Lecture slides for this course
have been prepared by Dr. Le Minh Huy,
EEE, Phenikaa University

Deep Learning

Chapter 1 Introduction

Dr. Van-Toi NGUYEN

EEE, Phenikaa University

Chapter 1: Course info & programming review

1. Course introduction and grades
2. History of Deep Learning
3. Deep learning applications
4. Materials

1. Course introduction and grades

Course introduction

Học phần “Học sâu” cung cấp những kiến thức cốt lõi của công nghệ học sâu, bao gồm: mô hình mạng nơ ron truyền thẳng; các kỹ thuật tổng quát hoá và tối ưu hóa các mô hình; mở rộng mô hình để làm việc với dữ liệu lớn; mạng CNN, RNN. Khóa học này trang bị các kỹ năng liên quan đến việc thiết kế, xây dựng và lập trình mô hình học sâu. Học phần cũng trang bị kiến thức cần thiết để học viên có thể sử dụng thư viện học sâu như Tensorflow để xây dựng một số ứng dụng cơ bản của công nghệ học sâu.

1. Course introduction and grades

Goals

- Tổng hợp lại kiến thức cơ bản về học sâu.
- Vận hành được các mô hình học sâu và ứng dựng trong một số bài toán thực tế.

Outcome requirements

- Phân tích được các kiến thức về mô hình mạng nơ ron truyền thẳng, các kỹ thuật tổng quát hoá và tối ưu hóa mô hình, mô hình học sâu CNN, RNN.
- Thiết kế được các mô hình học sâu CNN, RNN trong các bài toán thực tế với dữ liệu ảnh và dữ liệu chuỗi.

1. Course introduction and grades

Book

- Ian, Goodfellow; Yoshua, Bengio; Aaron, Courville (2016), Deep Learning, The MIT Press.

References

- Raschka, Sebastian (2019), Python Machine Learning :, Packt,, 9781789955750.
- Cs231n Stanford University
- Deep learning – deeplearning.ai

Grades

- Attendant + Homework: 10%
- Midterm project: 20%
- Final project: 70%

1. Course introduction and grades

Code on: Python, Google Colab, Tensorflow 2.0, Sklearn

MIT - Massachusetts Institute of Technology

Stanford University

Harvard University

1. Course introduction and grades

AI's Heroes

Yan Lecun

Yoshua Bengio

Geoffrey Hinton

Andrew Ng.

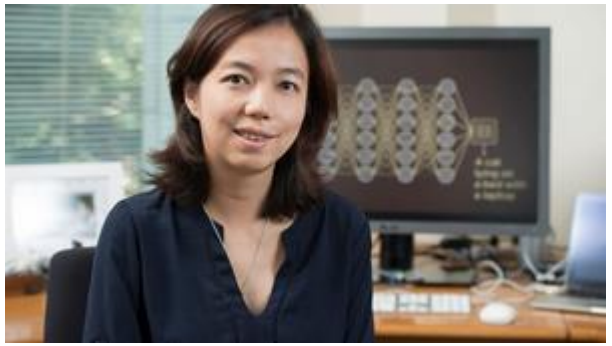

Fei Fei Li

1. Course introduction and grades

Chapter 1: Course Infor & Programming review - week 1

1. Course introduction and grades
2. History of Deep learning
3. Deep learning applications

Chapter 2: Building Neural Network from Scratch – week 2-7

1. Shallow neural network
2. Deep neural network
3. Building neural network: step-by-step (modulation)
4. Regularization
5. Dropout
6. Batch Normalization
7. Optimizers
8. Hyper-parameters
9. Practice

Midterm

Chapter 3: Convolutional Neural Network - week 8-10

1. Convolutional operator
2. History of CNN
3. Deep Convolutional Models
4. Layers in CNN
5. Applications of CNN
6. Practice

Midterm summary

Chapter 4: TensorFlow Library - week 11-13

1. Introduction to TensorFlow
2. Building a deep neural network with TensorFlow
3. Applications
4. Practice

Chapter 5: Recurrent Neural Network - week 14-15

1. Unfolding Computational Graphs
2. Building a Recurrent Neural Networks
3. Long Short-Term Memory
4. Vision with Language Processing
5. Application of RNN
6. Practice

1. Course introduction and grades

45 hours at Classes:

Theory + Coding practice

90 hours shelf-study at home:

Theory + Coding practice

2. History of Deep Learning

- Deep learning is a Subset of Machine Learning in which Artificial Neural Network adapt and learn from vast amounts of data

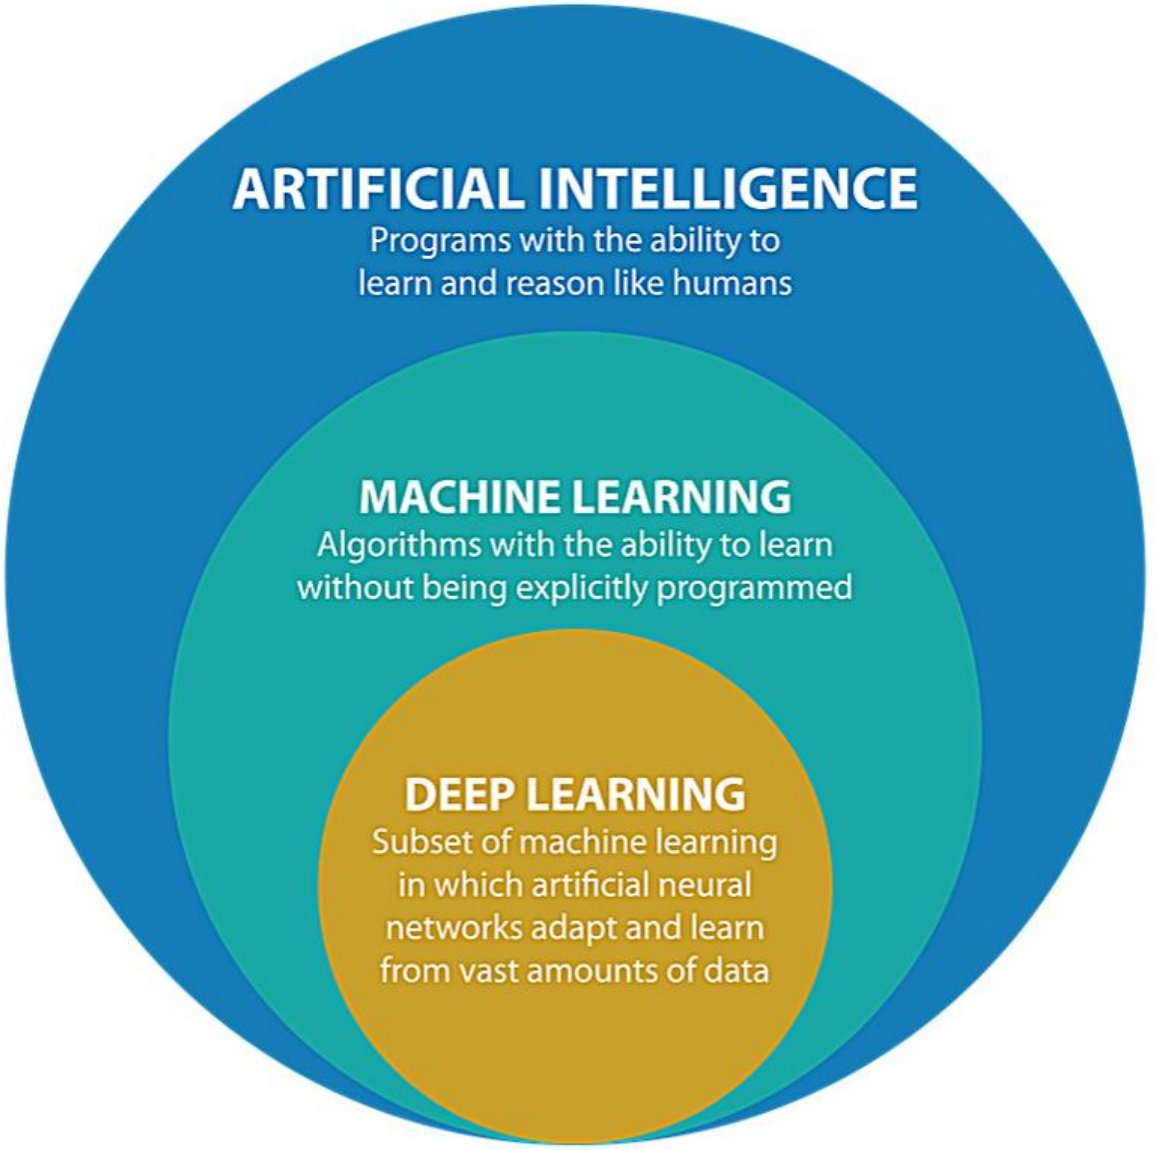

2. History of Deep Learning

- Deep learning is a Subset of Machine Learning in which Artificial Neural Network adapt and learn from vast amounts of data.

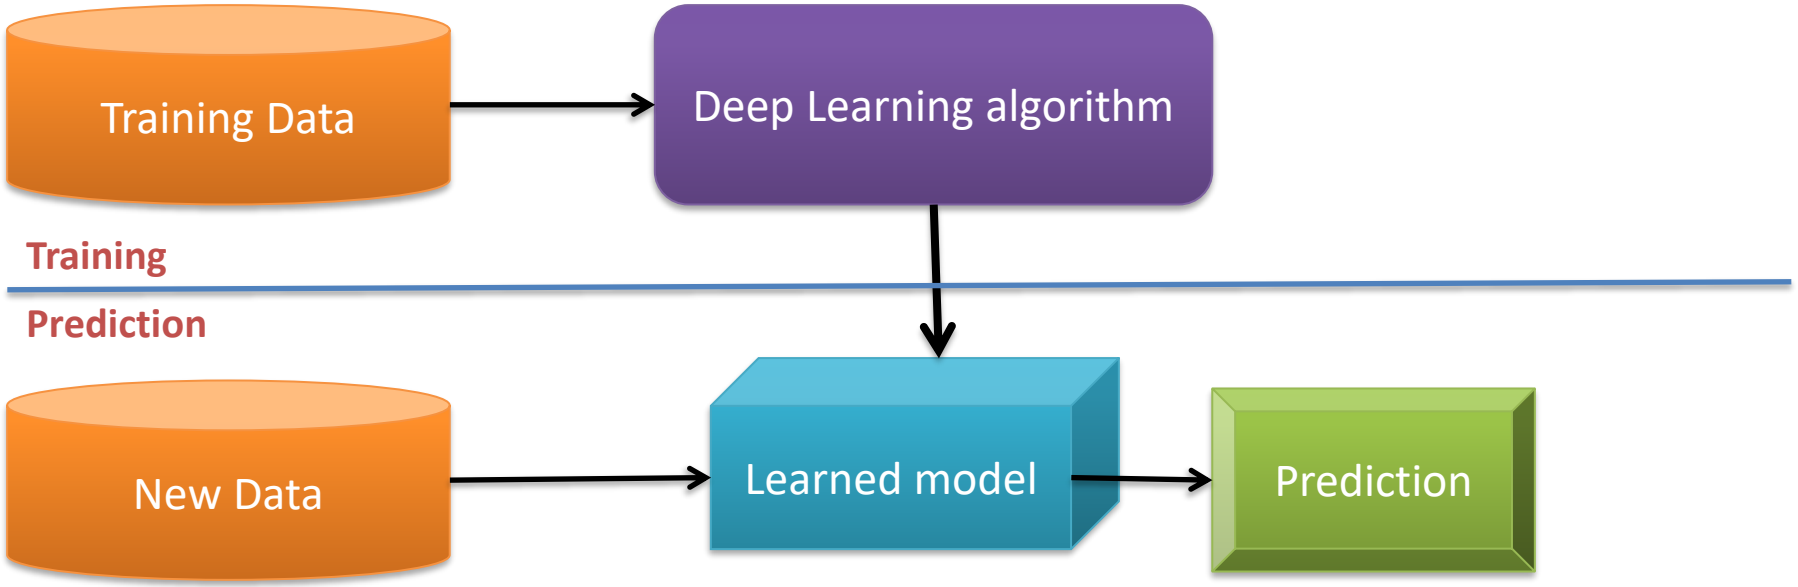

Methods that can learn from and make predictions on data

2. History of Deep Learning

Supervised: Learning with a **labeled training** set of data
Example: learn the *classification* of images based on image labels (dogs/cats, day time, numbers, etc.)

