading data...')
(x_train, y_train), (x_test, y_test) =
imdb.load_data(num_words=max_features)
print(len(x_train), 'train sequences')
print(len(x_test), 'test sequences')

print('Pad sequences (samples x time)')
x_train = sequence.pad_sequences(x_train, maxlen=maxlen)
x_test = sequence.pad_sequences(x_test, maxlen=maxlen)
print('x_train shape:', x_train.shape)
print('x_test shape:', x_test.shape)
```

# Keras IMDB LSTM

```
print('Build model...')
model = Sequential()
model.add(Embedding(max_features, 128))
model.add(LSTM(128, dropout=0.2, recurrent_dropout=0.2))
model.add(Dense(1, activation='sigmoid'))

# try using different optimizers and different optimizer configs
model.compile(loss='binary_crossentropy',
              optimizer='adam',
              metrics=['accuracy'])
```

# Keras IMDB LSTM

```
print('Train...')
model.fit(x_train, y_train,
          batch_size=batch_size,
          epochs=15,
          validation_data=(x_test, y_test))
score, acc = model.evaluate(x_test, y_test,

batch_size=batch_size)
print('Test score:', score)
print('Test accuracy:', acc)
```

# Keras IMDB LSTM

python imdb\_lstm.py  
Using TensorFlow backend.

Loading data...

25000 train sequences

25000 test sequences

Pad sequences (samples x time)

x\_train shape: (25000, 80)

x\_test shape: (25000, 80)

Build model...

Train...

Train on 25000 samples, validate on 25000 samples

Epoch 1/15

25000/25000 [=====] - 111s - loss: 0.4561 - acc: 0.7837 - val\_loss: 0.3892 - val\_acc: 0.8275

Epoch 2/15

25000/25000 [=====] - 112s - loss: 0.2947 - acc: 0.8792 - val\_loss: 0.4266 - val\_acc: 0.8353

Epoch 3/15

25000/25000 [=====] - 111s - loss: 0.2122 - acc: 0.9178 - val\_loss: 0.4133 - val\_acc: 0.8284

Epoch 4/15

25000/25000 [=====] - 112s - loss: 0.1461 - acc: 0.9450 - val\_loss: 0.4670 - val\_acc: 0.8260

Epoch 5/15

25000/25000 [=====] - 113s - loss: 0.1038 - acc: 0.9633 - val\_loss: 0.5580 - val\_acc: 0.8203

Epoch 6/15

25000/25000 [=====] - 113s - loss: 0.0739 - acc: 0.9749 - val\_loss: 0.6738 - val\_acc: 0.8174

Epoch 7/15

25000/25000 [=====] - 113s - loss: 0.0542 - acc: 0.9810 - val\_loss: 0.7463 - val\_acc: 0.8154

Epoch 8/15

25000/25000 [=====] - 113s - loss: 0.0428 - acc: 0.9856 - val\_loss: 0.8131 - val\_acc: 0.8157

Epoch 9/15

25000/25000 [=====] - 115s - loss: 0.0334 - acc: 0.9889 - val\_loss: 0.8566 - val\_acc: 0.8165

Epoch 10/15

25000/25000 [=====] - 114s - loss: 0.0248 - acc: 0.9920 - val\_loss: 0.9186 - val\_acc: 0.8165

Epoch 11/15

25000/25000 [=====] - 116s - loss: 0.0156 - acc: 0.9955 - val\_loss: 0.9016 - val\_acc: 0.8082

Epoch 12/15

25000/25000 [=====] - 117s - loss: 0.0196 - acc: 0.9942 - val\_loss: 0.9720 - val\_acc: 0.8124

Epoch 13/15

25000/25000 [=====] - 120s - loss: 0.0152 - acc: 0.9957 - val\_loss: 1.0064 - val\_acc: 0.8148

Epoch 14/15

25000/25000 [=====] - 121s - loss: 0.0128 - acc: 0.9961 - val\_loss: 1.1103 - val\_acc: 0.8121

Epoch 15/15

25000/25000 [=====] - 114s - loss: 0.0110 - acc: 0.9970 - val\_loss: 1.0173 - val\_acc: 0.8132

25000/25000 [=====] - 23s

Test score: 1.01734088922

Test accuracy: 0.8132

# Keras IMDB FastText

```
python imdb_fasttext.py
Using TensorFlow backend.
Loading data...
25000 train sequences
25000 test sequences
Average train sequence length: 238
Average test sequence length: 230
Pad sequences (samples x time)
x_train shape: (25000, 400)
x_test shape: (25000, 400)
Build model...
Train on 25000 samples, validate on 25000 samples
Epoch 1/5
25000/25000 [=====] - 14s - loss: 0.6102 - acc: 0.7397 - val_loss: 0.5034 - val_acc: 0.8105
Epoch 2/5
25000/25000 [=====] - 14s - loss: 0.4019 - acc: 0.8656 - val_loss: 0.3697 - val_acc: 0.8654
Epoch 3/5
25000/25000 [=====] - 14s - loss: 0.3025 - acc: 0.8959 - val_loss: 0.3199 - val_acc: 0.8791
Epoch 4/5
25000/25000 [=====] - 14s - loss: 0.2521 - acc: 0.9113 - val_loss: 0.2971 - val_acc: 0.8848
Epoch 5/5
25000/25000 [=====] - 14s - loss: 0.2181 - acc: 0.9249 - val_loss: 0.2899 - val_acc: 0.8855
Exception ignored in: <bound method BaseSession.__del__ of <tensorflow.python.client.session.Session object at
0x000001E3257DB438>>
Traceback (most recent call last):
  File "C:\Program Files\Anaconda3\lib\site-packages\tensorflow\python\client\session.py", line 587, in __del__
AttributeError: 'NoneType' object has no attribute 'TF_NewStatus'
```

# Keras IMDB CNN LSTM

