Machine Learning overview
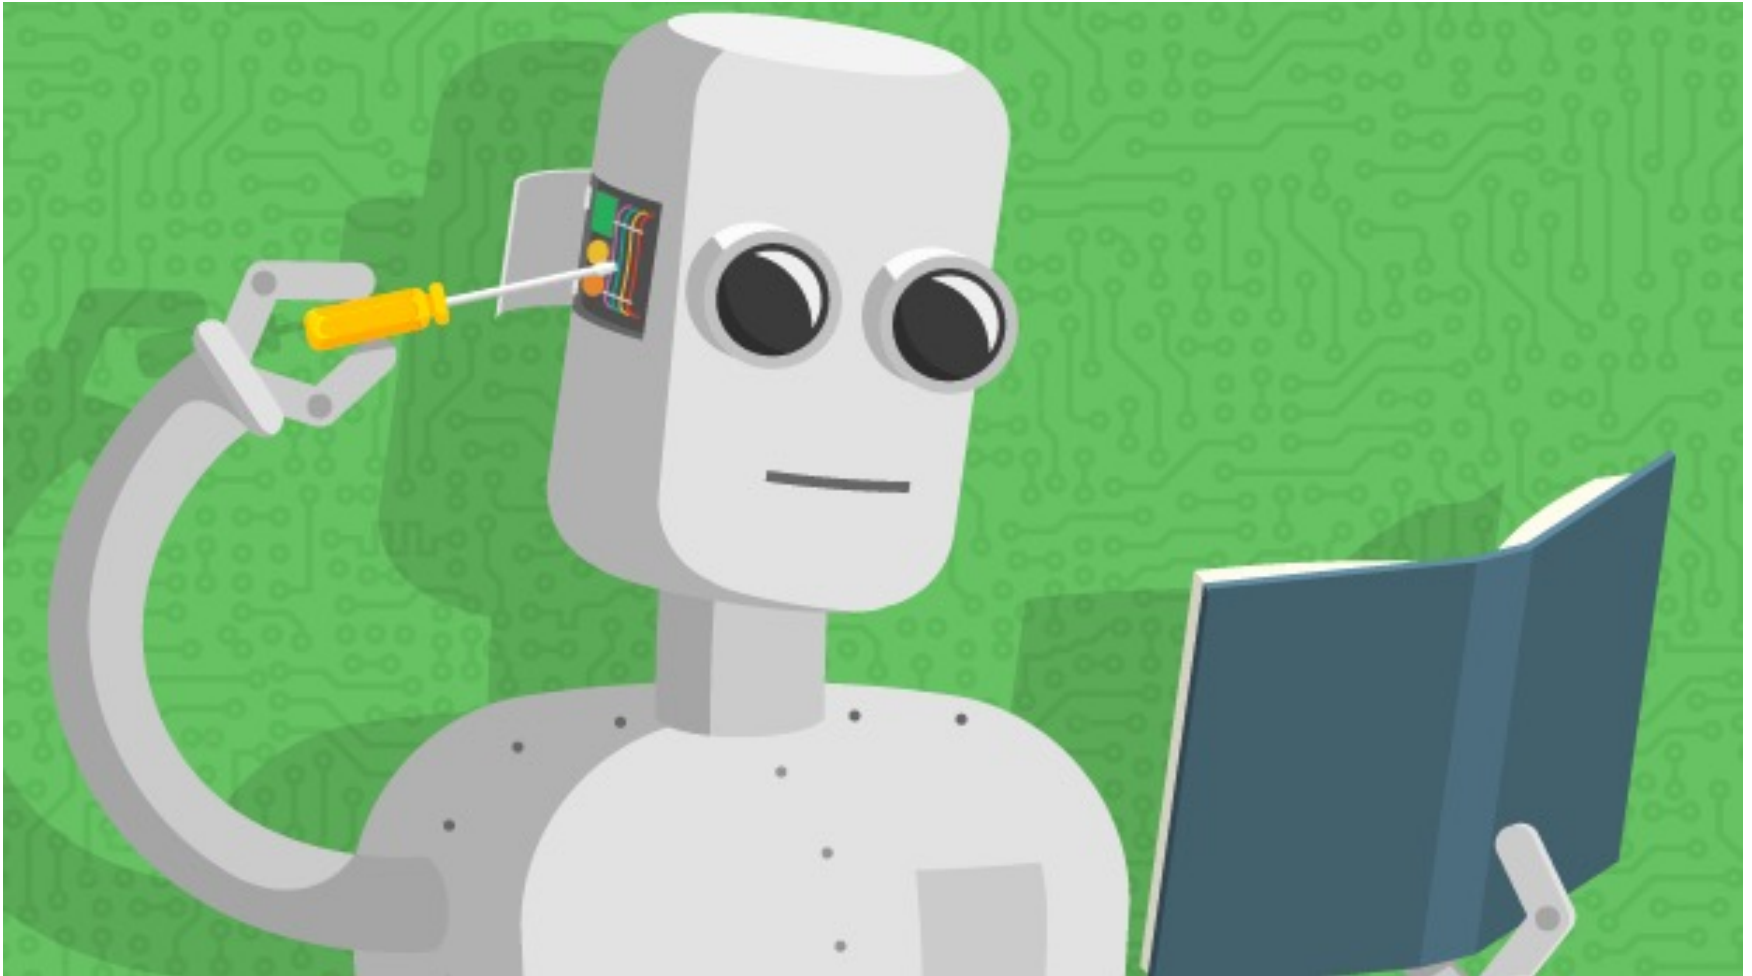## Chapter 19

What we will cover

- Some popular ML problems and algorithms
  - Take Machine Leaning, Data Science, NLP, Computer Vision for more
  - Use online resources & experiment on your own
- We will focus on when/how to use techniques and only touch on how/why they work
- Basic ML methodology and evaluation
- Use various platform for examples & demos (e.g., scikit-learn, Weka, TensorFlow, PyTorch)
  - Great for exploration and learning

What is learning?

- Learning denotes changes in a system that ... enable a system to do the same task more efficiently the next time – Herbert Simon

- Learning is constructing or modifying representations of what is being experienced – Ryszard Michalski

- Learning is making useful changes in our minds – Marvin Minsky

Why study learning?

- **Discover** new things or structure previously unknown
  - Examples: data mining, scientific discovery
- Fill in skeletal or **incomplete specifications** in a domain
  - Large, complex systems can't be completely built by hand & require dynamic updating to incorporate new info.
  - Learning new characteristics expands the domain or expertise and lessens the "brittleness" of the system
- Acquire **models directly from data** rather than by manual programming
- Build agents that can **adapt** to users, other agents, and their environment
- Understand and improve efficiency of **human learning**

AI and Learning Today

- 50s&60s: neural network learning popular

  Marvin Minsky did neural networks (NN) for his dissertation

- Mid 60s: replaced by paradigm of manually encoding & using symbolic knowledge

  Cf. [Perceptrons](#), Minsky & Papert book showed limitations of perceptron model of neural networks and helped kill off NN 🤔

- 90s: more data & processing power drove interest in statistical machine learning techniques & data mining

- Now: machine learning techniques & big data play biggest driver in almost all successful AI systems

  … and neural networks are the current favorite approach

seeAlso: [Timeline of machine learning](#)

