


Machine Learning overview

Chapter 19



What we will cover

- Some popular ML problems and algorithms
 - Take classes on Machine Learning, Data Science, NLP, Computer Vision for more
 - Use online resources & experiment on your own
- 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 for his dissertation (1954)
- Mid 60s: replaced by paradigm of manually encoding & using symbolic knowledge
Cf. [Perceptrons](#), Minsky & Papert book showed limitations of perceptron neural networks & helped kill off NN for decades 🤔
- 90s: new algorithms (e.g., [backtracking](#)) and 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

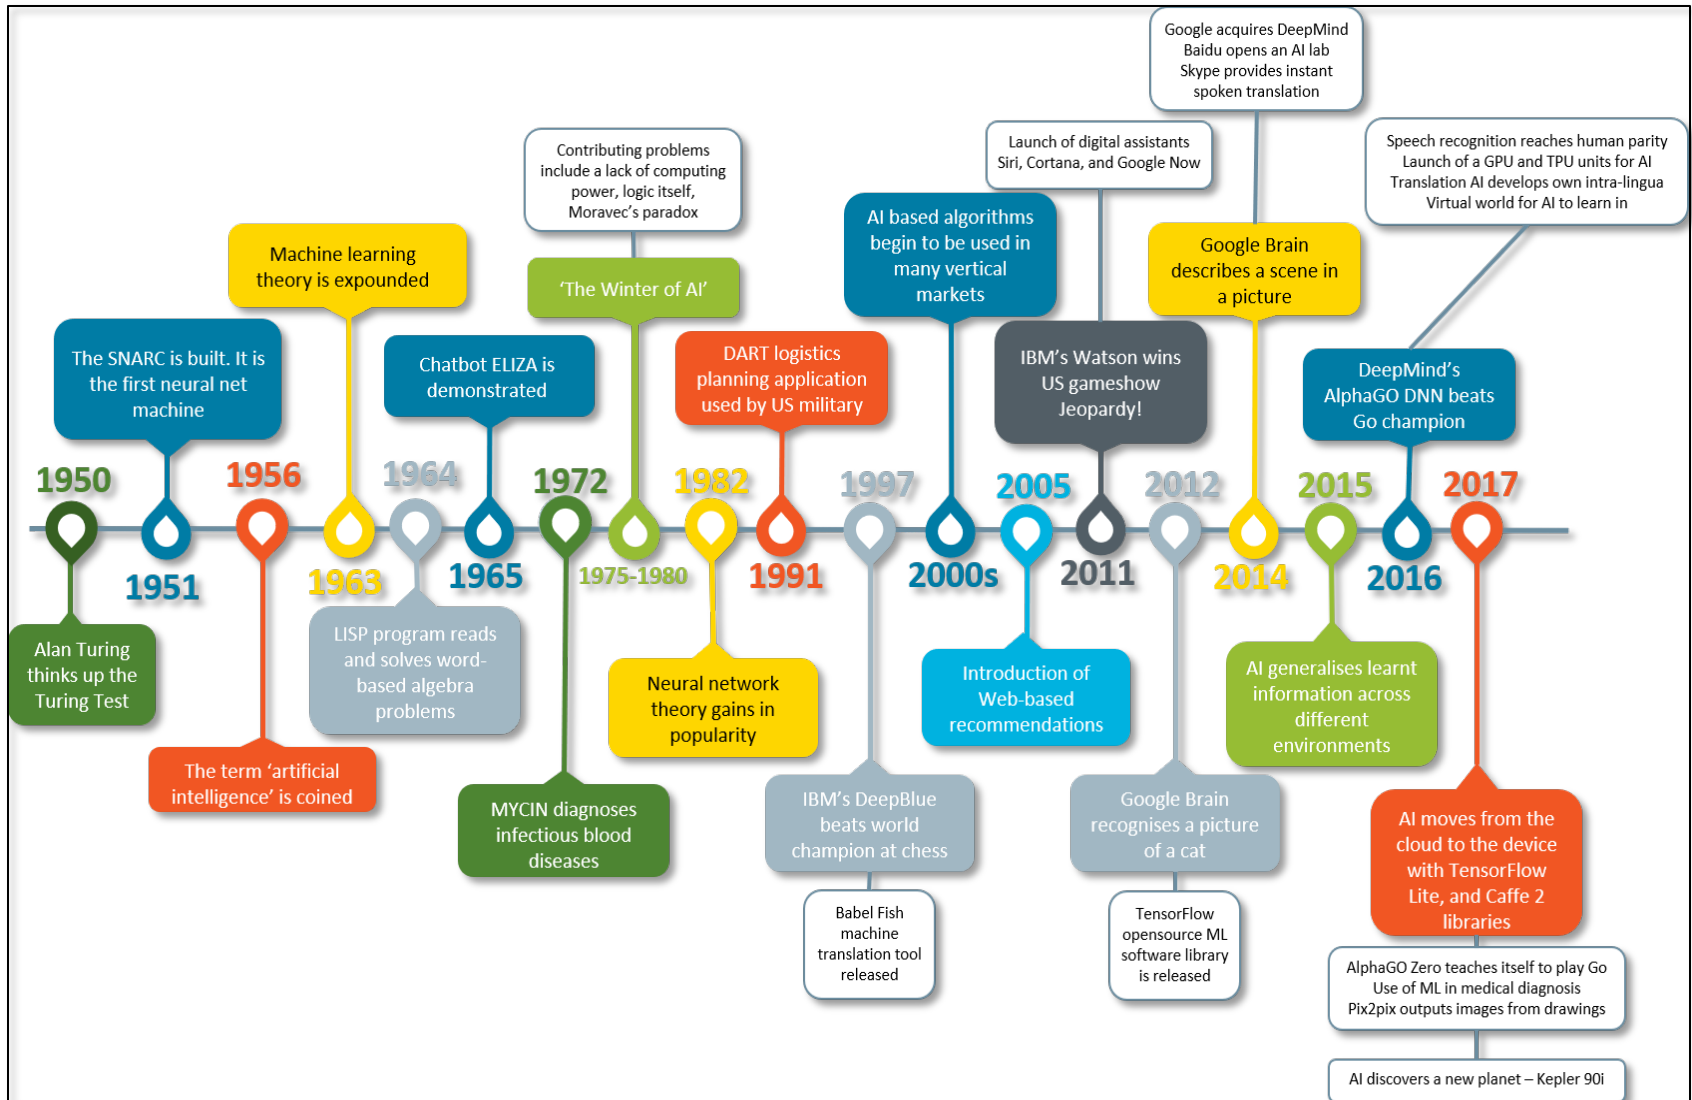
Marvin Minsky

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.



Early example of learning concepts from examples and non-examples (1970)

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



One of [many examples](#) you can find online

Neural Networks 2018-2022



Google's [AIY Vision Kit](#): an intelligent camera that can recognize objects, detect faces & 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.