Unsupervised: Discover **patterns** in **unlabeled** data
Example: *cluster* similar documents based on text

Reinforcement learning: learn to **act** based on **feedback/reward**
Example: learn to play Go, reward: *win or lose*

Classification

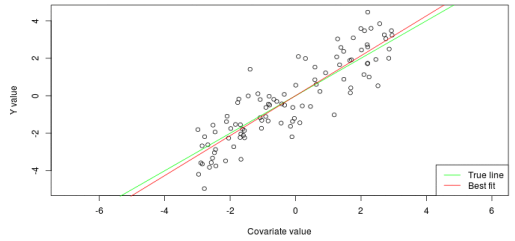

Regression

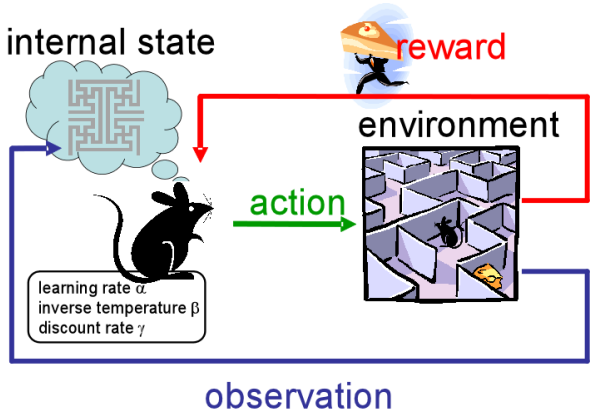

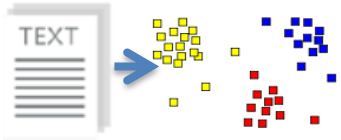

Clustering

Sources: *http://mbjoseph.github.io/2013/11/27/measure.html*
https://becominghuman.ai/the-very-basics-of-reinforcement-learning-154f28a79071

2. History of Deep Learning

Supervised: Learning with a **labeled training** set of data
Example: learn the *classification* of images based on image labels (dogs/cats, day time, numbers, etc.)

Unsupervised: Discover **patterns** in **unlabeled** data
Example: *cluster* similar documents based on text

Reinforcement learning: learn to **act** based on **feedback/reward**
Example: learn to play Go, reward: *win or lose*

Classification

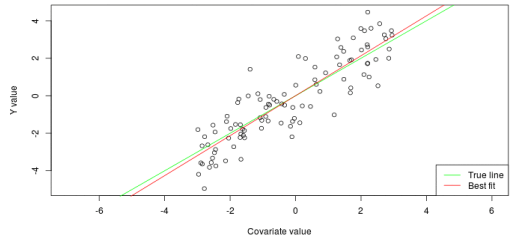

Regression

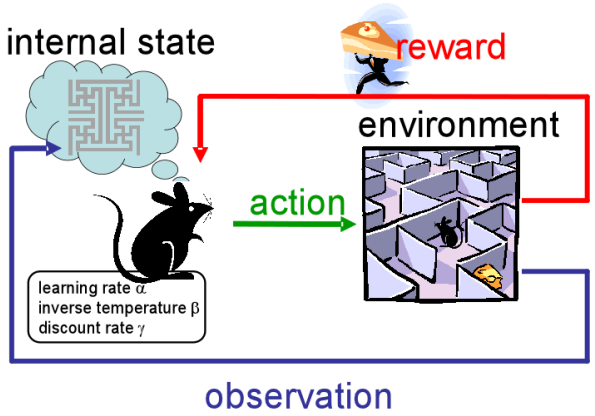

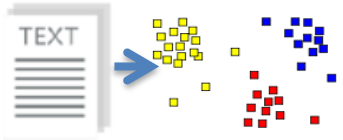

Clustering

Sources: *http://mbjoseph.github.io/2013/11/27/measure.html*
https://becominghuman.ai/the-very-basics-of-reinforcement-learning-154f28a79071

2. History of Deep Learning

Most deep learning methods work well because of **human-designed representations** and **input features**
DL becomes just **optimizing weights** to best make a final prediction

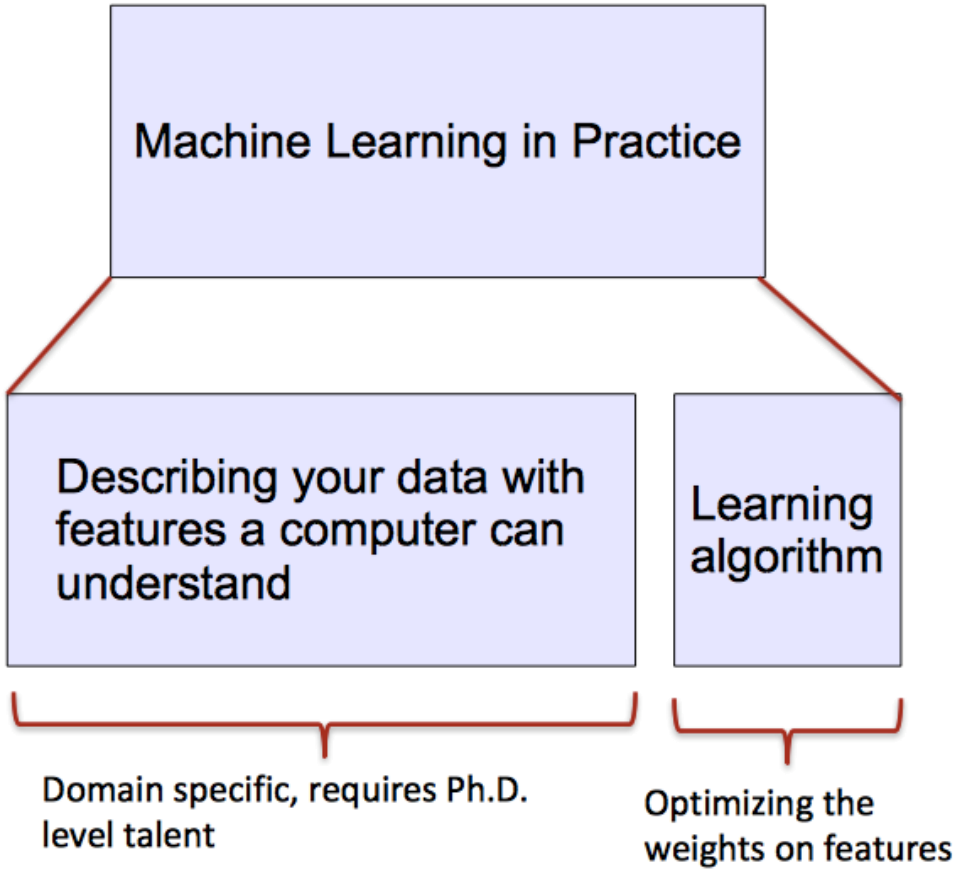

Feature	NER
Current Word	✓
Previous Word	✓
Next Word	✓
Current Word Character n-gram	all
Current POS Tag	✓
Surrounding POS Tag Sequence	✓
Current Word Shape	✓
Surrounding Word Shape Sequence	✓
Presence of Word in Left Window	size 4
Presence of Word in Right Window	size 4

NER: Named Entity Recognition
POS: Part of Speech

2. History of Deep Learning

- Hierarchy of representations with increasing levels of abstraction
- Image recognition
  - Pixel ➔ edge ➔ texton ➔ motif ➔ part ➔ object
- Text
  - Character ➔ word ➔ word group ➔ clause ➔ sentence ➔ story
- Speech
  - Sample ➔ spectral band ➔ sound ➔ … ➔ phone ➔ phoneme ➔ word

2. History of Deep Learning

- A sub-field of machine learning for learning **representations** of data.
- Exceptionally effective at **learning patterns**.
- Deep learning algorithms attempt to learn (multiple levels of) representation by using a **hierarchy of multiple layers**
- If you provide the system **tons of information**, it begins to understand it and respond in useful ways.

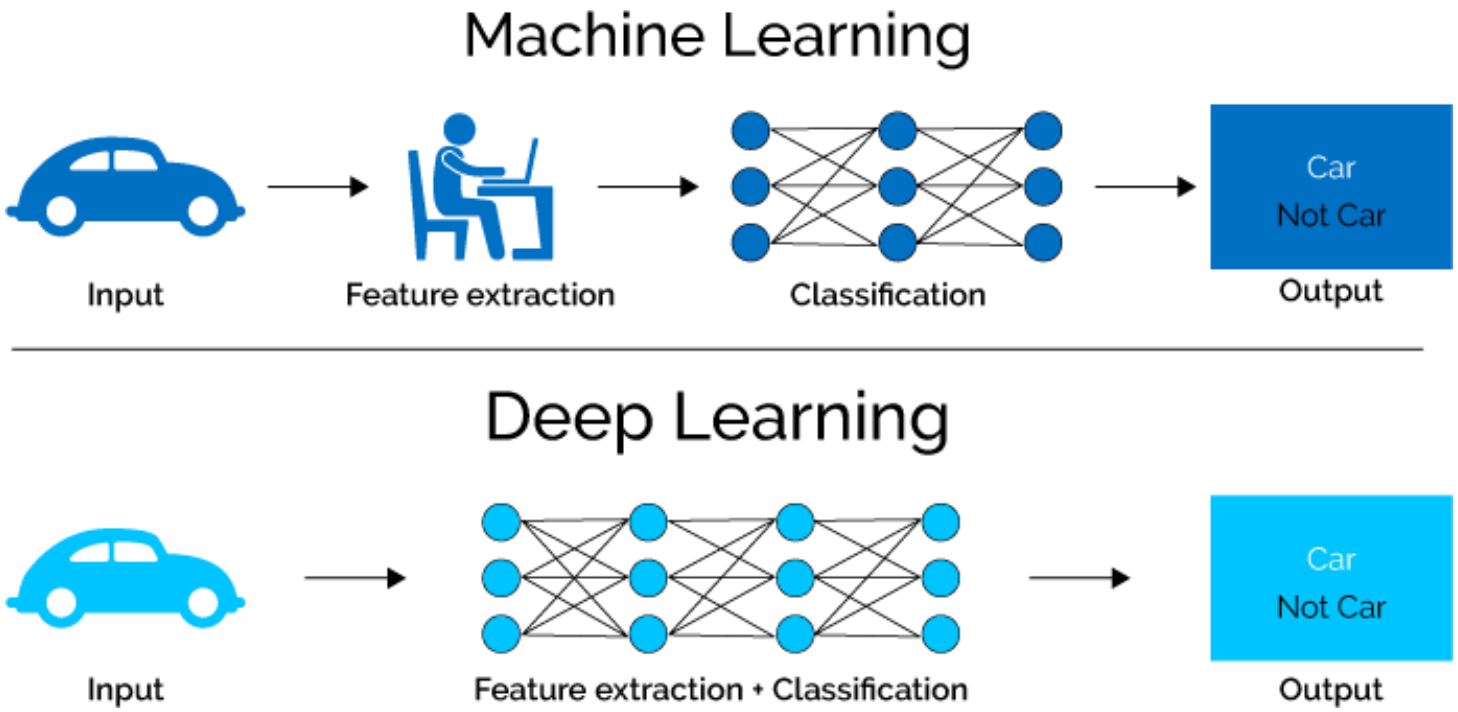

https://www.xenonstack.com/blog/static/public/uploads/media/machine-learning-vs-deep-learning.png

2. History of Deep Learning

Why is DL useful?

