Machine Learning

Sang-Hoon Oh

Division of Information Communication Convergence Engineering
Mokwon University

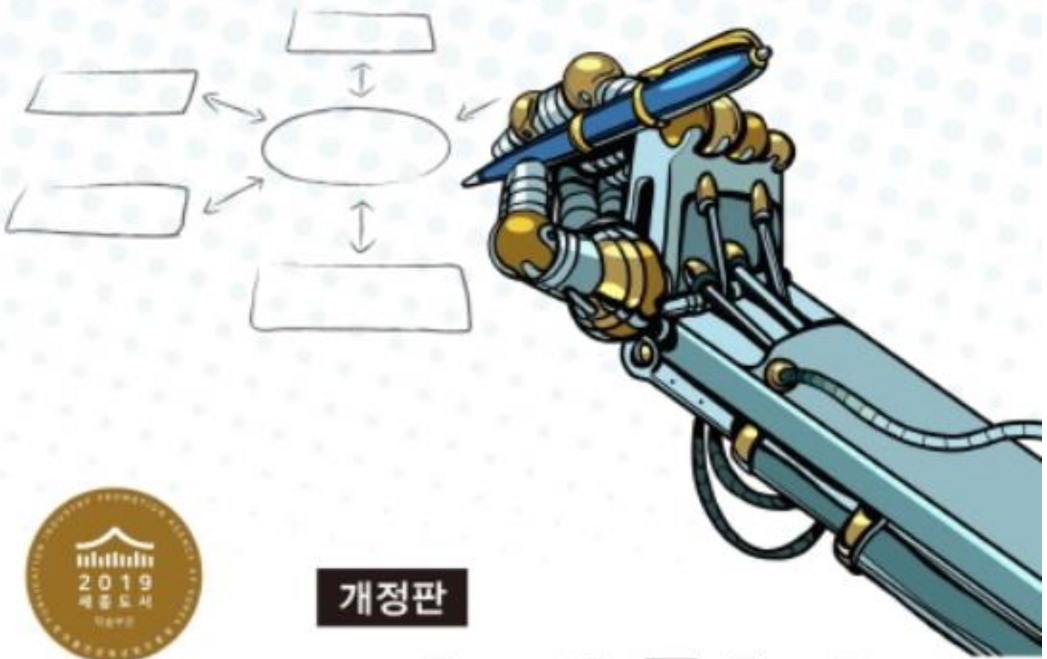

개정판

기계학습

수학적 이해에서 알고리즘까지

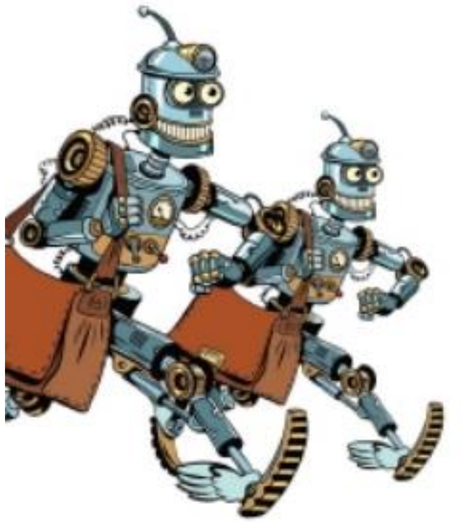

오상훈 지음

도서출판 홍릉

Contents

1.1. Artificial Intelligence and Machine Learning?

- 인공지능(Artificial Intelligence)?

- 사전적 정의(국립국어원 표준국어대사전)
  - 지능: 계산이나 문장 작성 따위의 지적 작업에서, 성취 정도에 따라 정하여지는 적응 능력. 지능 지수 따위로 수치화할 수 있다.

  - 인공지능: 『정보·통신』 인간의 지능이 가지는 학습, 추리, 적응, 논증 따위의 기능을 갖춘 컴퓨터 시스템. 전문가 시스템, 자연 언어의 이해, 음성 번역, 로봇 공학, 인공 시각, 문제 해결, 학습과 지식 획득, 인지 과학 따위에 응용한다.

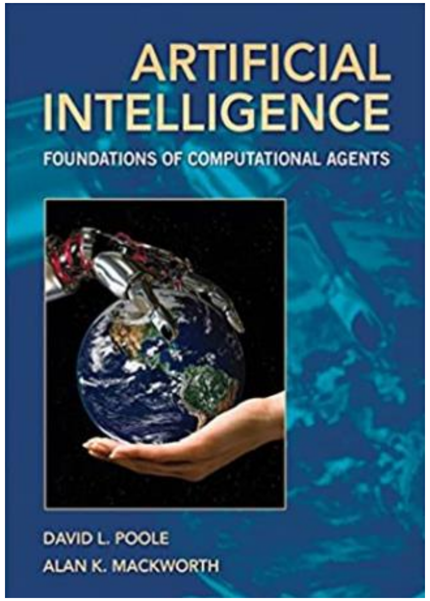

- Artificial Intelligence: the field that studies the synthesis and analysis of computational agents that act intelligently

- ANI(Artificial Narrow Intelligence): 특정 업무..Weak AI
- AGI(Artificial General Intelligence): 다양한 포괄적 업무..Strong AI
- ASI(Artificial Super Intelligence): 스스로 목표 설정 및 임무 수행

