Lecture 1: Course Introduction
CS480/680 Intro to Machine Learning

2023-1-10

Pascal Poupart
David R. Cheriton School of Computer Science

UNIVERSITY OF
WATERLOO

Outline

- Introduction to Machine Learning
- Course website and logistics

Instructor

- Pascal Poupart (Professor and CIFAR AI Chair)
  - 20+ years experience in Machine Learning

Machine Learning

- Traditional computer science
  - Program computer for every task

- New paradigm
  - Provide examples to machine
  - Machine learns to accomplish tasks based on examples

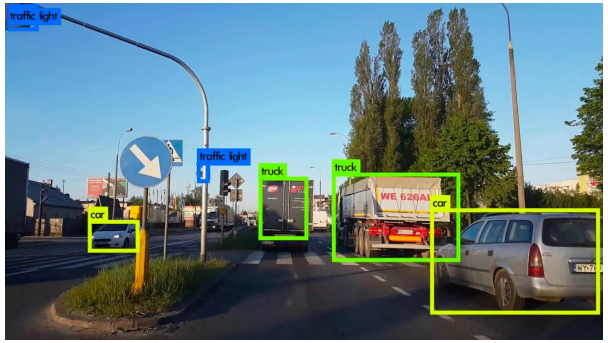

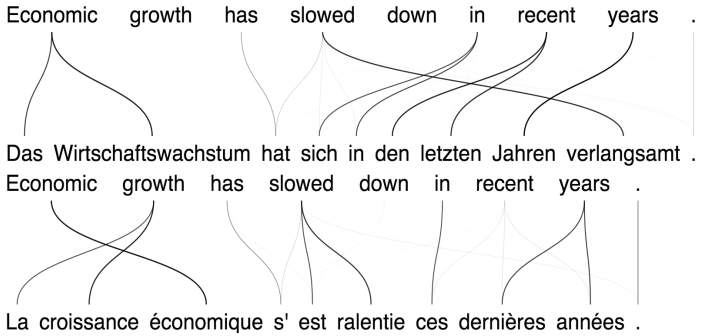

Definitions

- Arthur Samuel (1959): **Machine learning** is the field of study that gives computers the ability to learn without being explicitly programmed.

- Tom Mitchell (1998): A computer program is said to **learn** from experience E with respect to some class of tasks T and performance measure P, if its performance at tasks in T, as measured by P, improves with experience E.

Three Categories

Supervised learning

Reinforcement learning

Unsupervised learning

Supervised Learning

- Example: digit recognition (postal code)

- Simplest approach: memorization

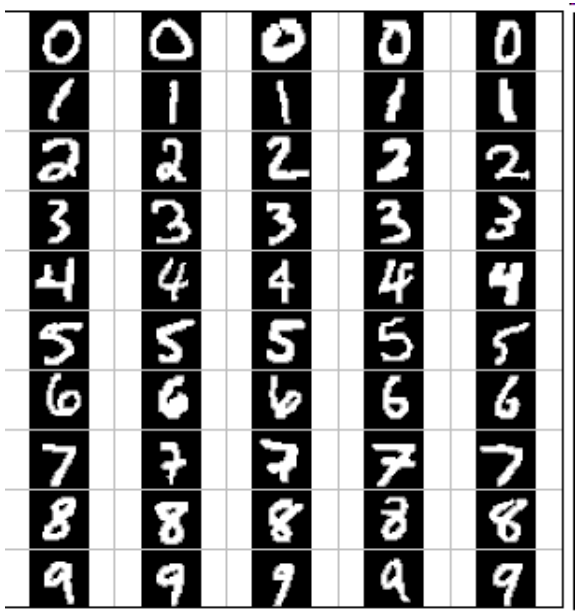

Supervised Learning

- Nearest neighbour:

More Formally

- Inductive learning (for supervised learning):
  - Given a training set of examples of the form $(x, f(x))$
    - x is the input, $f(x)$ is the output
  - Return a function h that approximates f
    - h is called the hypothesis

