Introduction to AI and Machine Learning

Sunghee Yun
Senior Applied Scientist

Amazon.com, Inc.

About the speaker

- Sunghee Yun
  - B.S., Electrical Engineering @ Seoul National University
  - M.S. & Ph.D., Electrical Engineering @ Stanford University
  - Samsung Electronics
    - CAE Team @ Semiconductor R&D Center of Samsung Electronics
    - Design Technology Team @ DRAM Development Lab. of Samsung Electronics
    - Memory Sales & Marketing Team @ Memory Business Unit of Samsung Electronics
    - Software R&D Center @ Samsung Electronics
  - Senior Applied Scientist @ Amazon
- Specialties
  - convex optimization
  - decentralized machine learning
  - deep reinforcement learning
  - recommendation systems

What is Artificial Intelligence (AI)?

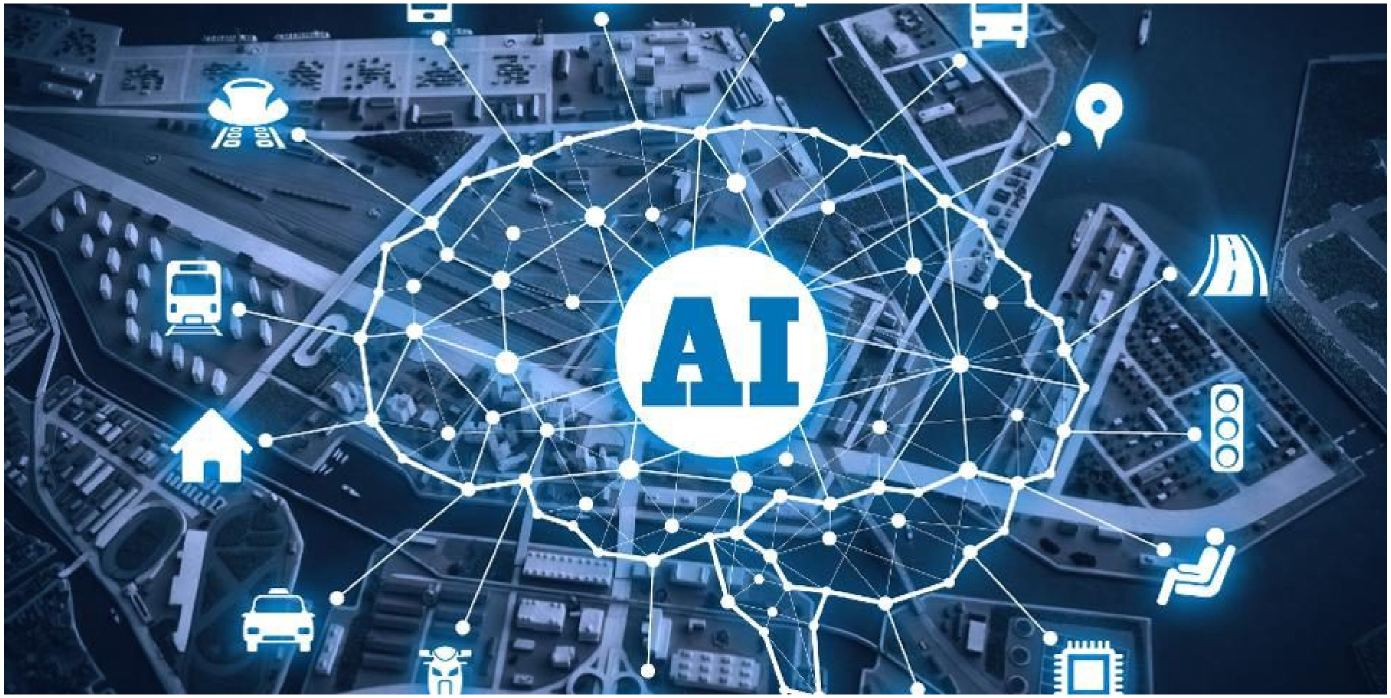

What is Artificial Intelligence (AI)?

