Practical Machine Learning

Workshop 1

Introduction to Machine Learning

Dr. Suyong Eum

OSAKA UNIVERSITY

Program for five workshops

Contents of the workshop	Time	Location
❑ Introduction to Machine Learning - Overview of Machine Learning Topics + tools - Discussion on the direction of the workshops during the semester	10am – 11:30am October 12 : Friday	C棟4F C401: Graduate School of Information Science and Technology
❑ Support Vector Machine (SVM) with Principal Component Analysis (PCA) - Lecture on SVM and PCA - Hand-on experience: data dimension reduction and classification	10am – 12:00am October 23 : Tuesday	
❑ Convolutional Neural Networks (CNN) - Lecture on Neural Network and CNN - Hand-on experience: Style transfer or Tiny ImageNet Challenge	10am – 12:00am November 6 : Tuesday	
❑ Recurrent Neural Network (RNN) - Lecture on RNN + LSTM + Seq-to-Seq and Attention mechanism - Hand-on experience: Character level language model + TACOTRON ?	10am – 12:00am November 20 :Tuesday	
❑ Reinforcement Learning (RL) - Lecture on RL + DQN + PG - Hand-on experience: CartPole game using OpenAI Gym	10am – 12:00am December 4 : Tuesday	

Website

- www.suyongeum.com/MLWS

- Materials will be available from the website.
- You can leave messages or questions and so they can be shared by all people.
- Private questions can be sent to me directly.
  - suyong@ist.osaka-u.ac.jp (Dr. Suyong Eum)
  - h-yang@ist.osaka-u.ac.jp (Dr. Hua Yang)

Pre-requisite for the workshops

- ❑ Good (?) Knowledge of Python
  - All examples will be shown in python.
  - Does anyone need python tutorials?

- ❑ Tutorials ….on demand
  - Python
  - Tensorflow
  - Google Cloud (virtual machine setup $300 free account for one year !)
  - Web programming (flask …. Django ….)

- ❑ A little bit of mathematics
  - Try to avoid as much as I can during the workshops

1st Workshop Outline

- ❑ Machine learning and its short history
- ❑ A typical process in the operation of machine learning algorithms with an example
- ❑ What you can do after this course
- ❑ Some tools that you need to prepare until next workshop

Machine Learning

A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P, if its performance at tasks in T, as measured by P, improves with experience E.

Machine Learning – Tom M. Mitchell, 1997

- Learning is a process to understand an underlying process through a set of observations.

vs

recognition

creation

action

History of Machine Learning

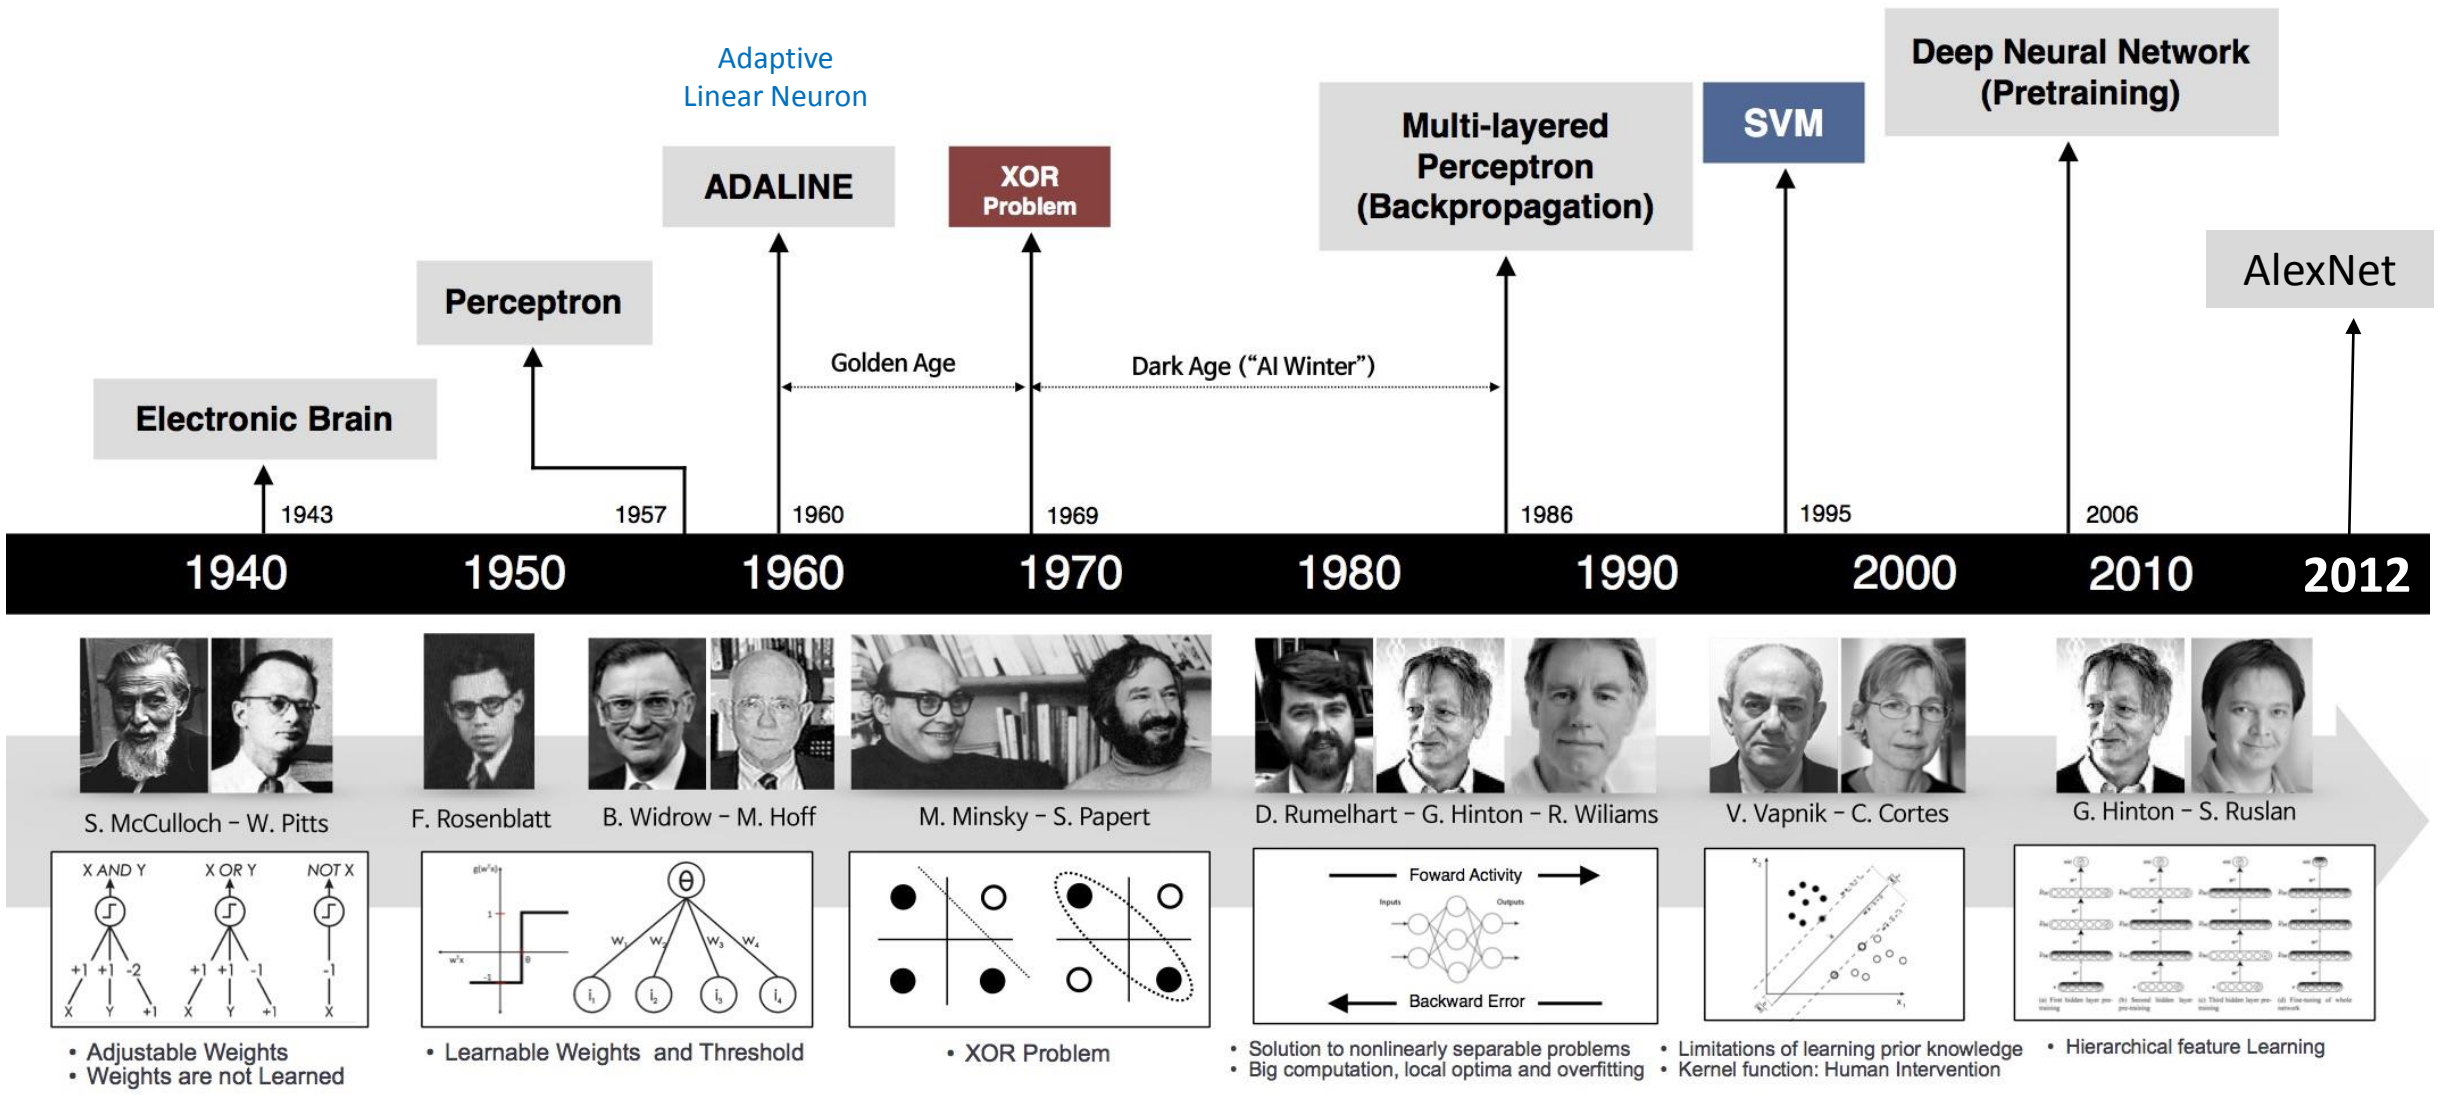

https://beamandrew.github.io/deeplearning/2017/02/23/deep_learning_101_part1.html