Turing Test

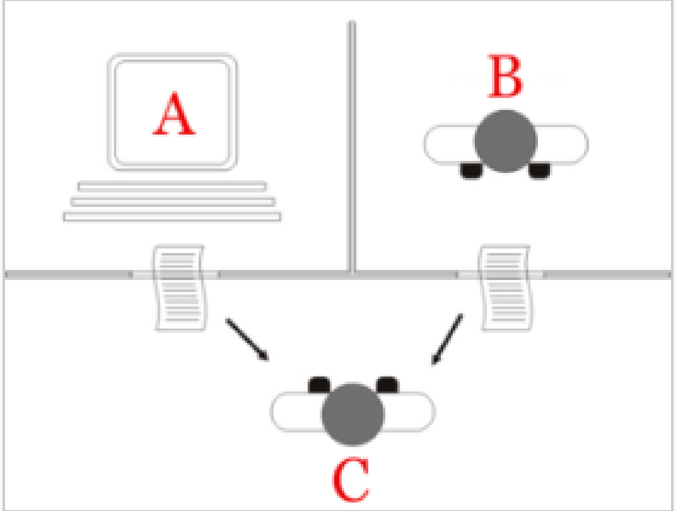

The "standard interpretation" of the Turing test, in which player C, the interrogator, is given the task of trying to determine which player – A or B – is a computer and which is a human. The interrogator is limited to using the responses to written questions to make the determination.[1]

- **A test of a machine's ability to exhibit intelligent behaviour equivalent to, or indistinguishable from, that of a human.**
- Turing proposed that a human evaluator would judge natural language conversations between a human and a machine designed to generate human-like responses. The evaluator would be aware that one of the two partners in conversation was a machine, and all participants would be separated from one another.
- The conversation would be limited to a text-only channel, such as a computer keyboard and screen, so the result would not depend on the machine's ability to render words as speech.
- If the evaluator could not reliably tell the machine from the human, the machine would be said to have passed the test.
- **The test results would not depend on the machine's ability to give correct answers to questions, only on how closely its answers resembled those a human would give.**

- John Searle: Chinese Room
  - 입출력에 따라 작동하는 정보처리 과정은 인간이 마음에 있어서 이해라는 심적 사건을 구현하지 못한다. 따라서, 인간의 마음은 컴퓨터의 정보처리 과정과 같지 않다.

Symbolic AI vs. Computational AI

- Symbolic AI (Rule-Based AI)
  - 문제를 해결하기 위한 규칙을 생성하고,
  - 그 규칙들을 프로그램 하여 문제를 해결하고자 함
  - 예) 전문가 시스템(Expert System)
  - 분석한 규칙에 어긋난 상황이 발생하면?

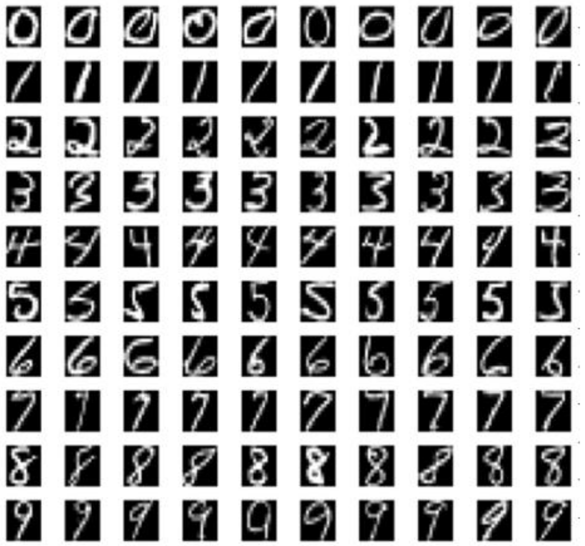

- Computational AI (Machine Learning)
  - Imitation of biological information processing
  - Artificial Neural Networks
  - Statistical Learning Theory
  - Learning from examples (데이터: 4차산업시대의 원유. 빅데이터)

  ⇒ Machine Learning

- Motivation for Machine Learning
- Advances in computer technology

  (store/access/process) :

  predictions, understanding and control.