- Manually designed features are often **over-specified**, **incomplete** and take a **long time to design** and validate
- Learned Features are **easy to adapt**, **fast** to learn
- Deep learning provides a very **flexible**, (almost?) **universal**, learnable framework for representing world, visual and linguistic information.
- Can learn in both unsupervised and supervised ways
- Effective **end-to-end** joint system learning
- Utilize large amounts of training data

Around 2010, DL started to outperform other ML techniques, first in speech and vision, then in Natural Language Processing (NLP)

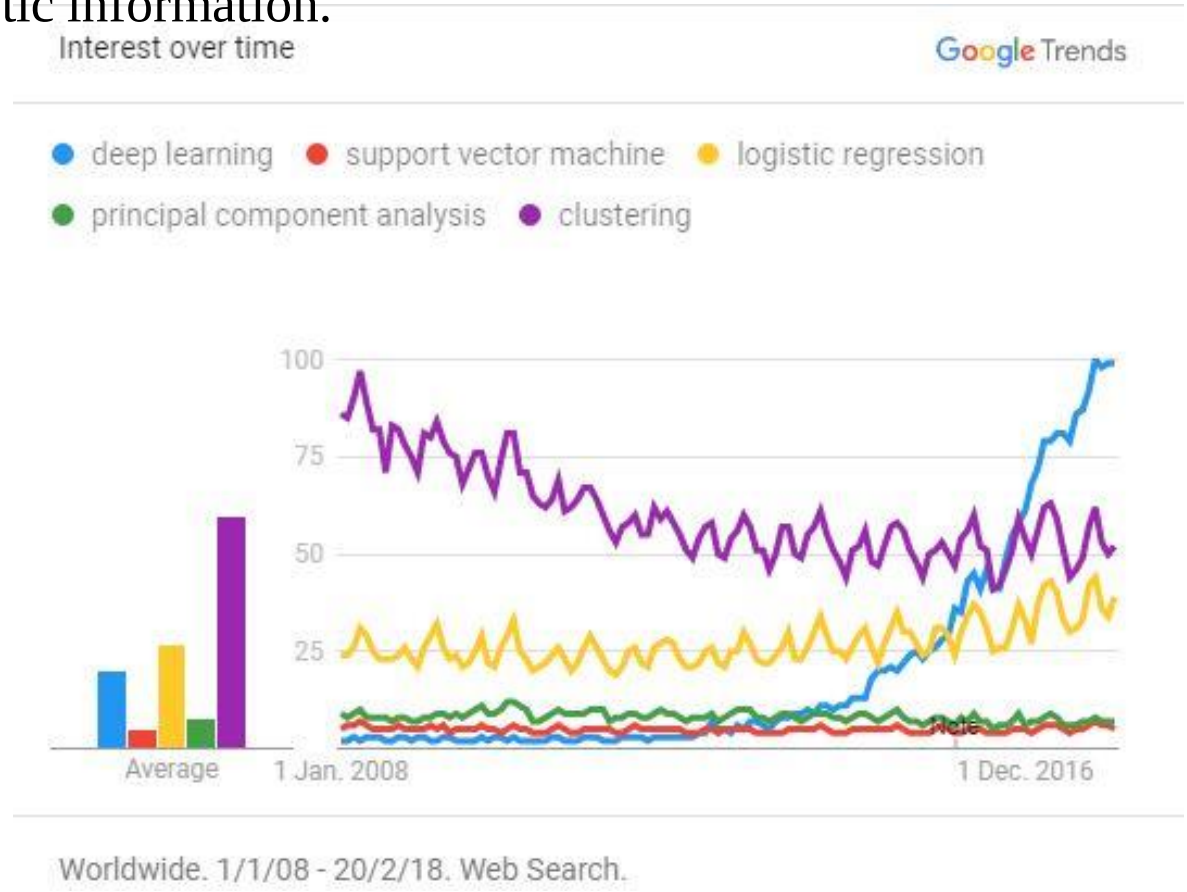

2. History of Deep Learning

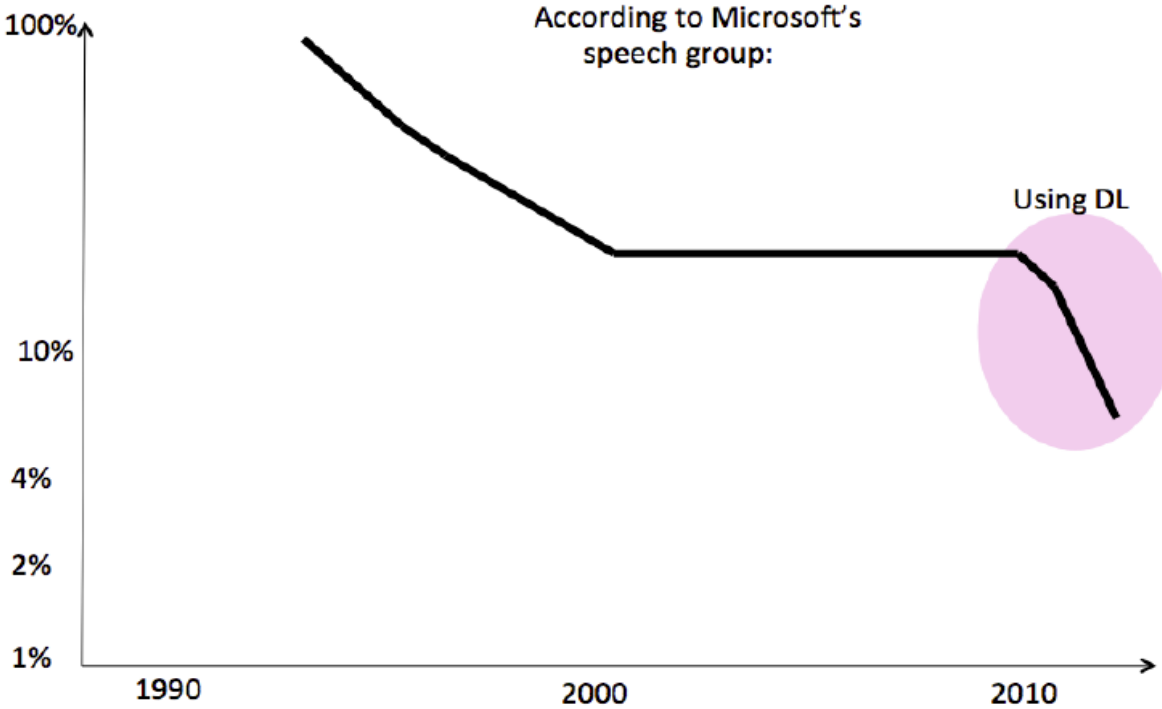

Deep Learning in Speech Recognition

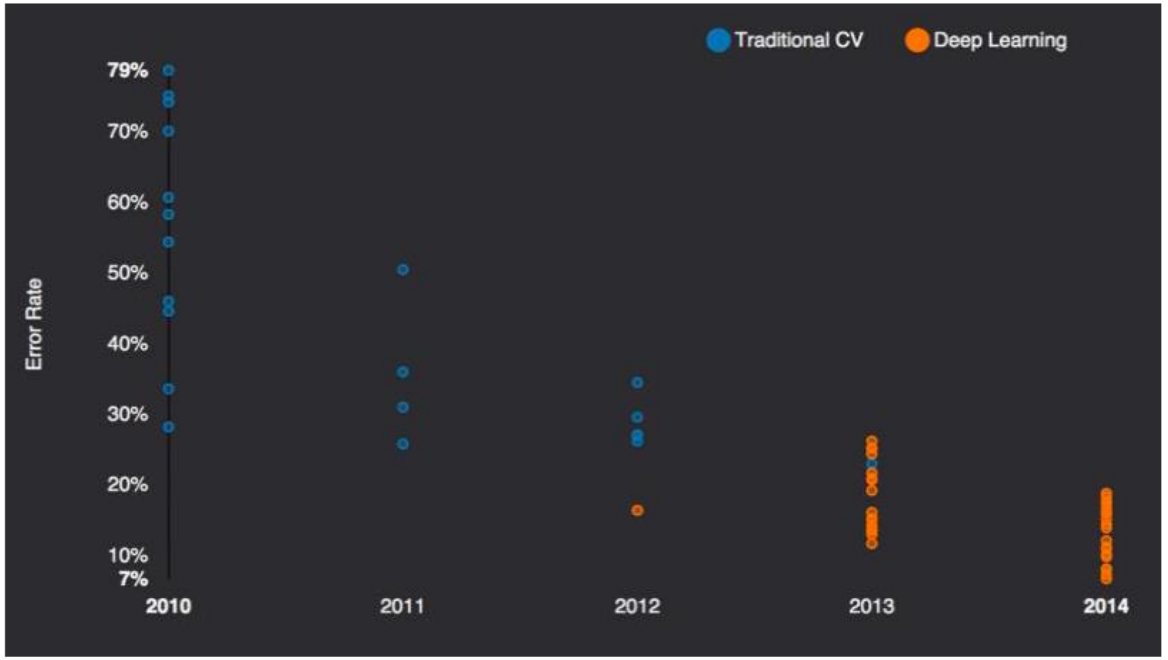

ImageNet: The "computer vision World Cup"

Several big improvements in recent years in NLP

- ✓ Machine Translation
- ✓ Sentiment Analysis
- ✓ Dialogue Agents
- ✓ Question Answering
- ✓ Text Classification …