```
python imdb_cnn_lstm_2.py
Using TensorFlow backend.
Loading data...
25000 train sequences
25000 test sequences
Pad sequences (samples x time)
x_train shape: (25000, 100)
x_test shape: (25000, 100)
Build model...
Train...
Train on 25000 samples, validate on 25000 samples
Epoch 1/2
25000/25000 [=====] - 64s - loss: 0.3824 - acc: 0.8238 - val_loss: 0.3591 - val_acc: 0.8467
Epoch 2/2
25000/25000 [=====] - 63s - loss: 0.1953 - acc: 0.9261 - val_loss: 0.3827 - val_acc: 0.8488
24990/25000 [=====>.] - ETA: 0s
Test score: 0.382728585386
Test accuracy: 0.848799994493
```

# imdb\_lstm\_2.py

```
from __future__ import print_function

from keras.preprocessing import sequence
from keras.models import Sequential
from keras.layers import Dense, Embedding
from keras.layers import LSTM
from keras.datasets import imdb

py_filename = 'imdb_lstm_2.py'
max_features = 20000
maxlen = 80 # cut texts after this number of words (among top max_features
most common words)
batch_size = 32
epochs = 20 #60

#%matplotlib inline
import matplotlib
import matplotlib.pyplot as plt
import numpy as np

import codecs
import datetime
import timeit
timer_start = timeit.default_timer()
#timer_end = timeit.default_timer()
#print('timer_end - timer_start', timer_end - timer_start)
```

# imdb\_lstm\_2.py

```
def getDateTimenow():
    strnow = datetime.datetime.now().strftime("%Y%m%d_%H%M%S")
    return strnow

def read_file_utf8(filename):
    with codecs.open(filename, 'r', encoding='utf-8') as f:
        text = f.read()
    return text

def write_file_utf8(filename, text):
    with codecs.open(filename, 'w', encoding='utf-8') as f:
        f.write(text)
        f.close()

def log_file_utf8(filename, text):
    with codecs.open(filename, 'a', encoding='utf-8') as f:
        #append file
        f.write(text + '\n')
        f.close()

log_file_utf8("logfile.txt", '***** ' + py_filename + ' *****')
log_file_utf8("logfile.txt", '***** Start DateTime: ' + getDateTimenow())

print('Start: ', datetime.datetime.now().strftime("%Y%m%d_%H%M%S"))
```

# imdb\_lstm\_2.py

```
print('Loading data...')
(x_train, y_train), (x_test, y_test) = imdb.load_data(num_words=max_features)
print(len(x_train), 'train sequences')
print(len(x_test), 'test sequences')

print('Pad sequences (samples x time)')
x_train = sequence.pad_sequences(x_train, maxlen=maxlen)
x_test = sequence.pad_sequences(x_test, maxlen=maxlen)
print('x_train shape:', x_train.shape)
print('x_test shape:', x_test.shape)

print('Build model...')
model = Sequential()
model.add(Embedding(max_features, 128))
model.add(LSTM(128, dropout=0.2, recurrent_dropout=0.2))
model.add(Dense(1, activation='sigmoid'))

# try using different optimizers and different optimizer configs
model.compile(loss='binary_crossentropy',
              optimizer='adam',
              metrics=['accuracy'])
```

# imdb\_lstm\_2.py

```
print('Train...')
print('model.fit: ', datetime.datetime.now().strftime("%Y%m%d_%H%M%S"))
history = model.fit(x_train, y_train,
                    batch_size = batch_size,
                    epochs = epochs,
                    validation_data = (x_test, y_test))

score, acc = model.evaluate(x_test, y_test,
                            batch_size=batch_size)

print('Test score:', score)
print('Test accuracy:', acc)
```

# imdb\_lstm\_2.py

```
timer_end = timeit.default_timer()
print('Timer: ', str(round(timer_end - timer_start, 2)), 's')
print('DateTime: ', datetime.datetime.now().strftime("%Y%m%d_%H%M%S"))
log_file_utf8("logfile.txt", 'Timer: ' + str(round(timer_end - timer_start, 2))
+ ' s')
log_file_utf8("logfile.txt", '***** End Datetime: ' +
datetime.datetime.now().strftime("%Y%m%d_%H%M%S"))

# summarize history for accuracy
#http://machinelearningmastery.com/display-deep-learning-model-training-history-in-keras/
print('history.history.keys():', history.history.keys())
print('history.history:', history.history)
log_file_utf8("logfile.txt", 'history.history:' + str(history.history))
```

# imdb\_lstm\_2.py

## # Deep Learning Training Visualization

```
plt.figure(figsize=(10, 8)) # make separate figure
ax1 = plt.subplot(2, 1, 1)
plt.plot(history.history['acc'])
plt.plot(history.history['val_acc'])
plt.title('model accuracy')
plt.ylabel('accuracy')
ax1.xaxis.set_major_locator(plt.NullLocator())
#plt.xlabel('epoch')
plt.legend(['train acc', 'test val_acc'], loc='upper left')
#plt.show()
ax2 = plt.subplot(2, 1, 2)
plt.plot(history.history['loss'])
plt.plot(history.history['val_loss'])
plt.title('model loss')
plt.ylabel('loss')
plt.xlabel('epoch')
plt.legend(['train loss', 'test val_loss'], loc='upper left')
plt.savefig("training_accuracy_loss_" + py_filename + "_" + str(epochs) +
".png", dpi= 300)
```

# imdb\_lstm\_2.py

```
#Log File for Deep Learning Summary Analysis
log_file_utf8("logfile.txt", 'DL_Summary:\t' + py_filename +
    '\tePOCHS\t' + str(epochs) +
    '\tscore\t' + str(score) +
    '\taccuracy\t' + str(acc) +
    '\tTimer\t' + str(round(timer_end - timer_start, 2)) +
    '\thistory\t' + str(history.history))
plt.show()
```

# python filename.py