The perfect storm

Data

Computational power

Algorithms

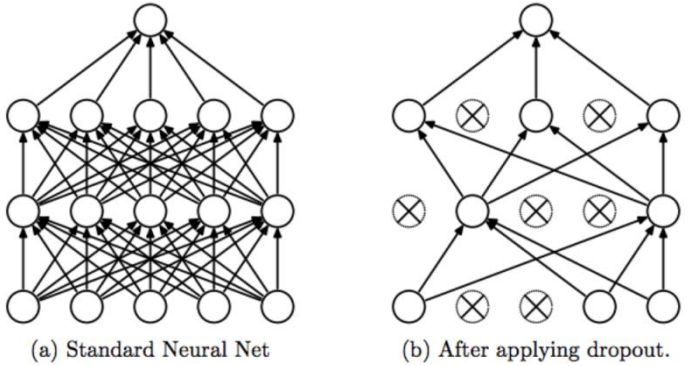

dropout

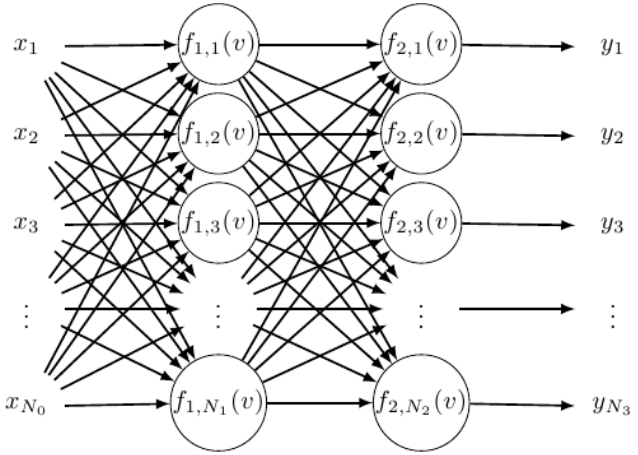

Backpropagation

https://beamandrew.github.io/deeplearning/2017/02/23/deep_learning_101_part1.html

Data

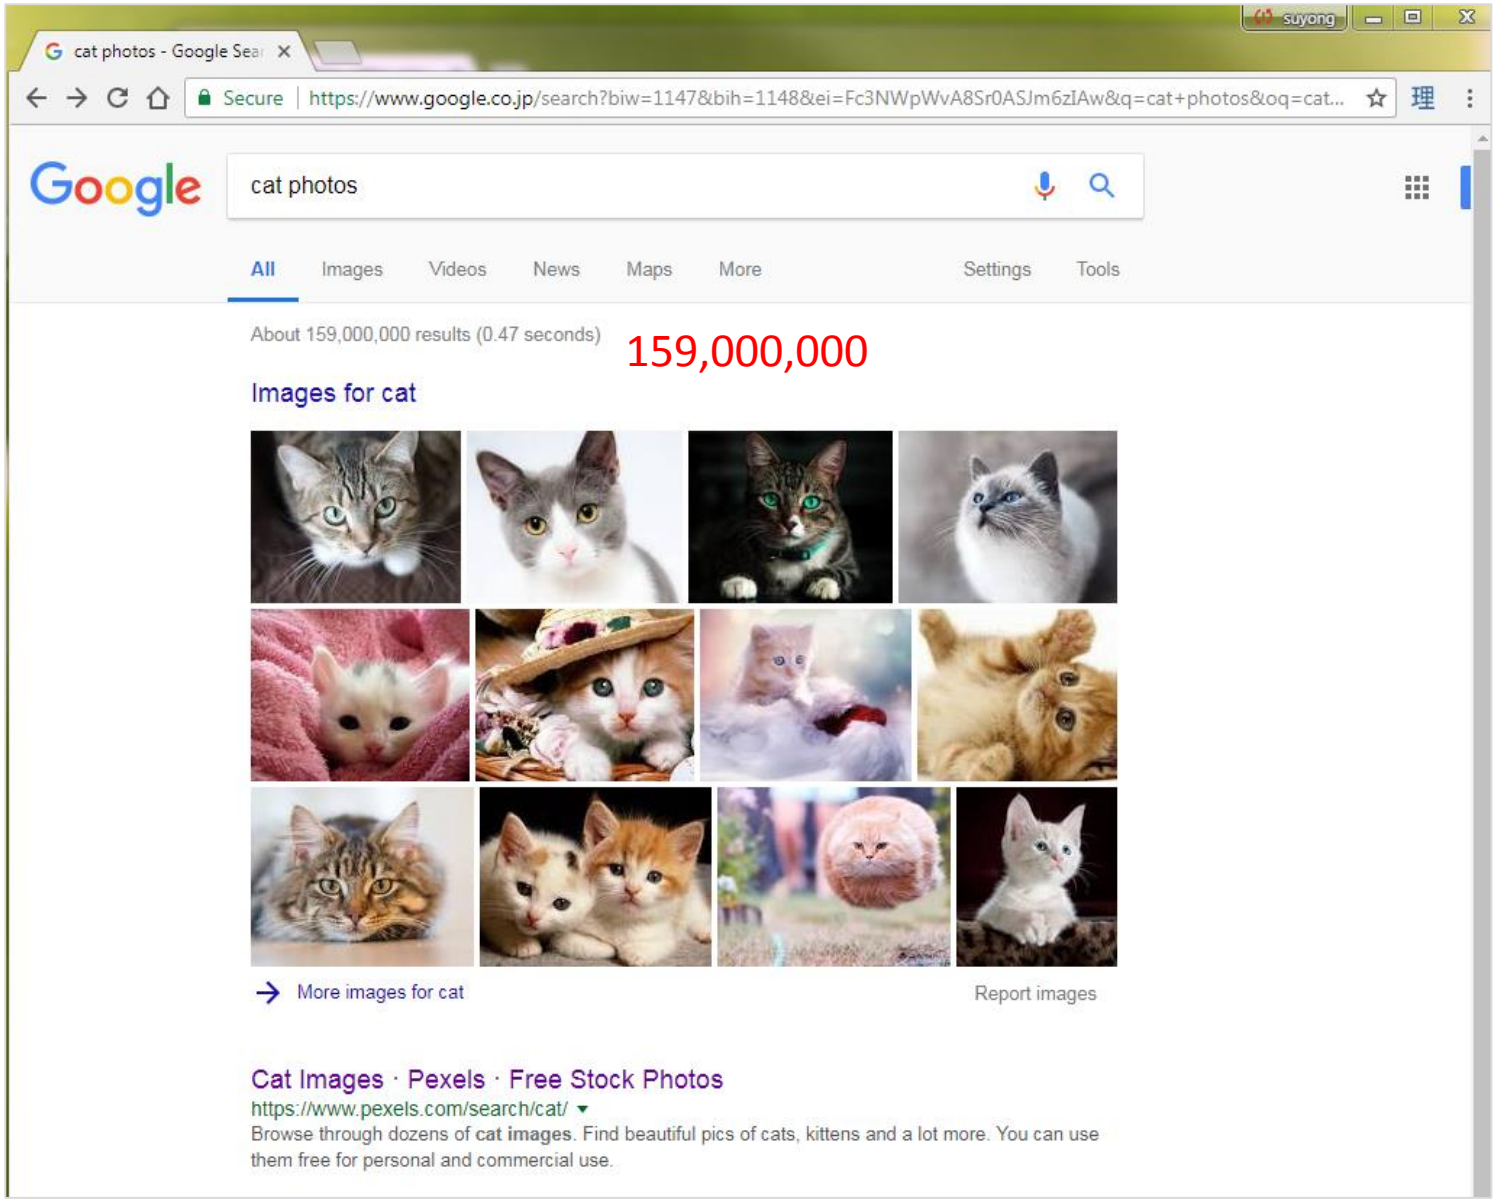
cat photos - Google Sea
Secure
https://www.google.co.jp/search?biw=1147&bih=1148&ei=Fc3NWpWvA8Sr0ASJm6zIAw&q=cat+photos&oq=cat...
Google
cat photos
All
Images
Videos
News
Maps
More
Settings
Tools
About 159,000,000 results (0.47 seconds)
159,000,000
Images for cat
More images for cat
Report images
Cat Images · Pexels · Free Stock Photos
https://www.pexels.com/search/cat/
Browse through dozens of cat images. Find beautiful pics of cats, kittens and a lot more. You can use them free for personal and commercial use.