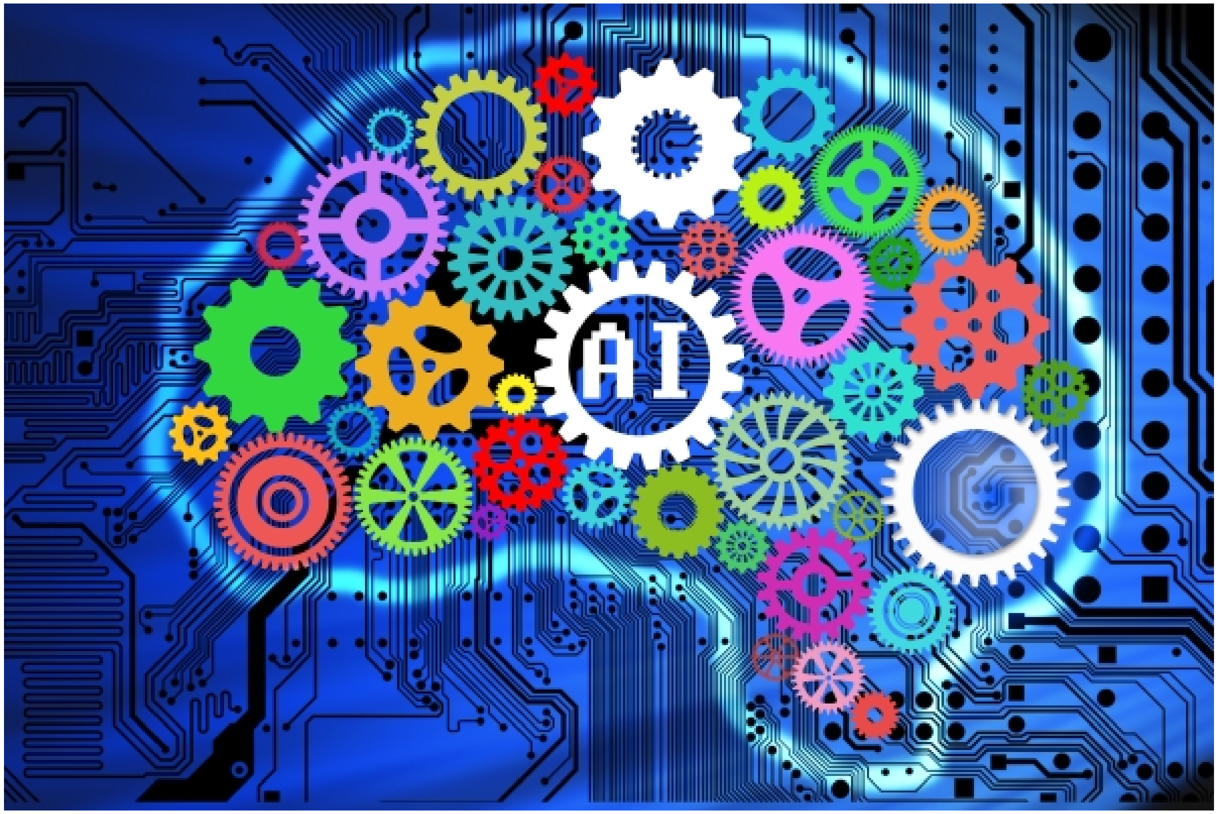

Areas AI makes great impacts on

- eCommerce: AI helps customers to find products fast in online retail store
- Healthcare: AI offers clinical decision support and document events electronically
- Logistics and Supply Chain: autonomous trucks and robitic picking system
- Medical Imaging: Deep learning can diagnose diseases very accurately
- Chatbot, self-driving car, biotechnology, robotics, advertising, finance, *etc.*