Leverage different levels of representation

- words & characters
- syntax & semantics

2. History of Deep Learning

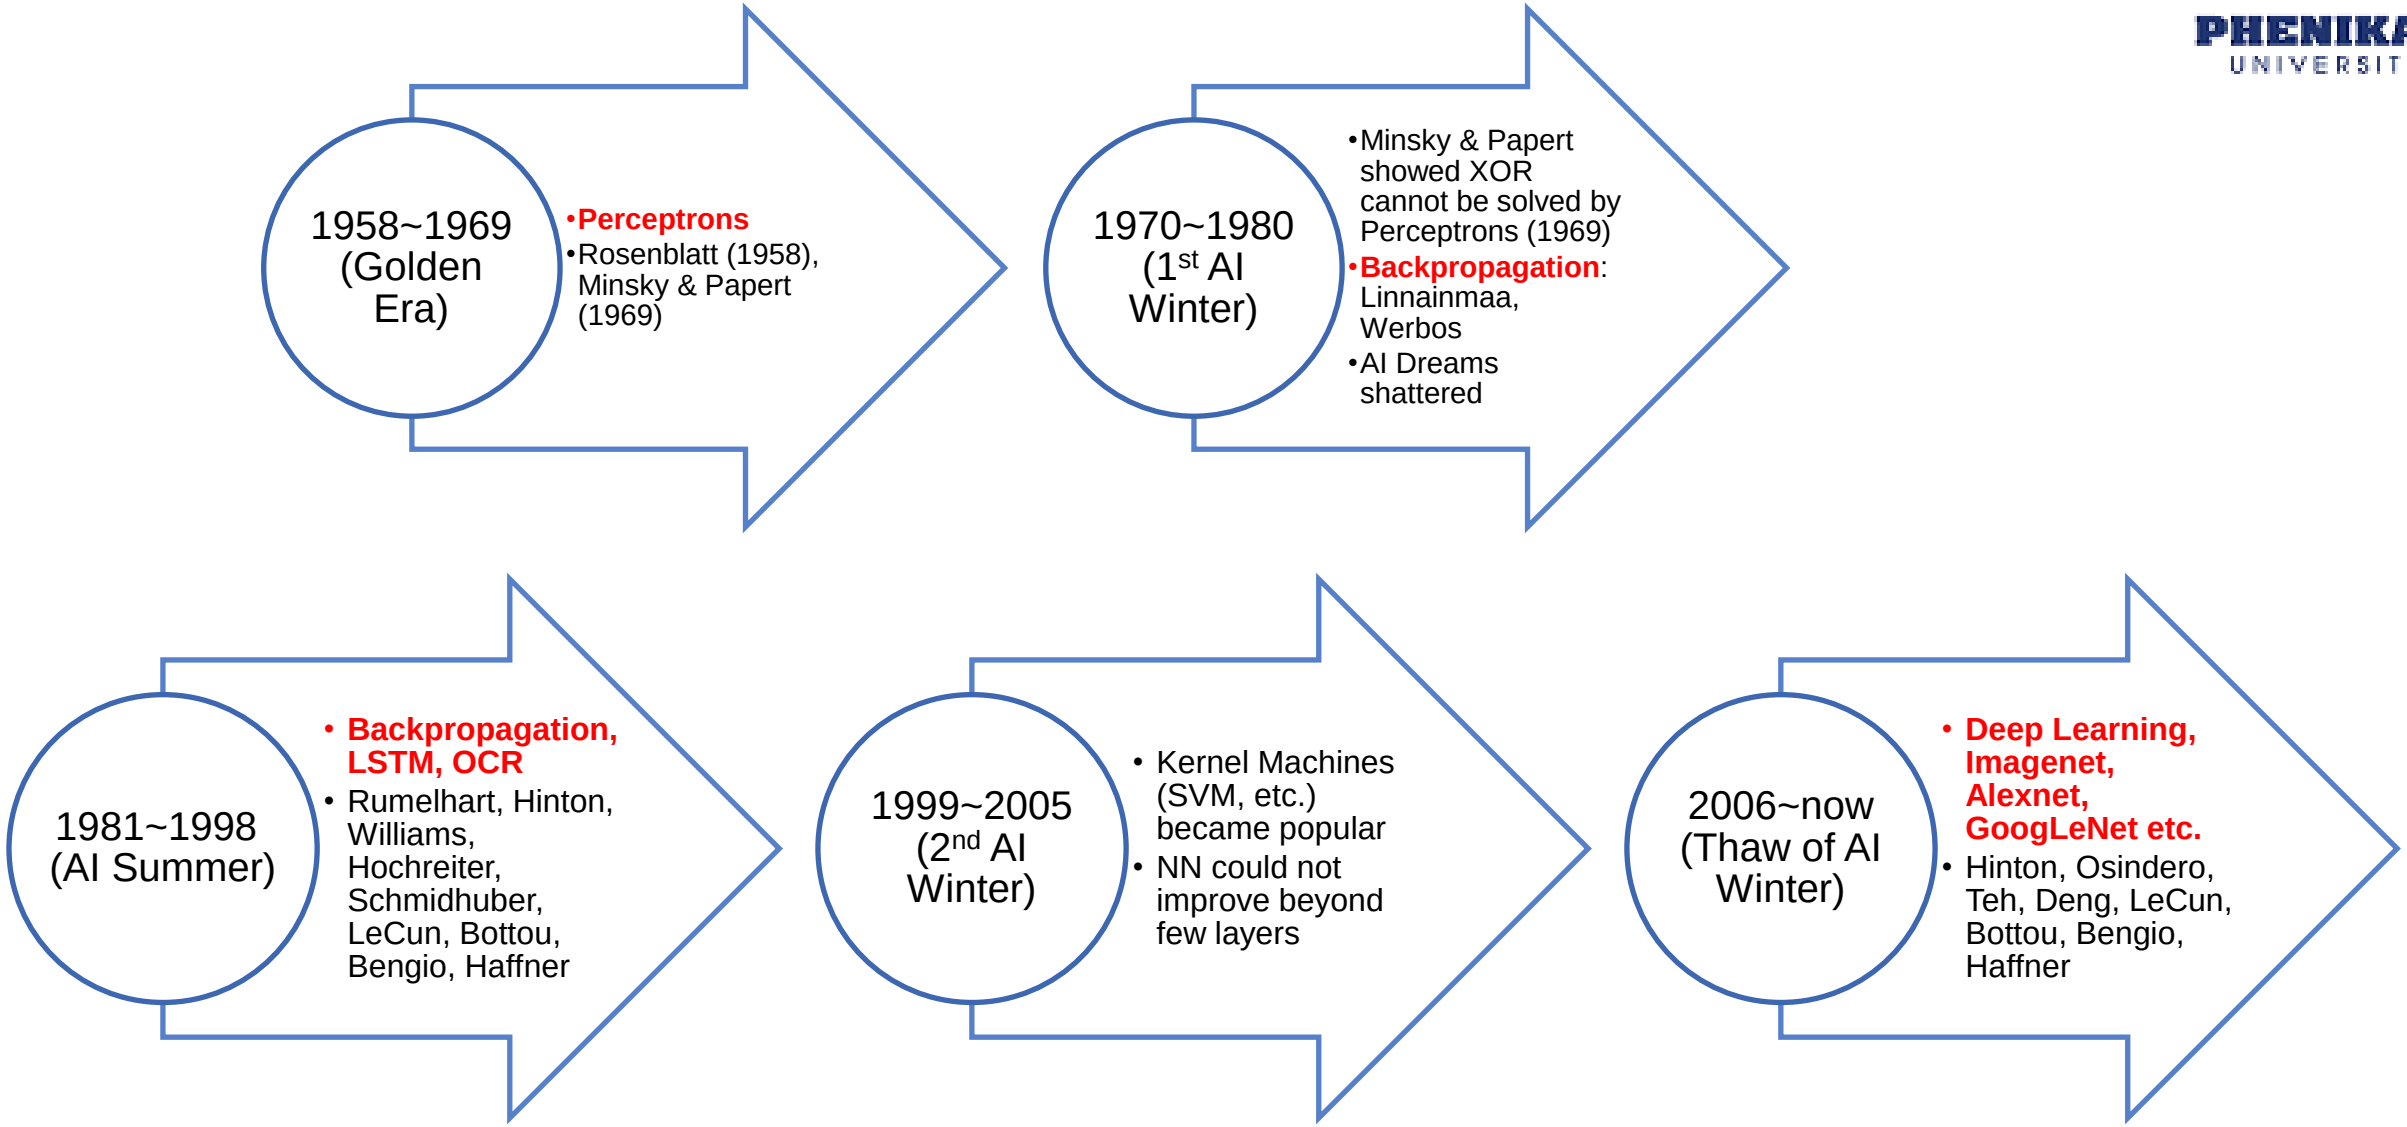

2. History of Deep Learning

1958-1969

- Rosenblatt proposed a machine for binary classifications
- Main idea
  - One weight w_i per input x_i
  - Multiply weights with respective inputs and add bias w_0
  - If result is larger than threshold δ, return 1, otherwise 0.

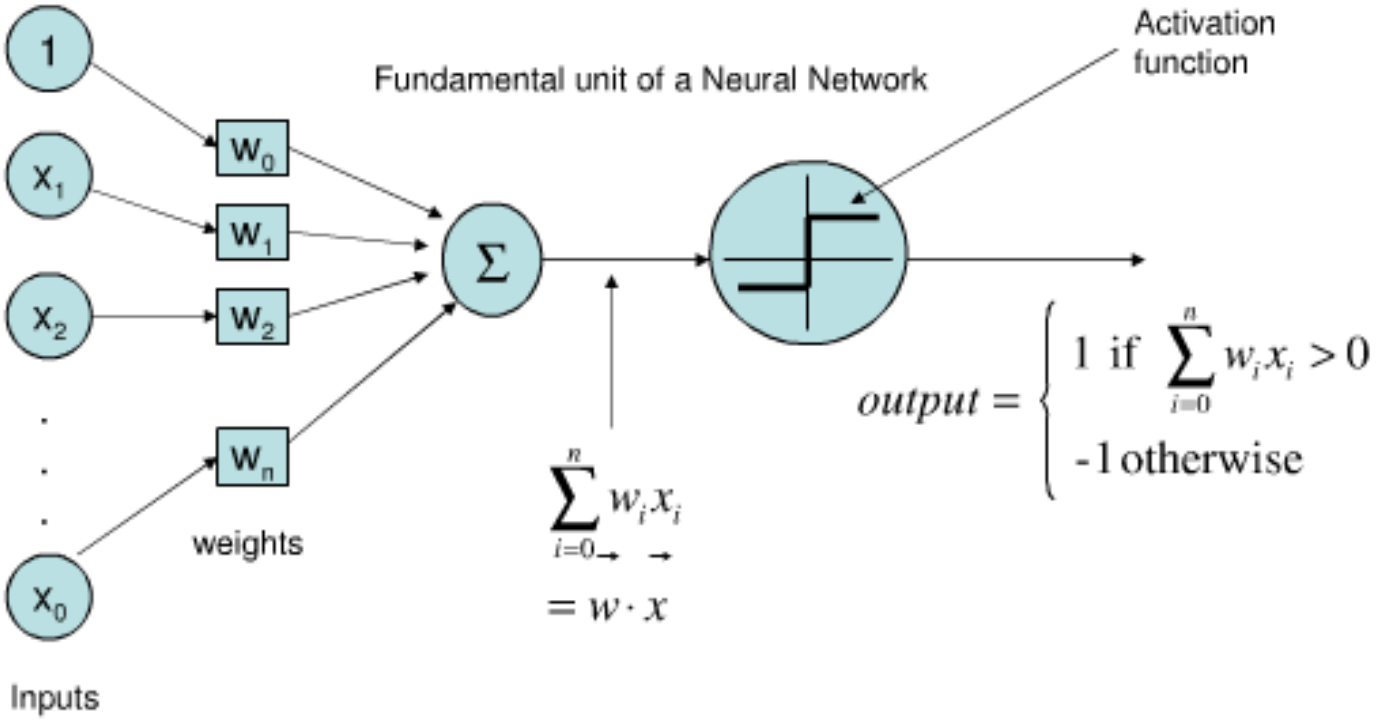

2. History of Deep Learning

Training a Perceptron

- **Learning Algorithm** by Rosenblatt
  - Initialize weights randomly
  - Take one sample x_i and predict y_i
  - For erroneous predictions, update weights
    - If the output was $\hat{y} = 0$ and $y_i = 1$, increase weights
    - If the output was $\hat{y} = 1$ and $y_i = 0$, decrease weights
  - Repeat until no errors are made

Video (18'54"): https://youtu.be/OVHc-7GYRo4

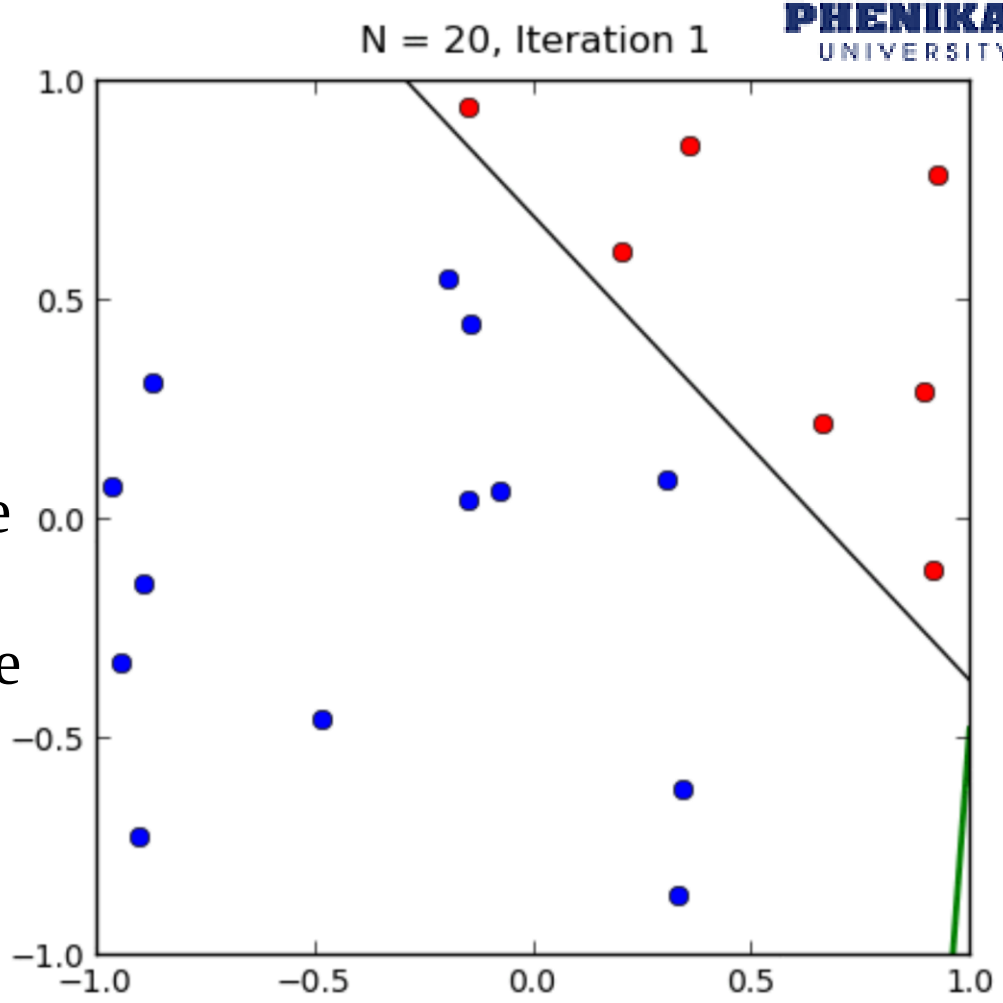

2. History of Deep Learning

Multi-layer Perceptron

- One perceptron = one decision
- Question: What about multiple decisions?
  - Eg. Digit classification
- Answer: Neural Network (NN) or Multi-Layer Perceptron (MLP)
  - Stack multiple perceptrons (neurons) into a single layer
  - Connect two or more layers by feeding output of one layer as input to the next layer