CPU vs GPU demo

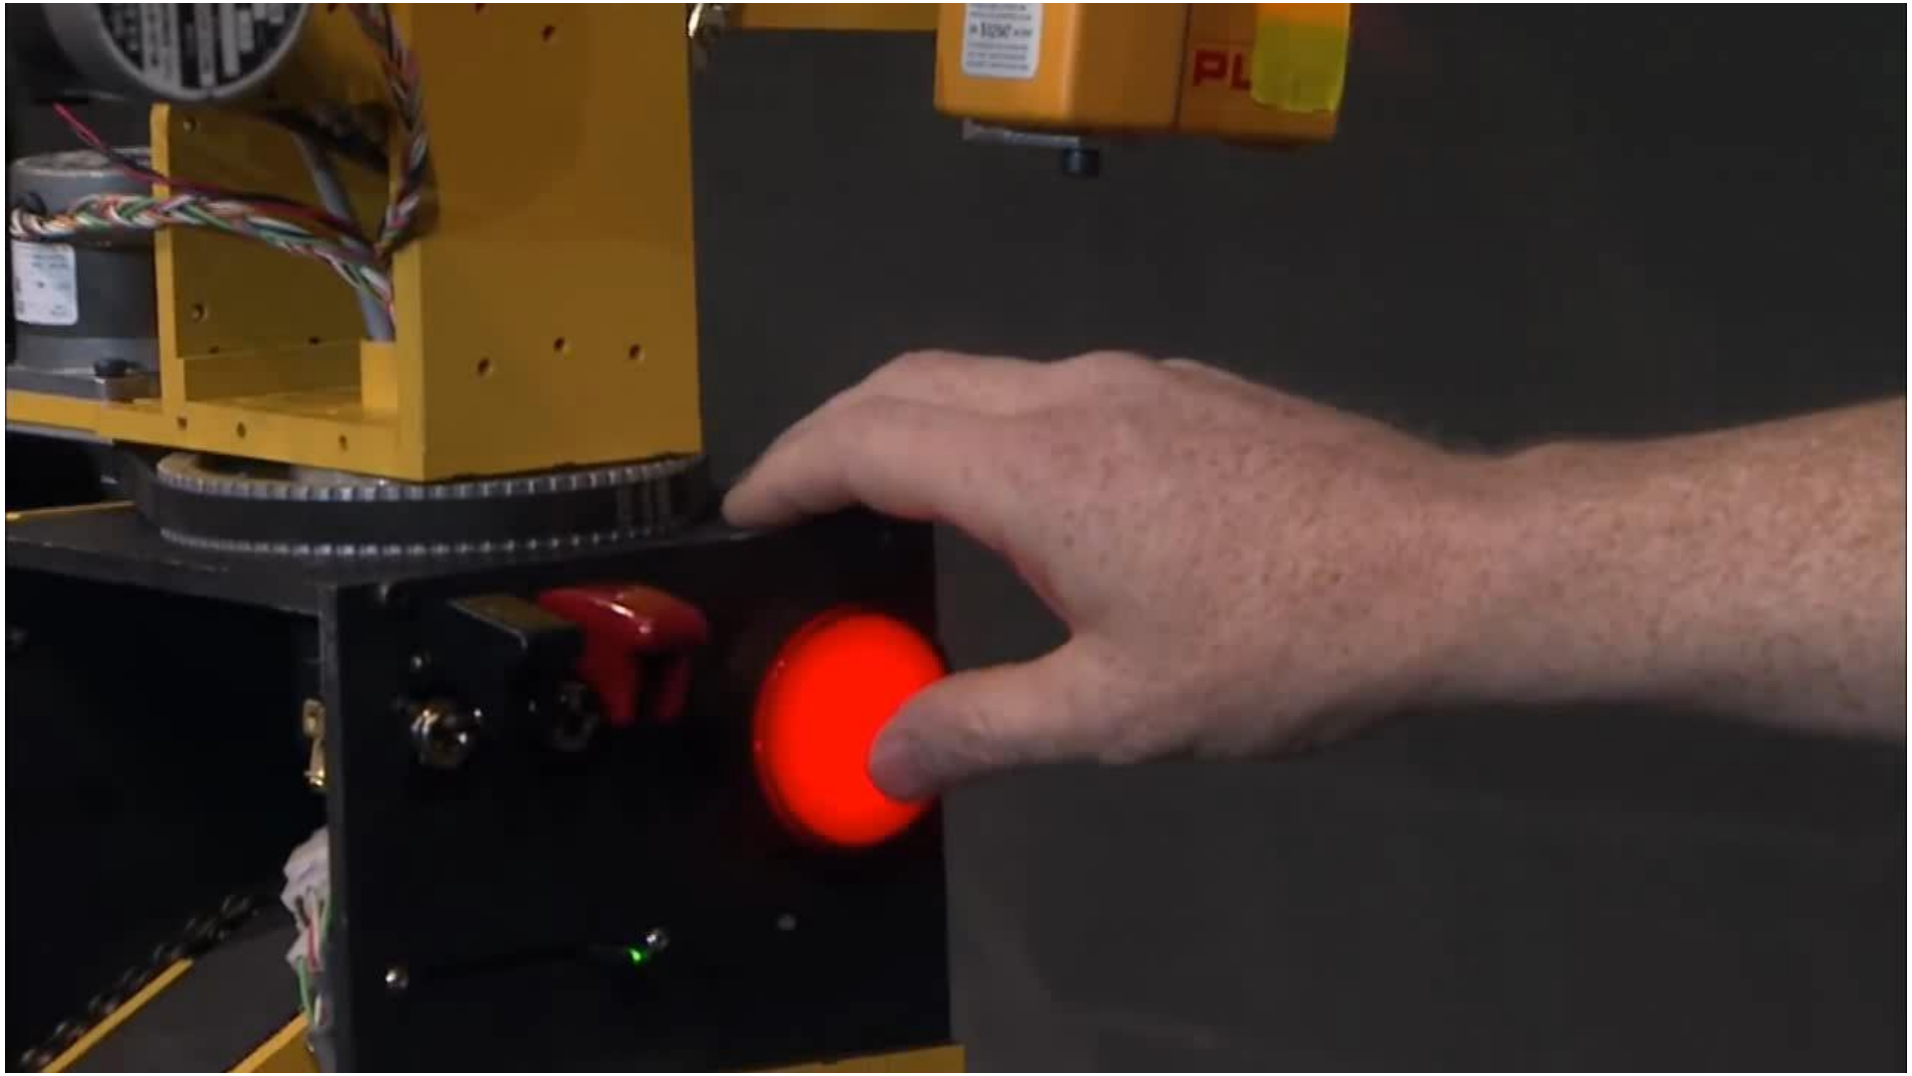

https://www.youtube.com/watch?v=-P28LKWTzrI

Types of machine learning algorithms

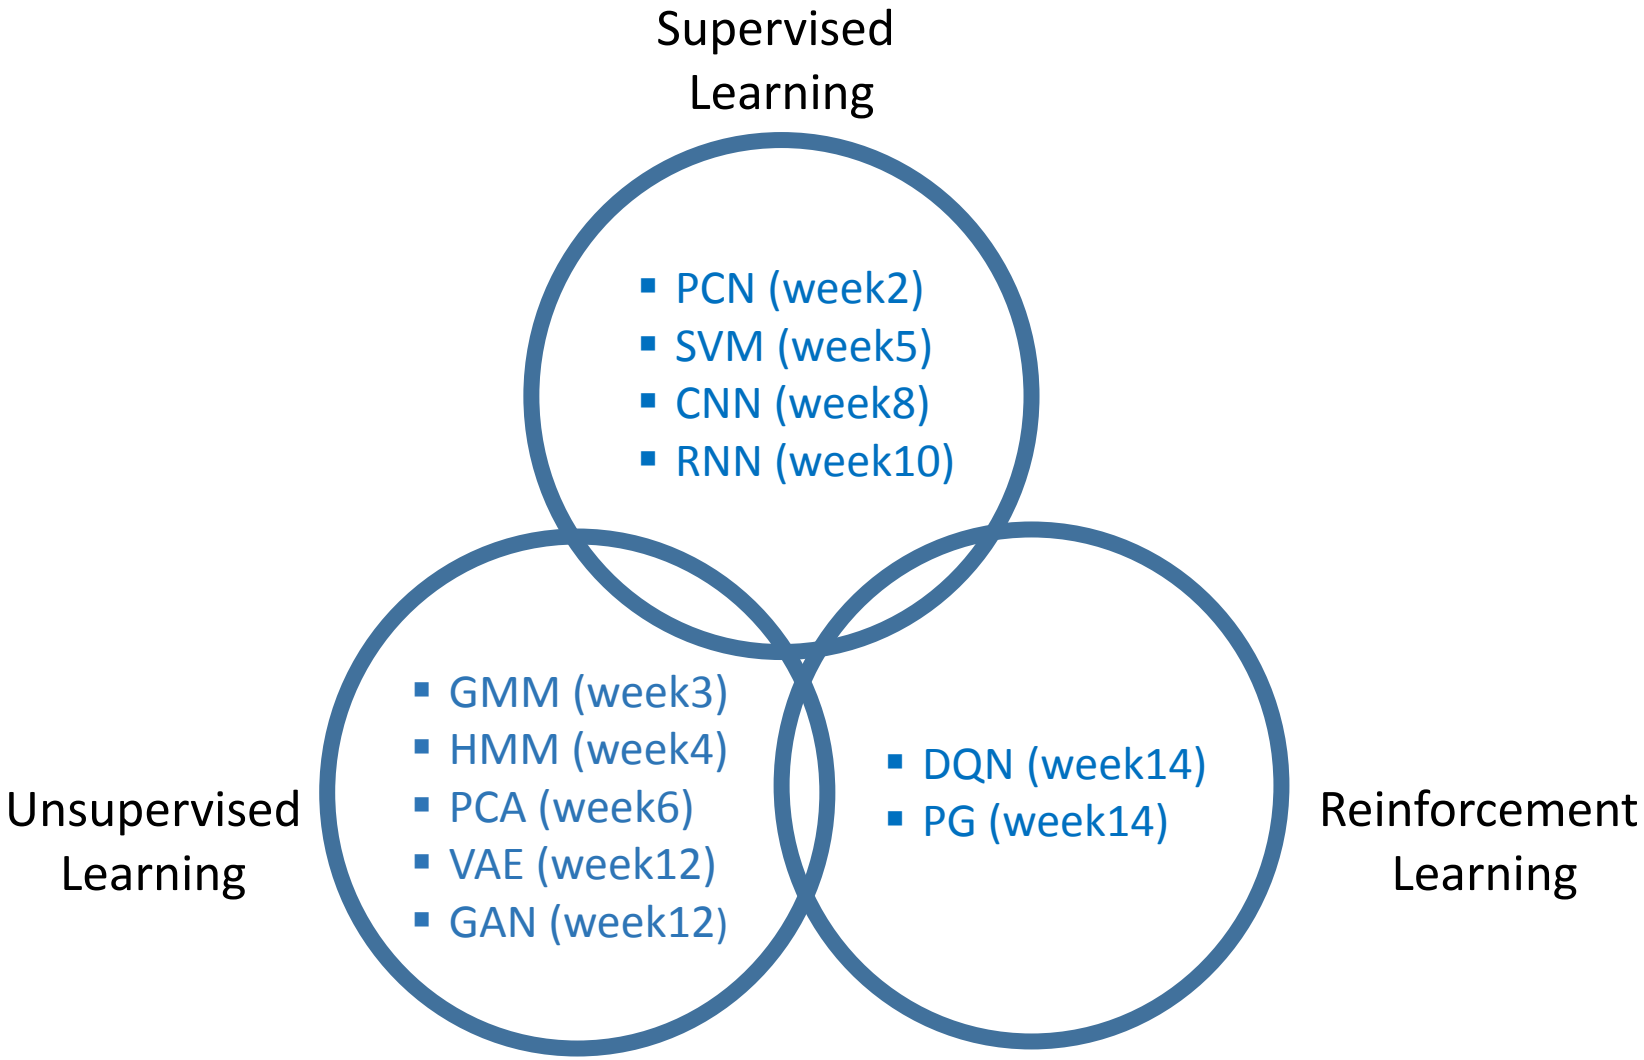

Supervised Learning

- Input + Output with Label
- Supervised learning is learning from by a knowledgeable external supervisor.

  - *Question -> Answer*

y = f(x)

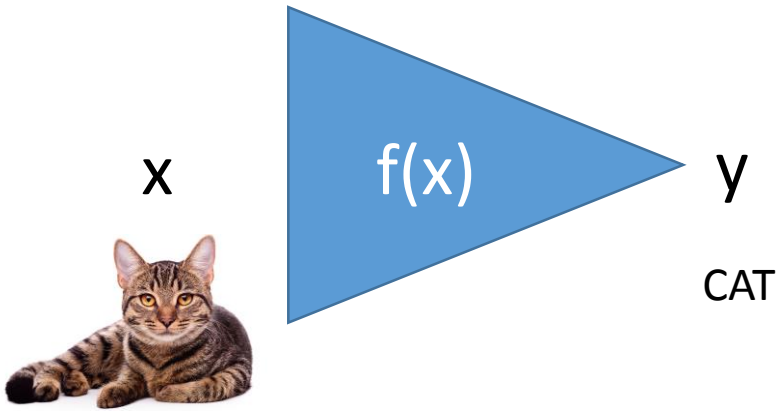

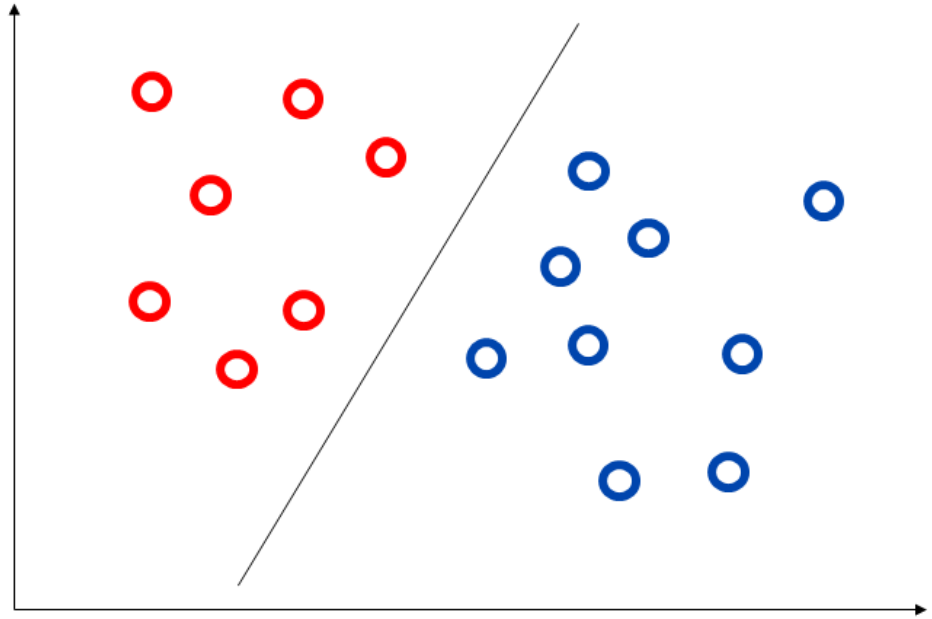

Unsupervised Learning

- ❑ Input + Output without Label
- ❑ Feature Learning

- *Question -> Question*

x = f(x)