Value Creation by AI

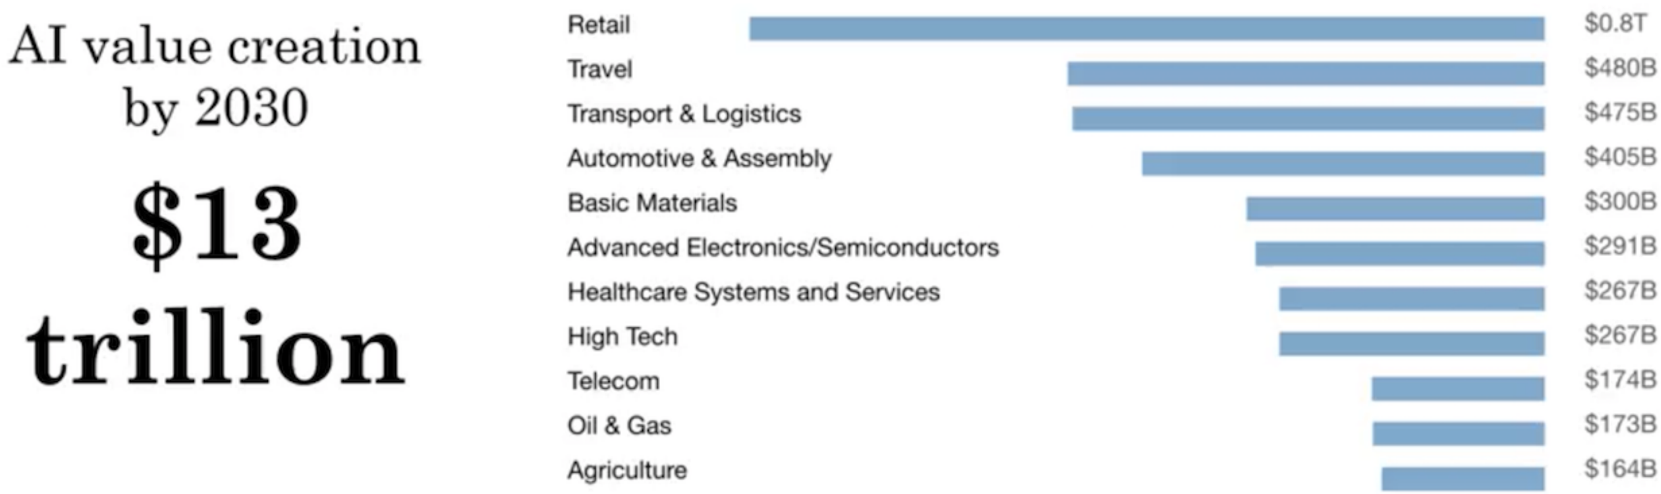

Source: McKinsey Global institute

Demystifying AI

- Artificial Intelligence (AI)
  - Artificial Narrow Intelligence (ANI)
    - smart speaker, self-driving car, web search, chatbot, personal assistant
  - Artificial General Intelligence (AGI)
    - does (almost) anything that a human do

Demystifying AI

- Artificial Intelligence (AI)
  - Artificial Narrow Intelligence (ANI)
    - smart speaker, self-driving car, web search, chatbot, personal assistant
  - Artificial General Intelligence (AGI)
    - does (almost) anything that a human do

AI, ML, DL, DM, DS?

- What are Artificial Intelligence (AI), Machine Learning (ML), Deep Learning (DL), Data Mining (DM), Data Science (DS)?
- How are they different?

AI, ML, DL, DM, DS?

- What are Artificial Intelligence (AI), Machine Learning (ML), Deep Learning (DL), Data Mining (DM), Data Science (DS)?
- How are they different?

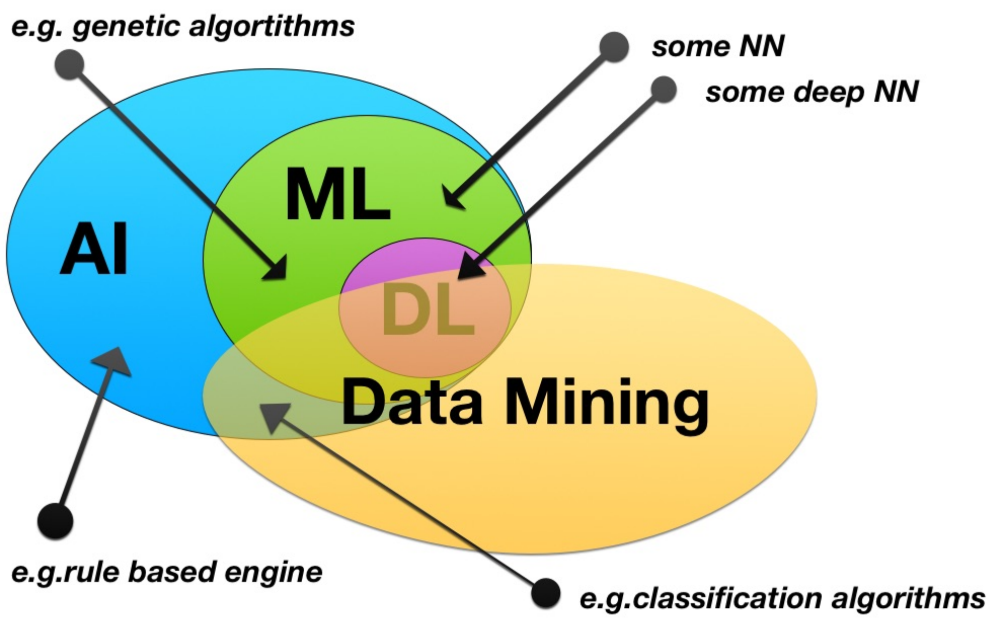

Machine Learning vs. Data Science

- Machine learning (ML)
  - "Field of study that gives computer the ability to learn wihtout being explicitly programmed" - Arthur Samuel (1959)
  - $e.g.$, software
- Data Science (DS)
  - Science of acquiring knowledge and insights from data
  - $e.g.$, slide deck

Three main ML methods

- Supervised learning
- Unsupervised learning
- Reinforcement learning

Supervised learning

- Data: $(x^{(i)}, y^{(i)}) \in \mathbf{R}^m \times \mathbf{R}^l$ $(i = 1, \ldots, N)$
- Goal: learn a function to predict y from x with parameters $\theta \in \mathbf{R}^n$

$$f(x; \theta) \sim y$$

  where $f : \mathbf{R}^m \times \mathbf{R}^n \to \mathbf{R}^l$

- Applications
  - classification
  - regression
  - object detection
  - semantic segmentation