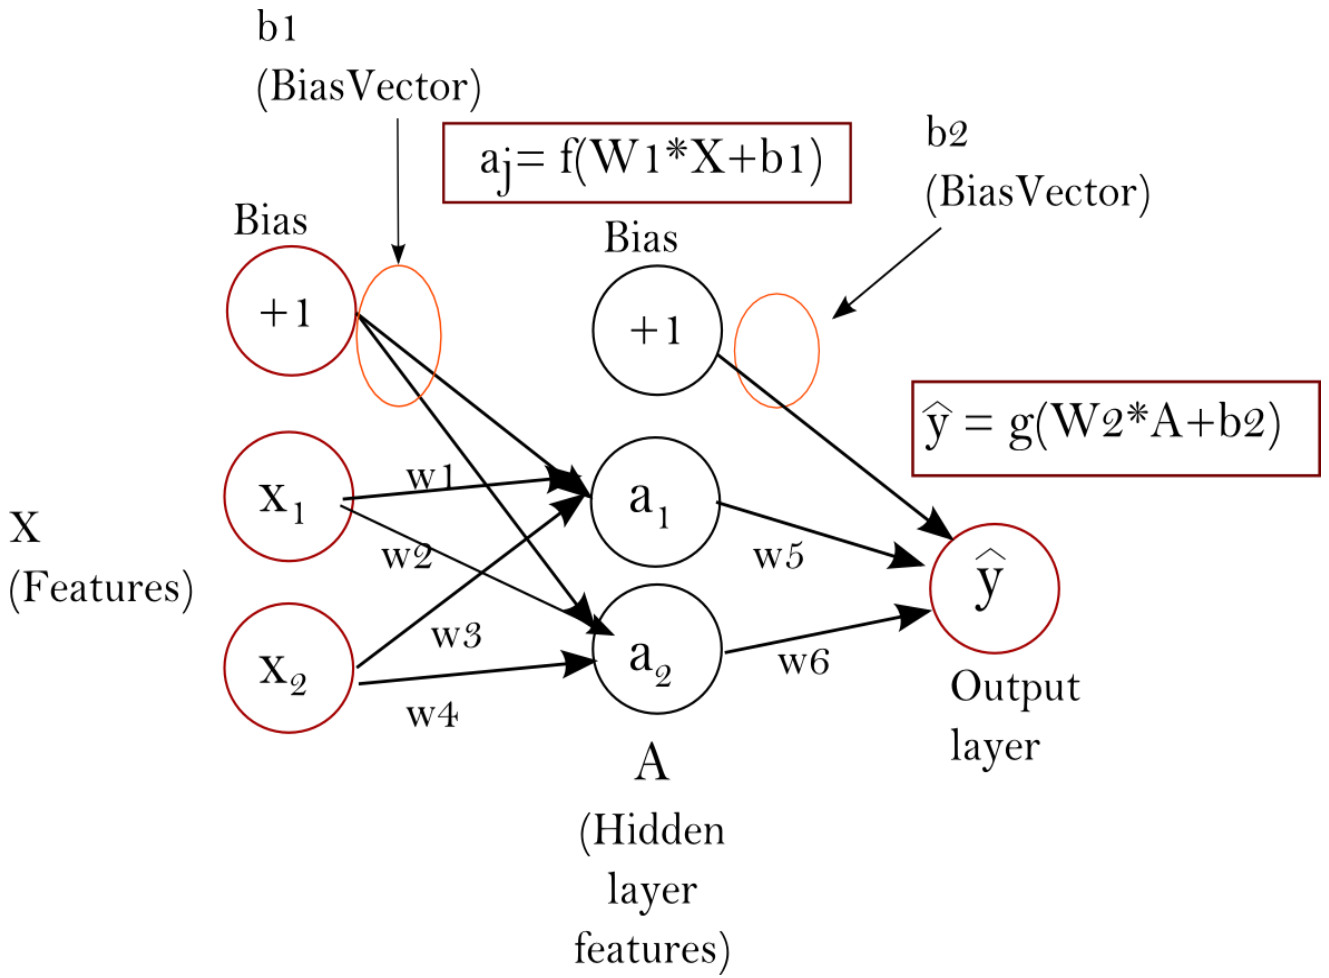

2. History of Deep Learning

1970~1980: 1st AI Winter

- XOR cannot be solved by Perceptron (Minsky)
- Perceptron training method cannot be applied to Neural Networks
- Funding slushed, Neural Networks were damned
- AI WINTER!!!
- Dreams shattered!
- Some significant results
  - Backpropagation: training method for NN (1970, 1974)

2. History of Deep Learning

1999~2005: 2nd AI Winter

- Kernel Machines (e.g. Support Vector Machines (SVM), etc.) became popular
  - Achieved similar accuracies
  - Included much fewer heuristics
  - Nice proofs on generalization
- Neural networks could not improve beyond a few layers
  - Lack of processing power (No GPUs)
  - Lack of data (No big, annotated datasets)
  - Overfitting (Models could not generalize)
  - Vanishing gradients (0.1*0.1*0.1*….*0.1 = 0.000000000001, too small for learning)
- AI community turned away from Neural Networks

2. History of Deep Learning

2006~now: Thaw of AI Winter

- Questions:
  - Are 1-2 hidden layers the best NN can do?
  - Or, is it the learning algorithm not really mature?
- Deep Learning (2006, Hinton, Osindero, Teh)
- Layer-by-layer training
  - Per-layer trained parameters initialize further training using contrastive divergence

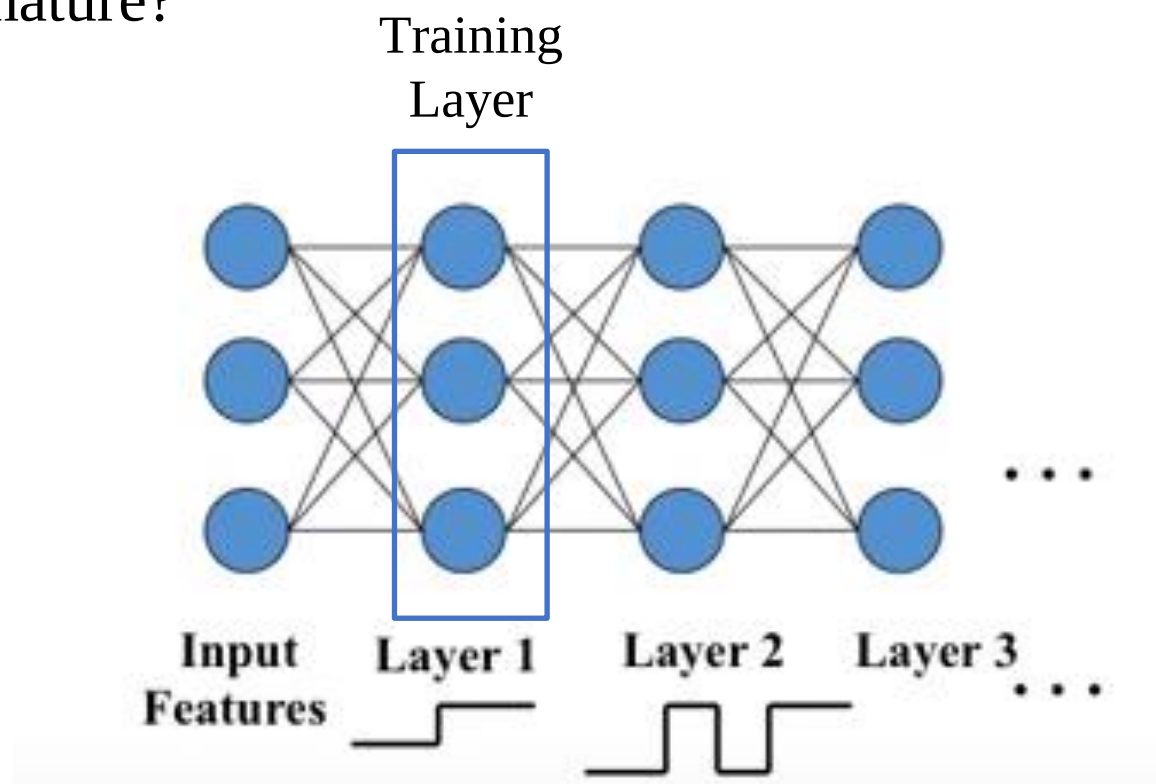

2. History of Deep Learning

Deep Learning is here …

- **ImageNet** dataset (Deng et al, 2009)
  - Collected images for each term of Wordnet (100,000 classes)
  - Tree of concepts organized hierarchically
    - "Ambulance", "Dalmatian dog", "Egyptian cat", …
  - **Imagenet Large Scale Visual Recognition Challenge** (ILSVRC)
    - 1 million images
    - 1,000 classes
    - Top-5 and top-1 error measured
      - Errors reduced drastically in the past 8 years (2010~2017): 28.2% ➔ 2.3%