Neural Networks 1960
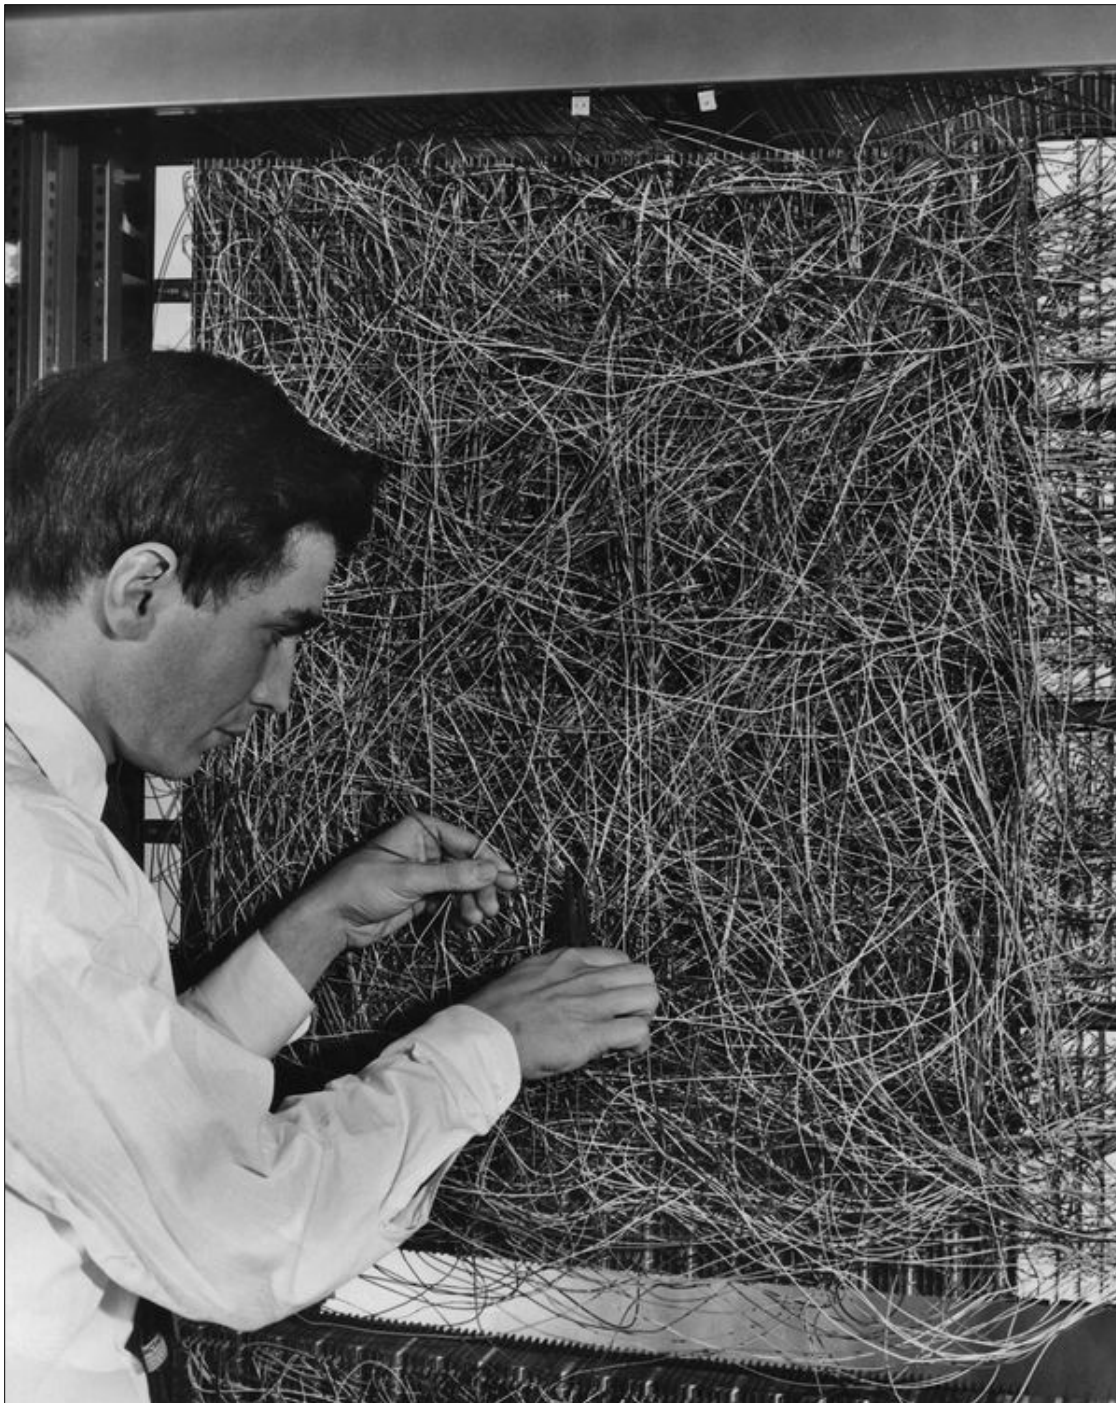
A man adjusting the random wiring network between the light sensors and association unit of scientist Frank Rosenblatt's Perceptron, or MARK 1 computer, at the Cornell Aeronautical Laboratory, Buffalo, New York, circa 1960. The machine is designed to use a type of artificial neural network, known as a perceptron.