```
python imdb_fasttext_2.py
```

```
python imdb_cnn_2.py
```

```
python imdb_lstm_2.py
```

```
python imdb_cnn_lstm_2.py
```

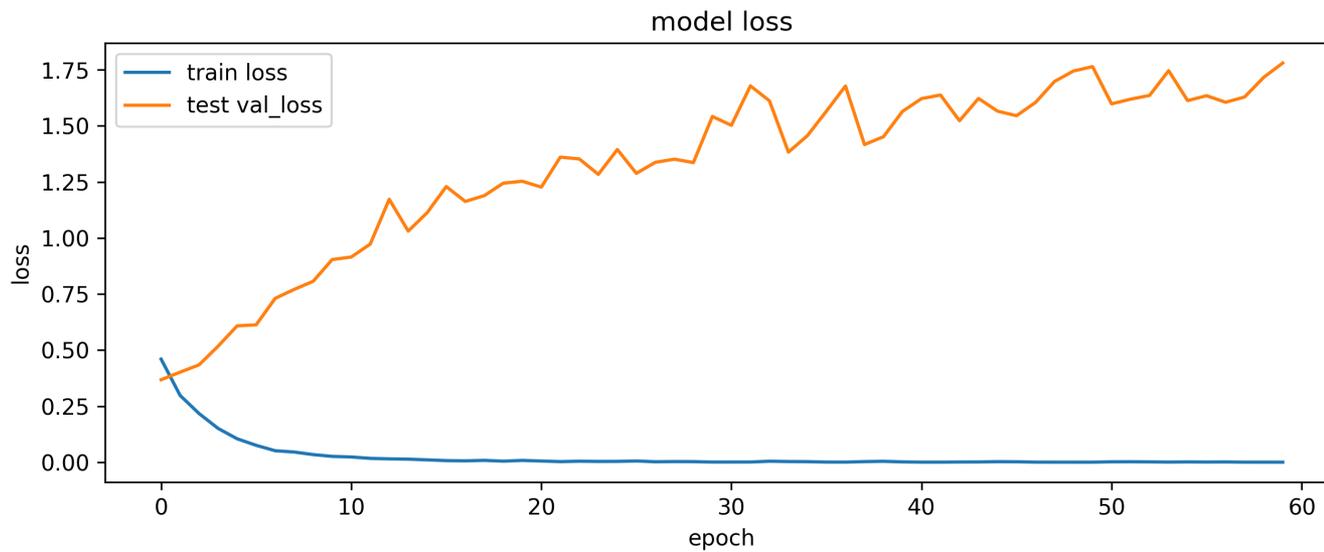
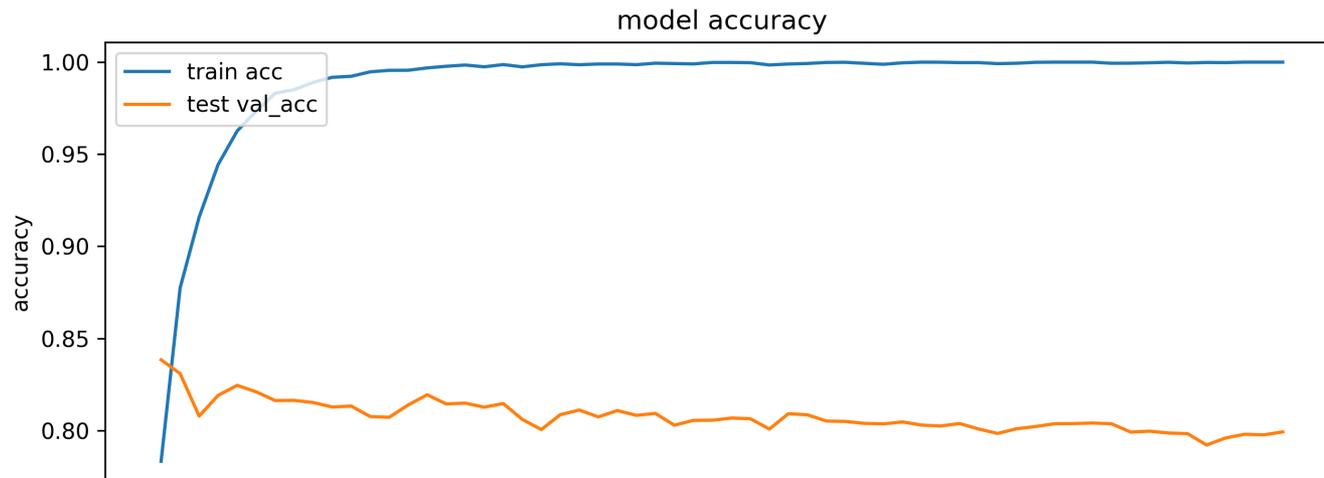
```
python imdb_bidirectional_lstm_2.py
```

# Deep Learning Summary

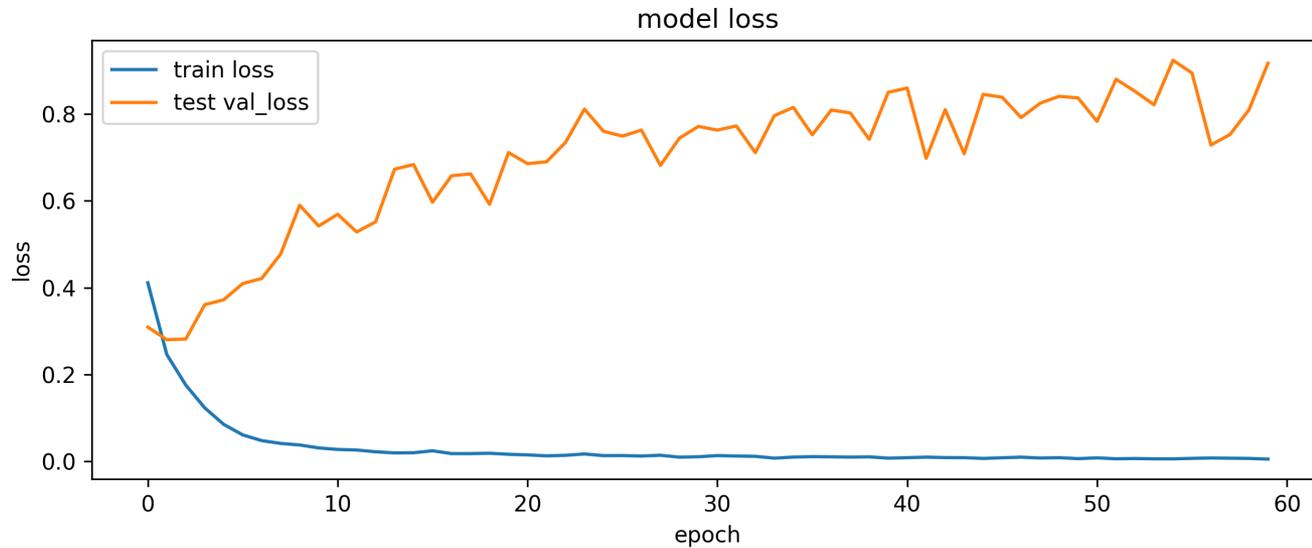
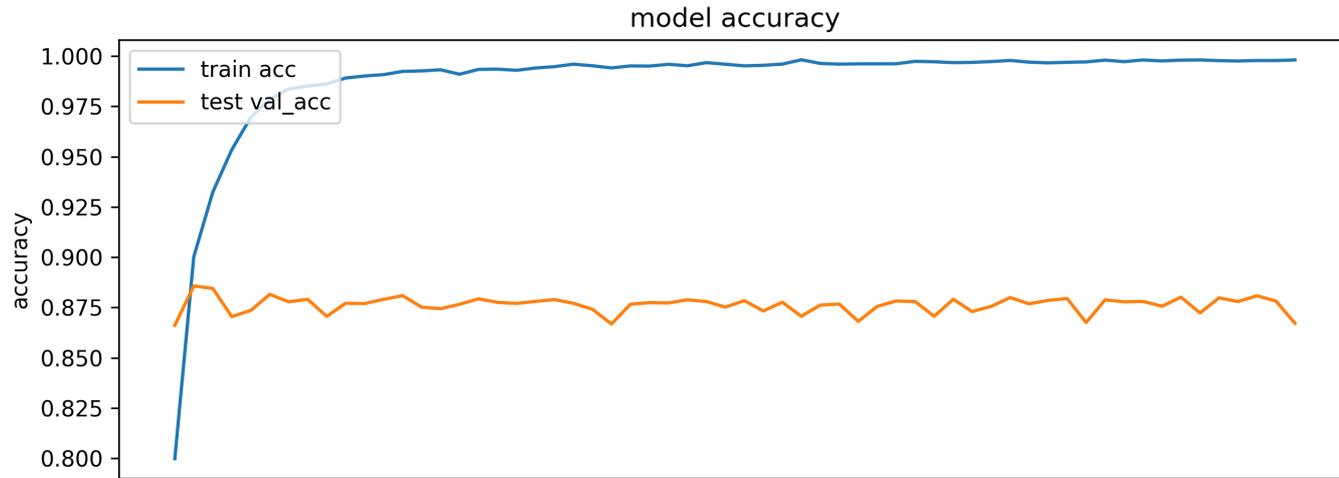
```
#Log File for Deep Learning Summary Analysis
log_file_utf8("logfile.txt", 'DL_Summary:\t' + py_filename +
'\tepochs\t' + str(epochs) +
'\tscore\t' + str(score) +
'\taccuracy\t' + str(acc) +
'\tTimer\t' + str(round(timer_end - timer_start, 2)) +
'\thistory\t' + str(history.history))
```

Model	epochs	Score	Accuracy	Timer (s)
imdb_lstm_2.py	30	0.6440	0.8540	<b>682.57</b>
imdb_cnn_2.py	30	0.7186	<b>0.8775</b>	4320.38