Sold for ~\$30 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

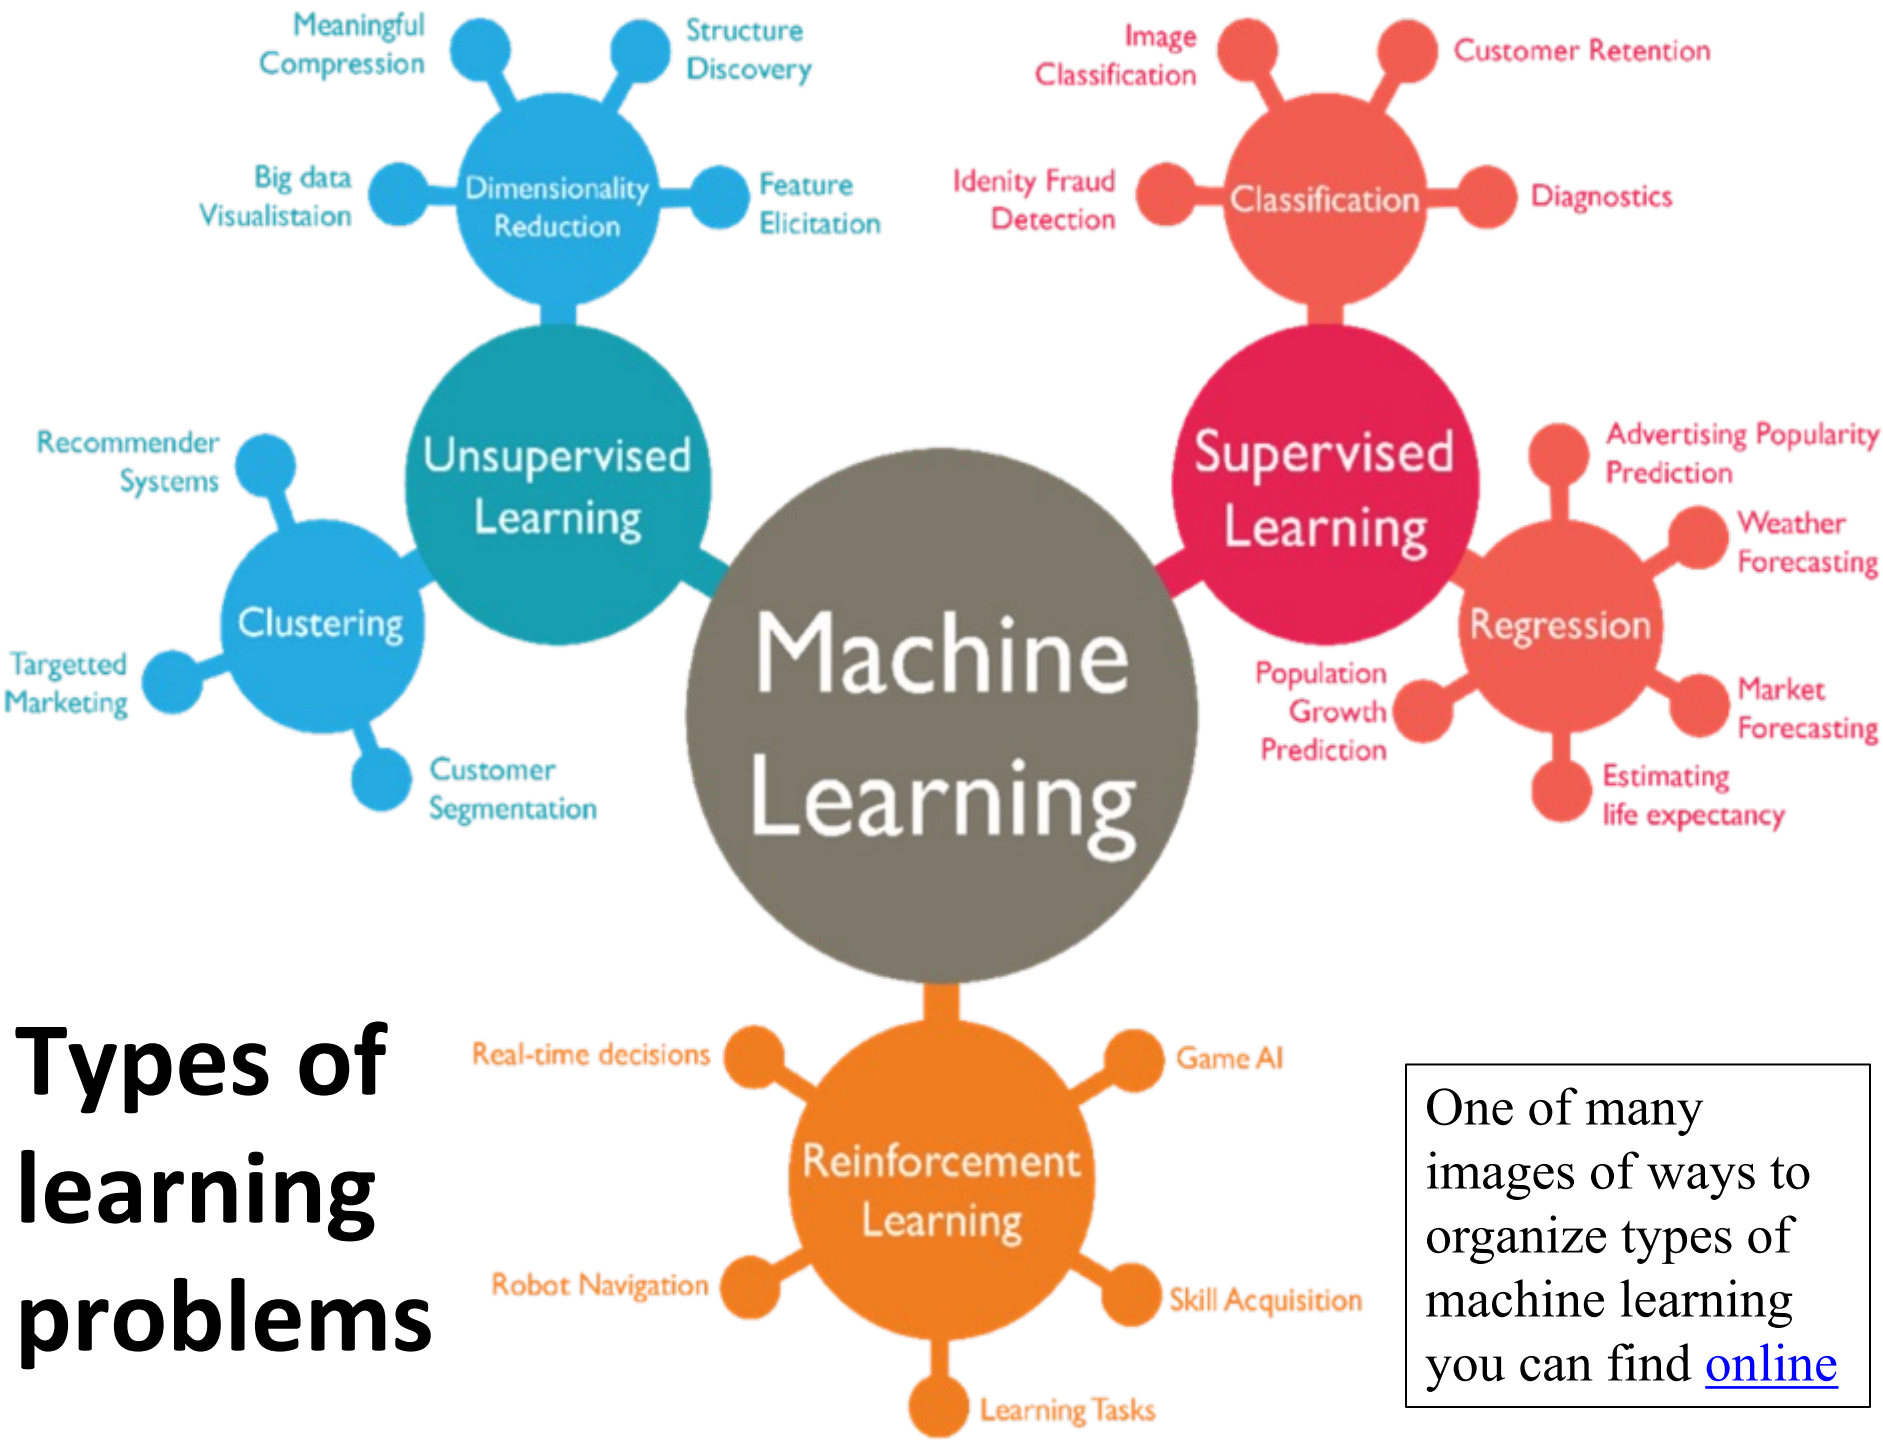
Many Machine learning paradigms (1)