Prediction

- Find function h that fits f at instances x

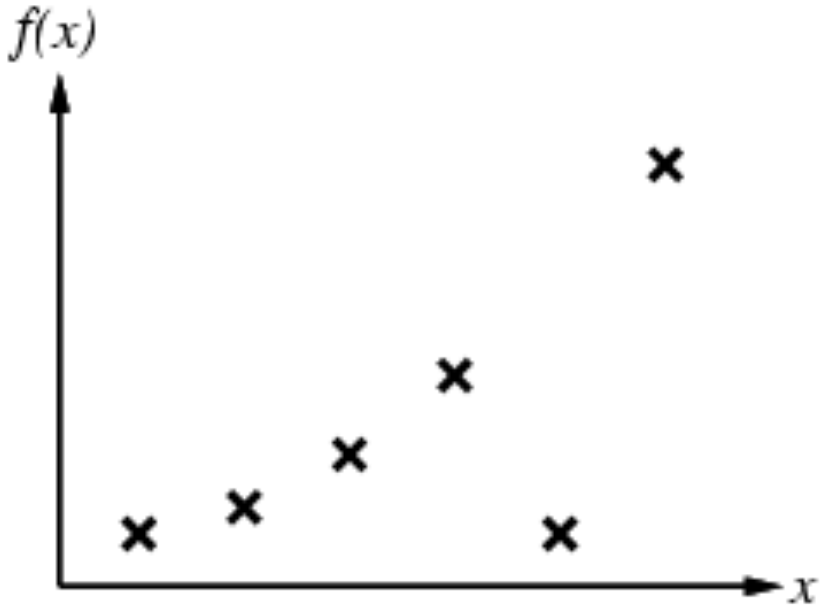

Prediction

- Find function h that fits f at instances x

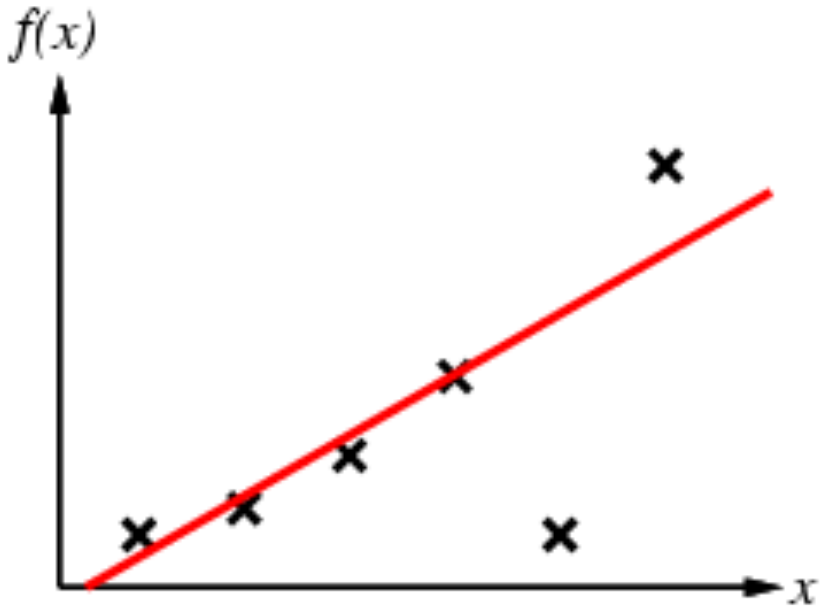

Prediction

- Find function h that fits f at instances x

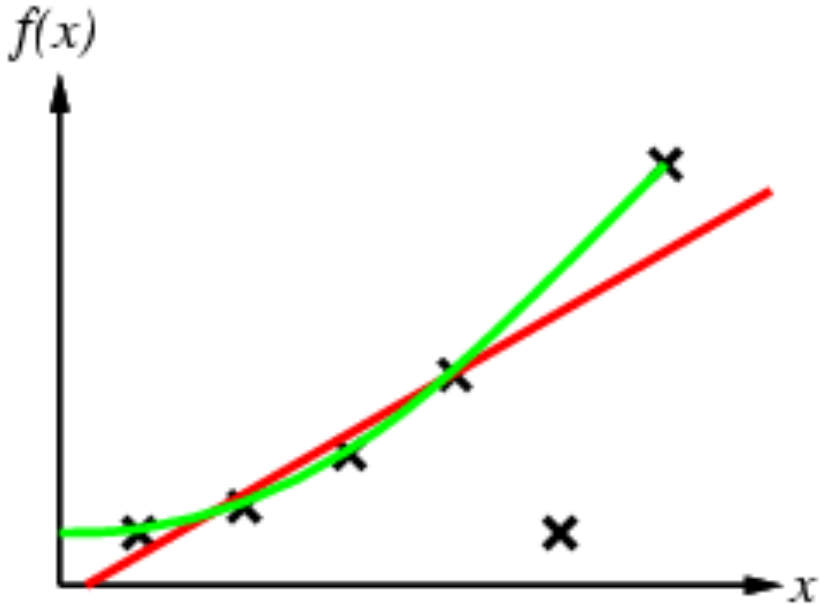

Prediction

- Find function h that fits f at instances x