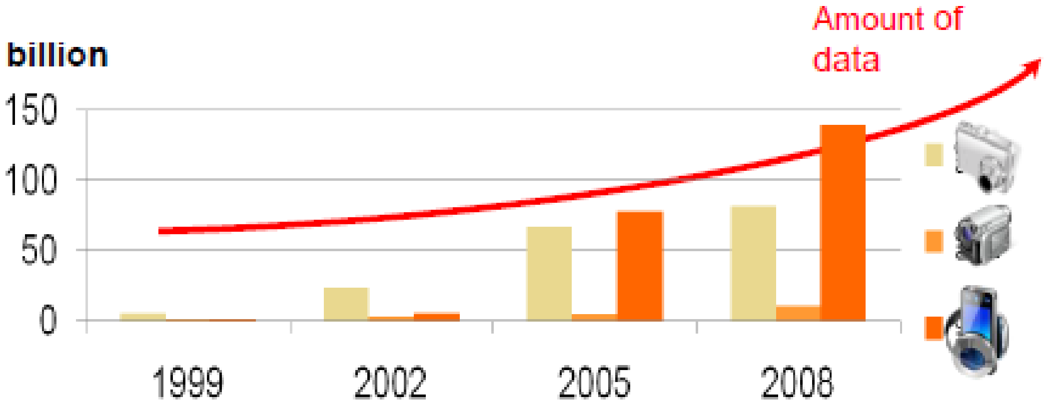

Data explosion

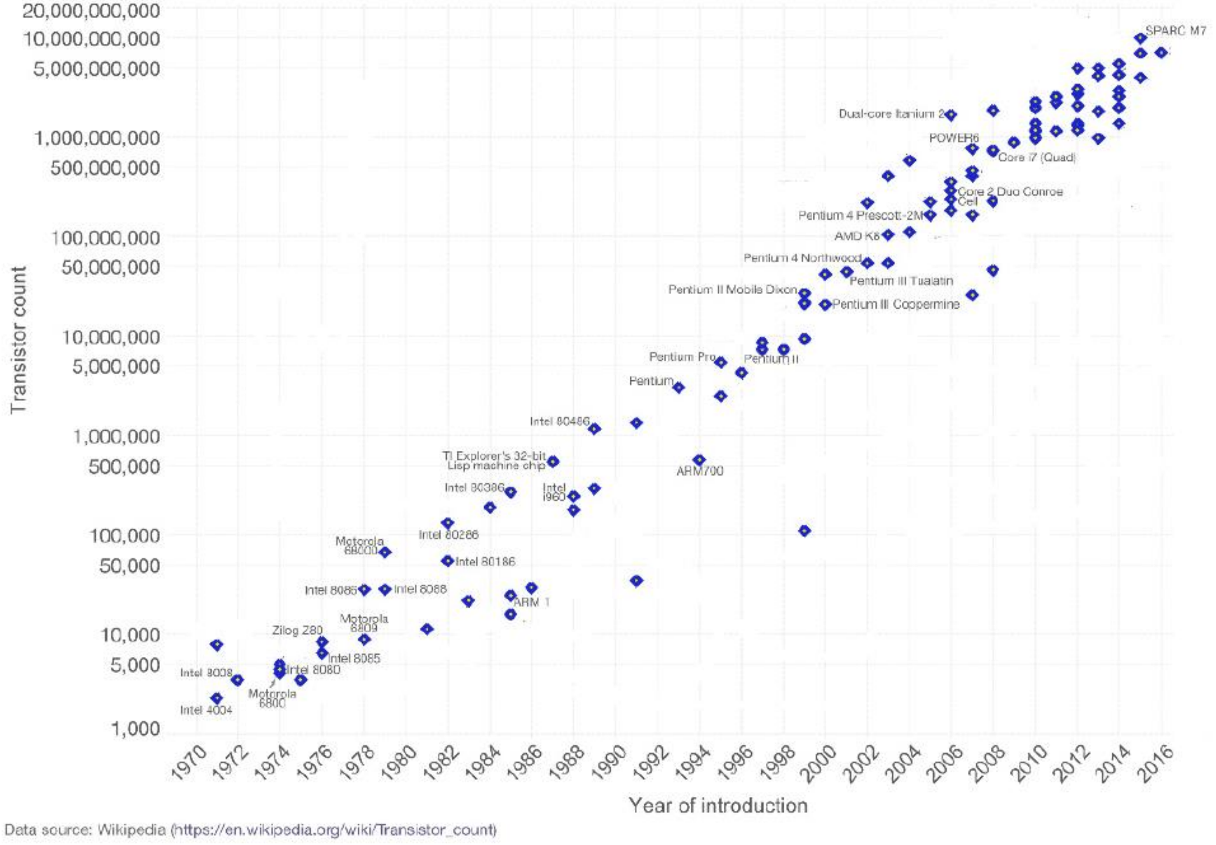

Moore's law

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 Artificial Intelligence(ML)

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

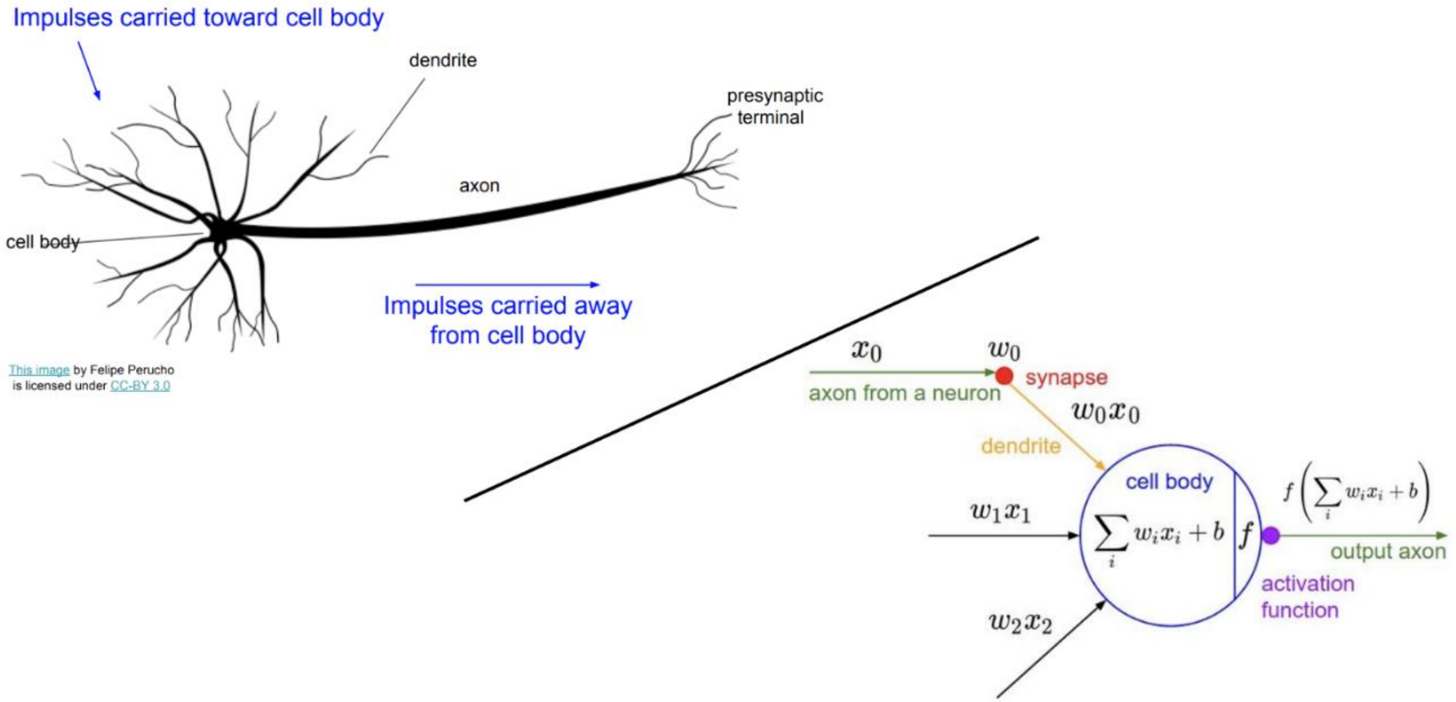

Side-by-side illustrations of biological and artificial neurons, via Stanford's CS231n. This analogy can't be taken too literally—biological neurons can do things that artificial neurons can't, and vice versa—but it's useful to understand the biological inspiration. See Wikipedia's description of biological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