2. History of Deep Learning

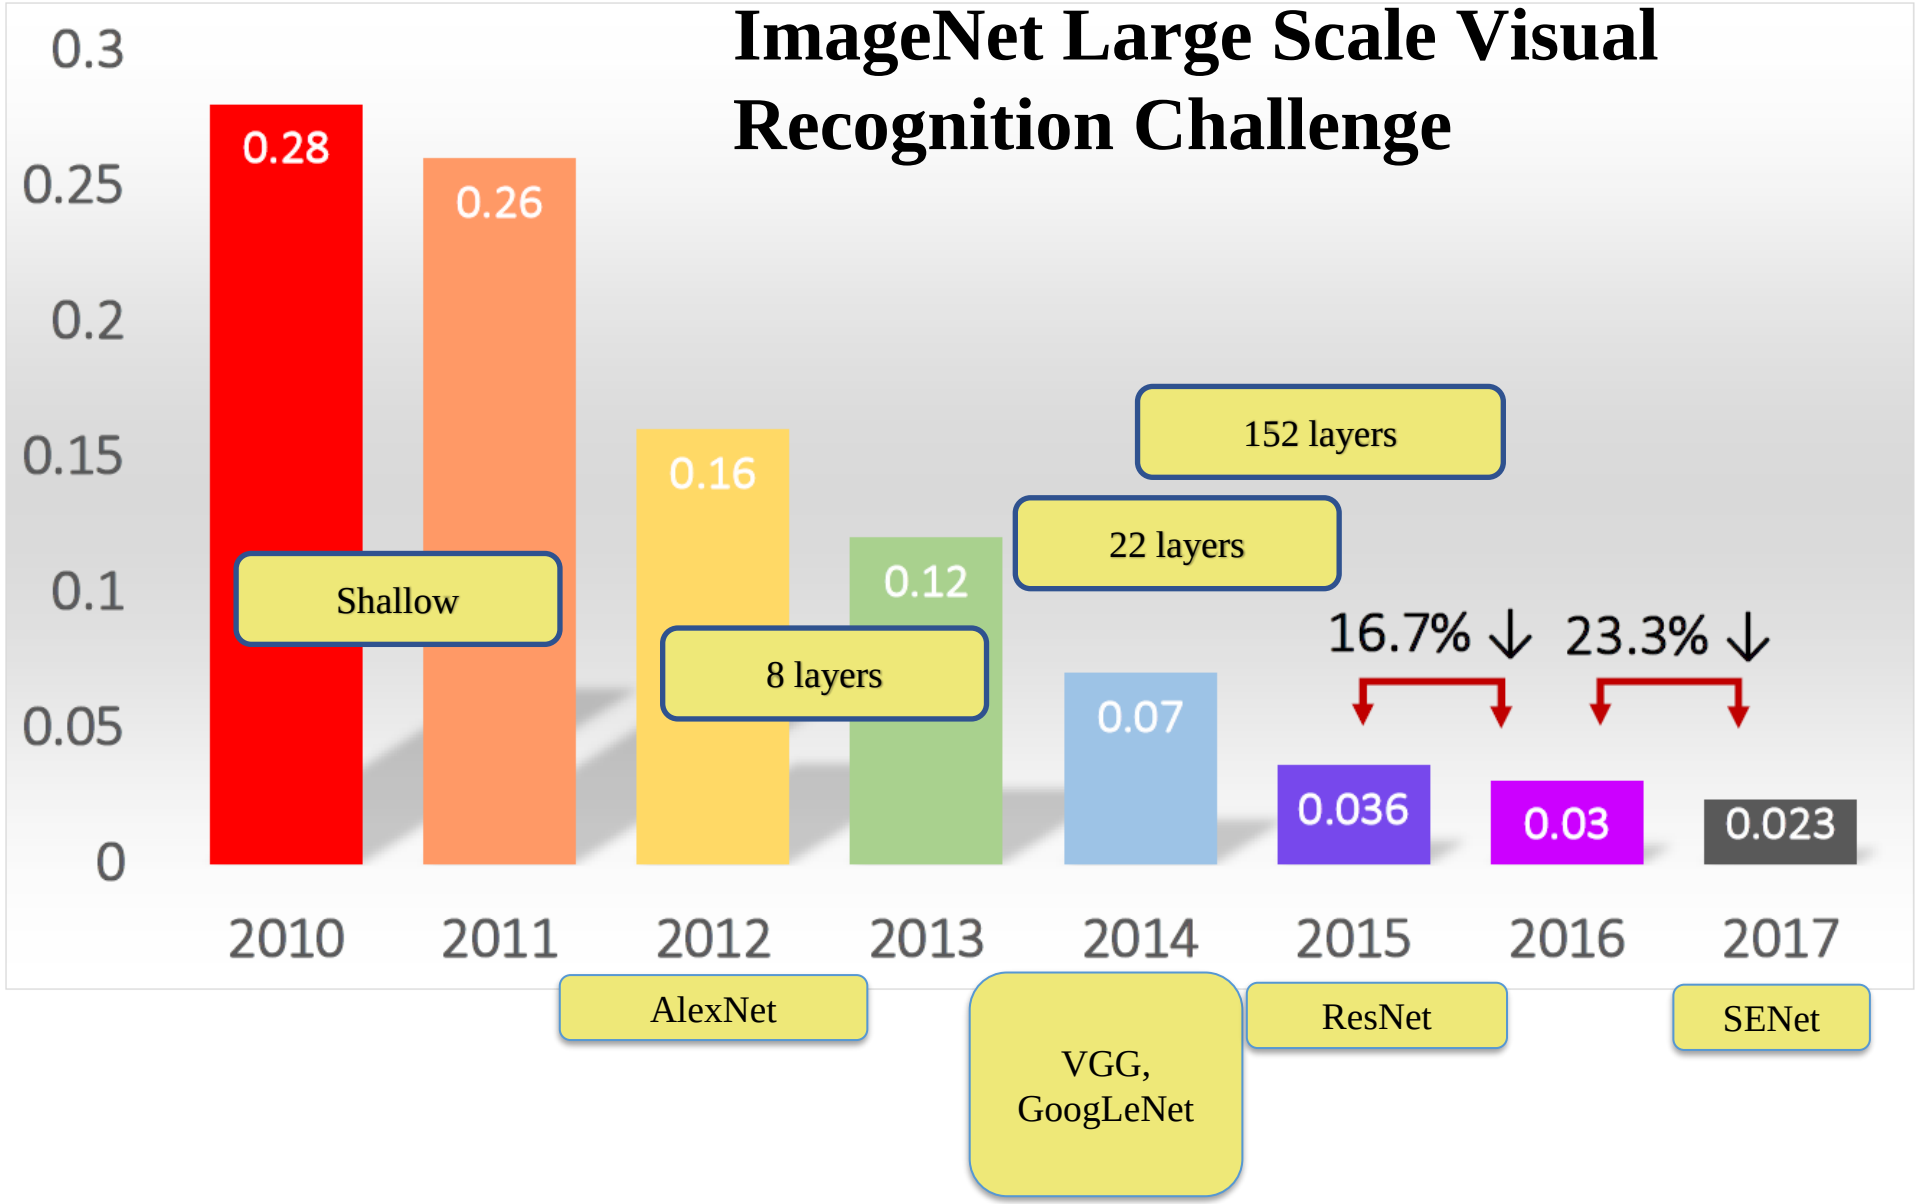

2. History of Deep Learning

Some FUN now …

- The Neural Network ZOO
  - Graphical notations for all kinds of neural networks
  - http://www.asimovinstitute.org/neural-network-zoo/
- A Neural Network Playground
  - An online interactive way to play with different network architectures
  - http://playground.tensorflow.org
- 8 Inspirational Applications of Deep Learning
  - Very interesting applications of deep learning
  - http://machinelearningmastery.com/inspirational-applications-deep-learning/

3. Applications of Deep Learning

(1/8) Automatic Colorization of B&W Images

- Large Convolutional Neural Networks (CNN)
- Website http://richzhang.github.io/colorizat/
- Video (5 s) http://whattogive.com/videoColou tion/

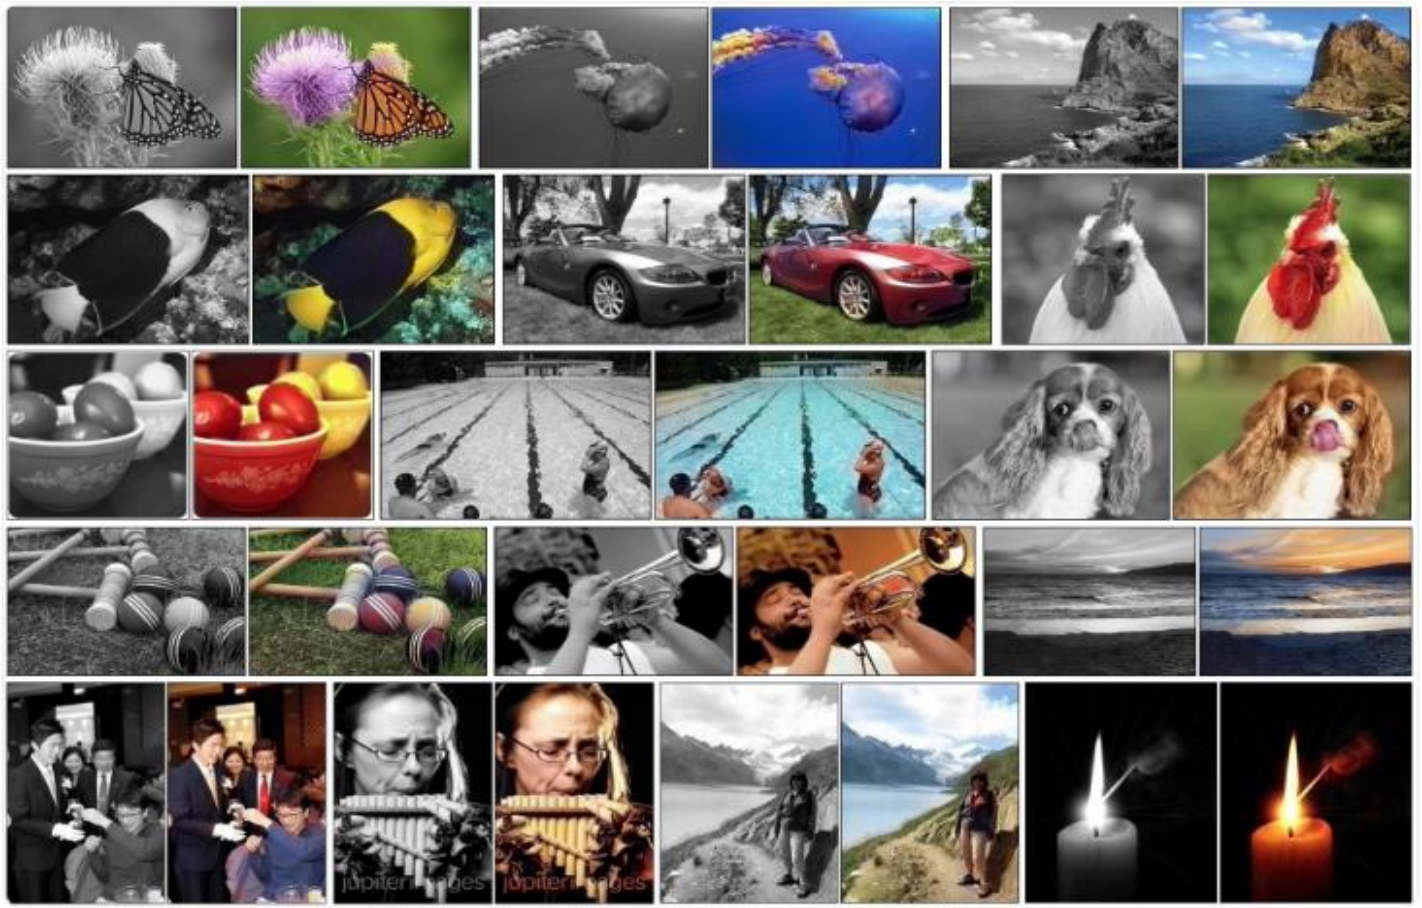

3. Applications of Deep Learning

(2/8) Automatically Adding Sounds

- Two types of NN
  - Large CNN for images
  - Large Long Short-Term Memory (LSTM) Recurrent Neural Networks (RNN) for sound
- News
  - http://news.mit.edu/2016/artificial-intelligence-produces-realistic-sounds-0613
- Visually Indicated Sounds (MIT)
  - http://vis.csail.mit.edu/
- Video (2.54 s)
  - https://youtu.be/0FW99AQmMc8

(3/8) Automatic Machine Translation

- Automatic Translation of Text
  - Large Long Short-Term Memory (LSTM) Recurrent Neural Networks
- Automatic Translation of Images
  - CNN + LSTM RNN

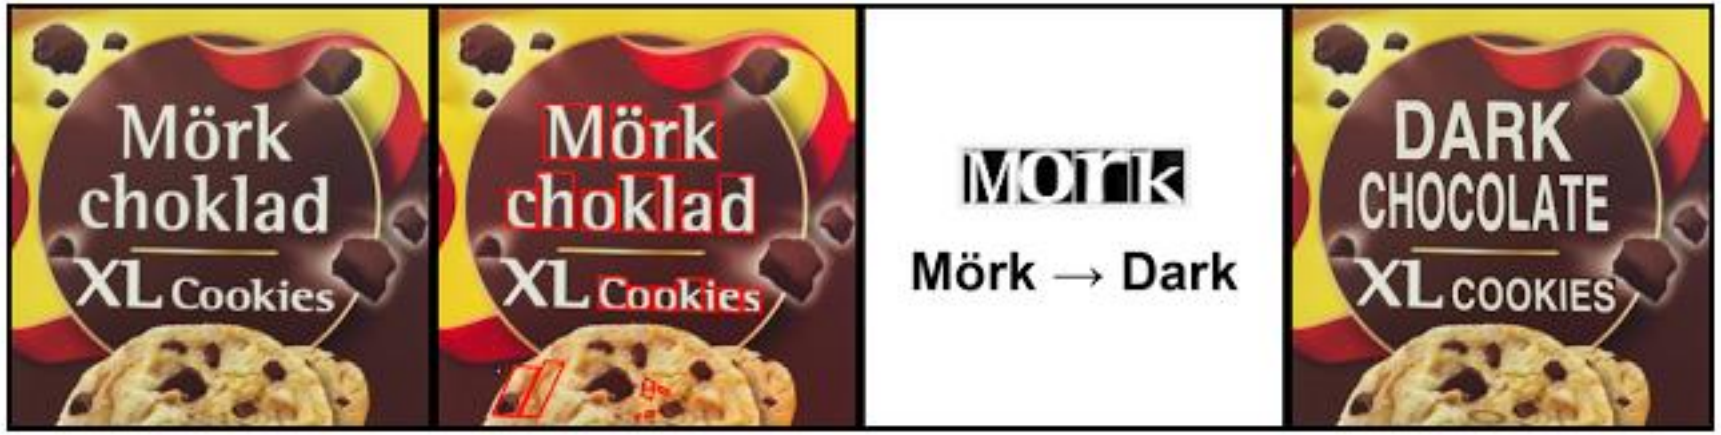

3. Applications of Deep Learning

(4/8) Object Classification & Detection

- Large deep CNN
- Paper on ImageNet Classification
  - http://www.cs.toronto.edu/~fritz/absps/imagenet.pdf

ConvNetJS: CIFAR-10 Demo

- ConvNetJS: CIFAR-10 Demo
  - http://cs.stanford.edu/people/karpathy/convnetjs/demo/cifar10.html
- Clarifai: 10,000 images
  - https://www.clarifai.com/