AI Learning in the 1970s
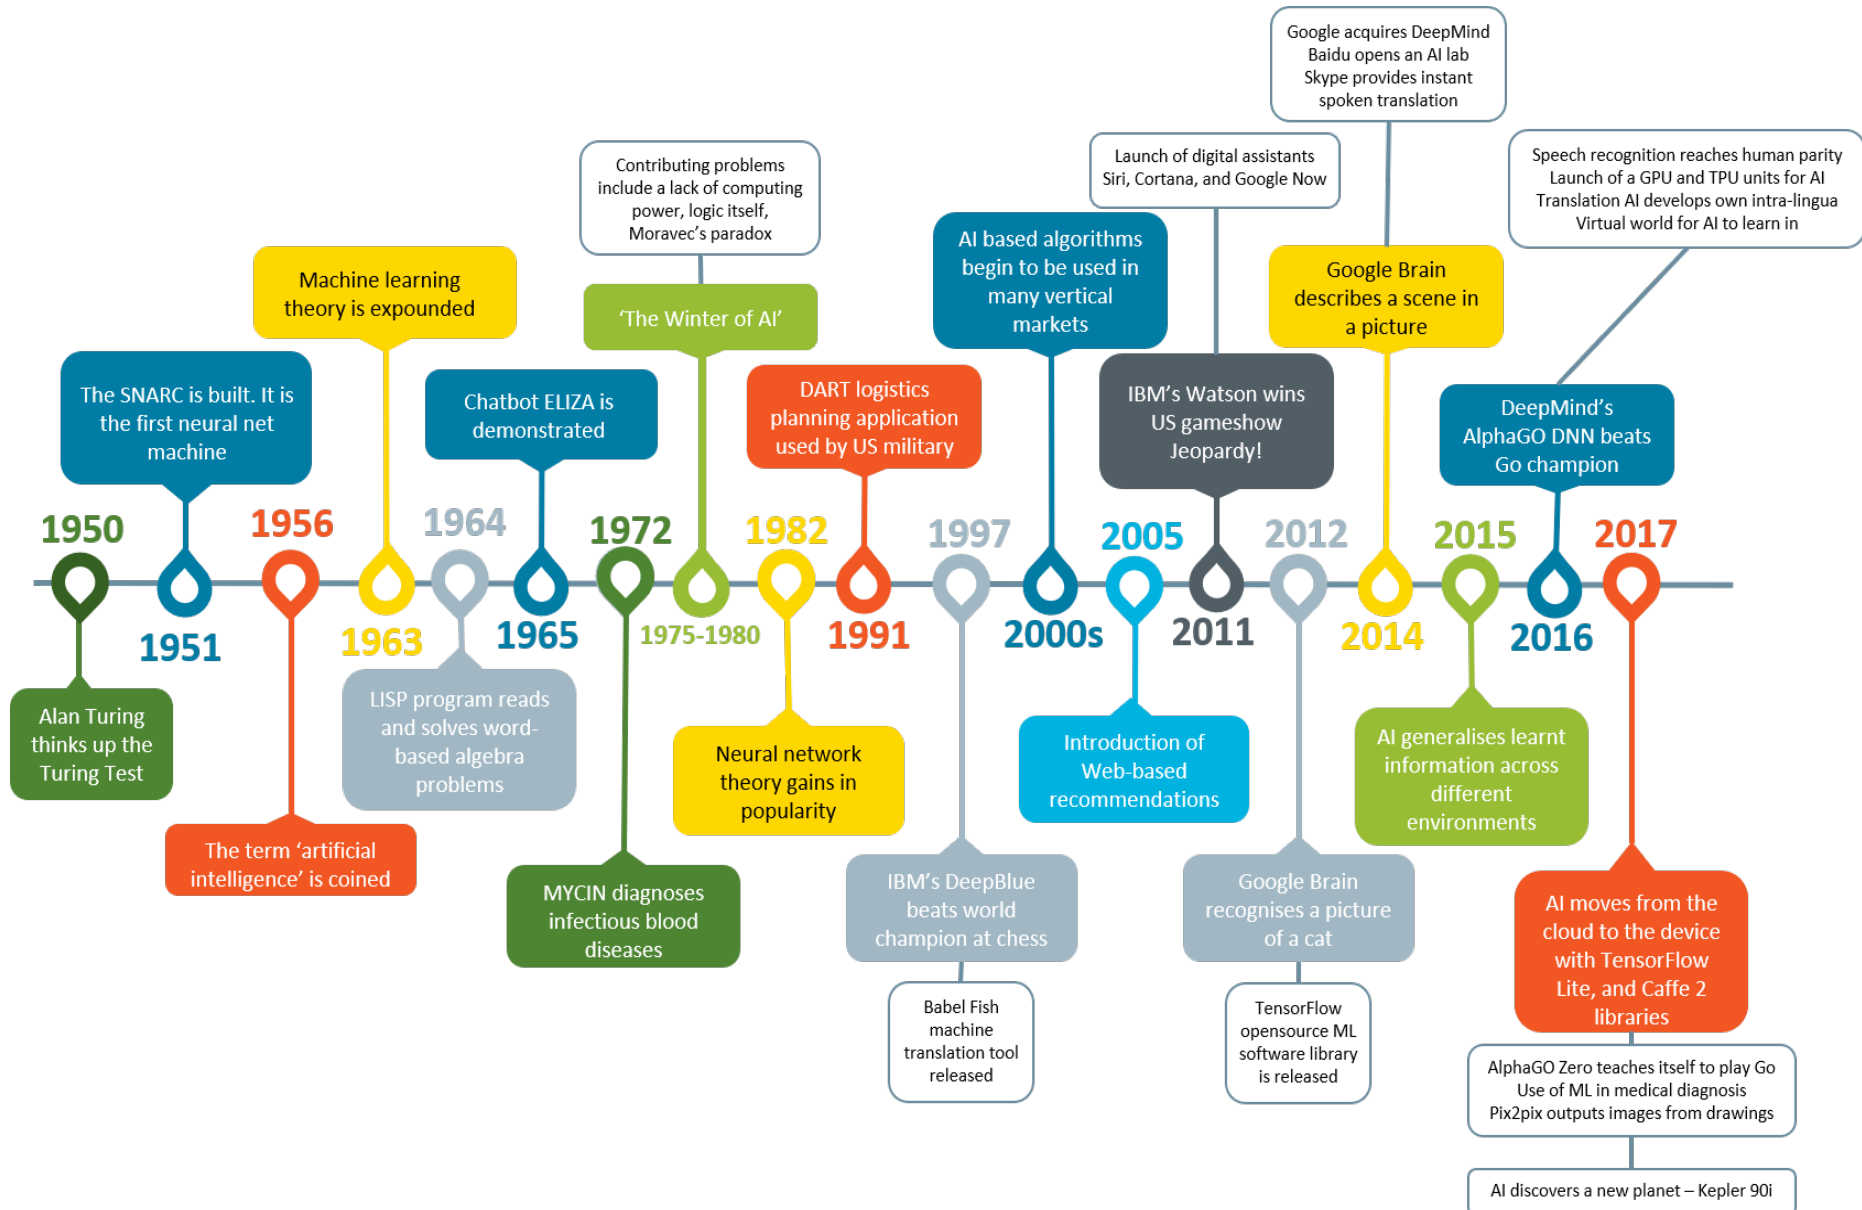
Winston's ARCH learning experiments …