Supervised learning

- give inputs (X), predict output (Y)
- examples
  - given an image, guess which objects are in the image
  - given texts, guess which words would follow the texts
  - given X-ray images, guess the probability of patient having some disease

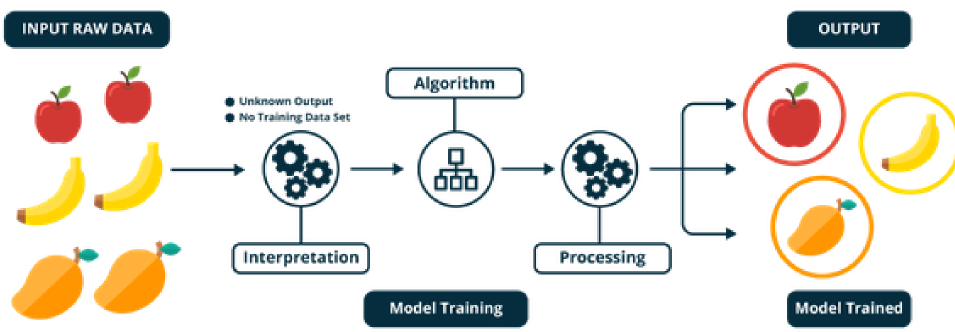

Unsupervised learning

- Data: $x^{(i)} \in \mathbf{R}^n$ $(i = 1, \ldots, N)$
- Goal: learn underlying hidden structure of x
- Applications
  - clustering
  - dimensionality reduction (matrix factorization)
  - featuring learning
  - density estimation
  - autoencoder

Unsupervised learning

- give inputs (X), find out structures of interest
- examples
  - given customer data, group customers into several categories ($e.g.$, for target marketing)
  - given data, estimate the probability distribution
  - given data, learn underlying structures

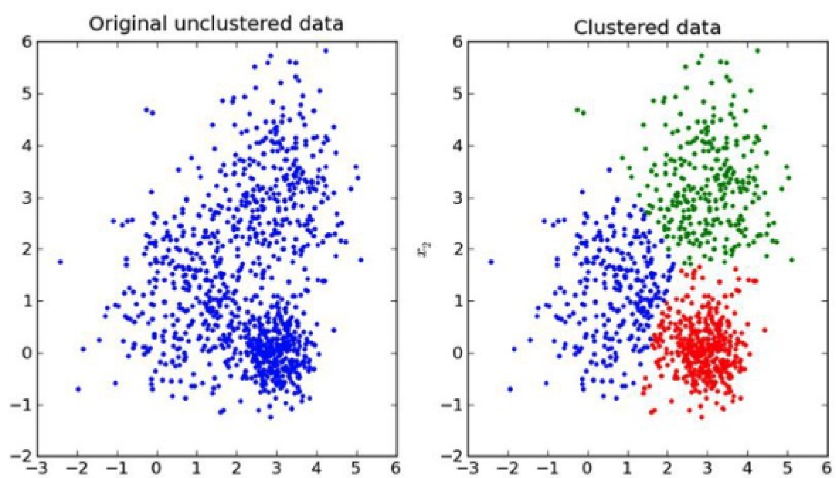

Reinforcement learning

- Agent actively interacts with environment to learn
  - unlike passive ways of learning, $e.g.$, supervised and unsupervised learnings
- Agent decides which actions to take based on history of actions and rewards
- Assumes that the environment reacts with uncertainty $\rightarrow$ stochastic formulation

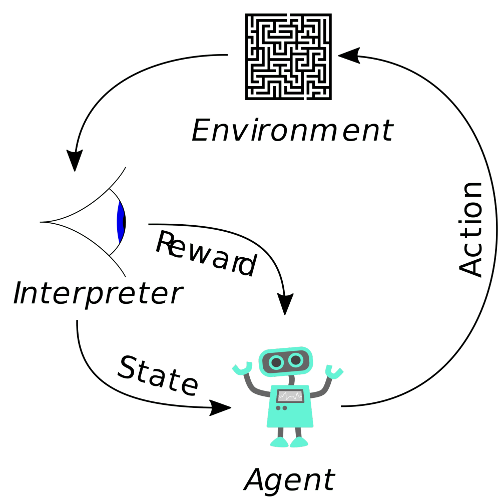

Reinforcement learning example: robot locomotion problem

- Objective: make the robot move forward
- State: angle and position of joints
- Action: torques applied to joints
- Reward: 1 if it's upright and moves forward @ each time step, 0 otherwise
- Sim-to-Real: Learning Agile Locomotion For Quadruped Robots