- **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

Many Machine learning paradigms (2)

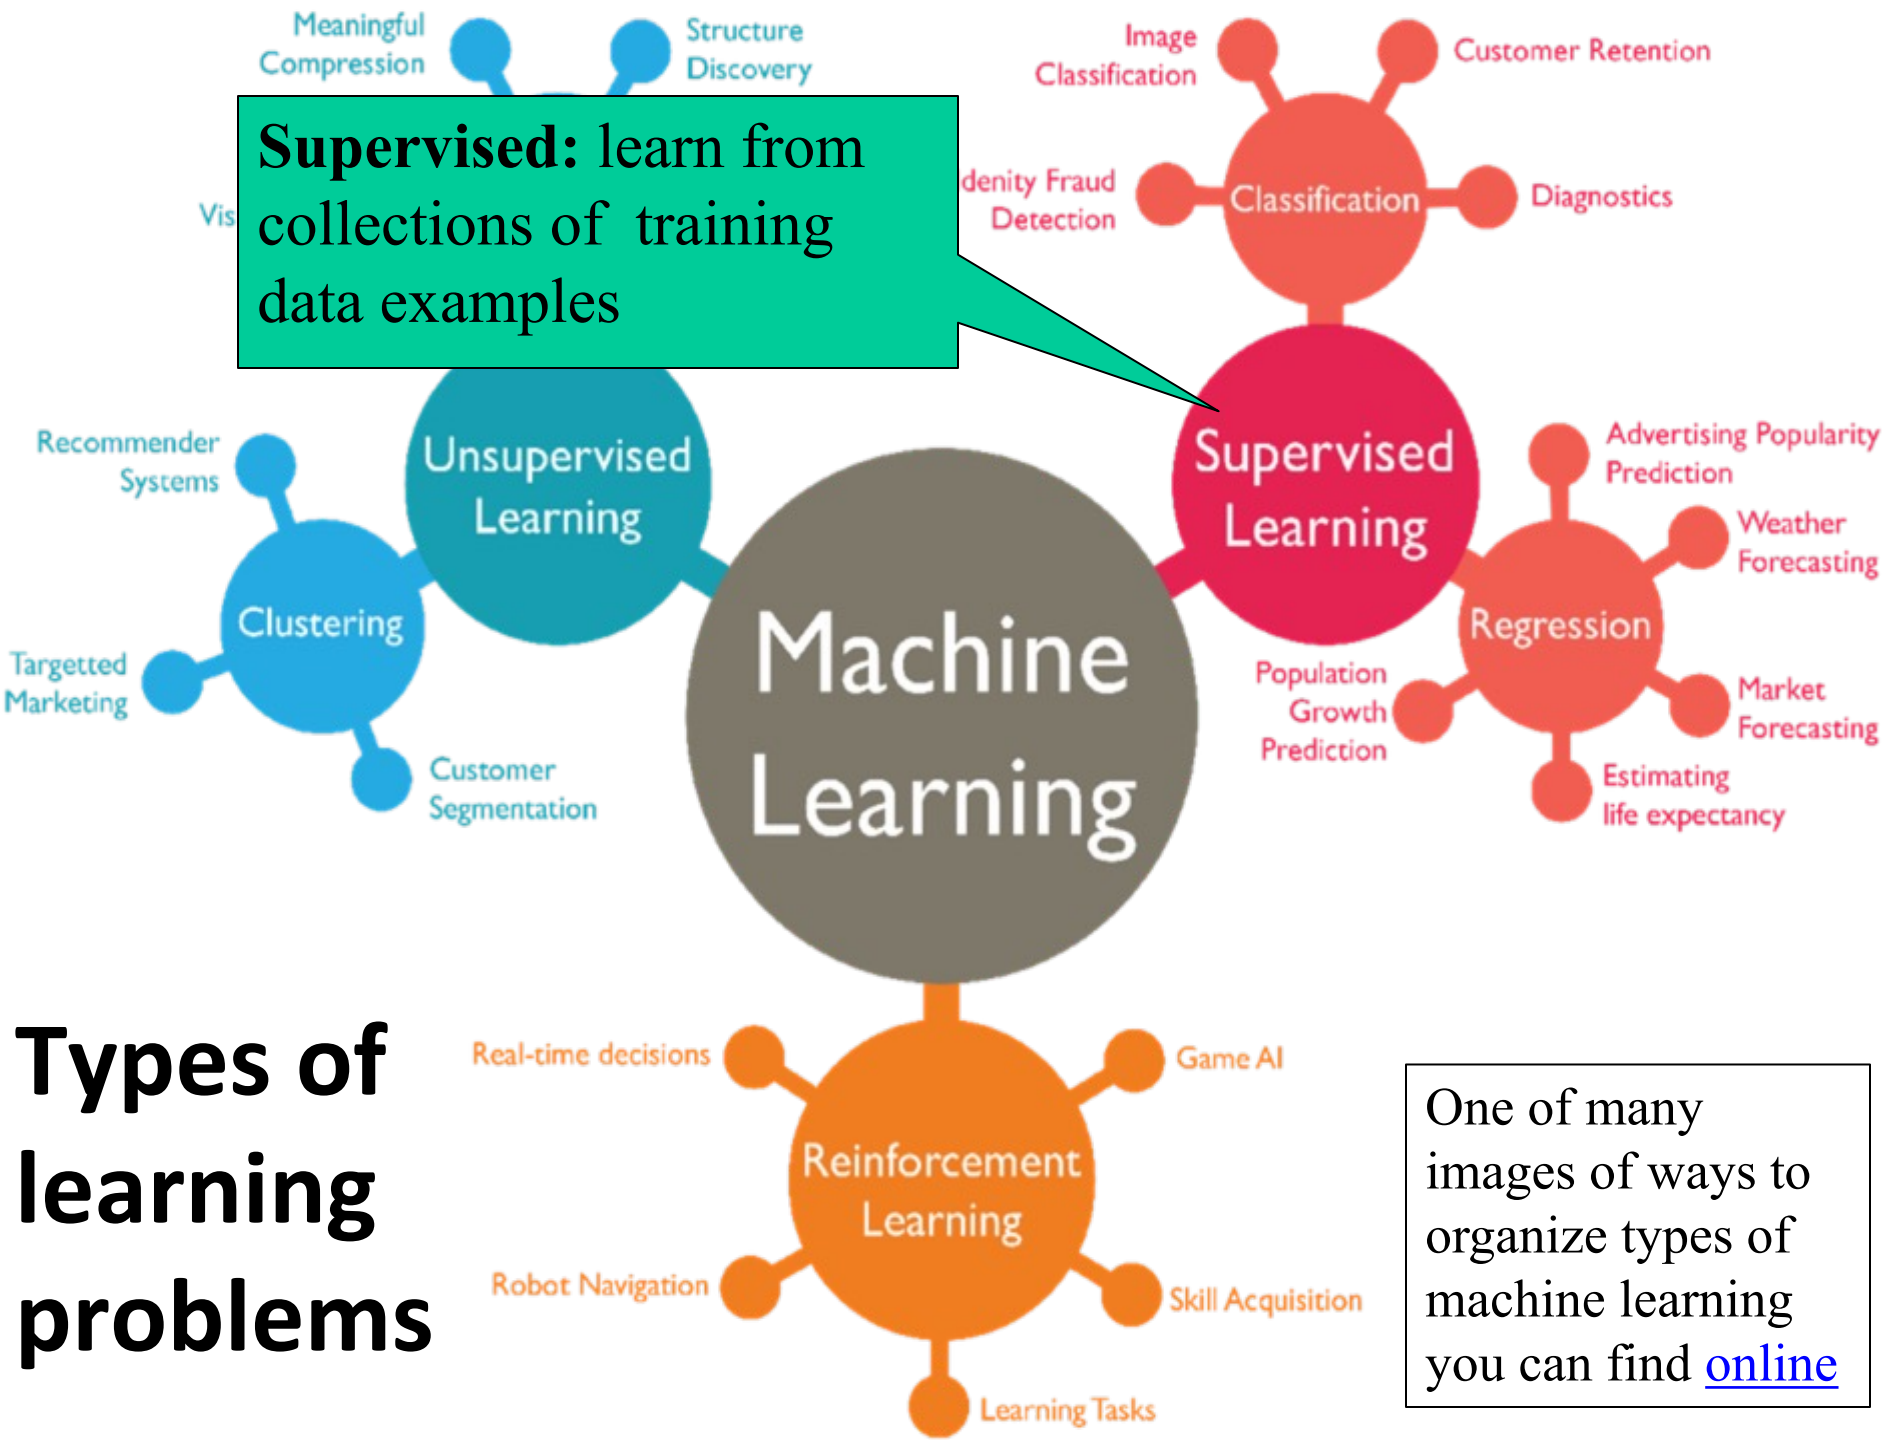
- **Analogy:** Find correspondence between different representations
- **Discovery:** Unsupervised, specific goal not given
- **Genetic algorithms:** *Evolutionary* search techniques, based on *survival of the fittest*
- **Reinforcement:** Feedback (positive or negative reward) given at the end of a sequence of steps
- **Deep learning:** *artificial neural networks* with *representation learning* for ML tasks

