Machine Learning

Sang-Hoon Oh

Division of Information Communication Convergence Engineering
Mokwon University

Contents

1.1. Machine Learning?

- Motivation for Machine Learning
- Advances in computer technology

  (store/access/process) :

  predictions, understanding and control.

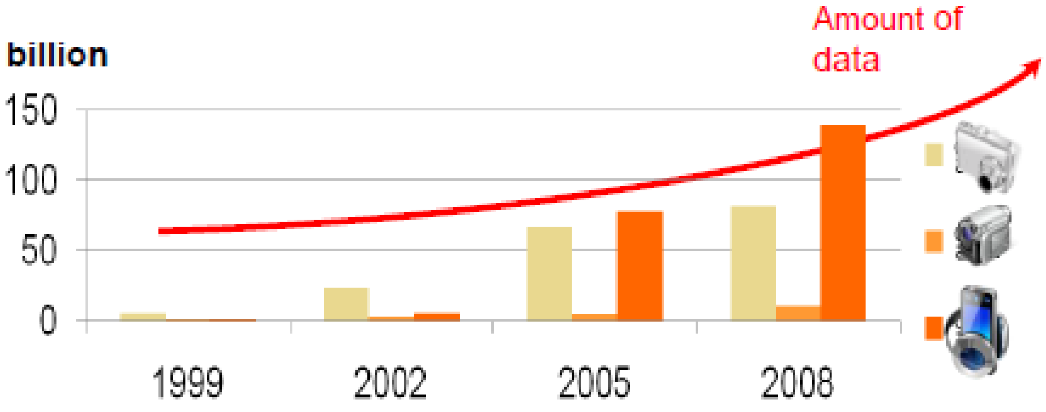

Data explosion

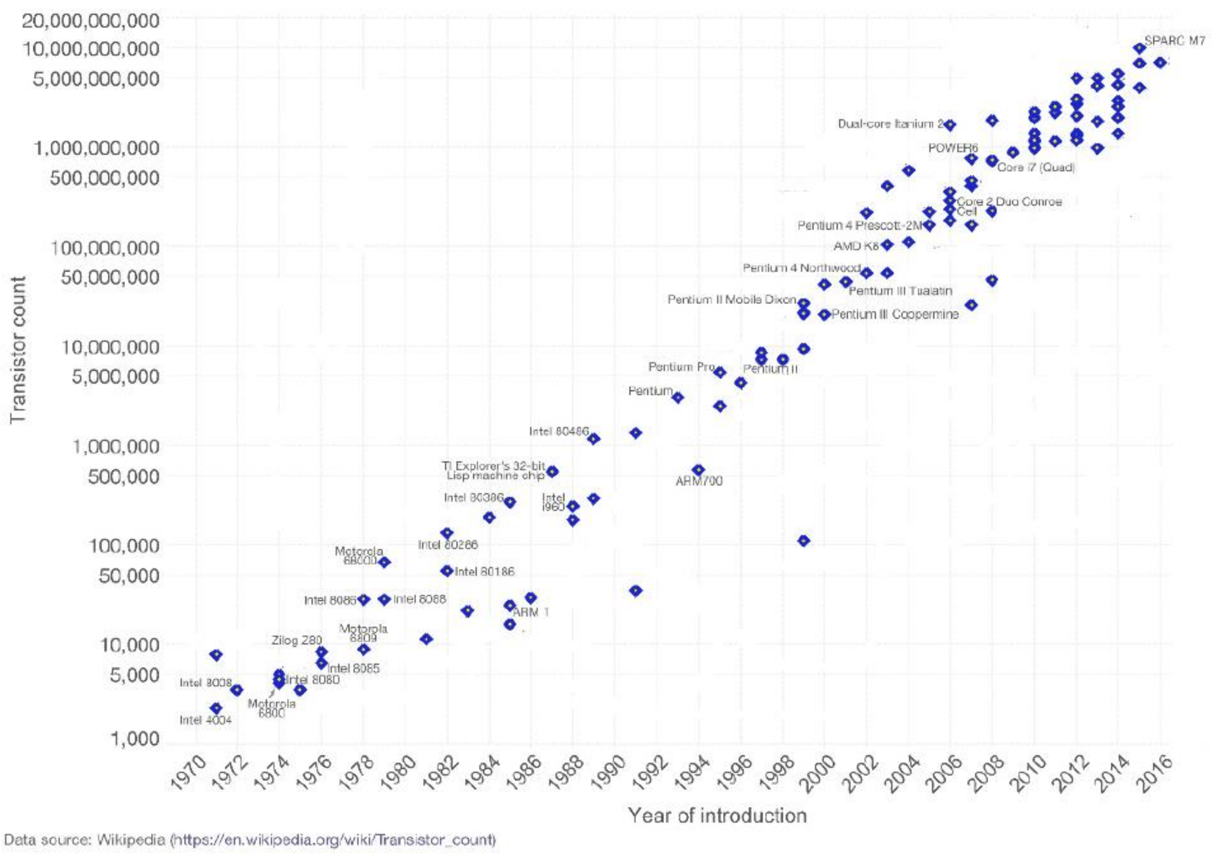

Moore's law

The number of transistors that can be inexpensively placed on an integrated circuit is increasing exponentially, doubling approximately every two years

- Very difficult to understand the details of the process that explains the generation of the data
- Need a technique to make a good and useful approximation in detecting certain patterns and regularities ⟹ Machine Learning!

1.2. Historical Background of Machine Learning

- Biological vs. Artificial Neuron
  - (1943) McCulloch and Pitts: Artificial neuron model
  - (1949) Hebbian learning
  - (1958) Perceptron (Rosenblatt)