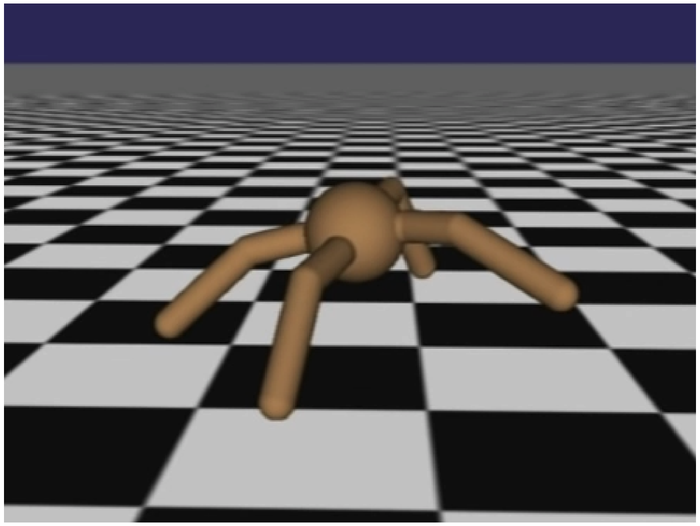

Reinforcement learning example: Atari games

- Objective: maximize score upon completion
- State: raw pixel inputs
- Action: game controls ($e.g.$, left, right, up, down)
- Reward: score increase or decrease @ each time step
- Google DeepMind's Deep Q-learning playing Atari Breakout

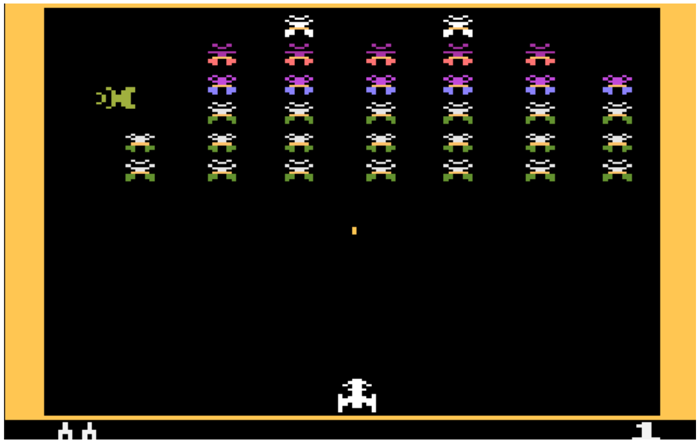

Reinforcement learning example: Go

- Objective: surround more territory (than the opponent)
- State: position of all pieces
- Action: where to put th next piece
- Reward: 1 if win at the end of the game, 0 otherwise

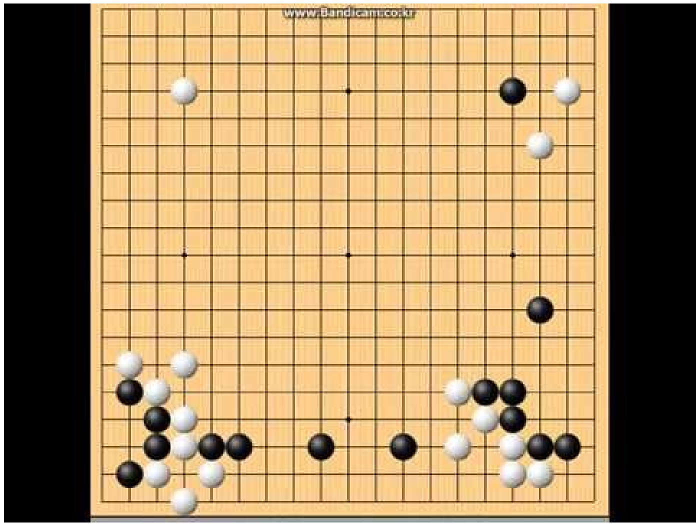

Genetic algorithm: learning to swing

- simulation example: Learn to swing by genetic algorithm

What do ML algorithms do?

- all ML algorithms essentially tries to reduce the difference between model output and measurements
- algorithms
  - stochastic gradient descent, momentum method, adaptive method, *etc.*