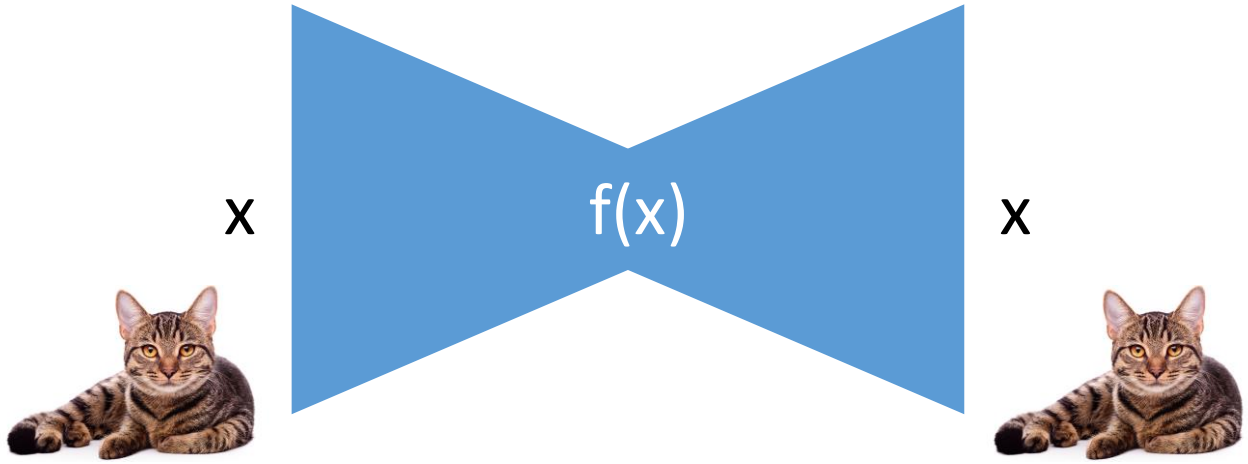

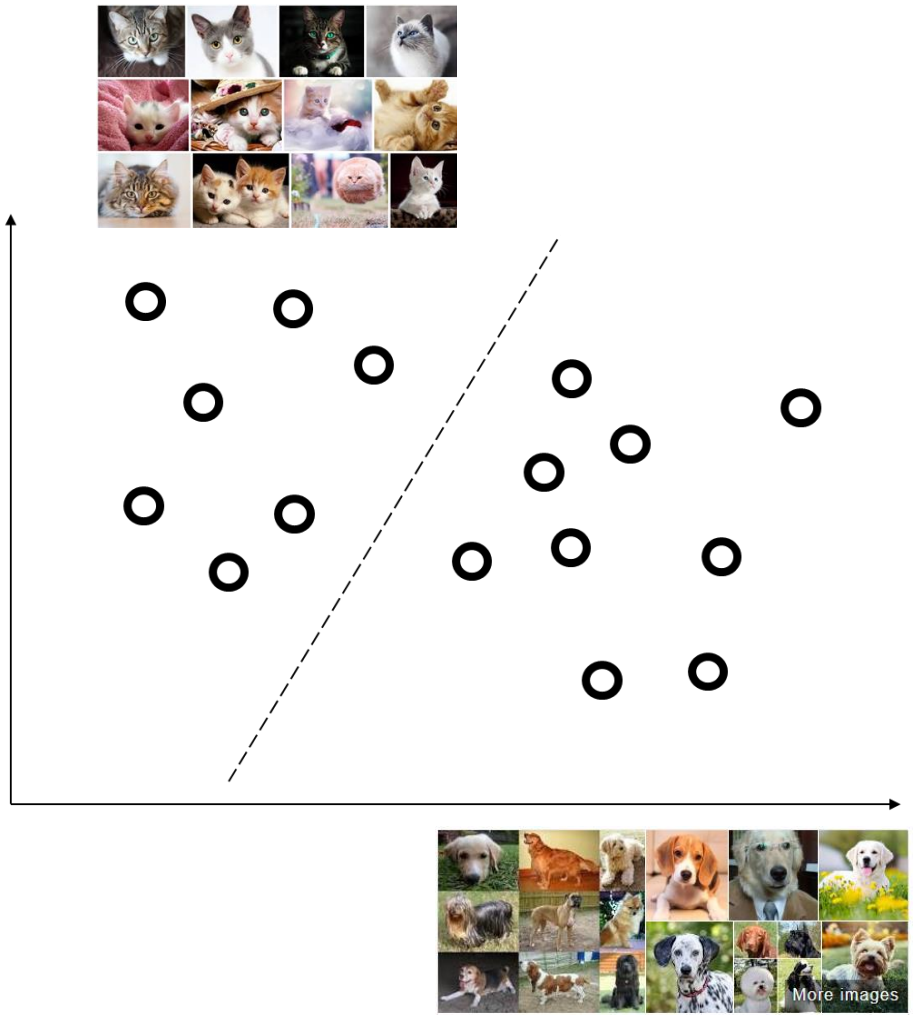

Reinforcement Learning

- Input + partial output with its quality: in some sense similar to supervised learning
- An action is rewarded/penalized to take a better action next time

- *Carrot and stick*

y = f(x)

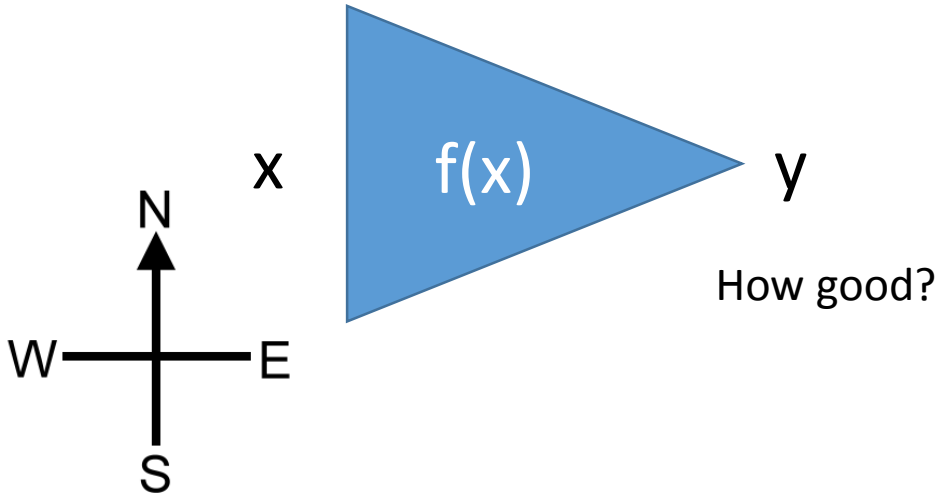

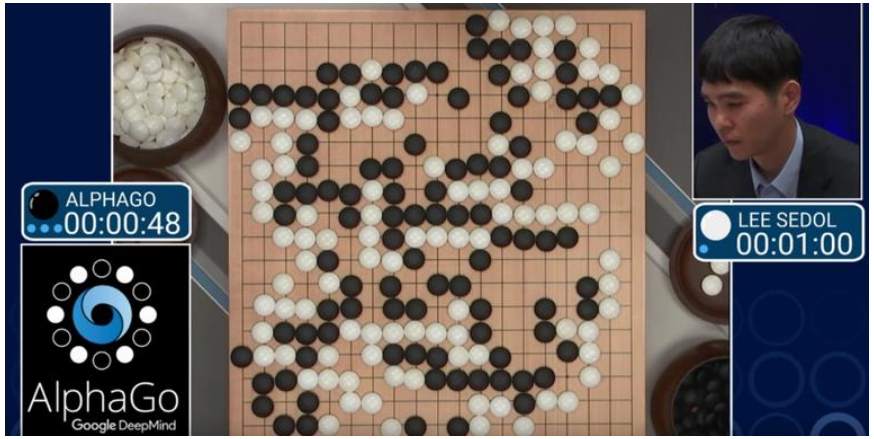

A typical process in the operation of machine learning algorithms with an example

Components of machine learning

- Raw data (including label)
- Input (features: dimension of a data point)
- Hypothesis (a function approximating a target function)
- Output
- Label

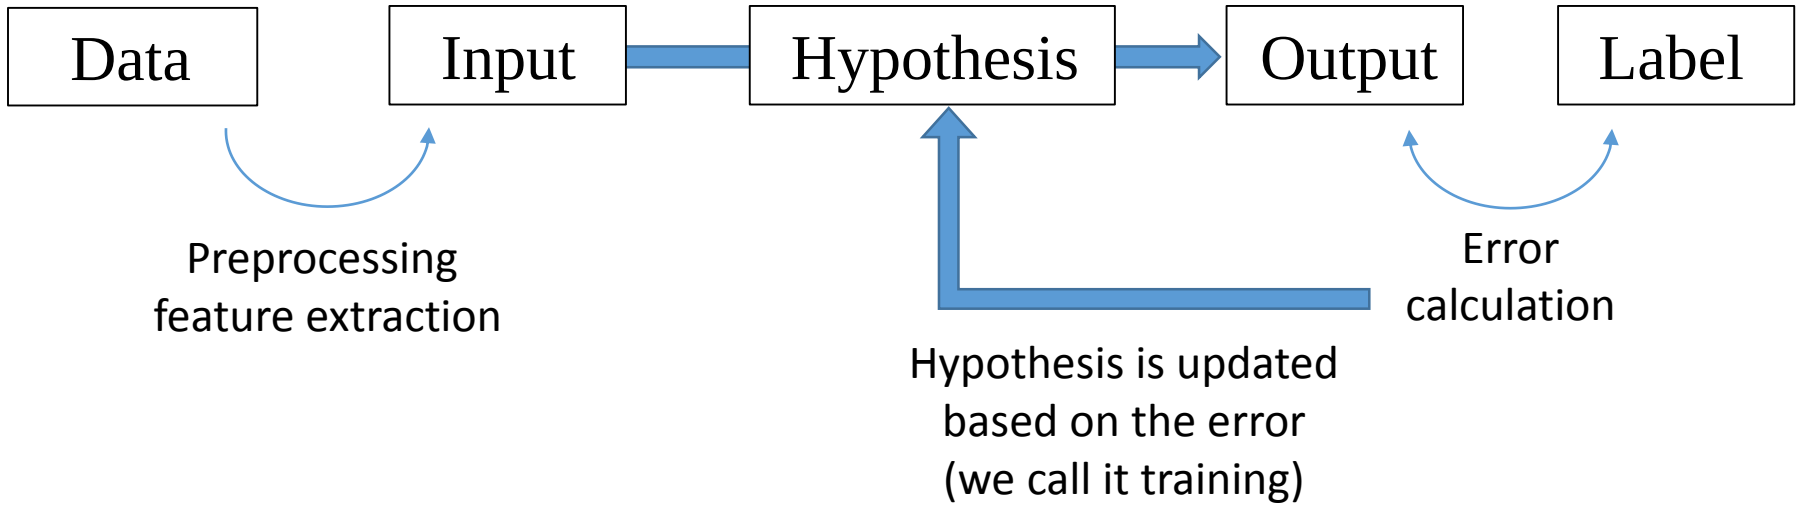

An example: male / female classification

❑ Given data per person

- Height: 170cm
- Weight: 52kg
- Foot size: 25cm
- Hand size: 20cm
- Nose height: 1.5cm
- Eye size: 2.5cm
- Hair length: 5cm