Marvin Minsky
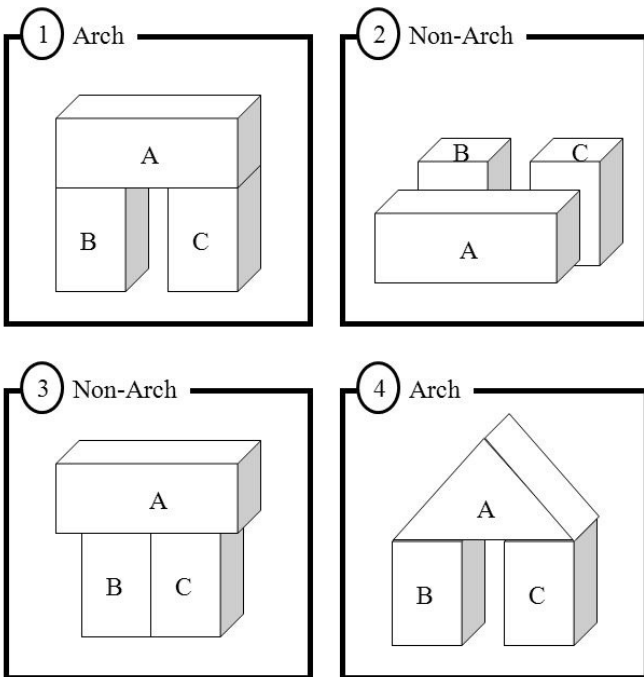
First we need to understand how to program machines to be intelligent in some way, then we can take on the task of getting the to learn how to do it.