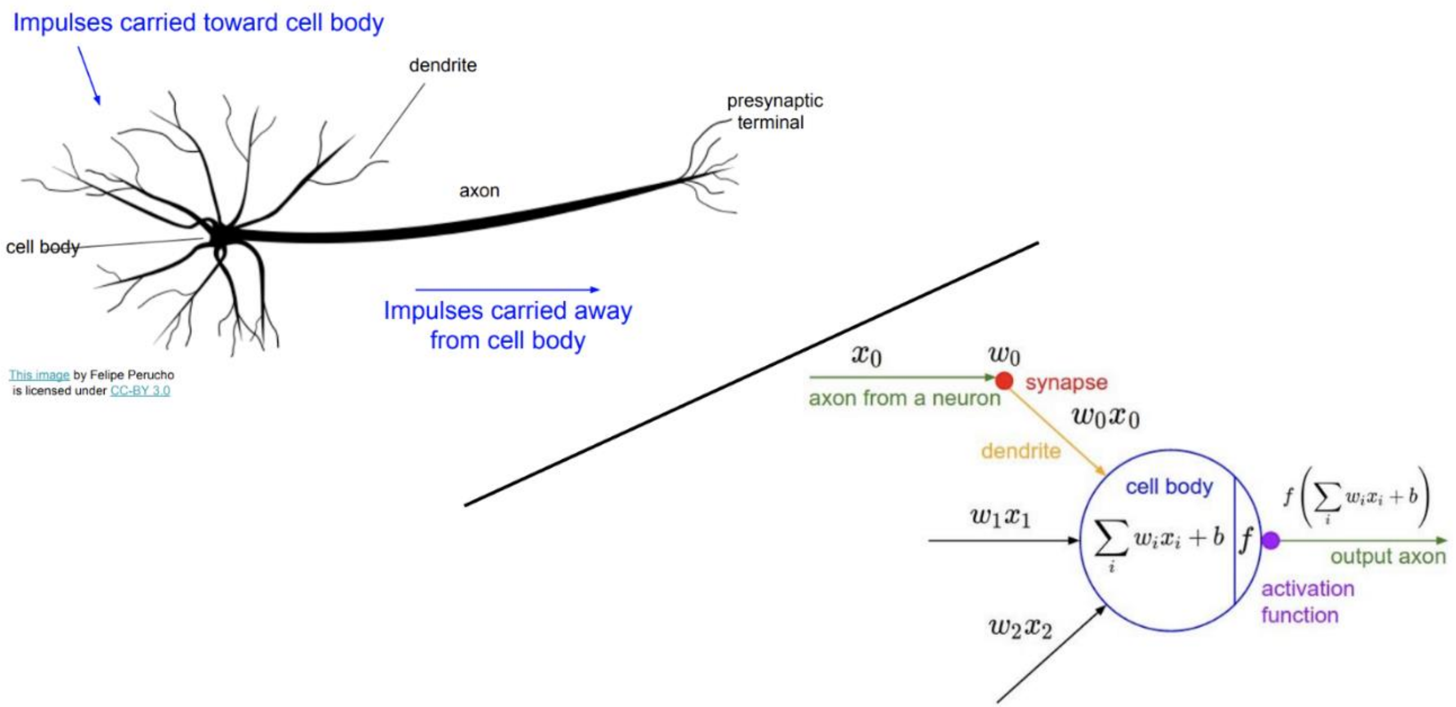

Side-by-side illustrations of biological and artificial neurons, via Stanford's CS231n. This analogy can't be taken too literally — biological neurons can do things that artificial neurons can't, and vice versa — but it's useful to understand the biological inspiration. See Wikipedia's description of biological vs. artificial neurons for more detail.

1.2. Historical Background of Machine Learning

1. Latency Period(1940~1980)

Latency Period (1940~1980)	
1943	McClloch & Pitts Neuron Model
1949	Hebbian Learning (Hebb)
1958	Perceptron (Rosenblatt)
1960	Delta Rule (Widrow & Hoff)
1967	Outstar Learning (Grossberg)
1969	Perceptron Book (Minsky & Papert)
1972	Associative Memory Neural Nets (Kohonen)
1973	Pattern Classification and Scene Analysis (Duda & Hart)
1977	Associative Memory Nets (Anderson)
1980	Neocognitron (Fukushima)

2. Quickening Period(1980~1995)

Quickening Period (1980~1995)	
1981	Parallel Models of Associative Memory (Hinton & Anderson)
1982	Self-Organizing Maps (Kohonen)
1982	Hopfield Neural Networks (Hopfield)
1983	Boltzmann Machine (Hinton & Sejnowski)
1985	Adaptive Resonance Theory (Carpenter & Grossberg)