- *Male*

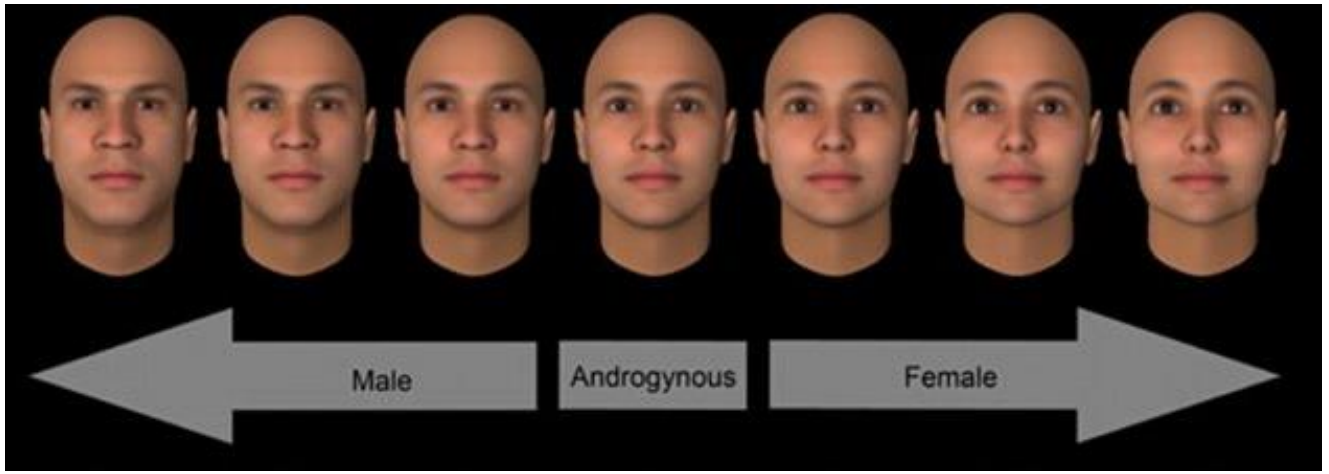

Components	Notation	Description
Data	$(x_1, y_1), (x_2, y_2), \ldots, (x_n, y_n)$	(a vector data, male/female)
Input	X	$x_n=\{x_1,x_2, \ldots, x_d\}$: d dimensions
Output	Y	Output data from hypothesis
Hypothesis	g: X → Y	(hypothesis) A model
Target function	f: X → Y	Unknown

An example: feature extraction

- Which features are important to tell that the given object is male or female?
- Assuming you chose two features and then you plot the data points

- Given data per person
  - Height: 170cm
  - Weight: 52kg
  - Foot size: 25cm
  - Hand size: 20cm
  - Nose height: 1.5cm
  - Eye size: 2.5cm
  - Hair length: 5cm
  - *Male*

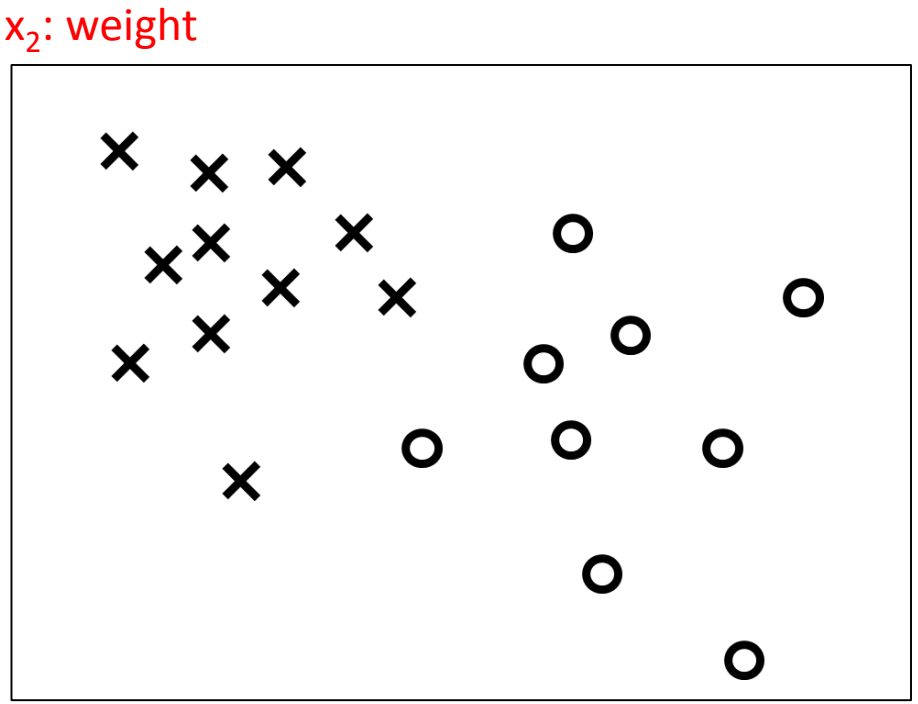

An example: hypothesis (model) selection

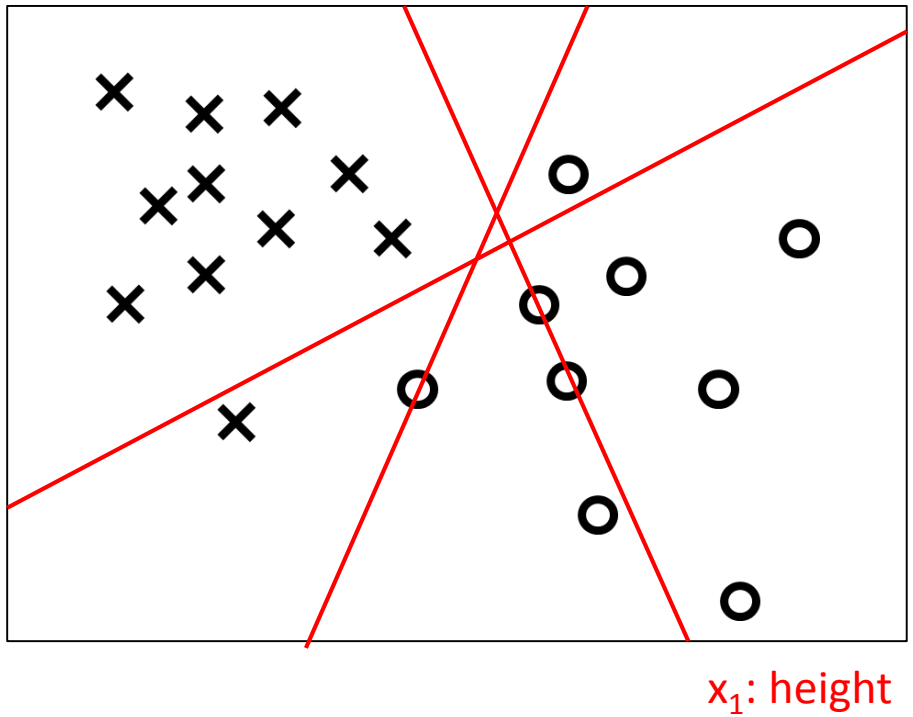

$$ax_1 + bx_2 + c = 0$$

An example: training the hypothesis to produce less error

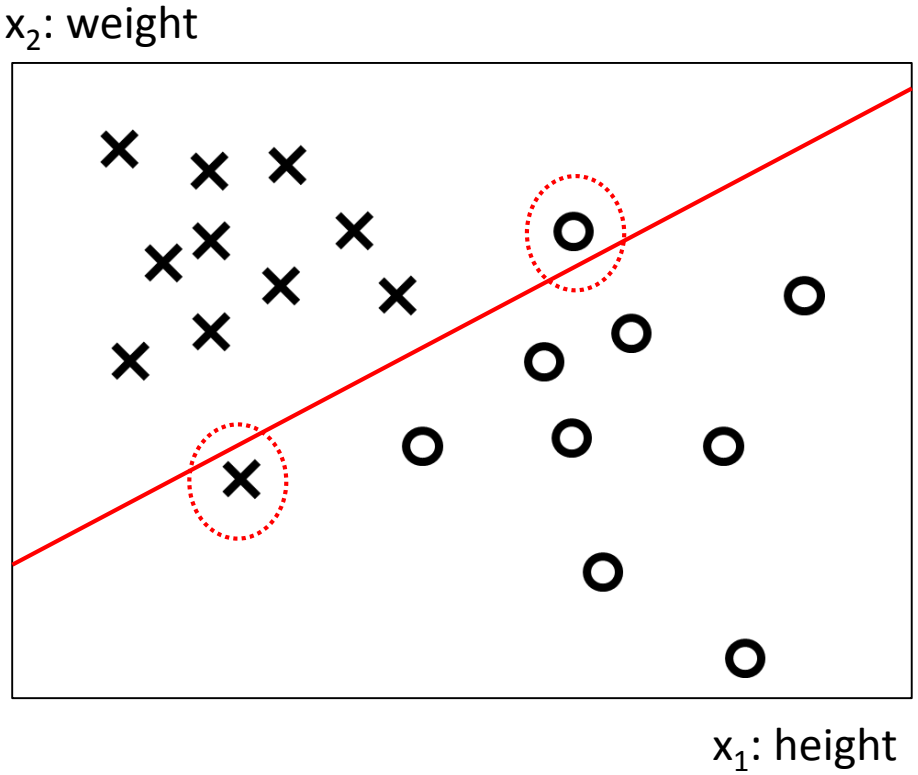

- Random selection of *W*
- Misclassified data points are found
- Update *W* in order to correctly classify the misclassified data points.
  - How? : depending on learning algorithm
    - Neural network: backpropagation?
    - Linear algebra: perceptron algorithm

An example: training the hypothesis to produce less error

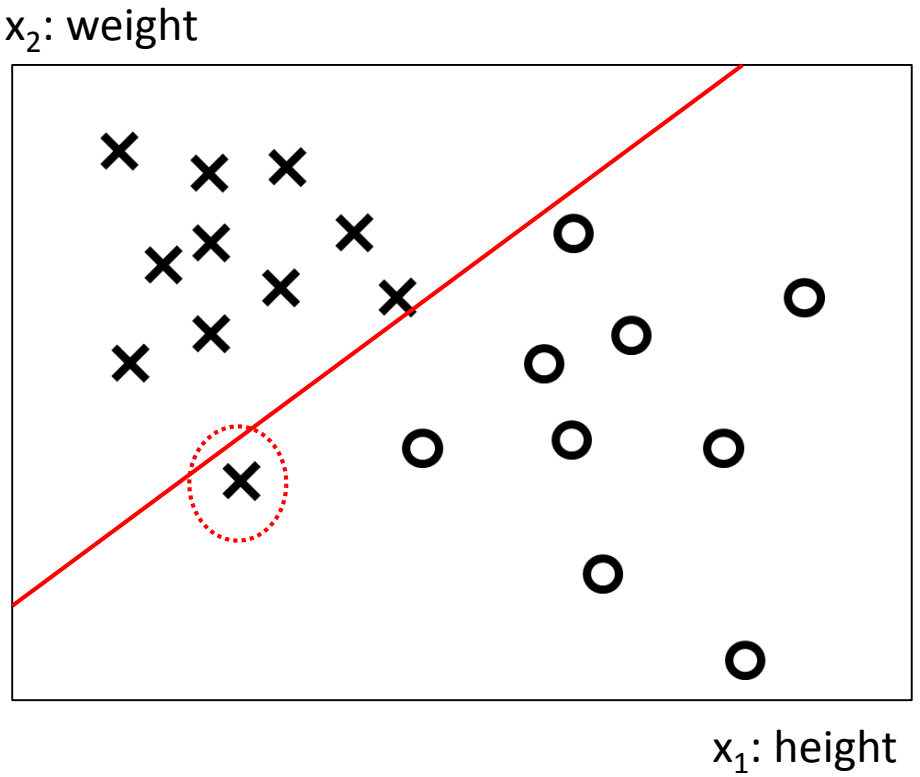

- Random selection of *W*
- Misclassified data points are found
- Update *W* in order to correctly classify the misclassified data points.
  - How? : depending on learning algorithm
    - Neural network: backpropagation?
    - Linear algebra: perceptron algorithm