AI timelines show Machine Learning beginning to dominate in the early 2000s

1950 — Alan Turing thinks up the Turing Test

1951 — The SNARC is built. It is the first neural net machine

1956 — The term 'artificial intelligence' is coined

1963 — Machine learning theory is expounded

1964 — LISP program reads and solves word-based algebra problems

1965 — Chatbot ELIZA is demonstrated

1972 — MYCIN diagnoses infectious blood diseases

1975-1980 — 'The Winter of AI'
Contributing problems include a lack of computing power, logic itself, Moravec's paradox

1982 — Neural network theory gains in popularity

1991 — DART logistics planning application used by US military

1997 — IBM's DeepBlue beats world champion at chess
Babel Fish machine translation tool released

2000s — Introduction of Web-based recommendations

2005 — AI based algorithms begin to be used in many vertical markets

2011 — IBM's Watson wins US gameshow Jeopardy!

2012 — Google Brain recognises a picture of a cat
TensorFlow opensource ML software library is released

2014 — Launch of digital assistants Siri, Cortana, and Google Now
Google acquires DeepMind
Baidu opens an AI lab
Skype provides instant spoken translation

2015 — Google Brain describes a scene in a picture
AI generalises learnt information across different environments

2016 — DeepMind's AlphaGO DNN beats Go champion
Speech recognition reaches human parity
Launch of a GPU and TPU units for AI
Translation AI develops own intra-lingua
Virtual world for AI to learn in

2017 — AI moves from the cloud to the device with TensorFlow Lite, and Caffe 2 libraries
AlphaGO Zero teaches itself to play Go
Use of ML in medical diagnosis
Pix2pix outputs images from drawings
AI discovers a new planet – Kepler 90i

Neural Networks 2018

Google's AIY Vision Kit is an intelligent camera that can recognize objects, detect faces and emotions. Download and use a variety of image recognition neural networks to customize the Vision Kit for your own creation. Included in the box: Raspberry Pi Zero WH, Pi Camera V2, Micro SD Card, Micro USB Cable, Push Button.