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)

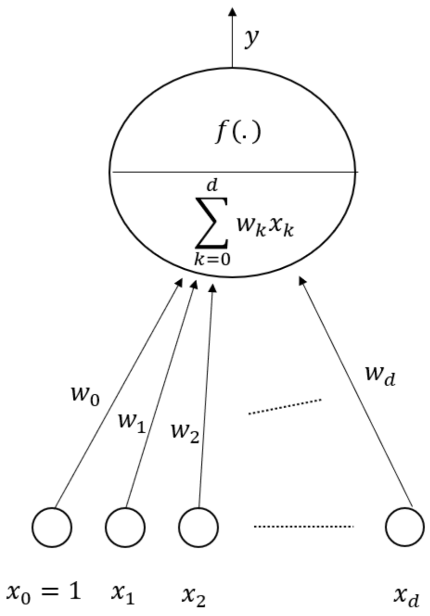

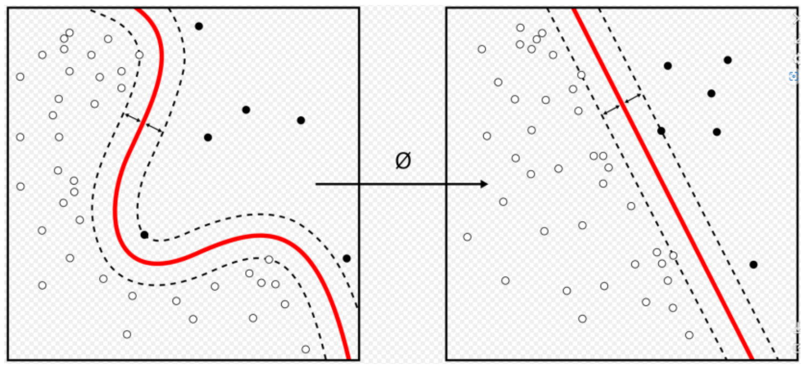

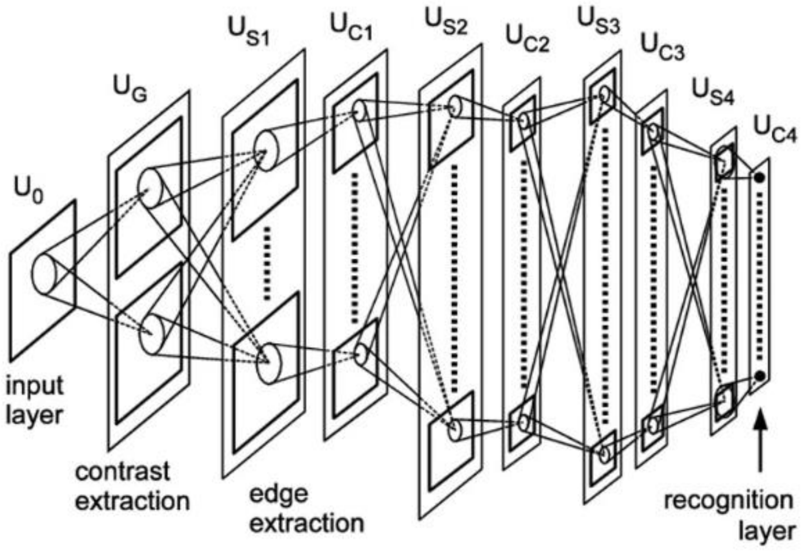

Neocognitiron

Historical Background of Machine Learning

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)

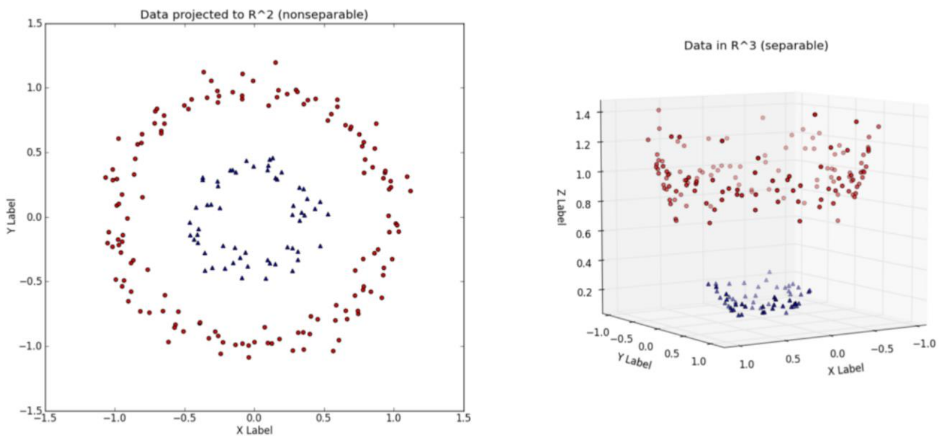

A nonseparable dataset in a two-dimensional space R^2, and the same dataset mapped onto threedimensions with the third dimension being x^2+y^2 (source: http://www.eric-kim.net/eric-kim-net/posts/1/kernel_trick.html)

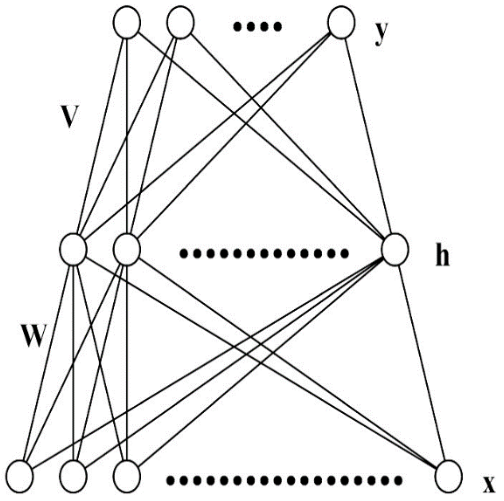

Multilayer Perceptron(MLP)