imdb_lstm_2.py	30	1.5716	0.8052	3958.93
imdb_cnn_lstm_2.py	30	1.3105	0.8240	2471.65
imdb_bidirectional_lstm_2.py	30	1.4083	0.8255	4344.36
imdb_fasttext_2.py	30	<b>0.6439</b>	0.8540	1117.78
imdb_fasttext_2.py	60	1.2335	0.8407	1297.02
imdb_cnn_2.py	60	0.9170	0.8672	8507.48
imdb_lstm_2.py	60	1.7803	0.7992	8039.67
imdb_cnn_lstm_2.py	60	1.4623	0.8137	4912.25
imdb_bidirectional_lstm_2.py	60	1.8975	0.8138	8589.17

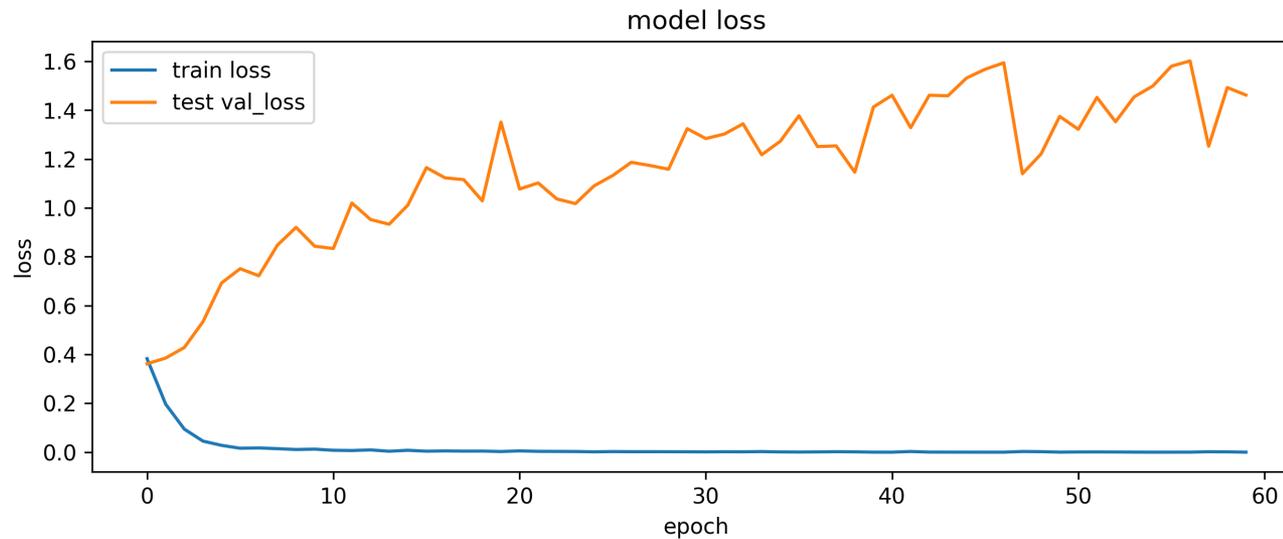
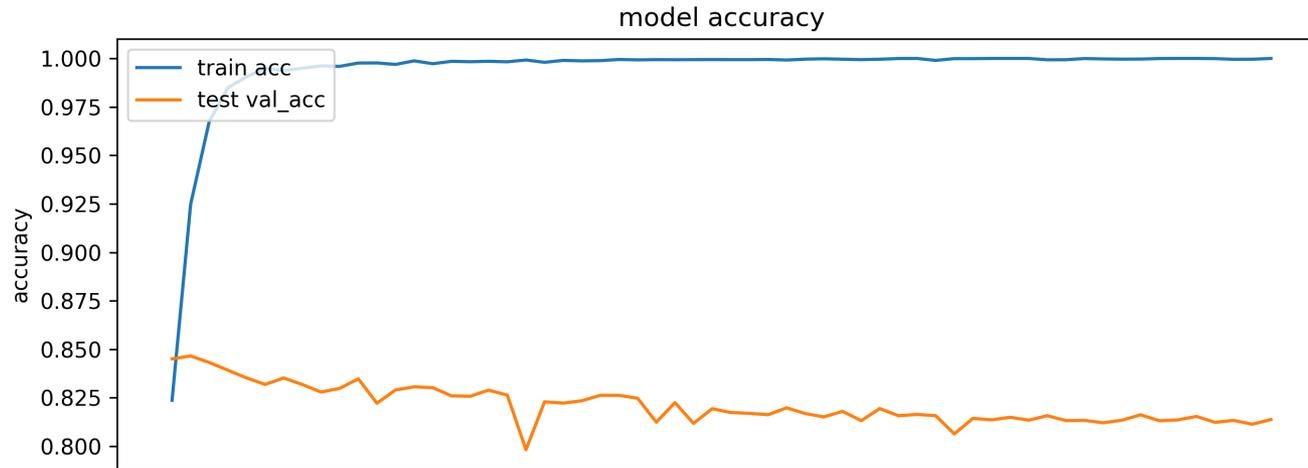
# imdb\_lstm\_2.py



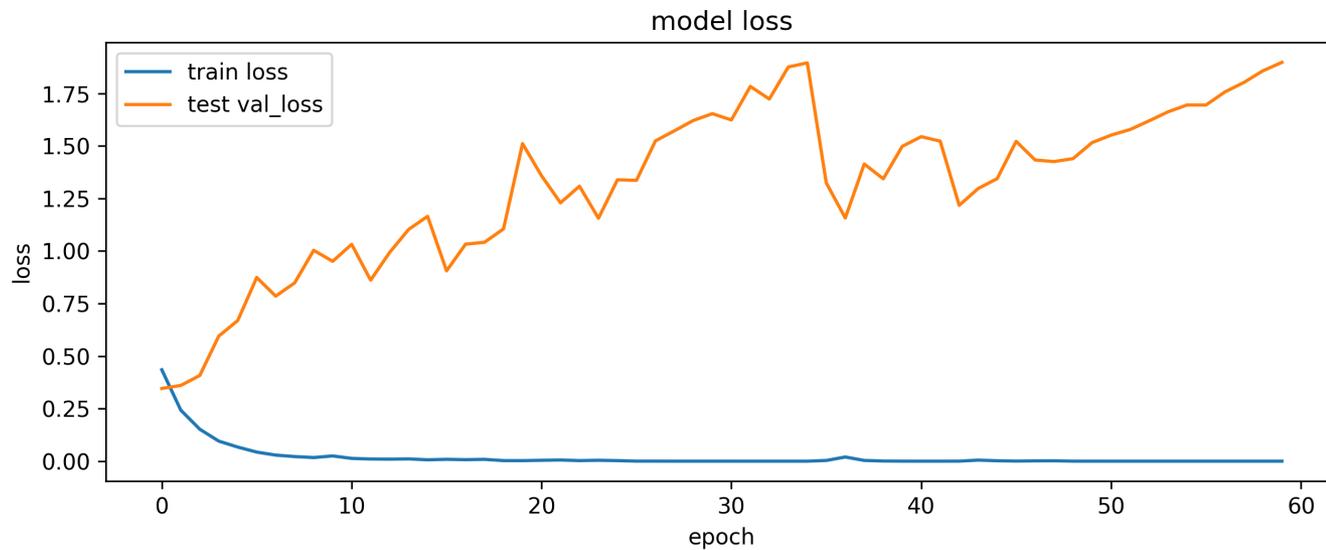
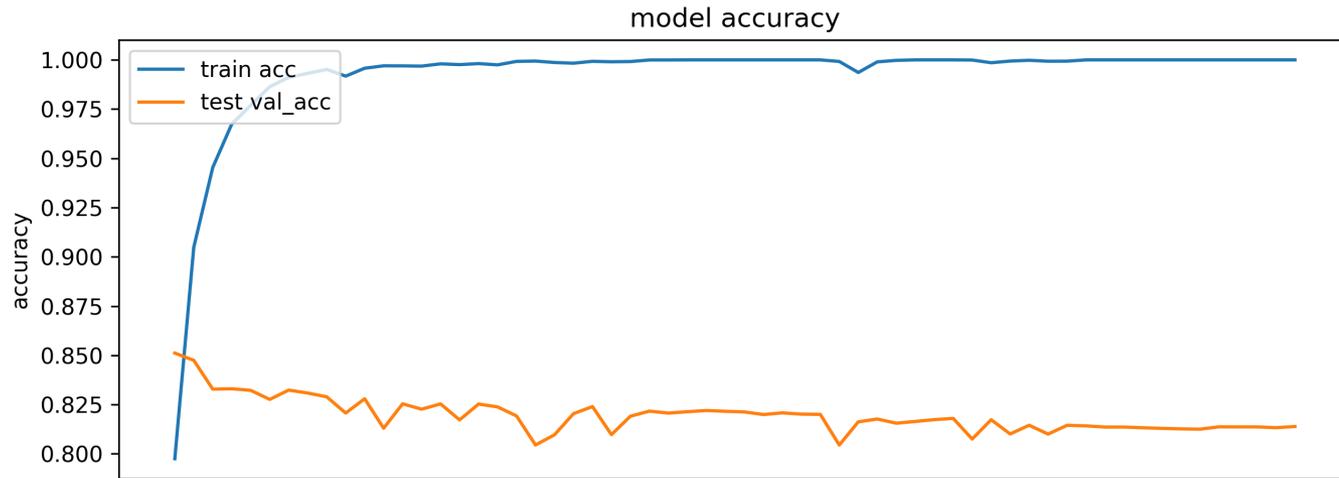
# imdb\_cnn\_2.py



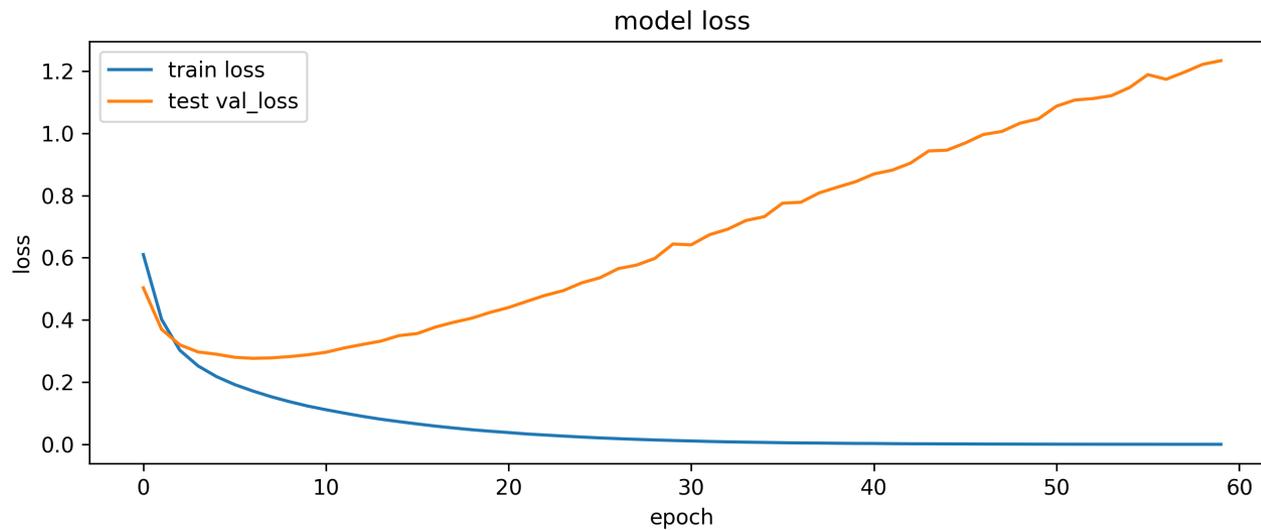
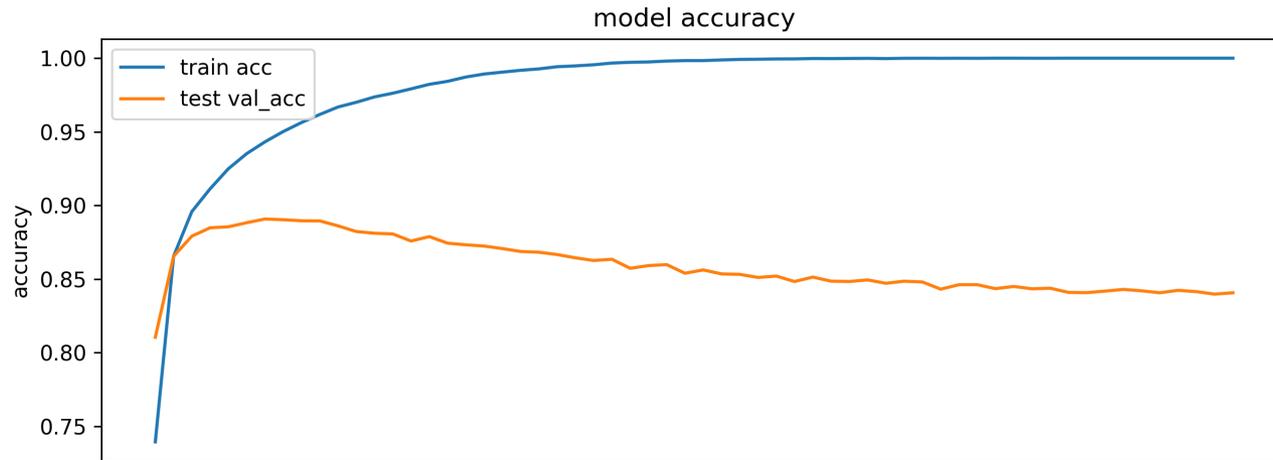
# imdb\_cnn\_lstm\_2.py



# imdb\_bidirectional\_lstm\_2.py



# imdb\_fasttext\_2.py



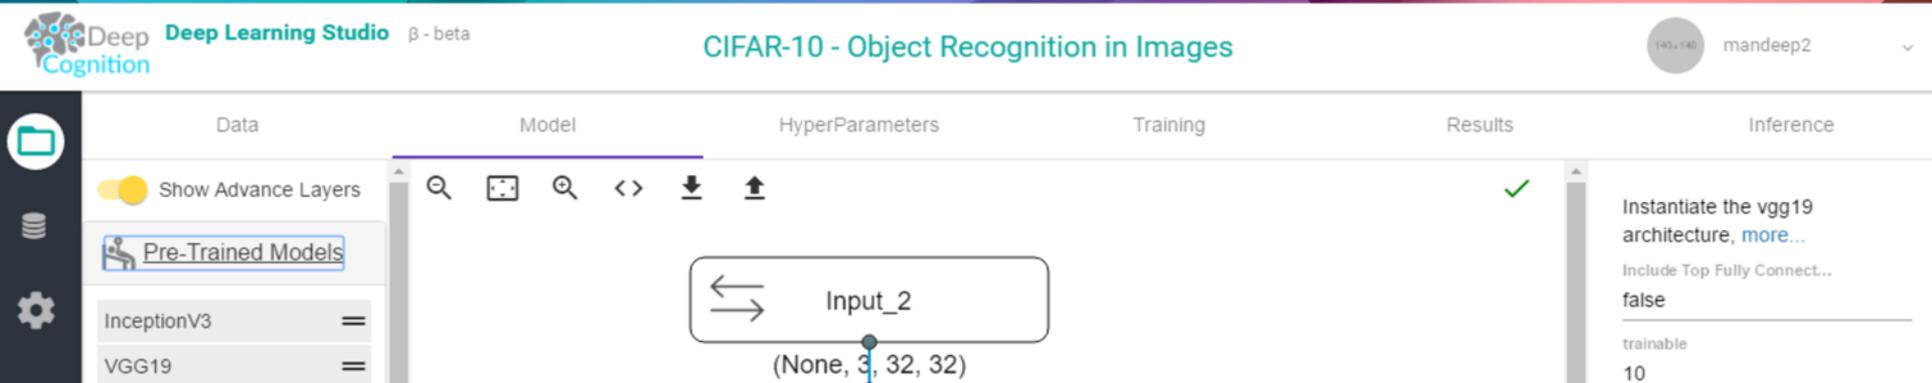