Currently $64.99 on Amazon

Machine Learning Successes

- Games: chess, go, poker
- Text sentiment analysis
- Email spam detection
- Recommender systems (e.g., Netflix, Amazon)
- Machine translation
- Speech understanding
- SIRI, Alexa, Google Assistant, …

- Autonomous vehicles
- Individual face recognition
- Understanding digital images
- Credit card fraud detection
- Showing annoying ads

Major Machine learning paradigms (1)

- **Rote**: 1-1 mapping from inputs to stored representation, learning by memorization, association-based storage & retrieval

- **Induction:** Use specific examples to reach general conclusions

- **Clustering**: Unsupervised discovery of natural groups in data

Major Machine learning paradigms (2)

- **Analogy:** Find correspondence between different representations

- **Discovery**: Unsupervised, specific goal not given

- **Genetic algorithms:** *Evolutionary* search techniques, based on *survival of the fittest*

- **Reinforcement:** Feedback (positive or negative reward) given at the end of a sequence of steps

- **Deep learning:** *artificial neural networks* with *representation learning* for ML tasks

Types of learning problems
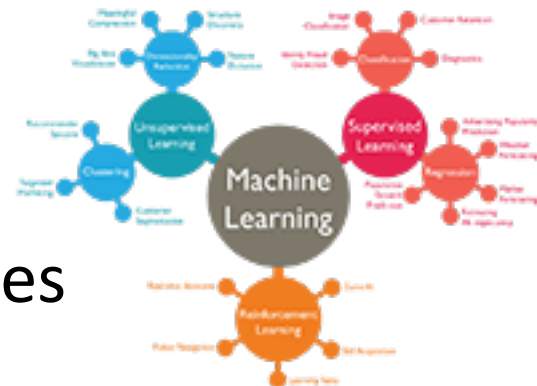
- **Supervised**: learn from training examples
  - Regression:
  - Classification: Decision Trees, SVM

- **Unsupervised**: learn w/o training examples
  - Clustering
  - Dimensionality reduction
  - Word embeddings

- **Reinforcement learning:** improve performance using feedback from actions taken

- Lots more we won't cover
  - Hidden Markov models, Learning to rank, Semi-supervised learning, Active learning, …

Machine Learning

Unsupervised Learning

Dimensionality Reduction
- Meaningful Compression
- Structure Discovery
- Big data Visualistaion
- Feature Elicitation

Clustering
- Recommender Systems
- Targetted Marketing
- Customer Segmentation

Supervised Learning

Classification
- Image Classification
- Customer Retention
- Idenity Fraud Detection
- Diagnostics

Regression
- Advertising Popularity Prediction
- Weather Forecasting
- Population Growth Prediction
- Market Forecasting
- Estimating life expectancy

Reinforcement Learning
- Real-time decisions
- Game AI
- Robot Navigation
- Skill Acquisition
- Learning Tasks

Machine Learning Problems

	Supervised Learning	Unsupervised Learning
Discrete	classification or categorization	clustering
Continuous	regression	dimensionality reduction

Supervised learning

- Given **training examples** of inputs & corresponding outputs, produce "correct" outputs for new inputs

- Two important scenarios:

  - **Classification:** outputs typically labels (goodRisk, badRisk); learn decision boundary to separate classes

  - **Regression:** aka *curve fitting* or *function approxima-tion*; Learn a *continuous* input-output mapping from examples, e.g., for a zip code, predict house sale price given its square footage

Unsupervised Learning

Given only *unlabeled* data as input, learn some sort of structure, e.g.:

- **Clustering**: group Facebook friends based on similarity of post texts and friends

- **Topic modeling**: Induce N topics and words most common in documents about each

- **Embeddings**: Find sets of words whose meanings are related (e.g., doctor, hospital)

- **Large Language Models:** Predict text that might follow a given test sequence (e.g., BERT, GPT-3)

Machine Learning

- ML's significance in AI has gone up and down over the last 75 years

  - Today it's **very** important for AI and data science

- Driving ML are three trends:

  - Cheaper and more powerful computing systems

  - Open-source ML tools & models (e.g., Weka, scikit-learn, TensorFlow, Huggingface, SpaCy, BERT …)

  - Availability of large amounts of data

- Understanding ML concepts and tools allow many to use them with success