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)

Stanford Racing and Victor Tango together at an intersection in the DARPA Urban Challenge Finals

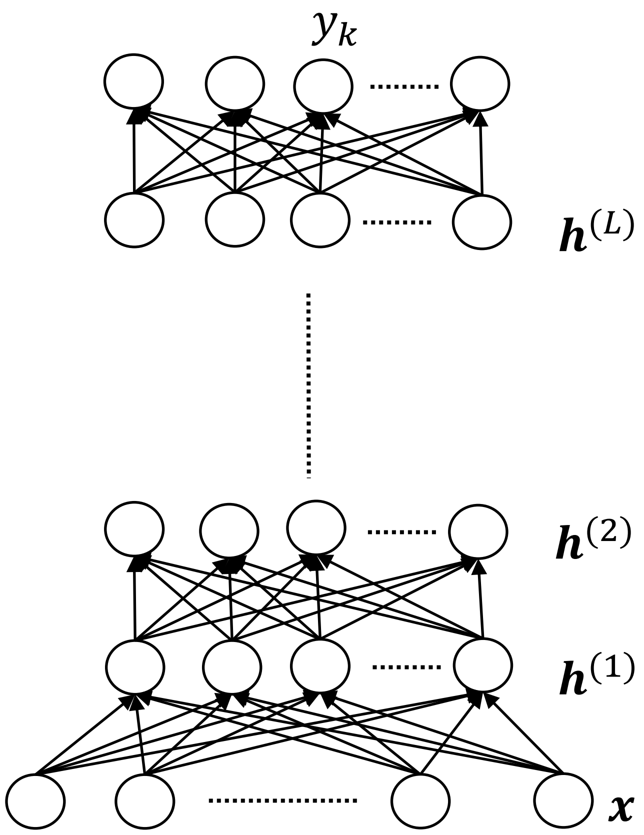

Deep Neural Network

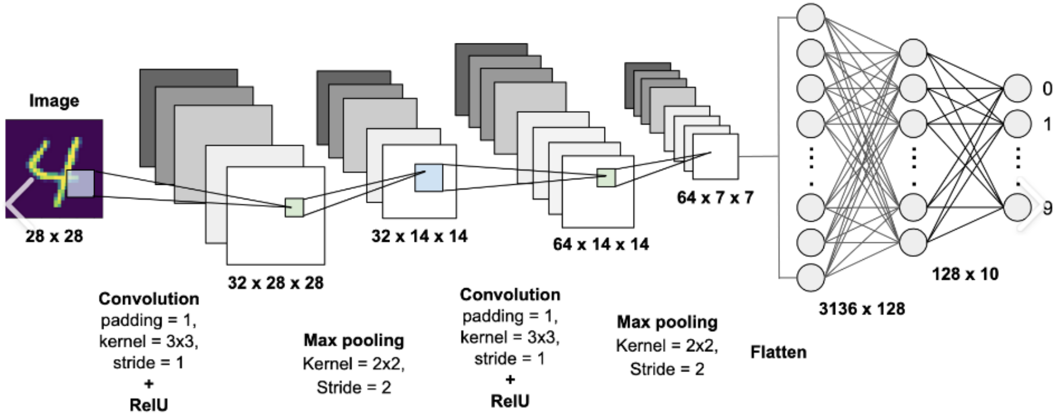

4. Leaping Period(2010~Current)

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
2020	Open AI: GPT-3
2022	Open AI: Dall-E-2 / ChatGPT
2022	Diffusion Model

Google Self-Driving Cars First Ride

Ken Jennings, Watson, and Brad Rutter in their *Jeopardy!* exhibition match.

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!
....