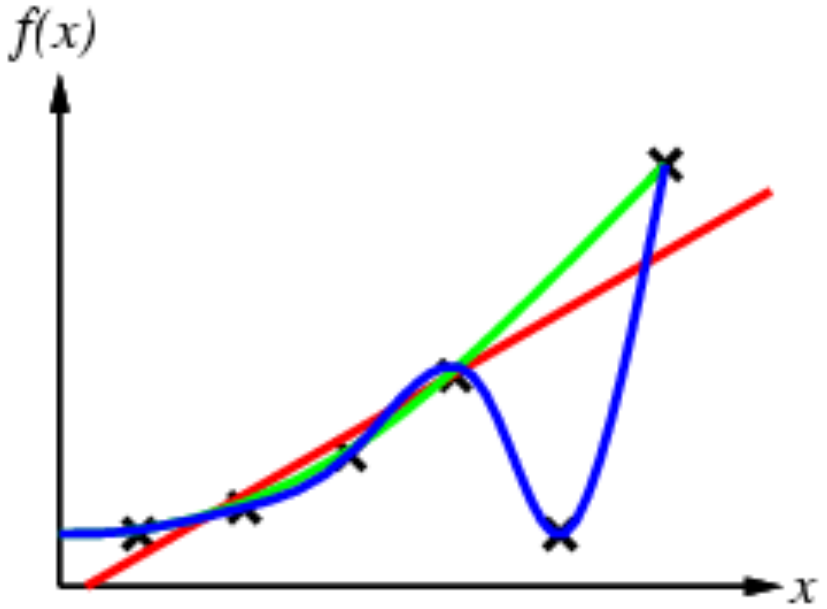

Prediction

- Find function h that fits f at instances x

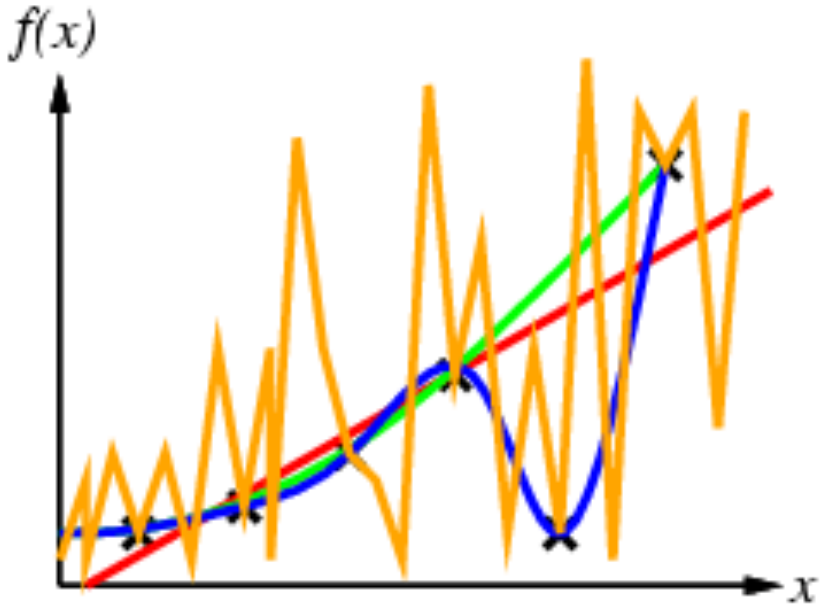

Generalization

- Key: a good hypothesis will generalize well (i.e., predict unseen examples correctly)

- Ockham's razor: prefer the simplest hypothesis consistent with data