1986	Error Back-Propagation Algorithm (Rumelhart, Hinton, &Williams)
1988	Bayesian Networks (Lauritzen, Spiegelhalter, & Pearl)
1992	Support Vector Machines (Boser, Guyon, & Vapnik)
1995	Statistical Learning Theory (Vapnik)

3. Growth Period(1995~2010)

Growth Period (1995~2010)	
1997	Independent Component Analysis (Bell & Sejnowski)
1998	Natural Gradient (Amari)
1998	Reinforcement Learning (Sutton & Barto)
1999	Learning in Graphical Model (Jordan)
1999	Kernel Machines (Schoelkopf & Smola)
2003	Boosting Algorithms (Freund & Shapire)
2005	DARPA Grand Challenge
2006	Restricted Boltzmann Machine/Deep Learning (Hinton & Sal akutdinov)
2009	Probabilistic Graphical Models (Koller & Friedman)

4. Leaping Period(2010~Current)

Leaping Period (2010~Current)	
2010	Google Car
2011	IBM Watson AI
2011	Apple Siri Personal Assistant
2012	AlexNet
2013	Allen Institute for Artificial Intelligence
2013	Facebook AI Research
2014	Baidu Deep Learning Institute
2014	Amazon Echo & Alexa
2014	Generative Adversial Networks (Goodfellow et al.)
2015	DARPA Robotics Challenge
2016	Google DeepMind AlphaGo
2016	Apple AI Lab
2016	Google Home & Assistant