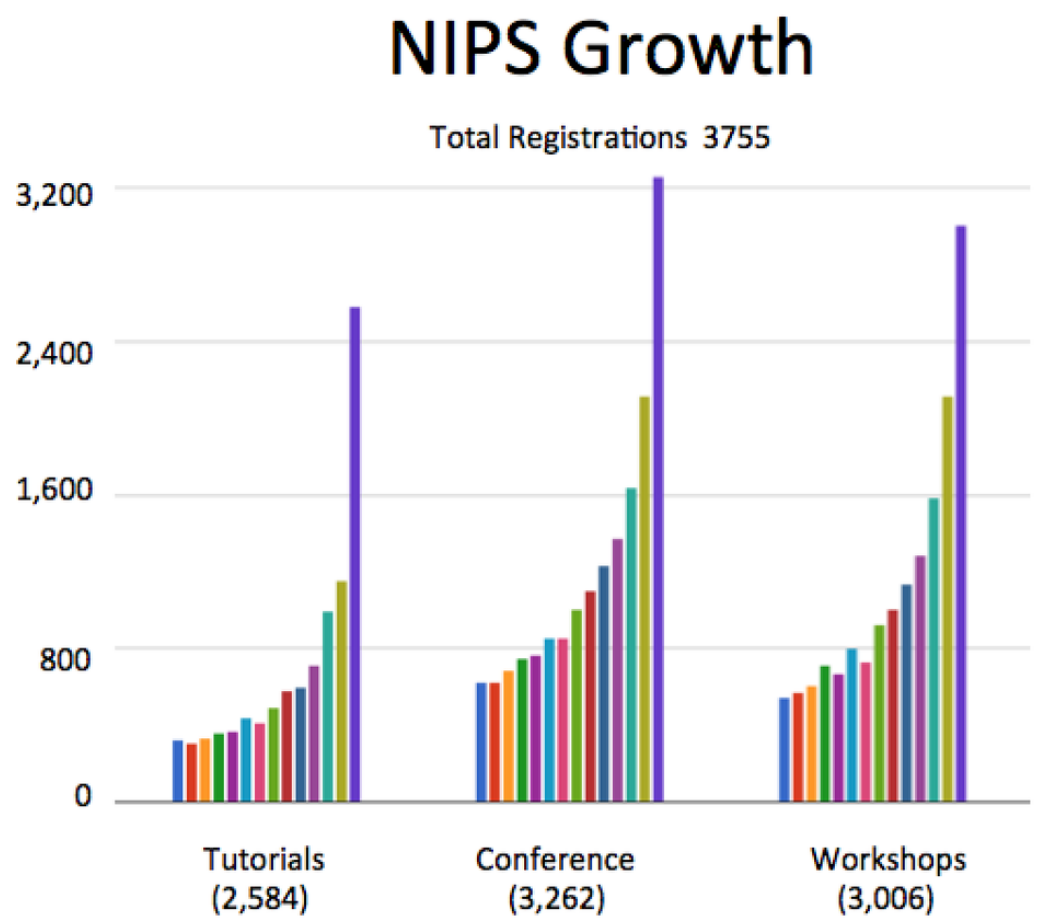

1.3. Types of Machine Learning

- Prediction
  - Supervised Learning: Given $\{(x_i, y_i)\}$ pairs, find a good understanding f: X→Y. Ex) Classification, Regression