3. Applications of Deep Learning

(5/8) Automatic Handwriting Generation

- Usage: used with forensic analysis

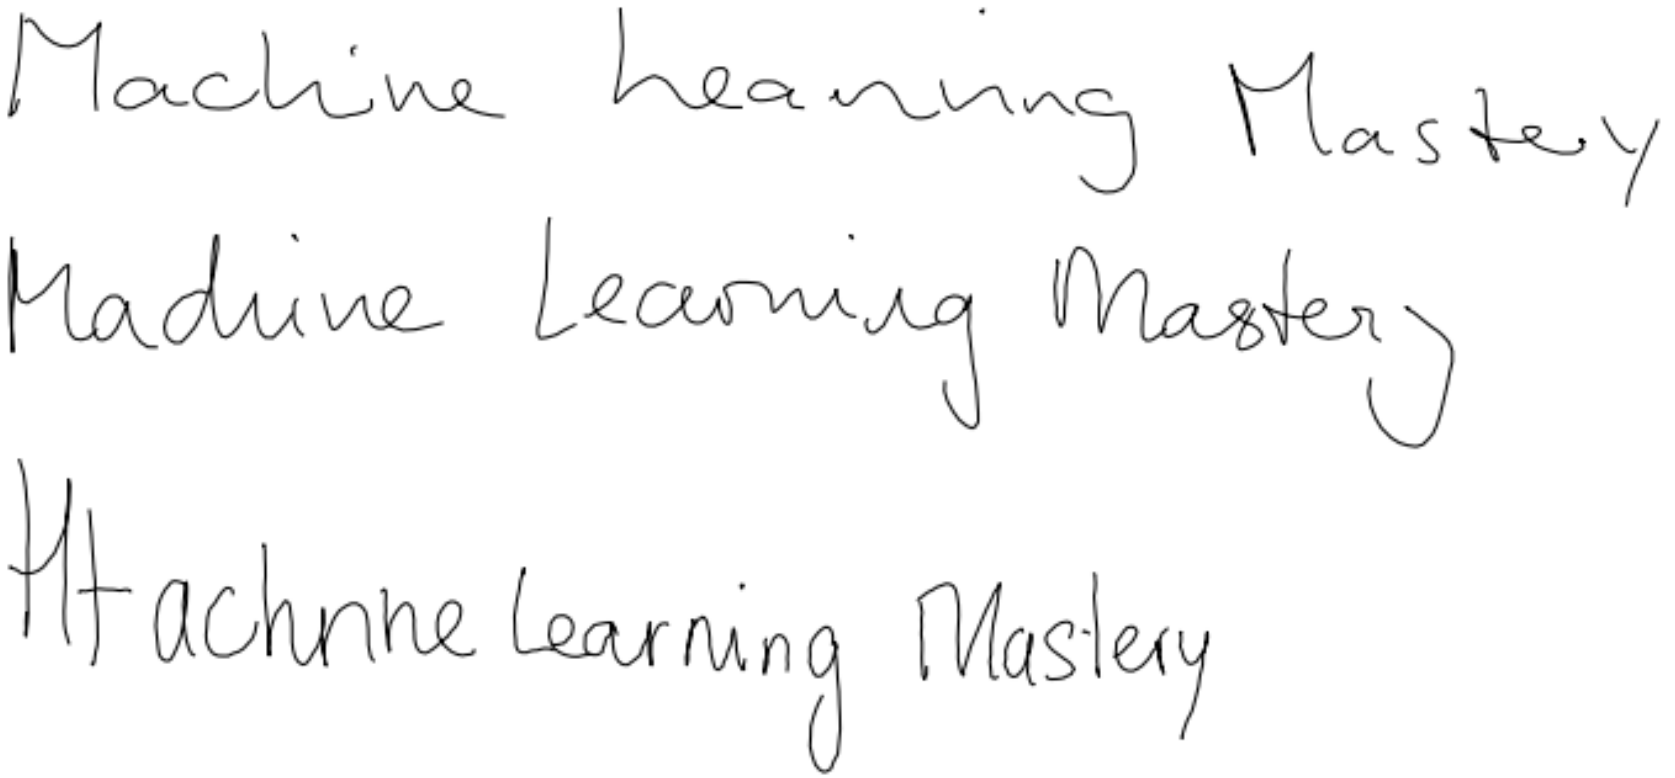

DEMO: http://www.cs.toronto.edu/~graves/handwriting.html

3. Applications of Deep Learning

(6/8) Automatic Text Generation

- Large RNN
- Code on Github
  - https://github.com/karpathy/char-rnn
- Paul Graham generator
- Shakespeare
- Wikipedia
- Algebraic Geometry (LaTeX)
- Linux source code
- Generating Baby Names

```
PANDARUS:
Alas, I think he shall be come approached and the day
When little srain would be attain'd into being never fed,
And who is but a chain and subjects of his death,
I should not sleep.

Second Senator:
They are away this miseries, produced upon my soul,
Breaking and strongly should be buried, when I perish
The earth and thoughts of many states.

DUKE VINCENTIO:
Well, your wit is in the care of side and that.

Second Lord:
They would be ruled after this chamber, and
my fair nues begun out of the fact, to be conveyed,
Whose noble souls I'll have the heart of the wars.

Clown:
Come, sir, I will make did behold your worship.

VIOLA:
I'll drink it.
```

Automatic Music Synthesis

- Using large RNN
- Authors comments:
  - This track was made using a RNN. Fed 500 mb guitar tabs in ASCII. It writes the tabs out in ASCII, I imported into GuitarPro, recorded the output, imported that into FL Studio, added some filters and a drum loop and got this. The notes and rhythms themselves are totally unedited.
- Music (5:20):
  - https://soundcloud.com/optometrist-prime/recurrence-music-written-by-a-recurrent-neural-network

3. Applications of Deep Learning

(7/8) Automatic Image Caption Generation

- Large CNN → Object Detection
- Large LSTM RNN → Caption Text Generation
- Deep Visual-Semantic Alignments for Generating Image Descriptions
  - http://cs.stanford.edu/people/karpathy/deepimagesent/

3. Applications of Deep Learning

(7/8) Automatic Image Caption Generation

- Demo
  - http://cs.stanford.edu/people/karpathy/deepimagesent/rankingdemo/

3. Applications of Deep Learning

(8/8) Automatic Game Playing

- Vision
- Decision Making
- Etc.