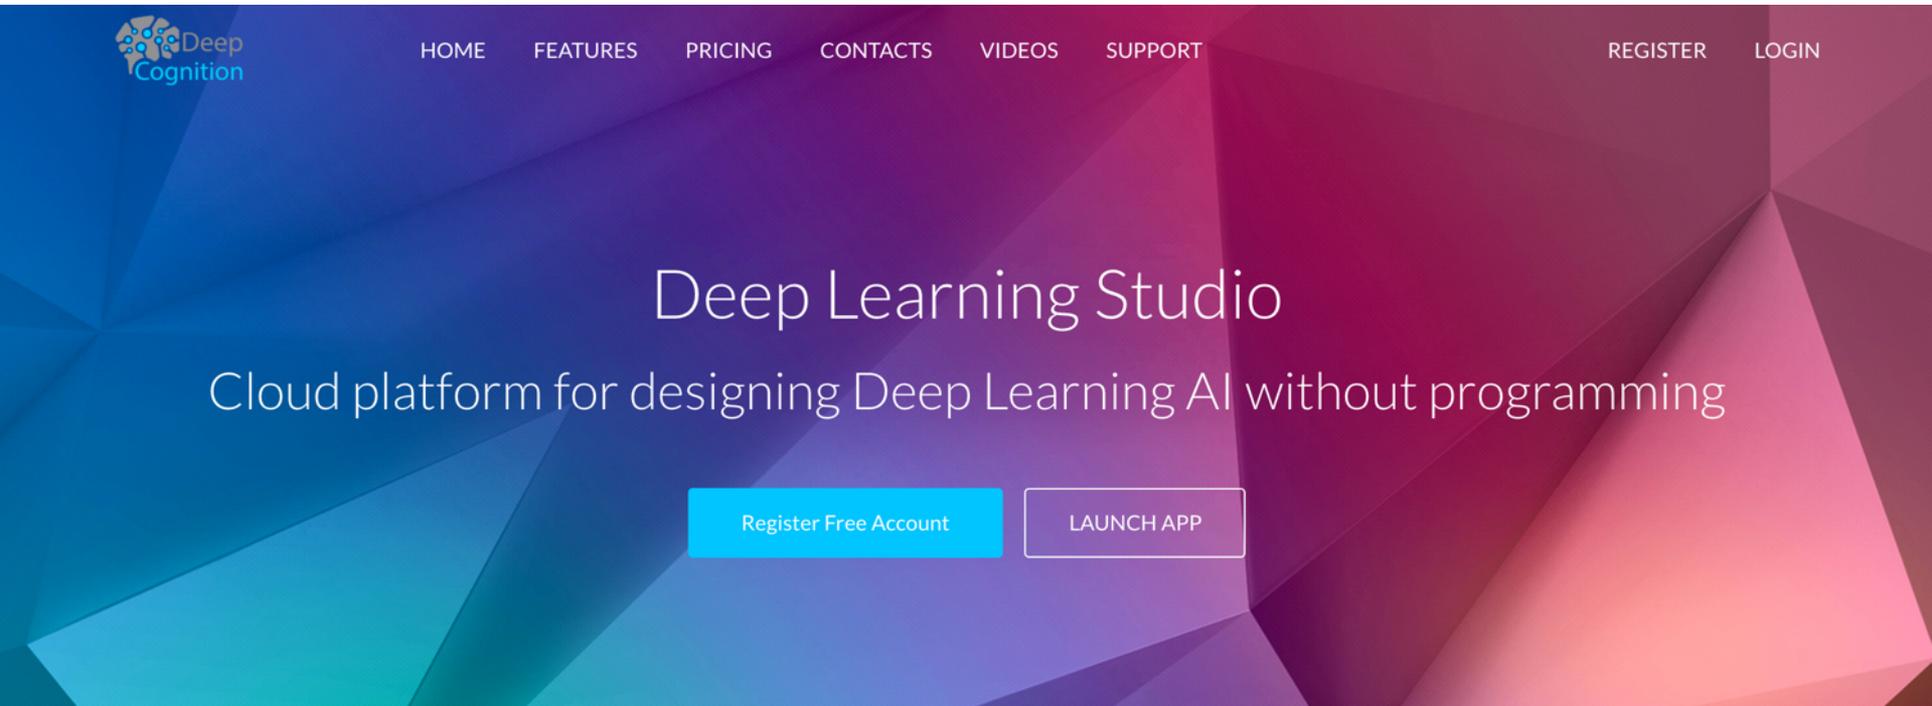
# Deep Learning with CPU vs. GPU

## Timings:

Hardware	Backend	Time / Epoch
CPU	TF	3 hrs
Titan X (maxwell)	TF	4 min
Titan X (maxwell)	TH	7 min

# Deep Learning Studio

Cloud platform for designing Deep Learning AI without programming



# Deep Learning Studio

Cloud platform for designing Deep Learning AI without programming

The screenshot displays the Deep Learning Studio interface for designing a neural network. The main workspace shows a flowchart of the model architecture:

- Input\_2**: (None, 3, 32, 32)
- VGG19\_1**: (None, 512, 1, 1)
- Flatten\_2**: (None, 512)
- Dense\_5**: (None, 100)
- Dropout\_1**: (None, 100)
- Dense\_3**: (None, 10)
- Output\_2**: (None, 10)