ImageNet Classification

- 1000 classes
- 1 million images

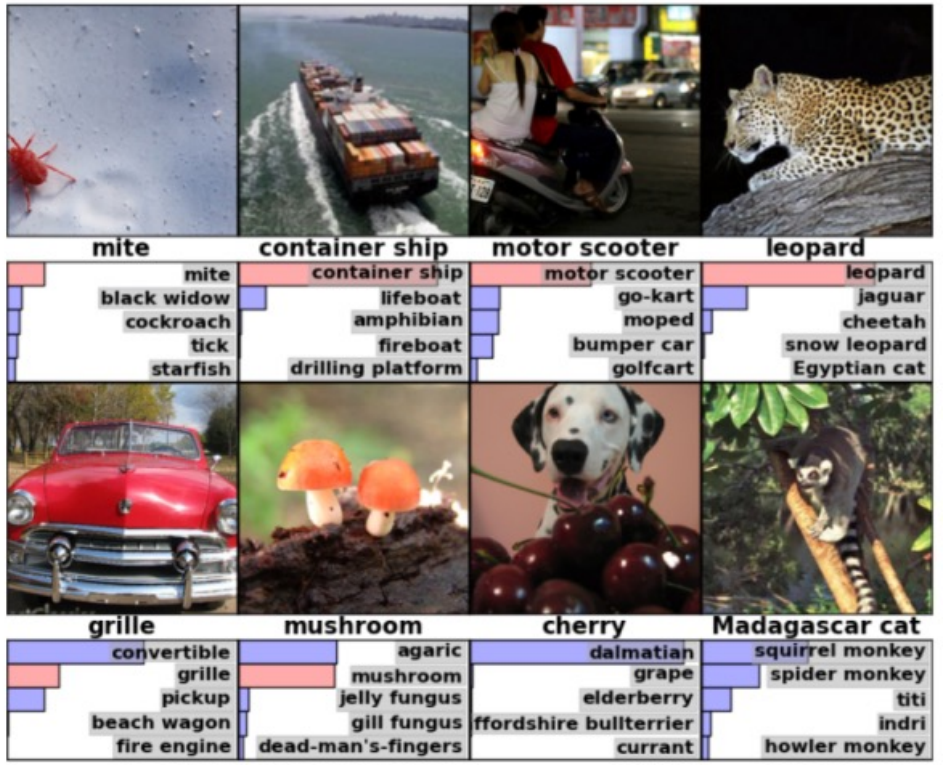

- Deep neural networks (supervised learning)

Unsupervised Learning

- Output is not given as part of training set
- Find model that explains the data
  - E.g. clustering, compressed representation, features, generative models