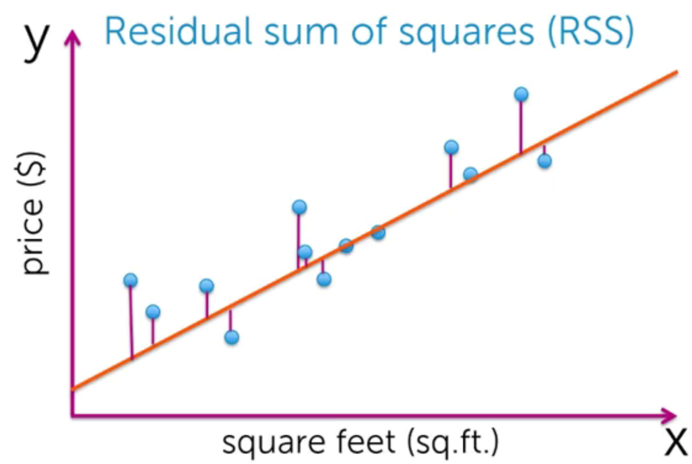

$$\mathrm{RSS}(w_0,w_1) = \sum_{i=1}^{N} (y_i - [w_0 + w_1 x_i])^2$$

Mathematical formulation for ML

- given training set, $\{(x^{(1)}, y^{(1)}), \ldots, (x^{(m)}, y^{(m)})\}$, where $x^{(i)} \in \mathbf{R}^p$ and $y^{(i)} \in \mathbf{R}^q$
- want to find function $g_\theta : \mathbf{R}^p \to \mathbf{R}^q$ with learning parameter, $\theta \in \mathbf{R}^n$
  - $g_\theta(x)$ desired to be as close as possible to y for future $(x, y) \in \mathbf{R}^p \times \mathbf{R}^q$
  - $i.e.$, $g_\theta(x) \sim y$
- define a loss function $l : \mathbf{R}^q \times \mathbf{R}^q \to \mathbf{R}_+$
- solve the optimization problem:

$$\begin{array}{ll} \text{minimize} & f(\theta) = \frac{1}{m} \sum_{i=1}^m l(g_\theta(x^{(i)}), y^{(i)}) \\ \text{subject to} & \theta \in \Theta \end{array}$$

ML example: regression

- problem: what is a reasonable price for a house?
  - what would a rational (or rather normal) human being do?
  - ML approach:
    - collect data: x: size, y: price
    - train model: draw a line to represent (typical) trend
    - predict a price from the line

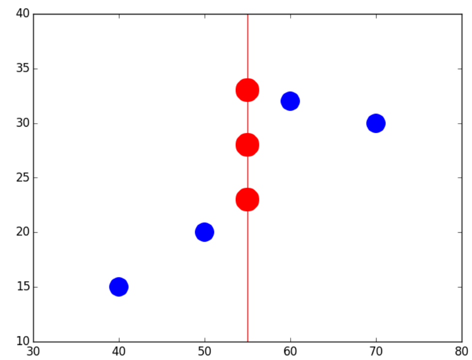

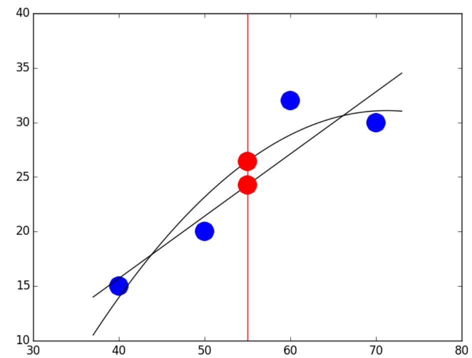

ML example: multi-variate regression

- what if we have more than one x? or rather more than two x's?

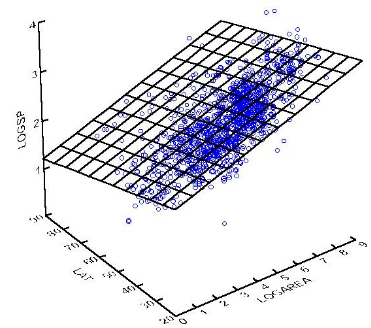

- what if highly nonlinera and nonconvex fitting function is needed?

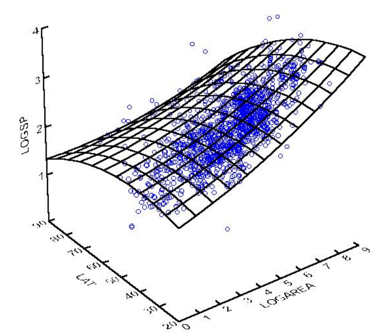

- ALL ML algorithms try to solve these difficult problems in (smart) ways!

Deep learning

- train the neural network with many layers by tweaking weights on connections
  - universal approximation theorem: feed-forward network with a single hidden layer can approximate any continuous functions
  - Bayesian inference: the more data it sees, the smarter it gets