A Brief Review of NIPS 2015
February 2016
Article Authors: Zak Shafran

The twenty ninth Conference on Neural Information Processing Systems (NIPS), a single track machine learning and computational neuroscience conference, was held in Montreal, during a relatively balmy winter week, spanning December 7th to 12th. The conference saw a record number of attendees. As illustrated in the graph, during the last 14 years, there has been a steady increase in the number of attendees in all three categories -- tutorials, conference, and workshops; and biggest growth this year was in tutorials where the attendance doubled compared to the previous year!
……………………..

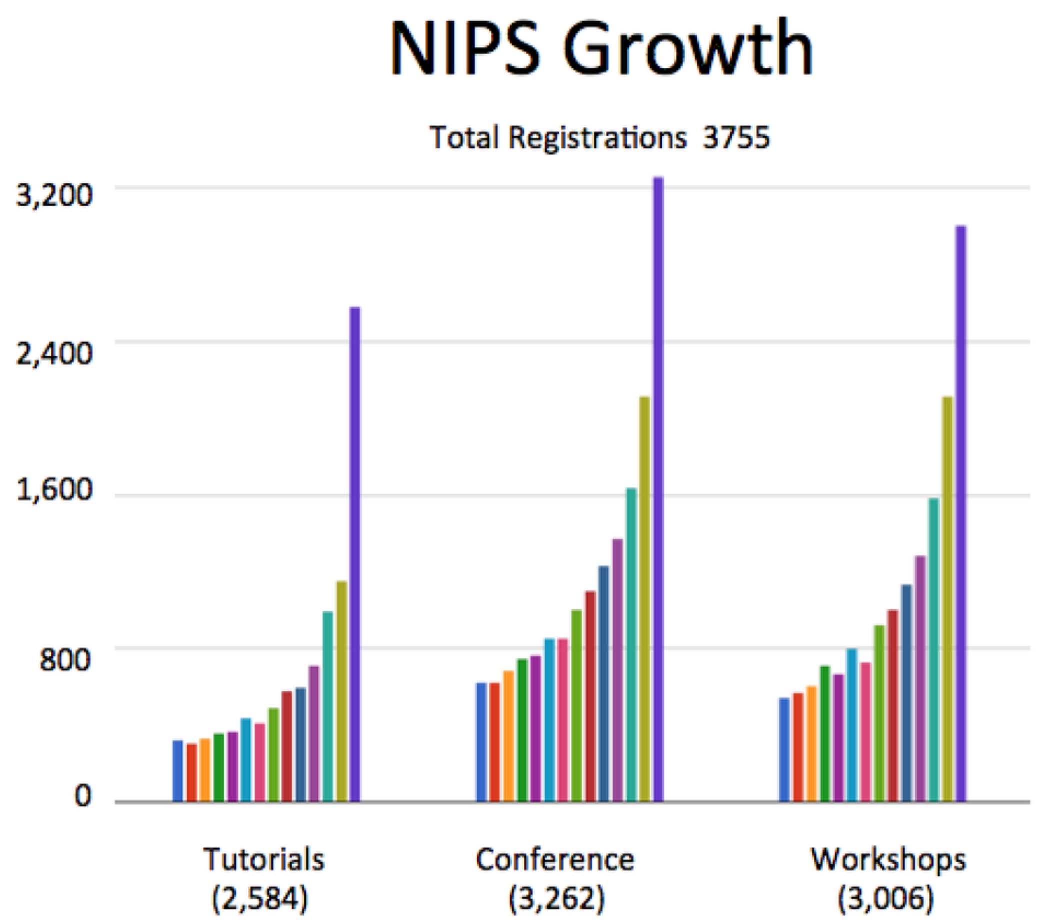

1.3. Types of Machine Learning

- Prediction
  - Supervised Learning: Given $\{(x_i, y_i)\}$ pairs, find a good understanding f: X→Y. Ex) Classification, Regression
  - Active Learning : Find f: X→Y, but $y_1, \dots, y_N$ are initially hidden and there is a charge for each label you want revealed. By intelligent adaptive querying, you can get away with significantly fewer labels than those needed in a regular supervised learning framework
  - Semi-supervised learning: Given both small number of labels and large number of unlabeled data, find f: X→Y.
- Understanding
  - Unsupervised Learning: Given $\{x_i\}$ find something interesting. Ex) Clustering, novelty-detection
- Policy(Control)
  - Reinforcement Learning: find f: P→A to max R