An example: training the hypothesis to produce less error

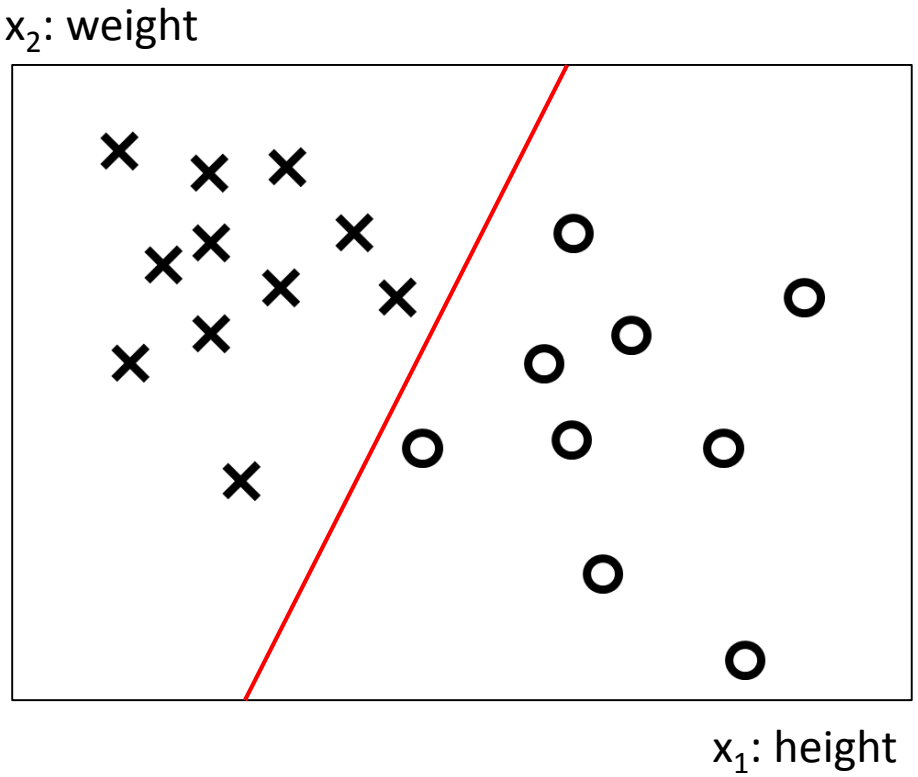

- Random selection of *W*
- Misclassified data points are found
- Update *W* in order to correctly classify the misclassified data points.
  - How? : depending on learning algorithm
    - Neural network: backpropagation?
    - Linear algebra: perceptron algorithm

Cross validation (CV)

- You are generally given one big training data set.
- How to verify goodness of your model?

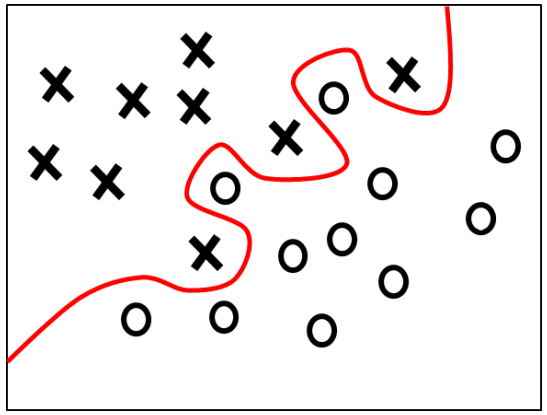

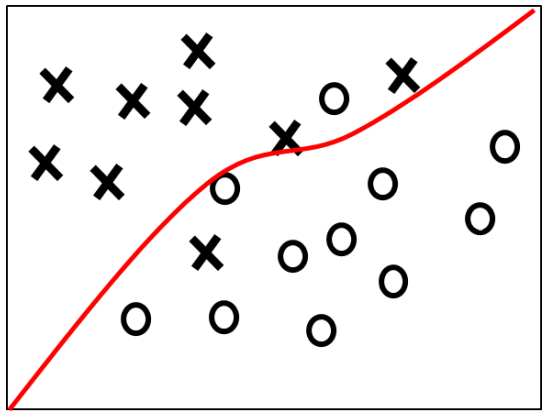

Occam's razor

$$y(x,w) = w_0 + w_1 x + w_2 x^2 + \ldots + w_M x^M = \sum_{j=0}^{M} w_j x^j$$

What you can do after this course

Image classification

Large Scale Visual Recognition Challenge (ILSVRC)
- 1000 class objects
- around 1.4 million images

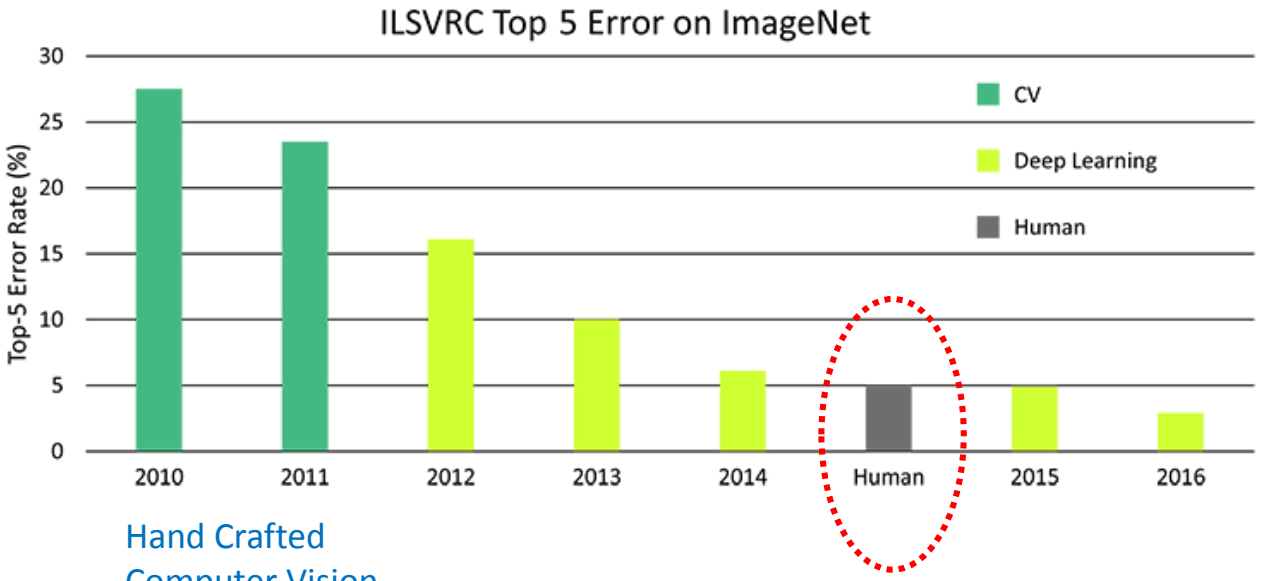

Hand Crafted Computer Vision approaches

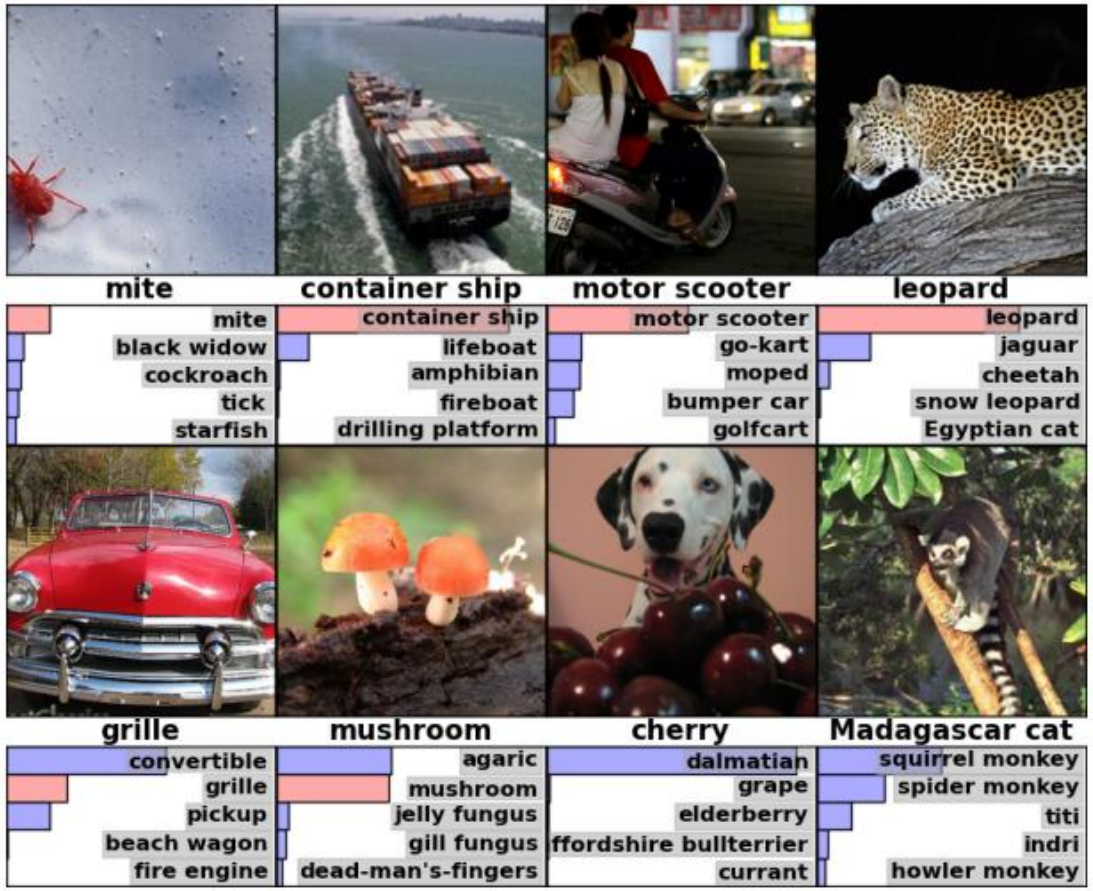

https://www.dsiac.org/resources/journals/dsiac/winter-2017-volume-4-number-1/real-time-situ-intelligent-video-analytics
http://www.cs.toronto.edu/~fritz/absps/imagenet.pdf

Image detection

2016

Mean Average Precision

Real-Time Detectors	Train	mAP
100Hz DPM [31]	2007	16.0
30Hz DPM [31]	2007	26.1
Fast YOLO	2007+2012	52.7
YOLO	2007+2012	**63.4**
Less Than Real-Time		
Fastest DPM [38]	2007	30.4
R-CNN Minus R [20]	2007	53.5
Fast R-CNN [14]	2007+2012	70.0
Faster R-CNN VGG-16[28]	2007+2012	73.2
Faster R-CNN ZF [28]	2007+2012	62.1
YOLO VGG-16	2007+2012	66.4