Types of learning problems



One of many images of ways to organize types of machine learning you can find [online](#)

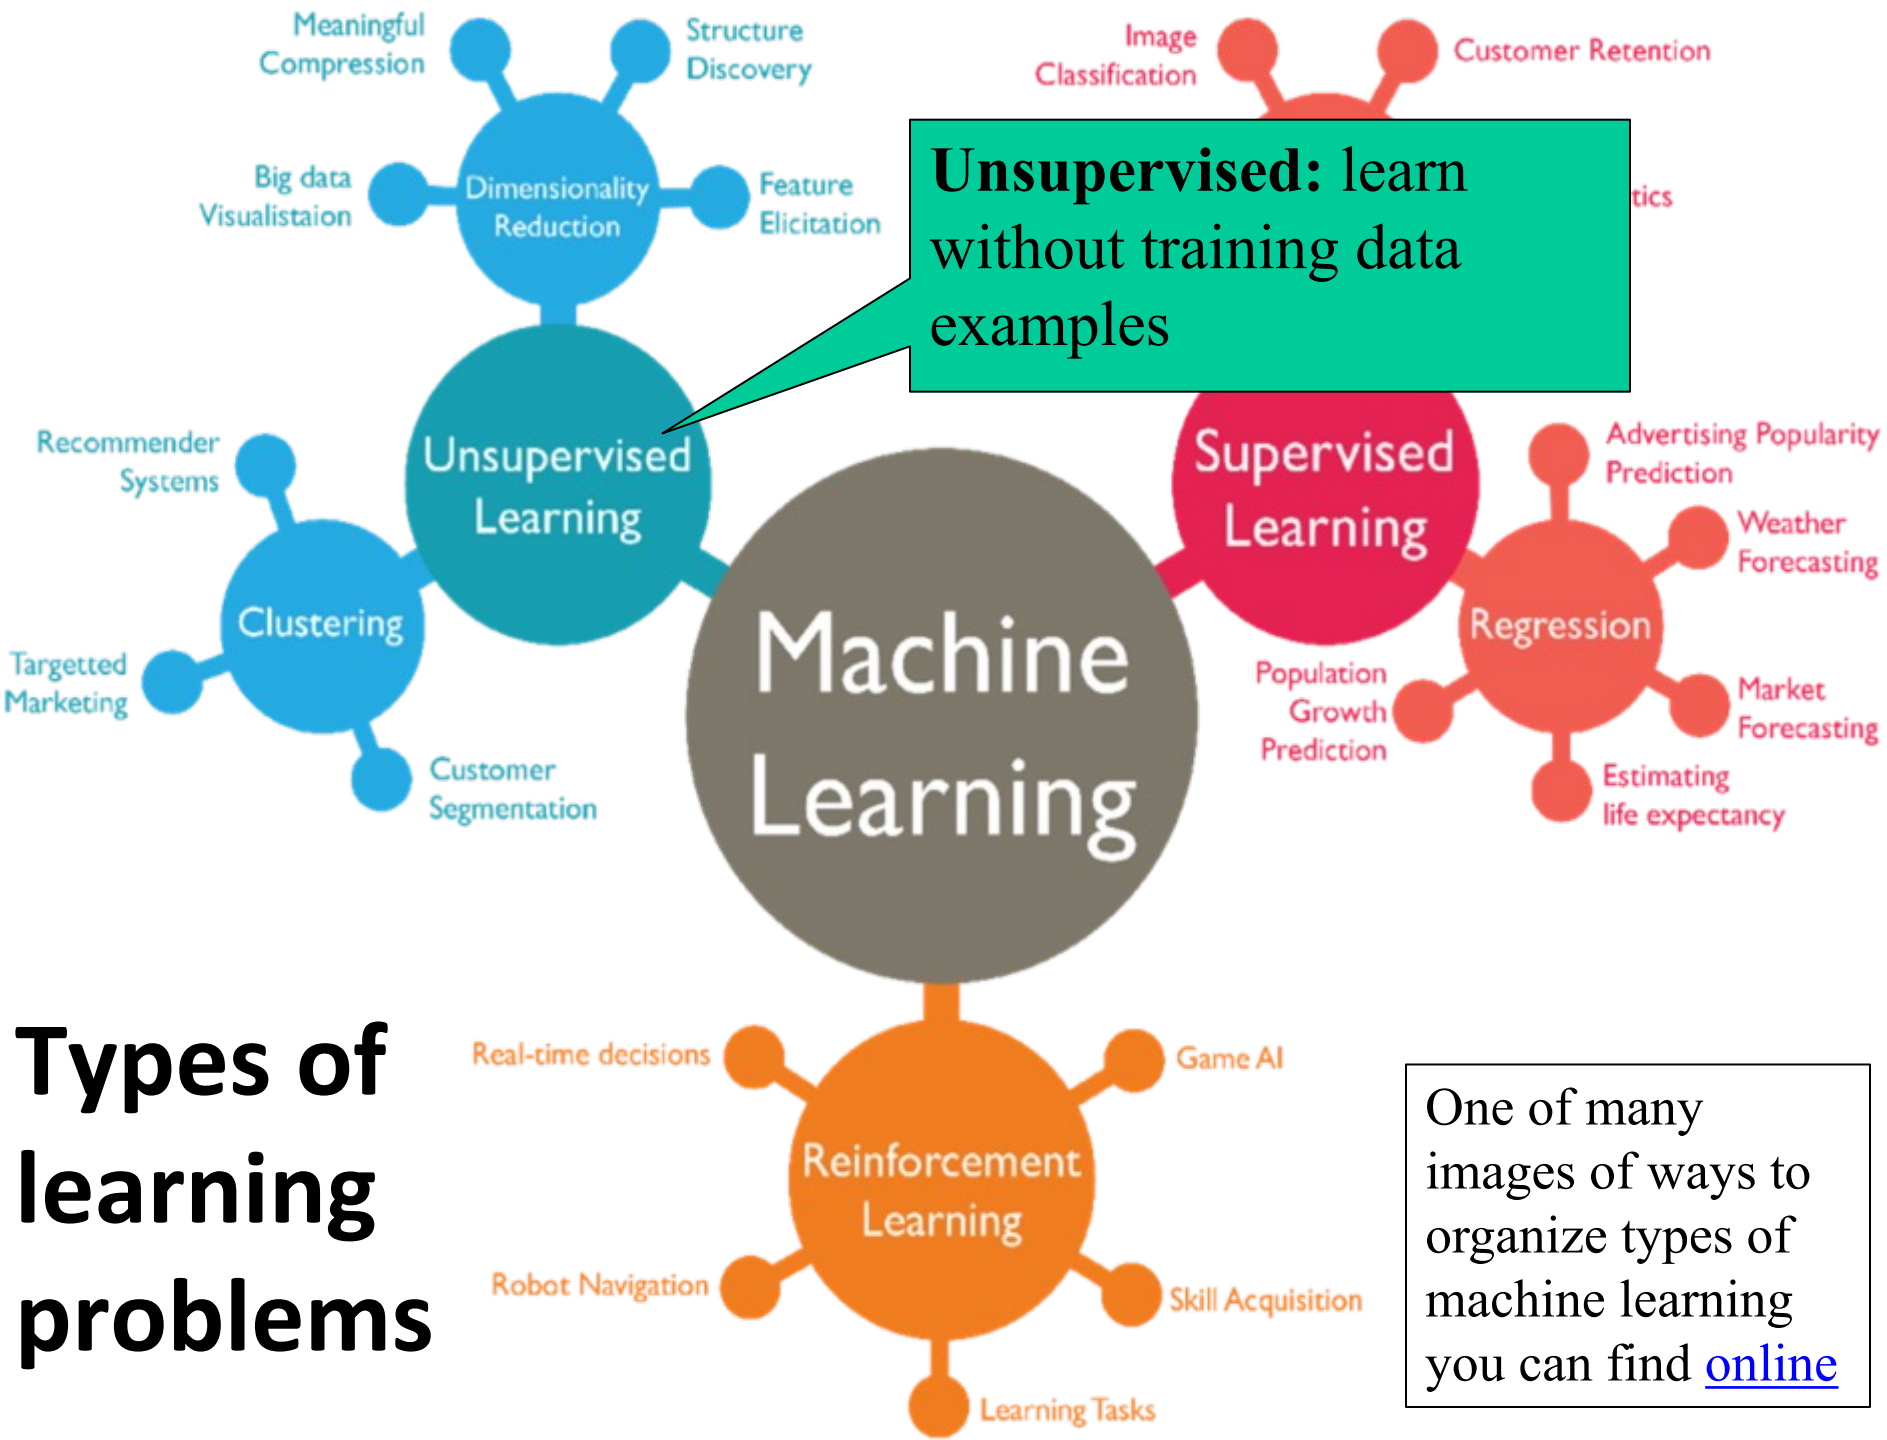
Supervised: learn from collections of training data examples