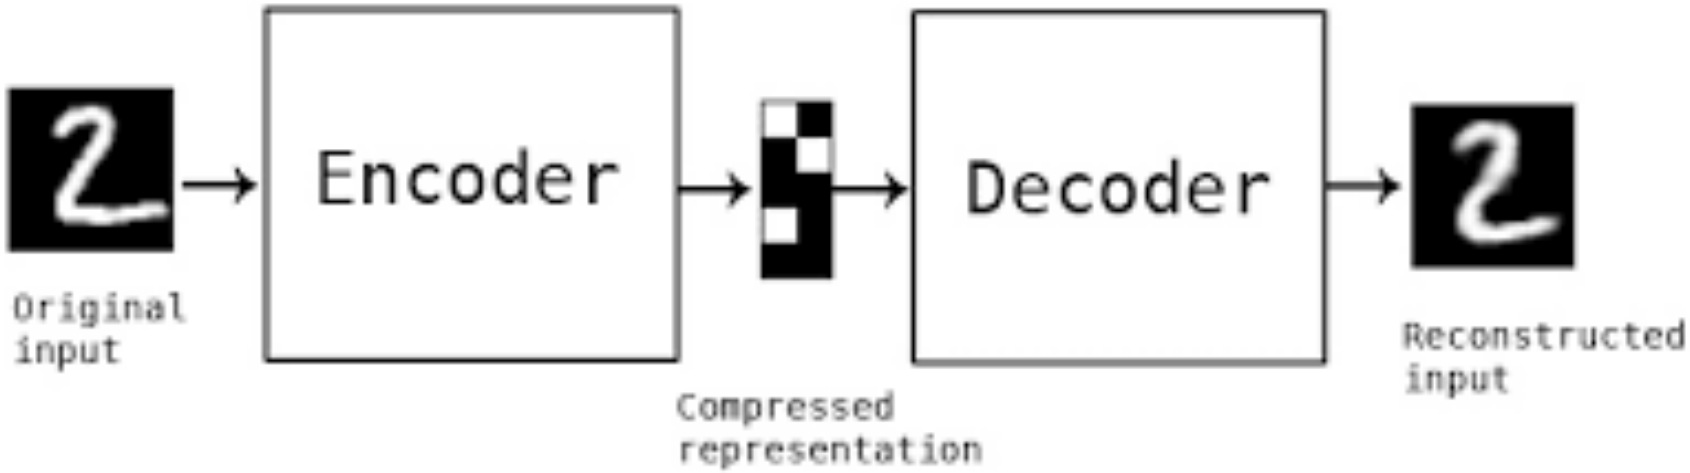

Unsupervised Feature Generation

- Encoder trained on large number of images

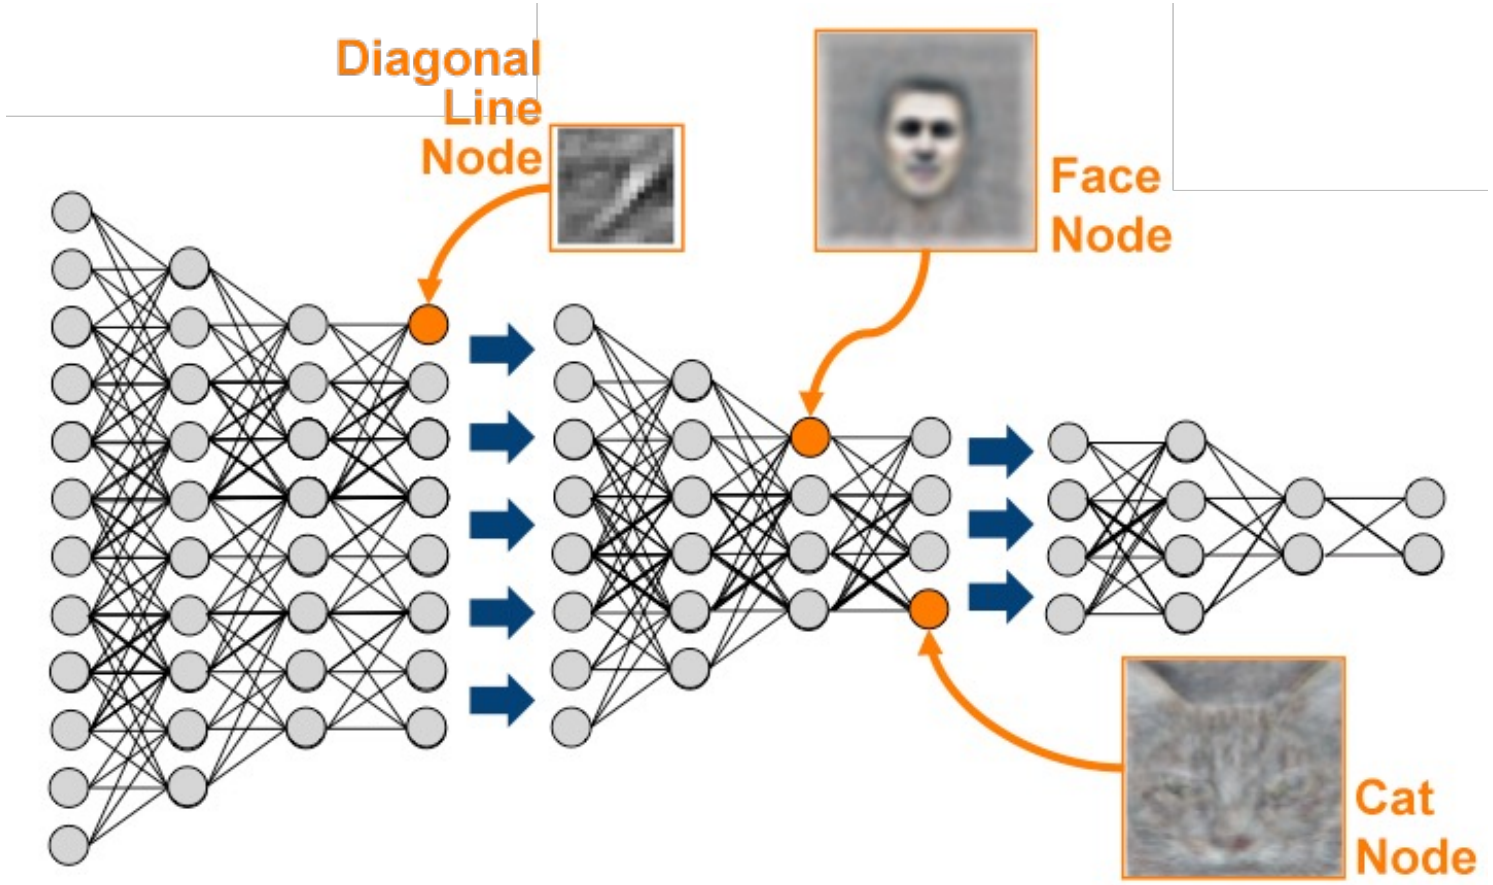

Unsupervised Image Generation

- Which images are real? And which ones are fake?

Real

CelebA (Liu et al., 2015)

StyleGAN2 (Karras et al., 2020)

Fake!

- Image generation: variational autoencoders, generative adversarial networks, diffusion models