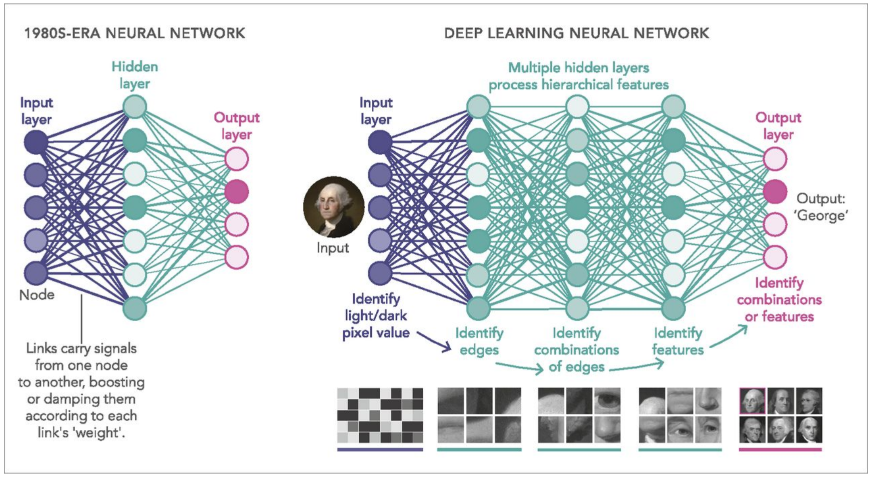

Deep learning example

- problem: predict (or estimate proper) flight price
- inputs: origin/destination airports, departure date, airline
- output: price

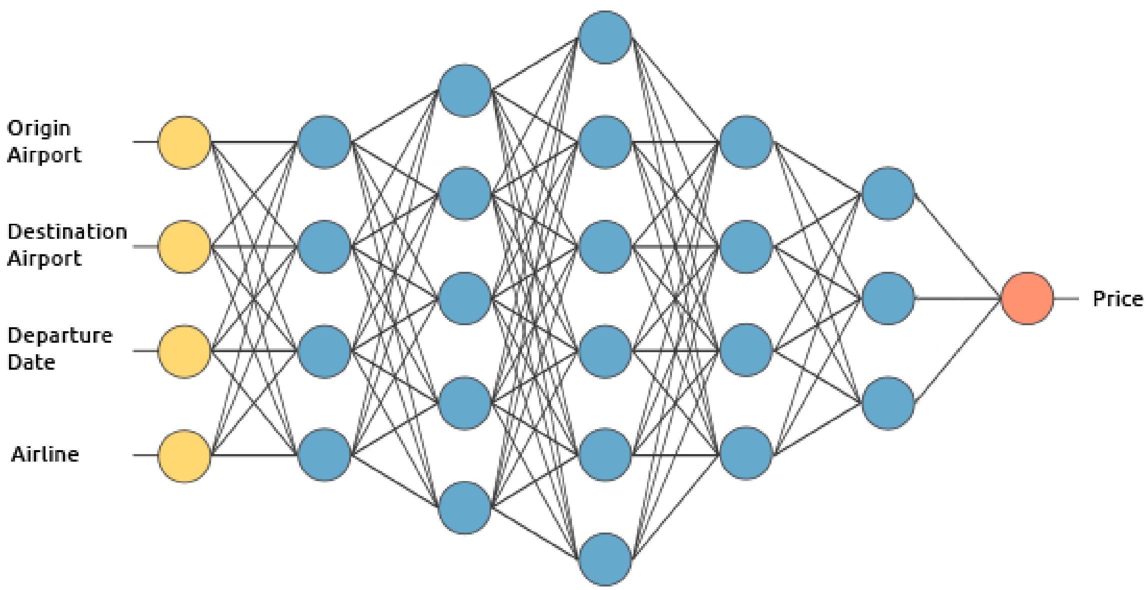

Why now?

- enormous data: eCommerce, multi-media, digital data
- computation power: Moore's laws, cloud computing, GPU

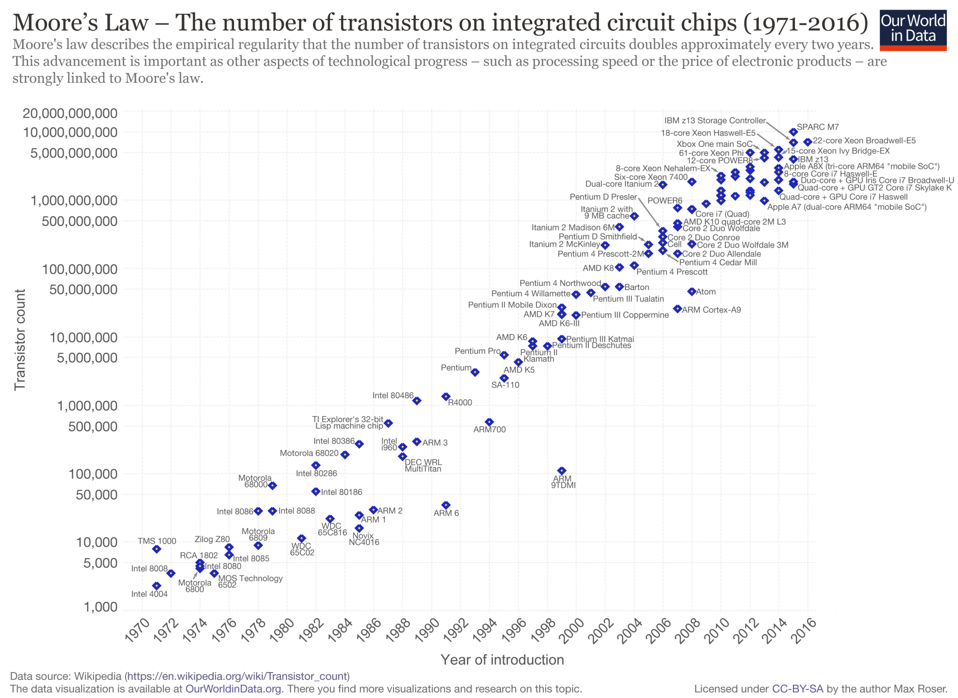

How can we do better?