  (P: history of perception, A: action. R: long term reward)
- Generative Model
  - Generative Adversarial Network: The field of unsupervised learning which aims to study algorithms that learn the underlying structure of the given data, without specifying a target value.

2.2. Classification by *k*-Nearest Neighbors Algorithm

- Classification
  - Ex) Credit scoring
    - Low-risk and high-risk customers from their incomes and savings

  - Class membership
  - Other examples
    - Handwritten digit recognition
    - Speech Recognition
    - Fraud detection

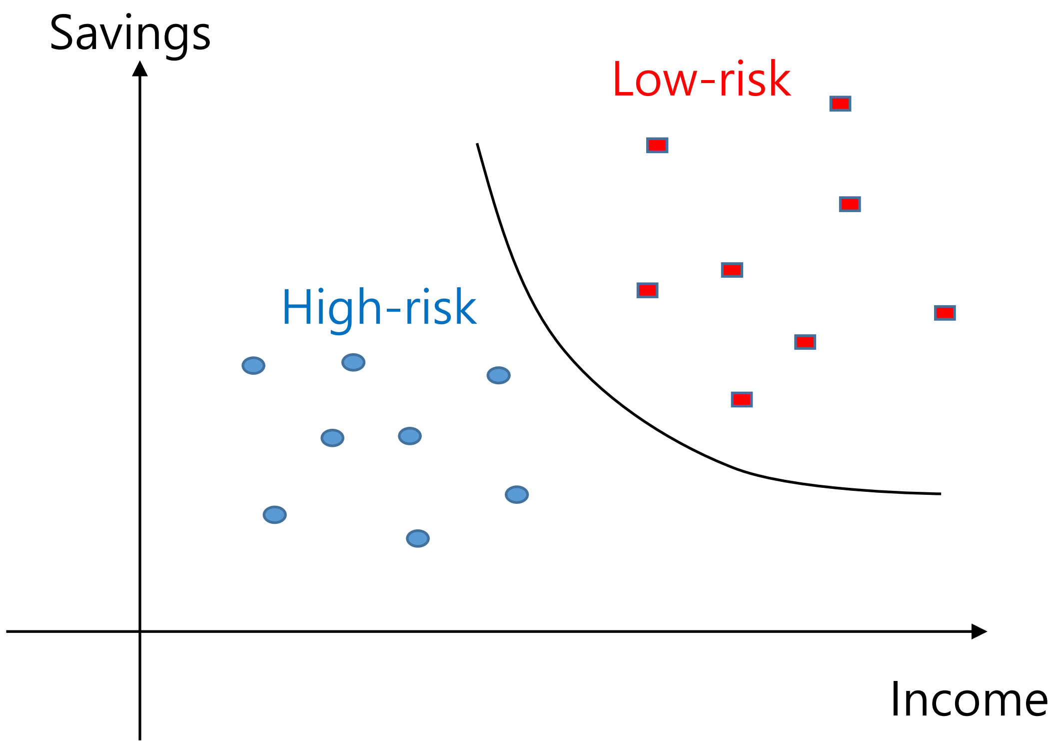

2.3. Regression by *k*-Nearest Neighbors Algorithm

- Regression
  - Ex) The price of a used car

  - Output values: continuous(regression) vs. discrete (classification)

  - Other examples
    - Stock price prediction
    - Oil price prediction
    - Water level prediction