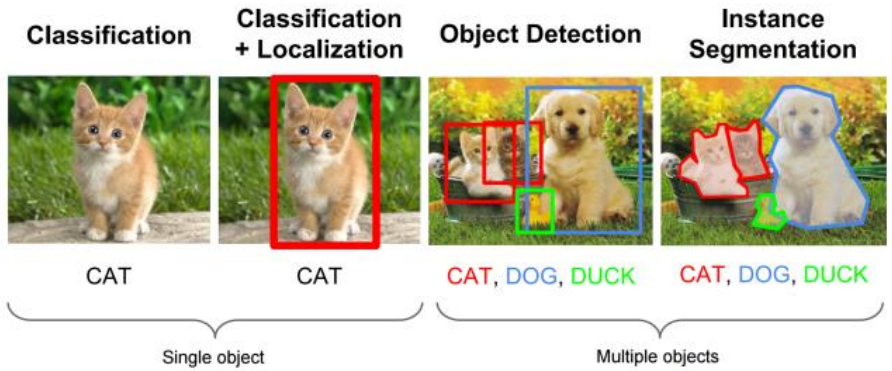

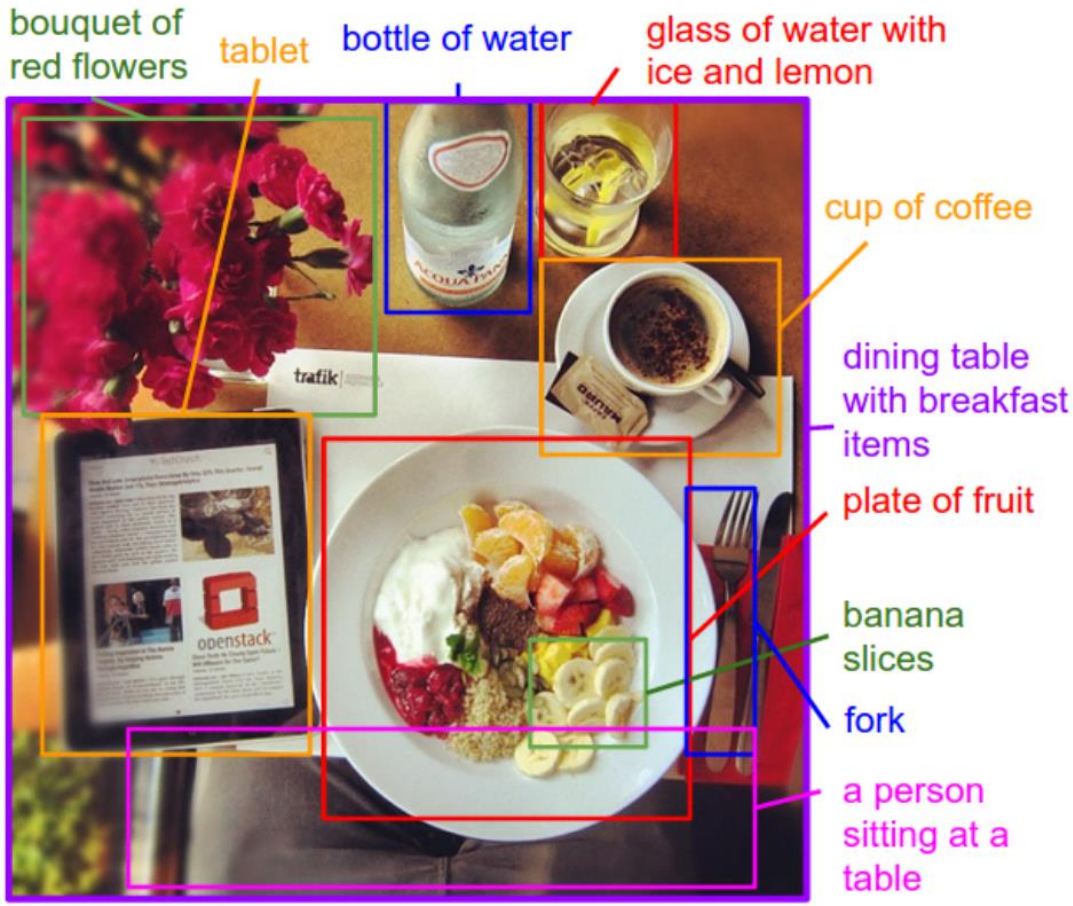

- Classification
- Localization
- Detection
- Segmentation

https://arxiv.org/pdf/1506.02640.pdf
https://arxiv.org/pdf/1412.2306v2.pdf

Image style transfer

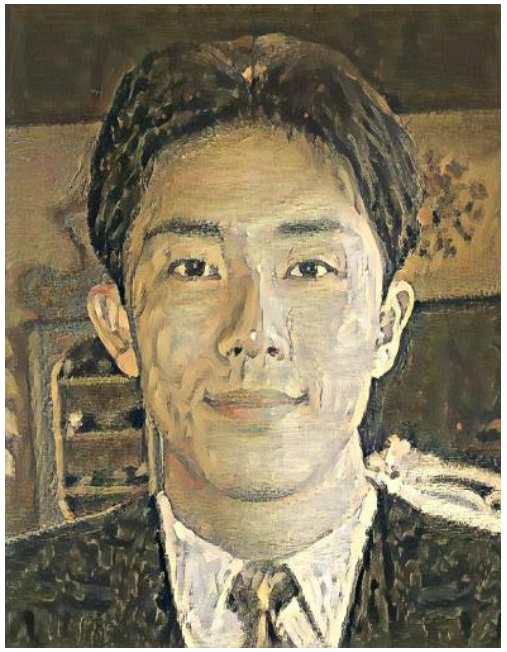

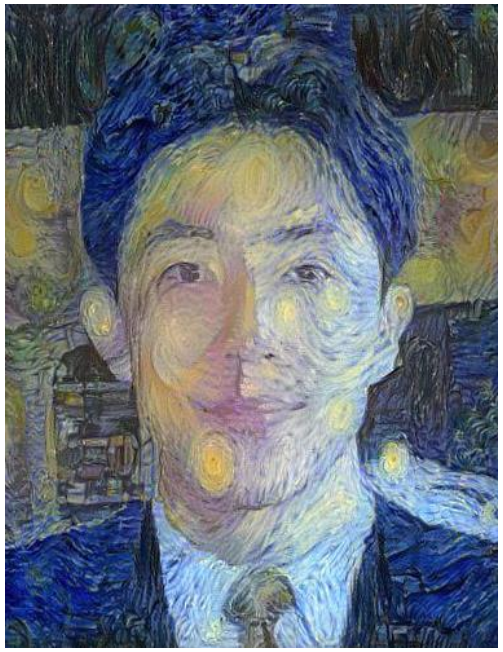

Image caption

a man is riding a motorcycle on a street
logprob: -8.65

a bus is parked on the side of the road
logprob: -7.19

a zebra standing in a field of grass
logprob: -7.88

a woman is standing in front of a store
logprob: -11.40

a woman holding a teddy bear in front of a mirror
logprob: -9.65

a baby laying on a bed with a stuffed bear
logprob: -8.85

https://cs.stanford.edu/people/karpathy/deepimagesent/generationdemo/

Creation: writing

```
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.
```

Writing new episodes of Friends is easy if you use a neural network

"Chandler: Well, I proposed to my shoe..."

By James Vincent | @jjvincent | Jan 21, 2016, 4:03am EST

SHARE

For $\bigoplus_{n=1,\ldots,m}$ where $\mathcal{L}_{m_\bullet} = 0$, hence we can find a closed subset $\mathcal{H}$ in $\mathcal{H}$ and any sets $\mathcal{F}$ on X, U is a closed immersion of S, then $U \to T$ is a separated algebraic space.

Proof. Proof of (1). It also start we get

$$S = \mathrm{Spec}(R) = U \times_X U \times_X U$$

and the compariciely in the fibre product covering we have to prove the lemma generated by $\coprod Z \times_U U \to V$. Consider the maps M along the set of points Sch_{fppf} and $U \to U$ is the fibre category of S in U in Section, ?? and the fact that any U affine, see Morphisms, Lemma ??. Hence we obtain a scheme S and any open subset $W \subset U$ in $Sh(G)$ such that $\mathrm{Spec}(R') \to S$ is smooth or an

$$U = \bigcup U_i \times_{S_i} U_i$$

which has a nonzero morphism we may assume that f_i is of finite presentation over S. We claim that $\mathcal{O}_{X,x}$ is a scheme where $x, x', s'' \in S'$ such that $\mathcal{O}_{X,x'} \to \mathcal{O}'_{X',x'}$ is separated. By Algebra, Lemma ?? we can define a map of complexes $\mathrm{GL}_{S'}(x'/S'')$ and we win. □

To prove study we see that $\mathcal{F}|_U$ is a covering of $\mathcal{X}'$, and $\mathcal{T}_i$ is an object of $\mathcal{F}_{X/S}$ for $i > 0$ and $\mathcal{F}_p$ exists and let $\mathcal{F}_i$ be a presheaf of $\mathcal{O}_X$-modules on $\mathcal{C}$ as a $\mathcal{F}$-module. In particular $\mathcal{F} = U/\mathcal{F}$ we have to show that

$$\widetilde{M}^\bullet = \mathcal{I}^\bullet \otimes_{\mathrm{Spec}(k)} \mathcal{O}_{S,s} - i_X^{-1}\mathcal{F})$$

is a unique morphism of algebraic stacks. Note that

$$\mathrm{Arrows} = (Sch/S)^{opp}_{fppf}, (Sch/S)_{fppf}$$

and

$$V = \Gamma(S, \mathcal{O}) \longmapsto (U, \mathrm{Spec}(A))$$

is an open subset of X. Thus U is affine. This is a continuous map of X is the inverse, the groupoid scheme S.

Proof. See discussion of sheaves of sets. □

The result for prove any open covering follows from the less of Example ??. It may replace S by $X_{spaces,étale}$ which gives an open subspace of X and T equal to S_{Zar}, see Descent, Lemma ??. Namely, by Lemma ?? we see that R is geometrically regular over S.

http://karpathy.github.io/2015/05/21/rnn-effectiveness/
https://www.theverge.com/2016/1/21/10805398/friends-neural-network-scripts