- more data, stronger computation power, larger neural net, the better!

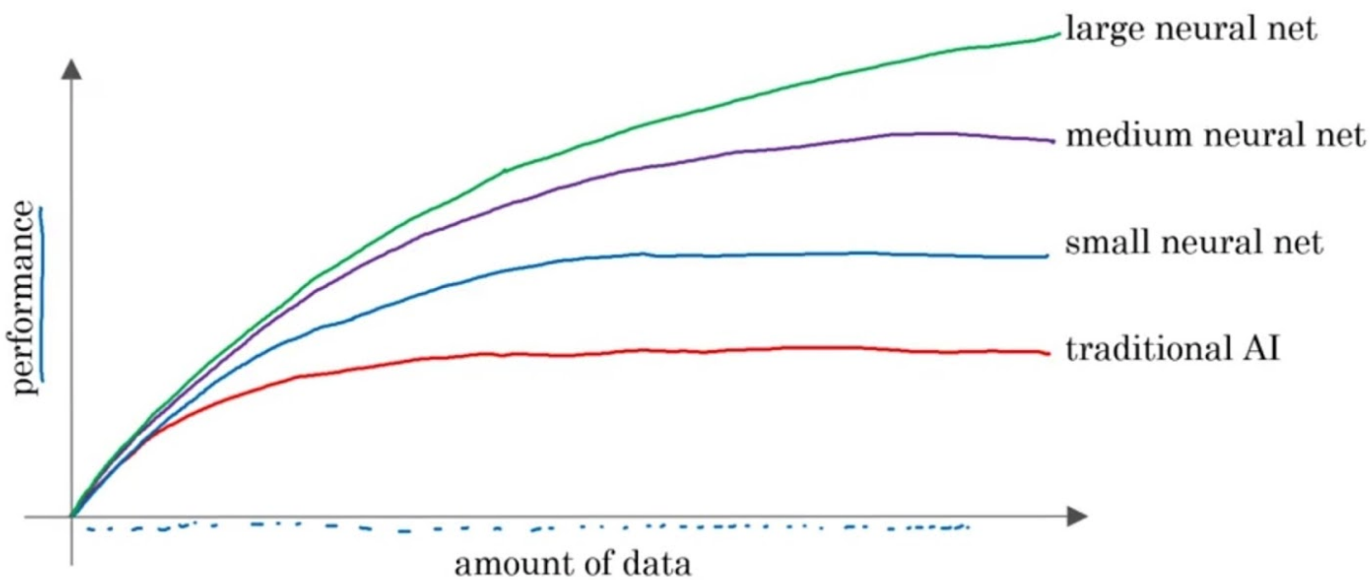

Demystifying AI

- ML today can do
  - relatively simple tasks
  - what humans can do within seconds
- ML today cannot do
  - quite complicated tasks requiring human intelligence
  - what takes long time for humans to do

What ML today can do

- when a customer sends e-mail
  - “The toy arrived two days late, so I was not able to give it to my niece for her birthday. Can I return it?”
- humans would say
  - “Oh, I’m sorry to hear that”, “I hope your niece had a good birthday”, “Yes, we can help with . . . ”
- AI could only say
  - “Thank you for your e-mail”

Self-driving car

- ML can recognize objects ($e.g.$, semantic segmentation)

- ML cannot recognize people's intentions

Advantages of ML

- However, machines
  - never get tired or sleeps
  - never complain about their pay
  - do not increase errors because they repeatedly do the same task
  - have perfect memory and precise computation ability
- for examples, for 24 hours a day and 7 days a week,
  - Amazon recommendation system learns model with data from hundreds of millions of customers
  - Google photos learns modles with trillions of photo images

Things to discuss

- Would singularity come?
- Could machines have consciousness?
- Is Skynet plausible?
- Would humans lose many jobs?

Conclusion

- AI has changed the world; its impact on our world is significant
- ML algorithms try to reduce errors
  - neural network has amazing capability
  - development of high performance computer has enabled many difference
- we do not know the future, but most likely
  - the singularity would not come (soon)
  - machines cannot have consciousness
  - many jobs will be lost, but humans will find a way (as they have)