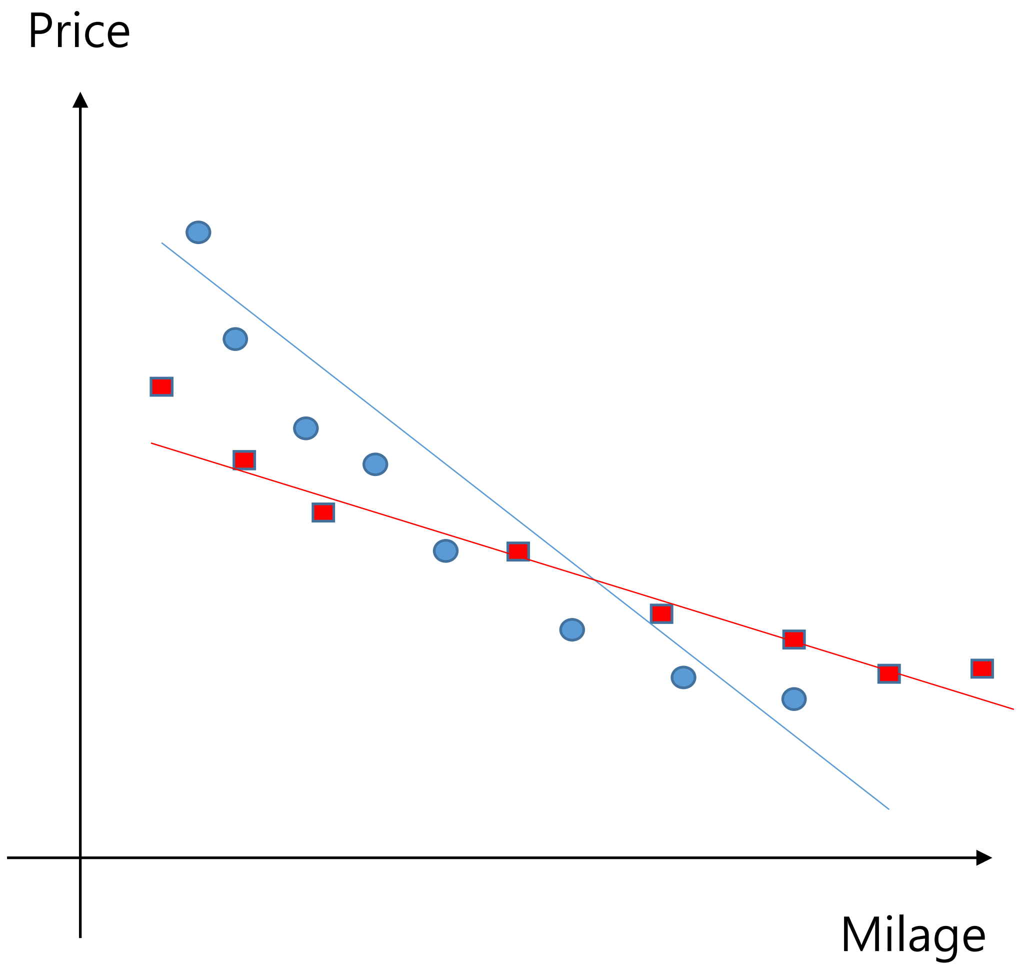

Applications

○**Medicine/Biology**

- (Predict) Possibility of second attack (patient hospitalized due to heart attack)
- (Predict) Glucose in the blood of diabetic person
- (Predict) Predict cancer
- Bioinformatics.

○**Economy/finance**

- (Predict) Predict stock price
- (Predict) Credit card fraud detection

○**IT**

- (Predict) Handwritten ZIP code recognition
- (Predict) Spam mail
- (Predict) Speech Recognition
- (Understanding) Call patterns for network optimization

○**Astronomy/Physics**

- (Understanding) Discovering astronomical features

○**Intelligence**

- (Policy) Flying helicopters upside down
- (Policy) Game playing
- (Policy) Robot navigation