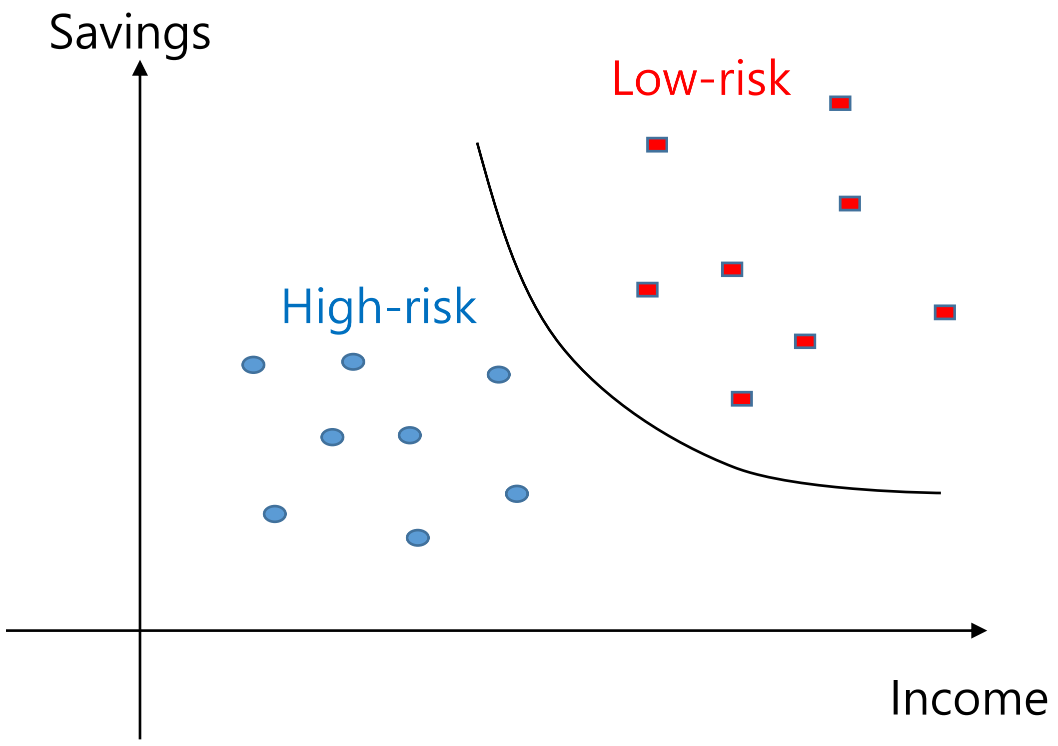

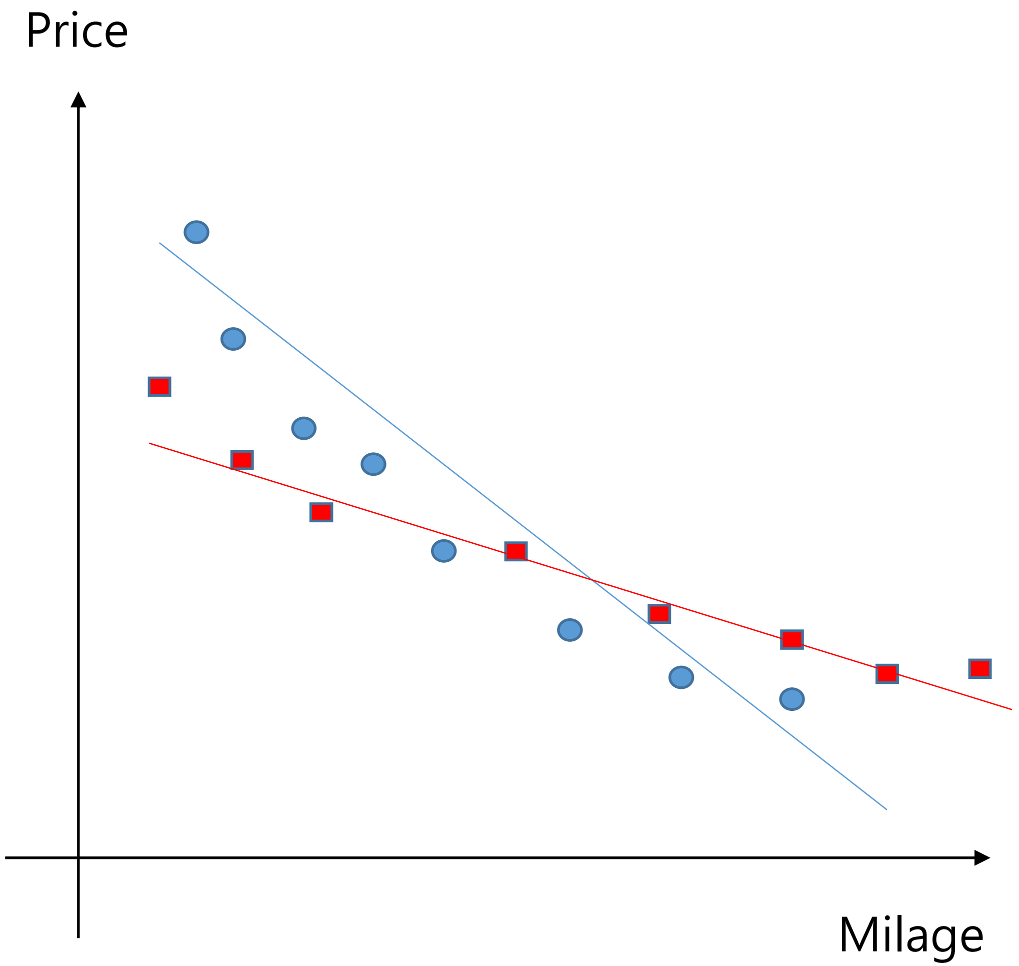

- Prediction
  - Supervised Learning: Given $\{(x_i, y_i)\}$ pairs, find a good understanding f: X→Y. Ex) Classification, Regression
  - Active Learning : Find f: X→Y, but $y_1, \ldots, y_N$ are initially hidden and there is a charge for each label you want revealed. The hope is that by intelligent adaptive querying, you can get away with significantly fewer labels than you would need 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.

Types of Machine Learning

- Understanding
  - Unsupervised Learning: Given $\{x_i\}$ find something interesting. Ex) Clustering, novelty-detection

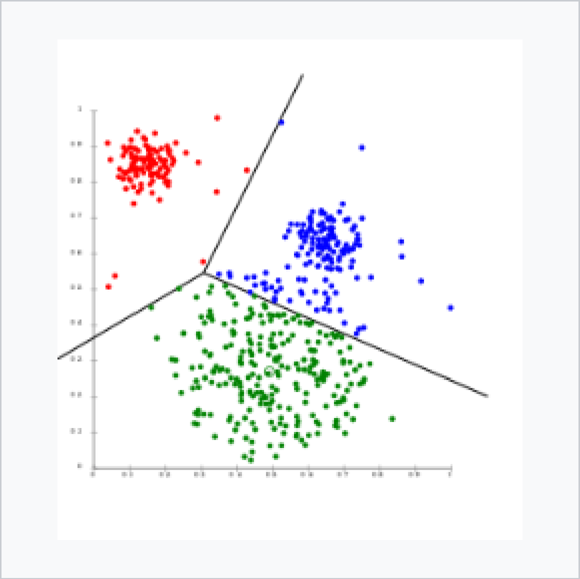

k-means separates data into Voronoi cells, which assumes equal-sized clusters (not adequate here)

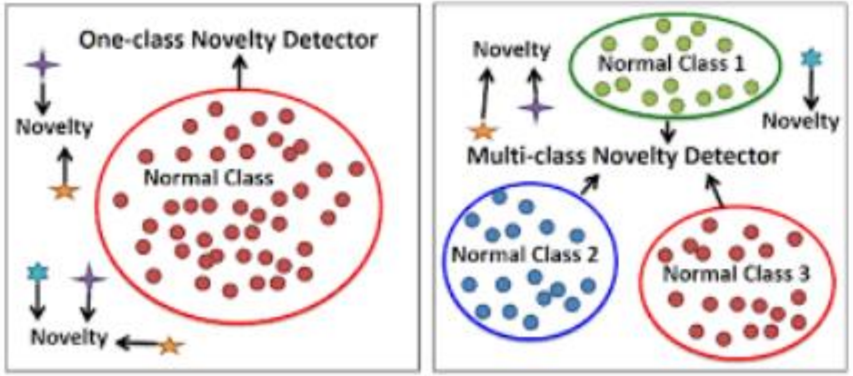

Source

Types of Machine Learning

- Policy(Control)
  - Reinforcement Learning: find f: P→A to max R (P: history of perception, A: action. R: long term reward)

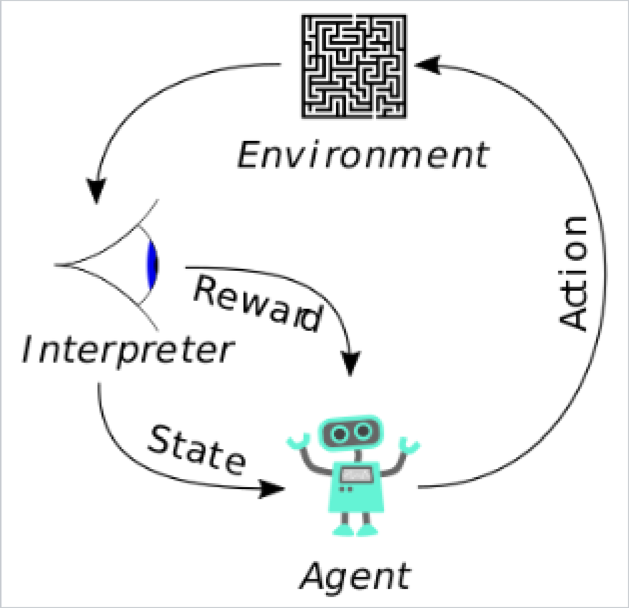

The typical framing of a Reinforcement Learning (RL) scenario: an agent takes actions in an environment, which is interpreted into a reward and a representation of the state, which are fed back into the agent.

Types of Machine Learning

- 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.