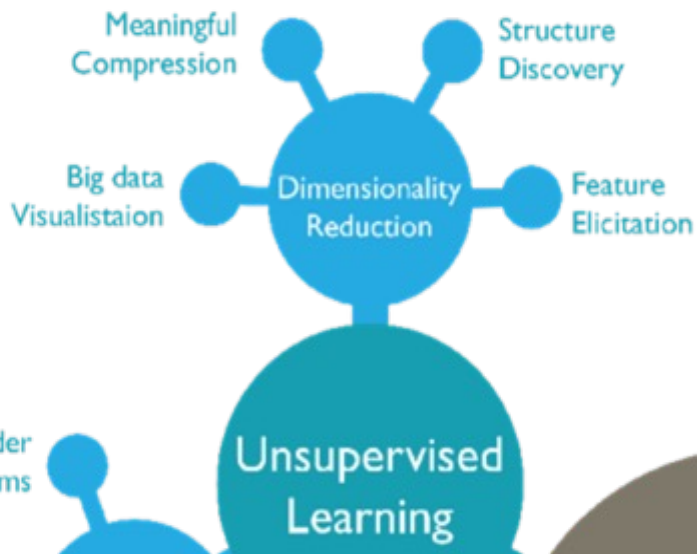


Types of learning problems

One of many images of ways to organize types of machine learning you can find [online](#)

Types of learning problems





Reinforcement learning:
improve performance
using feedback from
actions taken

Types of learning problems

One of many images of ways to organize types of machine learning you can find [online](#)

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
 - Still need other AI techniques for human-like intelligence
- Driving ML are three trends:
 - Cheaper and more powerful computing systems and availability of large amounts of data
 - Open-source ML tools & models (e.g., Weka, scikit-learn, TensorFlow, Huggingface, SpaCy, BERT ...)
 - Success of [foundation models](#) like GPT-4 for supporting useful applications
- Understanding ML concepts and tools allow many to use them with success