https://youtu.be/TmPfTpjtdgg

4. Materials

- Python Programming Language
- Deep Learning Frameworks
- Calculus

4. Materials

- Python
- Jupyter Notebook
- Google Colab

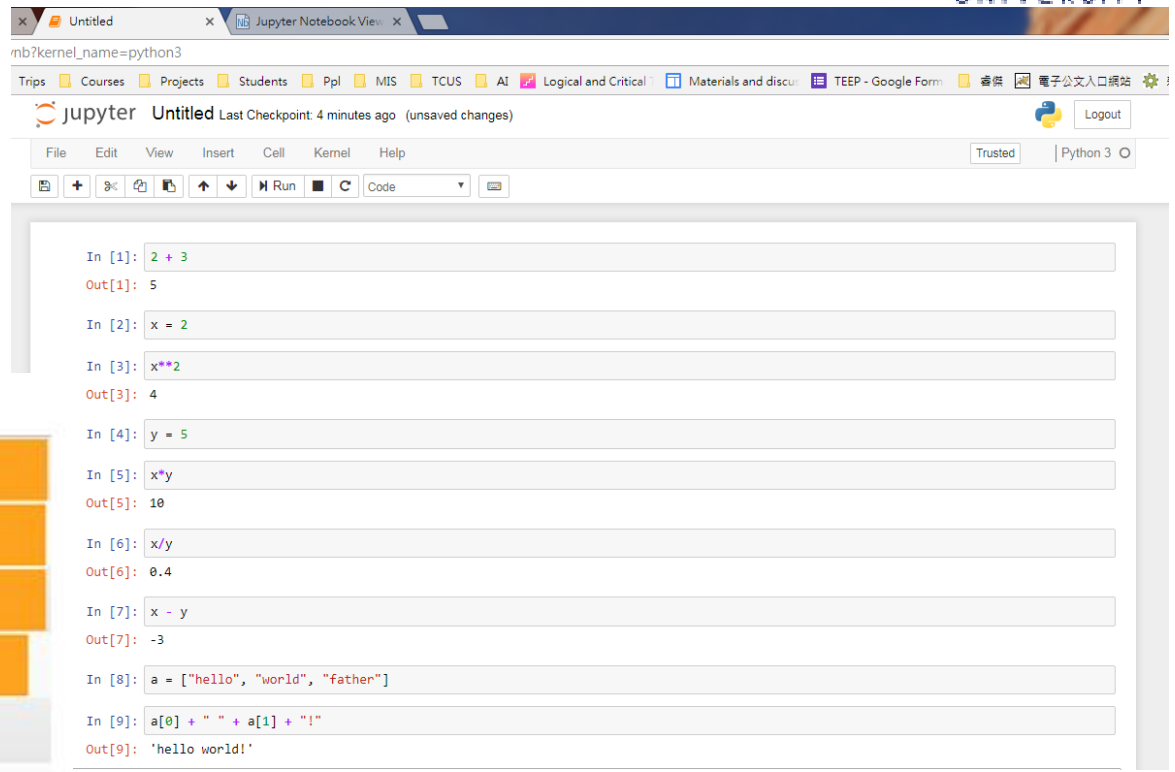

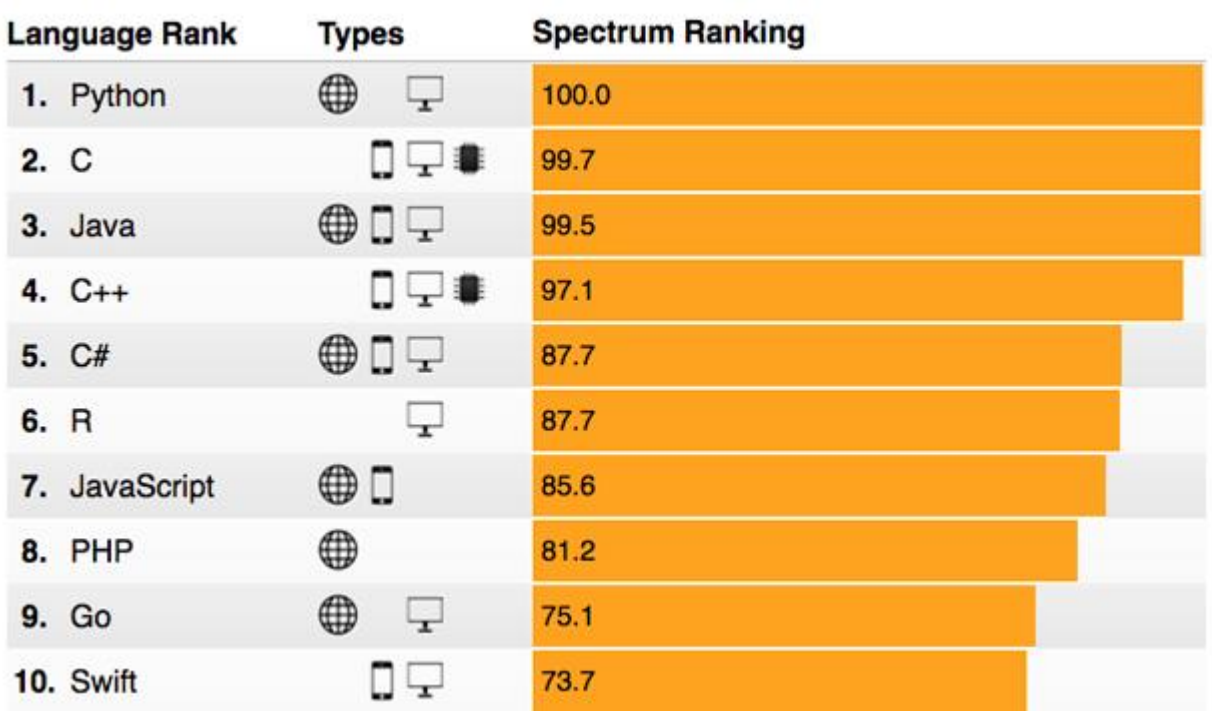

4. Materials

Deep learning Frameworks

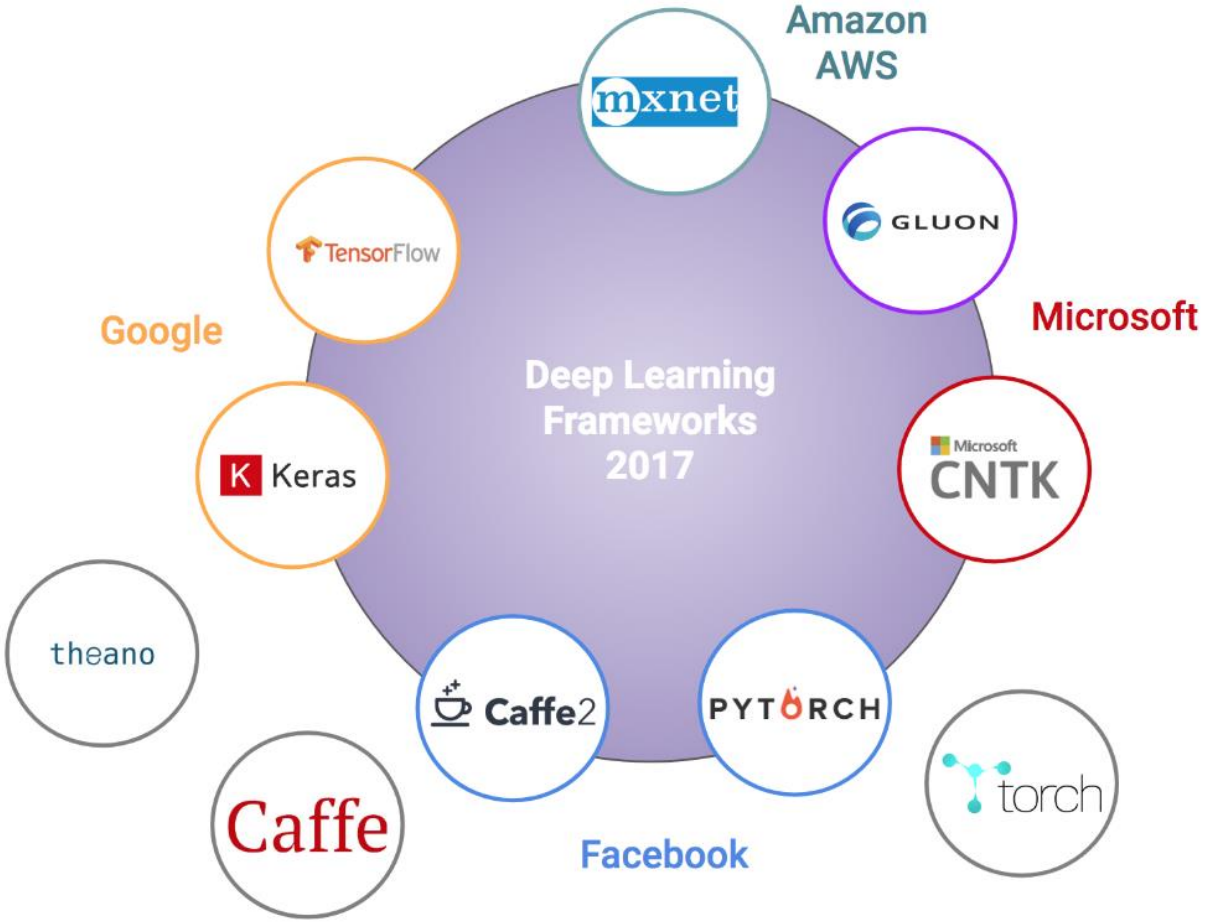

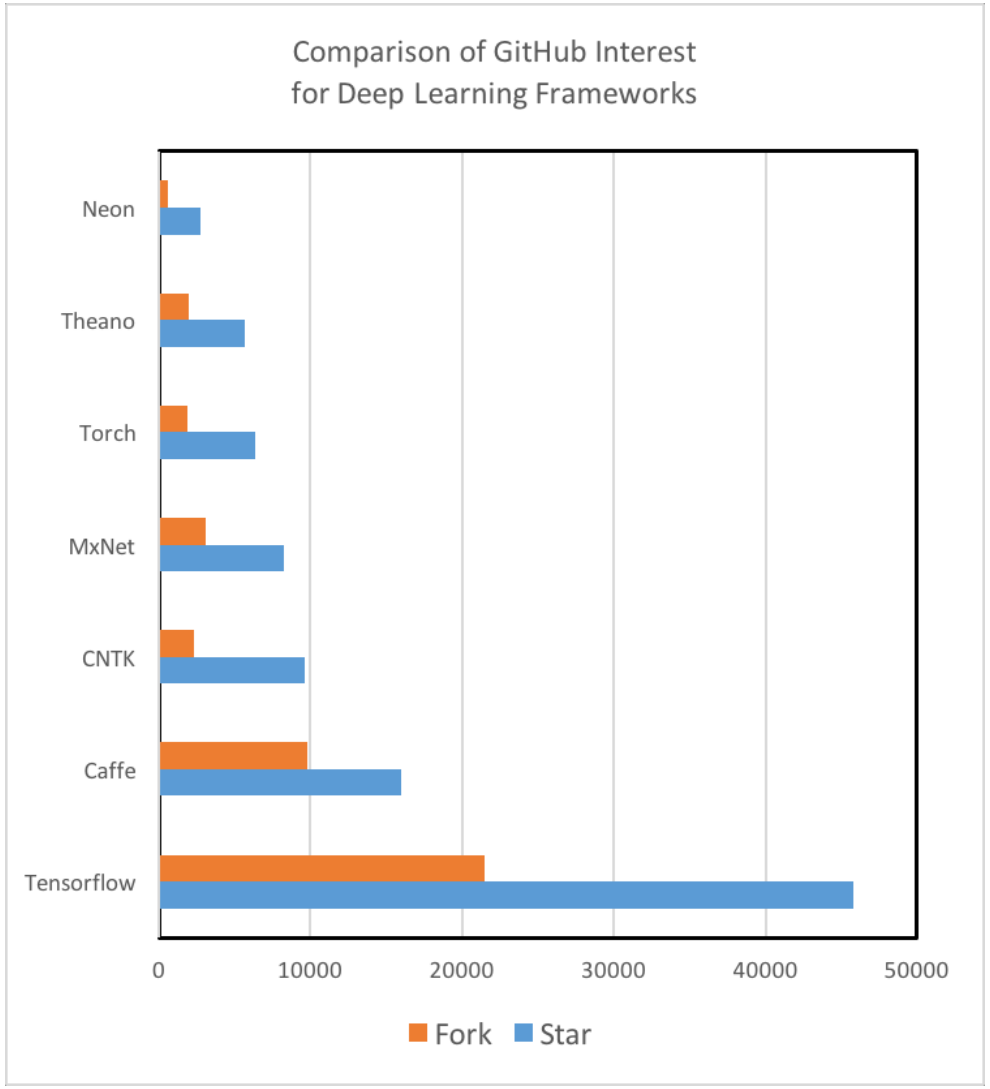

4. Materials

Deep learning Frameworks

https://www.kdnuggets.com/2017/03/getting-started-deep-learning.html

	Languages	Tutorials and training materials	CNN modeling capability	RNN modeling capability	Architecture: easy-to-use and modular front end	Speed	Multiple GPU support	Keras compatible
Theano	Python, C++	++	++	++	+	++	+	+
Tensor-Flow	Python	+++	+++	++	+++	++	++	+
Torch	Lua, Python (new)	+	+++	++	++	+++	++	
Caffe	C++	+	++		+	+	+	
MXNet	R, Python, Julia, Scala	++	++	+	++	++	+++	
Neon	Python	+	++	+	+	++	+	
CNTK	C++	+	+	+++	+	++	+	

4. Materials

PHENIKAA UNIVERSITY

Linear Algebra Review and Reference

Zico Kolter (updated by Chuong Do)

September 30, 2015

Contents

4. Materials

- Python programming lecture of mine
- CS229 Python & NumPy

Chapter PYTHON

Dr. Minhhuy Le
606-A4, EEE, Phenikaa Uni.
huy.leminh@phenikaa-uni.edu.vn

Outline

Chapter 4

1. Python Programming Development
2. Characters, List, Files
3. Loops structures and Booleans, compare to MATLAB
4. Function and Class
5. Plots

CS229 Python & Numpy

Jingbo Yang

How is python related to with others?

Python 2.0 released in 2000

(Python 2.7 "end-of-life" in 2020)

Python 3.0 released in 2008

(Python 3.6+ for CS 229)

Can run interpreted, like MATLAB

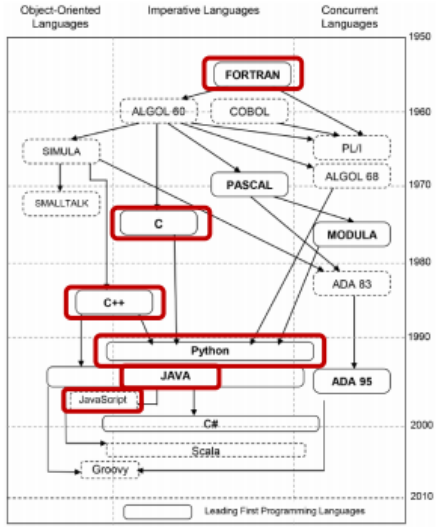

4. Materials

- **Google Colab: Free GPU & CPU** https://colab.research.google.com/

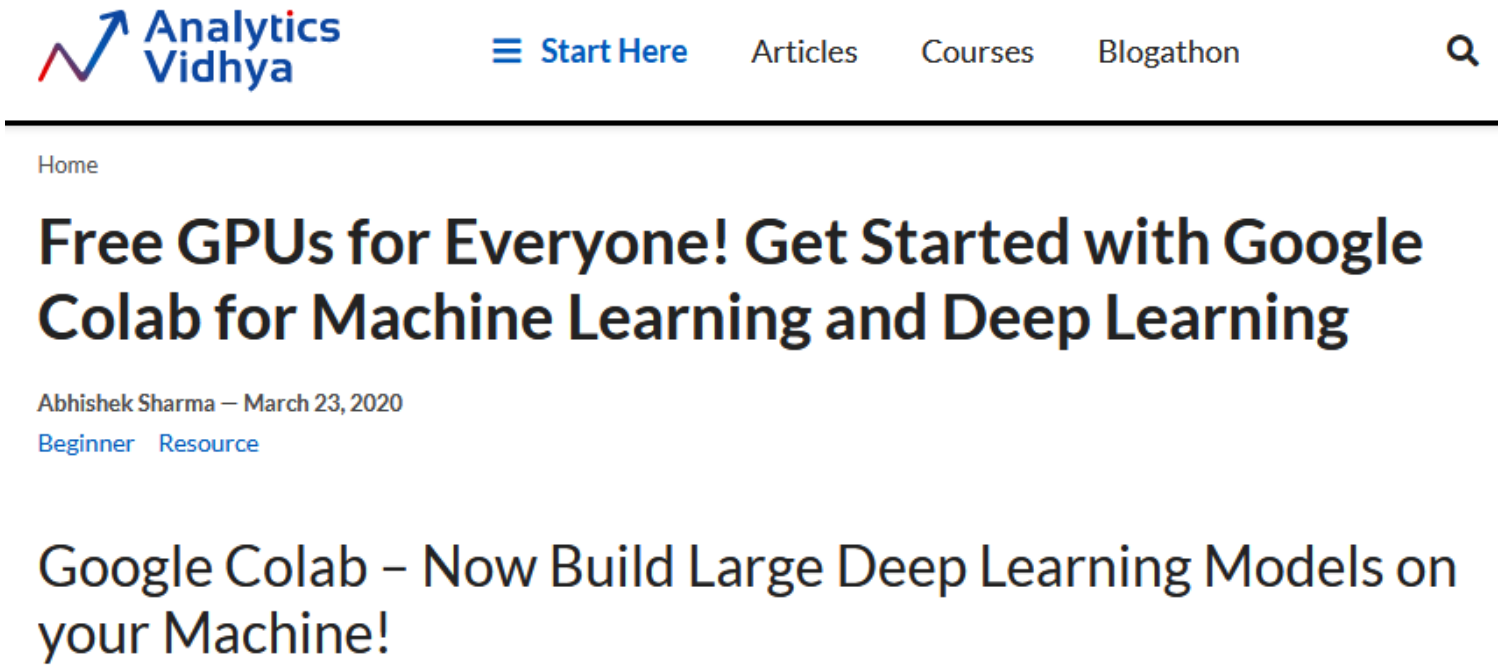

https://www.analyticsvidhya.com/blog/2020/03/google-colab-machine-learning-deep-learning/#1

5. Conclusions

- Remind: Python programming & Maths
- More self-study is required
- Coding from scratch
- Project based exam