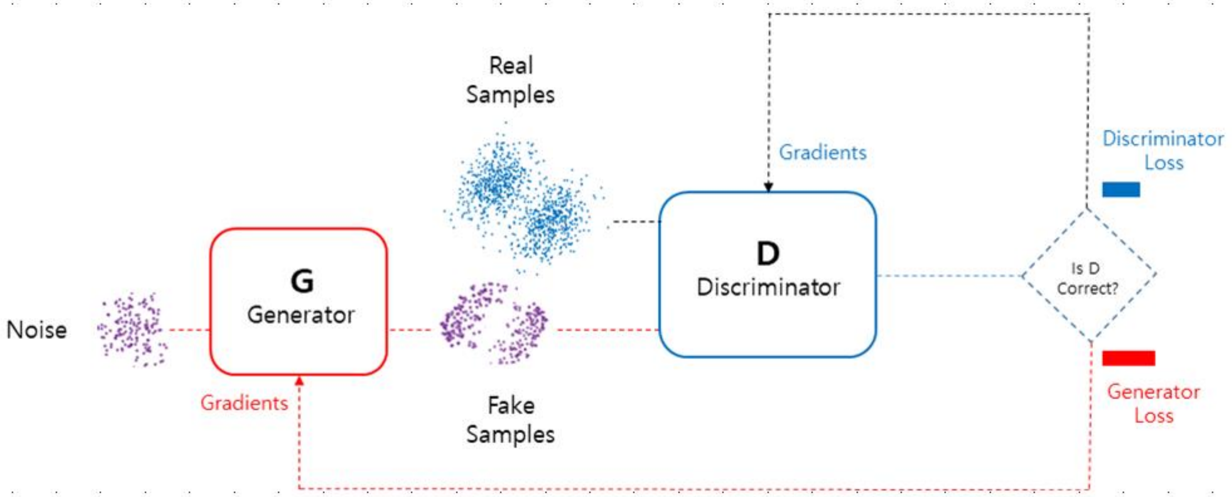

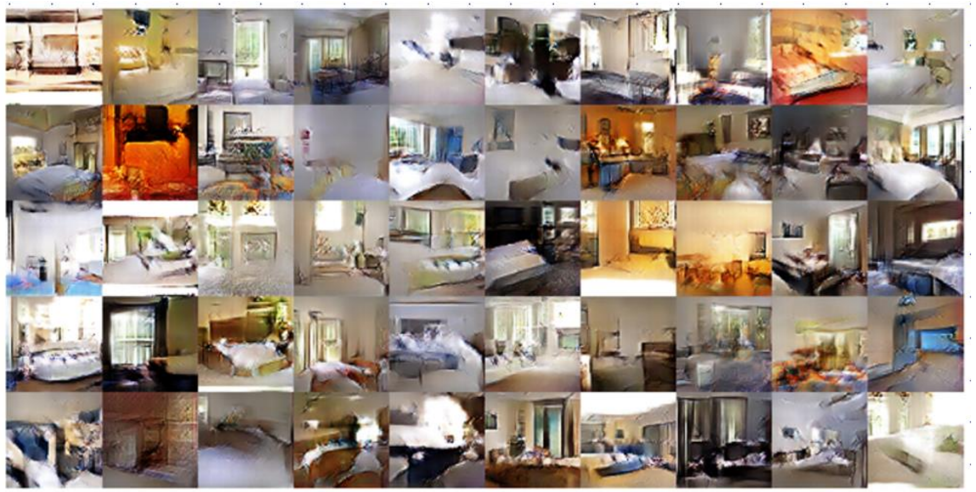

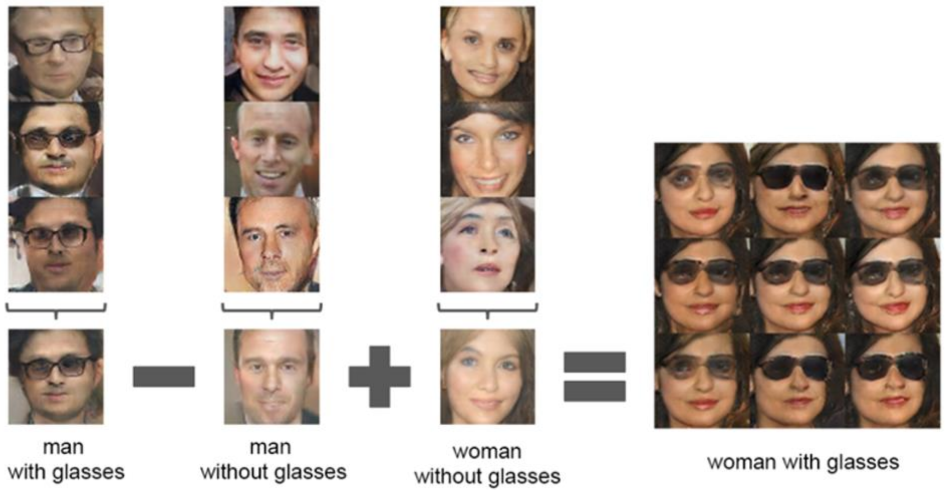

Applications

- **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

Applications

- 음성신호처리/음성인식
- 영상신호처리/패턴인식
- 디지털신호처리
- 자연어처리
- 로보틱스

Applications

- **인공비서: Amazon Echo**
- **가상 인간 (Virtual Human)**
- **의료인공지능 (AI in Healthcare)**
- **인공지능교육 (AI in Education)**
- **자율주행: Google Waymo**
- **스마트 홈 (Smart Home or Home Automation)**
- **스마트 그리드 (Smart Grid)**
- **스마트 공장 (Smart Factory or Smart Manufacturing)**
- **스마트 농장 (Smart Farm)**
- **인공지능과 법률 (AI in Law and Legal Practice)**
- **인공지능과 예술 (AI in Arts)**
- **휴머노이드 로봇 (Humanoid Robot AI)**
- **…….**