Creation: music composition

https://magenta.tensorflow.org/performance-rnn

Generation: image generation

Generated by a machine

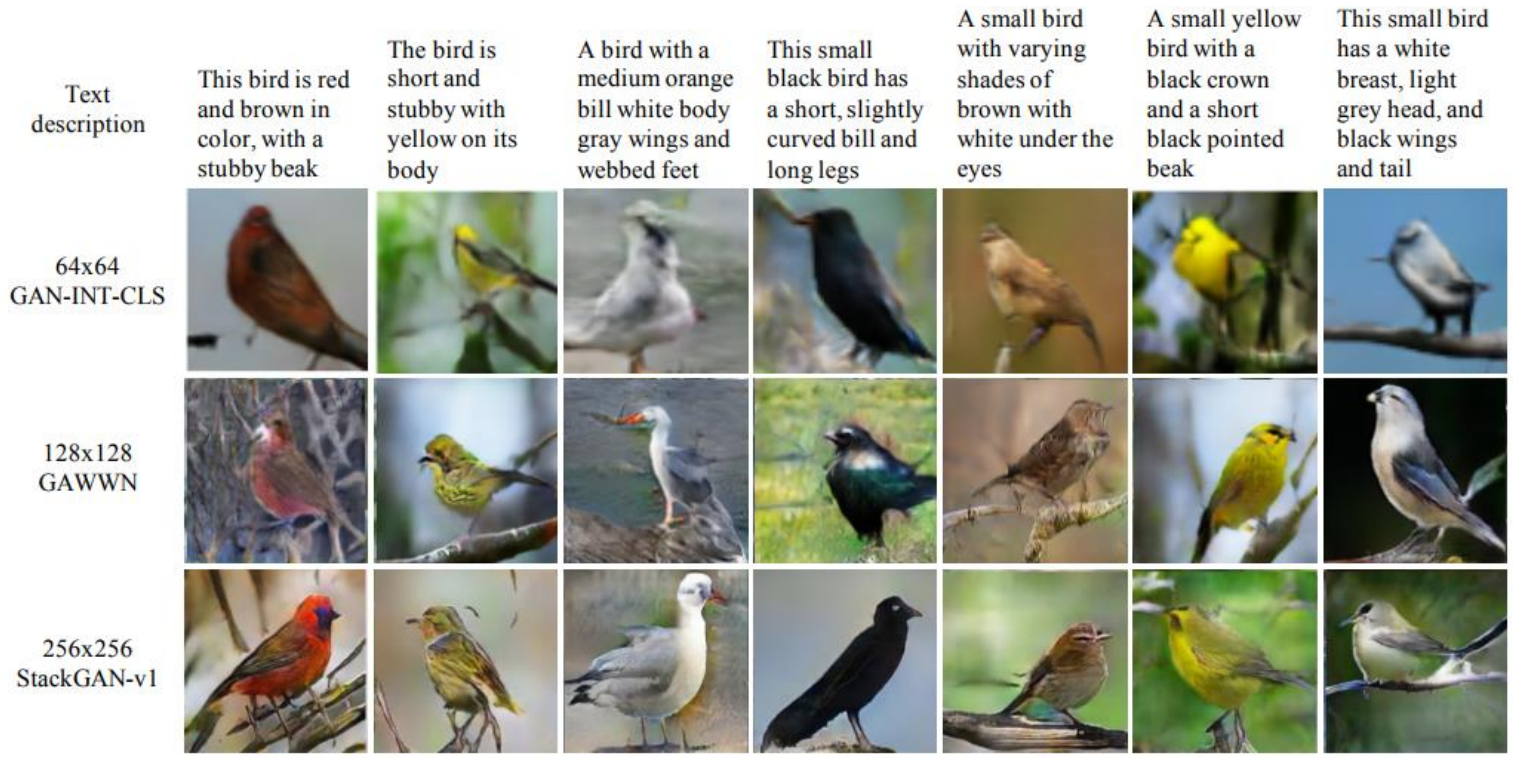

Generated by a machine based on given text

https://arxiv.org/pdf/1703.10717.pdf
https://arxiv.org/pdf/1710.10916.pdf

Playing Atari game

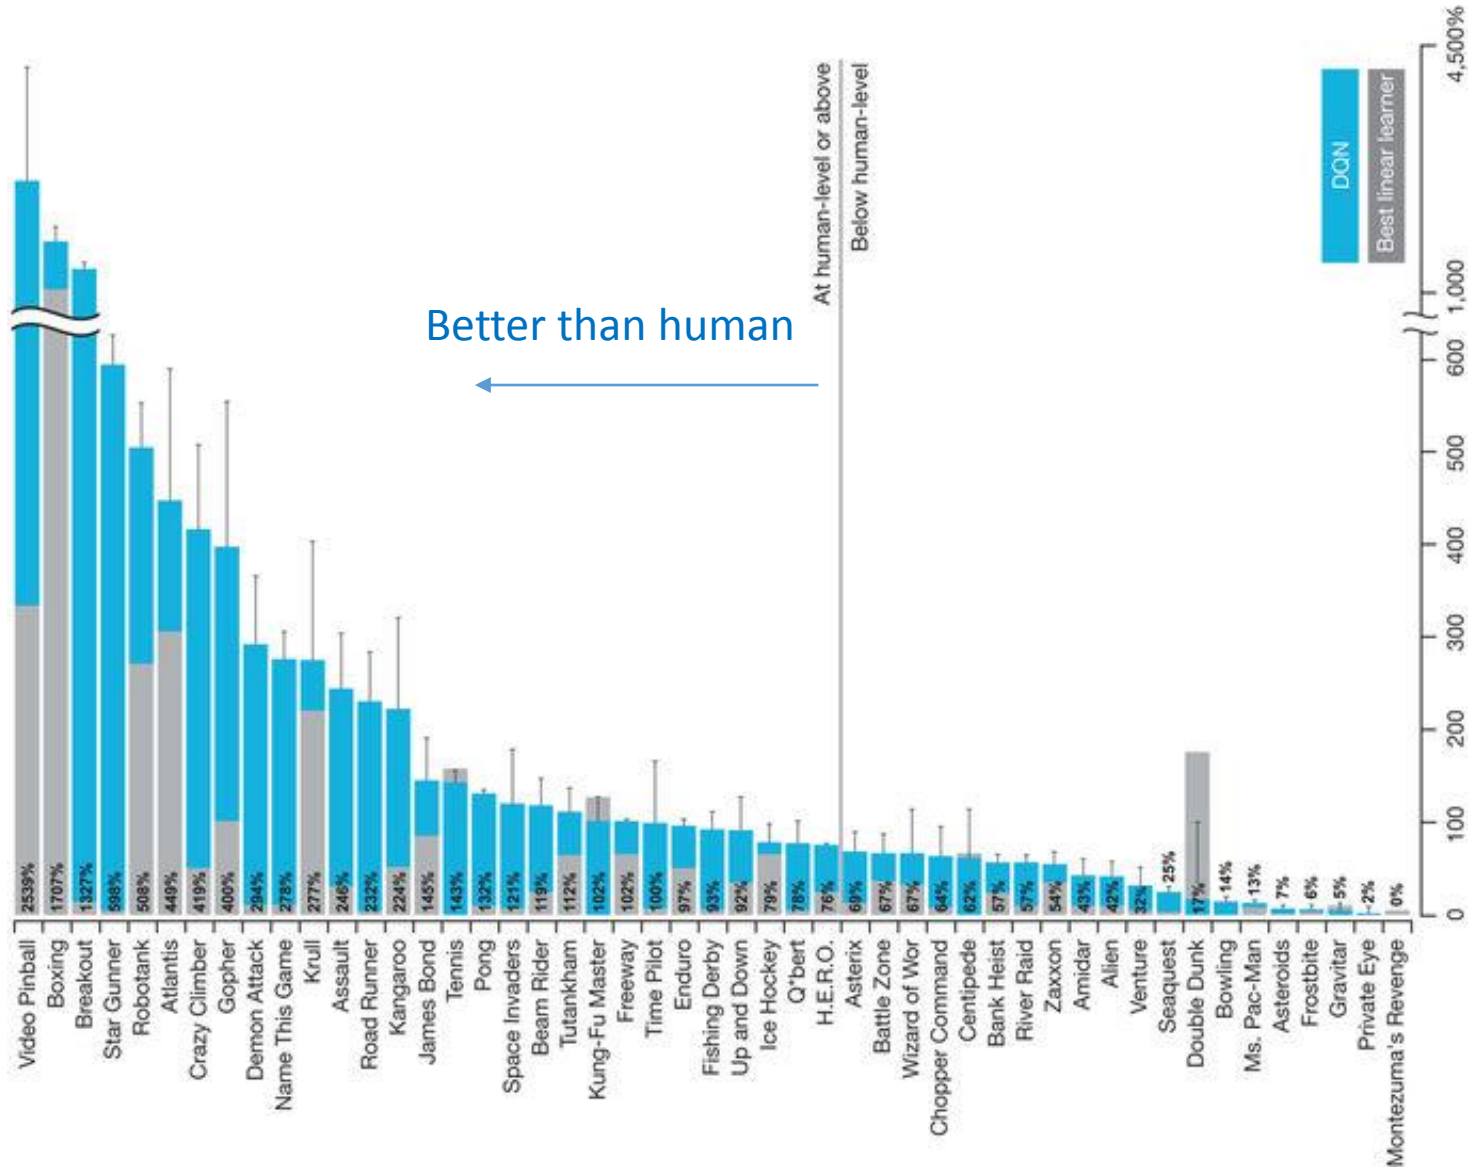

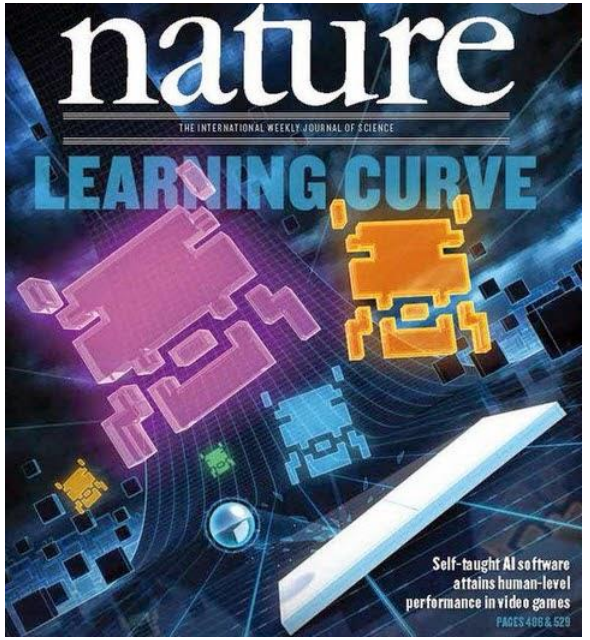

https://www.nature.com/articles/nature14236 nature, 2015

Alpha GO (2015)

- Game Go: 10^{170} state space
- Beat European Champion: October 2015
- Beat World Champion: March 2016

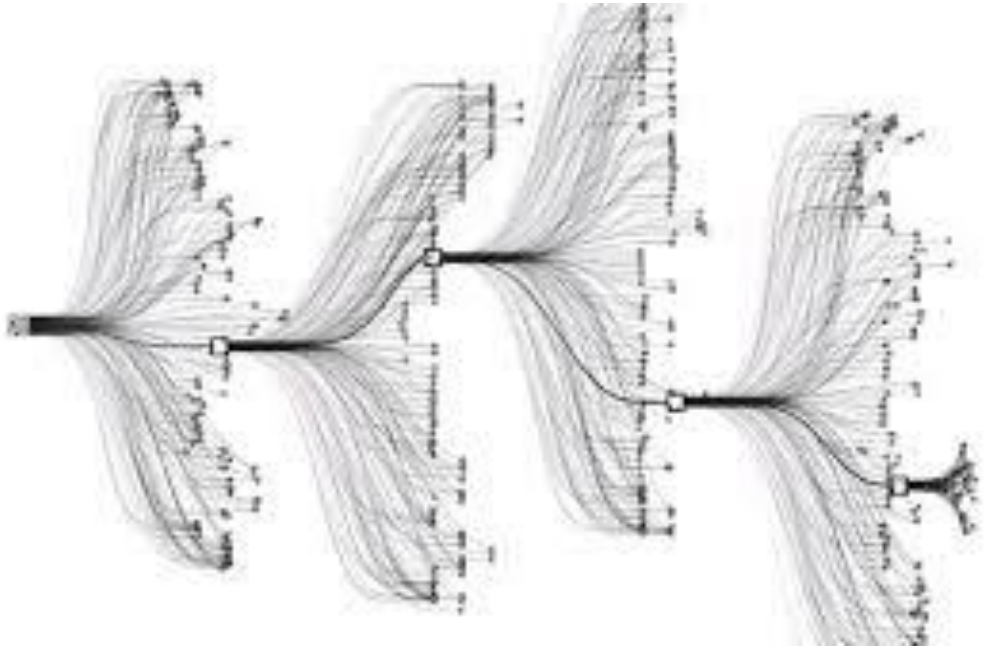

https://deepmind.com/research/alphago/

Automatic Driving

- 1.3 million people die every year in car accidents.
- 94% of those accidents involve human error.
- 70% of the manned Taxis is related to labor cost.

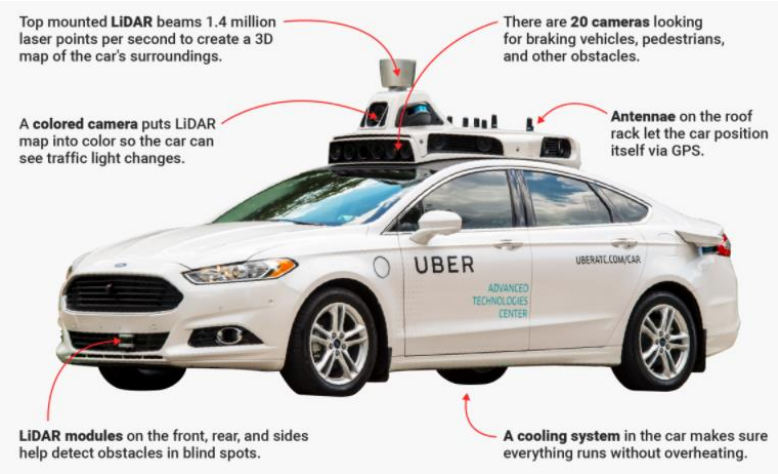

https://newsroom.uber.com/us-pennsylvania/new-wheels/

Tools that we are going to use during the workshops in this semester

Working environment during the workshop

- Online
  - Colab (https://colab.research.google.com/)
    - Need a google account
  - Google Cloud (https://cloud.google.com/)
    - Need a google account
    - $300 free account for one year - credit card information required to use some extra functions.

- Offline
  - Installation of the tools in your notebook
    - Anaconda (https://www.anaconda.com/download/)
      - Python 3.7 version
      - Then, you can install almost other packages with “conda” installation
    - Jupyter Notebook

Backup slides