The left sidebar contains a list of pre-trained models: InceptionV3, VGG19, VGG16, and ResNet50. Below this are categories for Special Functions, Convolutional Layers, Core Layers, Pooling Layers, Recurrent Layers, Advanced Activations Layers, Convolutional Recurrent Layers, and Noise Layers.

The right-hand panel shows configuration options for the VGG19 architecture, including a checkbox for "Show Advance Layers" (checked), a search icon, and a "Show Advance Options" toggle (unchecked). The configuration details include:

- Instantiate the vgg19 architecture, [more...](#)
- Include Top Fully Connect... **false**
- trainable **10**



# Deep Learning Studio

Cloud platform for designing Deep Learning AI without programming

eta

## MNIST Handwritten Digits Classifier

Model      HyperParameters      Training      Results

Dataset Source: Testing ▾

Training Run: Run0 ▾    Start Inference    or    Download Trained Model

Digit Label	Image	predictions
• 9		• 9
• 1		• 1
• 1		• 1
• 5		• 3
• 0		• 0
• 5		• 5
• 1		• 1
• 2		• 6
• 2		• 2
• 3		• 3

« Previous    1    2    3    4    5    ...    351    Next »

Download Results

# References

- Sunila Gollapudi (2016), Practical Machine Learning, Packt Publishing
- Sebastian Raschka (2015), Python Machine Learning, Packt Publishing
- Martin Gorner (2017), TensorFlow and Deep Learning without a PhD, Part 1 (Google Cloud Next '17), <https://www.youtube.com/watch?v=u4aIGiomYP4>
- Martin Gorner (2017), TensorFlow and Deep Learning without a PhD, Part 2 (Google Cloud Next '17), <https://www.youtube.com/watch?v=fTUwdXUffI8>
- Martin Gorner (2017), TensorFlow and Deep Learning without a PhD, <https://goo.gl/pHeXe7>, <https://codelabs.developers.google.com/codelabs/cloud-tensorflow-mnist>
- Deep Learning Basics: Neural Networks Demystified, <https://www.youtube.com/playlist?list=PLiaHhY2iBX9hdHaRr6b7XevZtgZRa1PoU>
- Deep Learning SIMPLIFIED, <https://www.youtube.com/playlist?list=PLjJh1vISEYgvGod9wWiydumYl8hOXixNu>
- TensorFlow: <https://www.tensorflow.org/>
- Theano: <http://deeplearning.net/software/theano/>
- Keras: <http://keras.io/>
- Deep Learning Studio: Cloud platform for designing Deep Learning AI without programming, <http://deepcognition.ai/>
- <http://p.migdal.pl/2017/04/30/teaching-deep-learning.html>
- <https://github.com/leriomaggio/deep-learning-keras-tensorflow>