Reinforcement Learning Problem

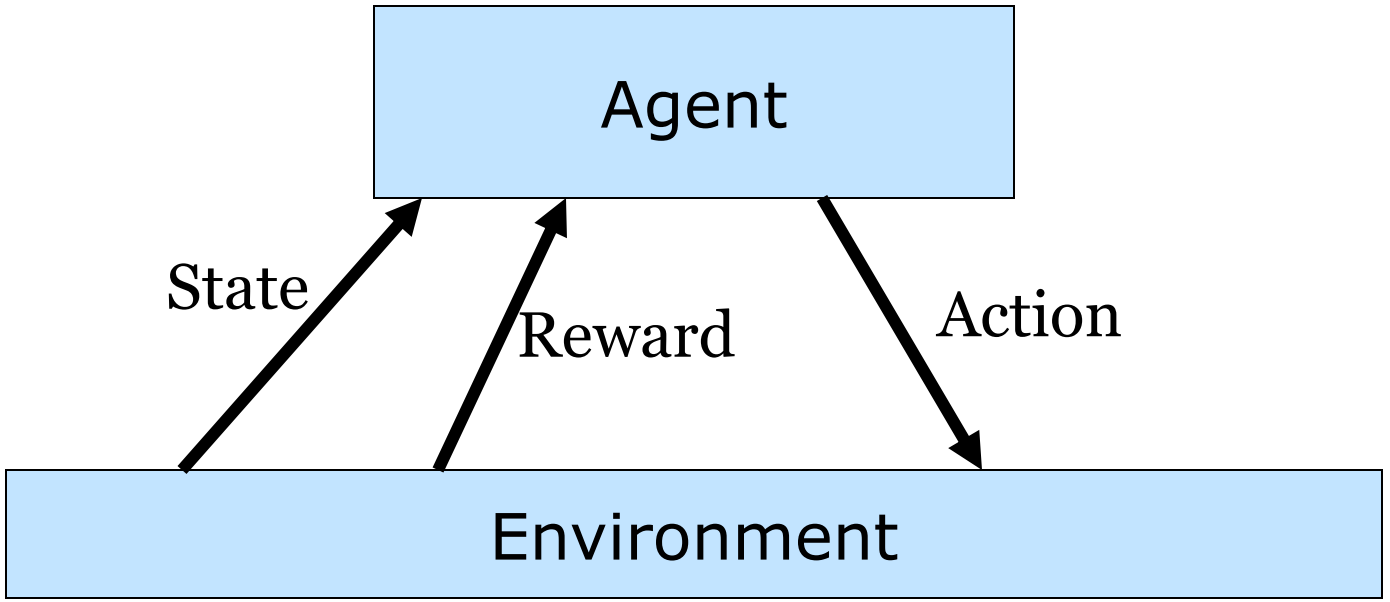

Goal: Learn to choose actions that maximize rewards

Animal Psychology

- Negative reinforcements:
  - Pain and hunger
- Positive reinforcements:
  - Pleasure and food
- Reinforcements used to train animals

- Let's do the same with computers!

Game Playing

- Example: Go (one of the oldest and hardest board games)

- **Agent:** player
- **Environment:** opponent
- **State:** board configuration
- **Action:** next stone location
- **Reward:** +1 win / -1 loose

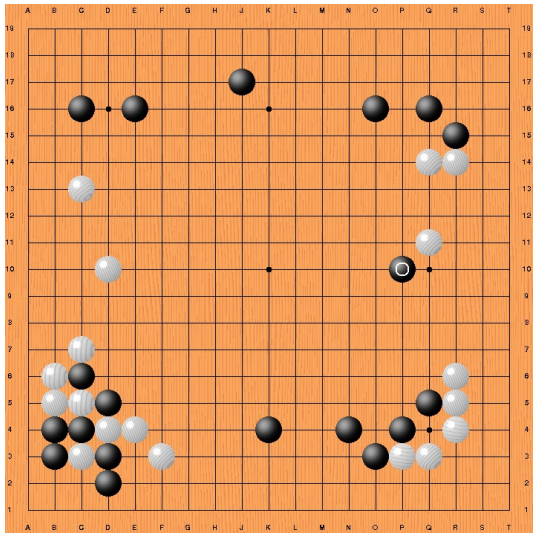

2016: AlphaGo defeats top player Lee Sedol (4-1)
Game 2 move 37: AlphaGo plays unexpected move (odds 1/10,000)

Combining Unsupervised, Supervised and Reinforcement Learning

- Modern systems:
  - Phase 1: unsupervised feature extraction (no labels)
  - Phase 2: supervised training (exploit labels)
  - Phase 3: fine tune by reinforcement learning (exploit reinforcements)

- Alpha Go: supervised + reinforcement learning
- Sentiment analysis with BERT: unsupervised + supervised learning
- ChatGPT: supervised + reinforcement learning

Applications of Machine Learning

- Speech recognition: Siri, Cortana
- Natural Language Processing: Machine translation, dialog systems
- Computer vision: Image and video analysis
- Robotic Control: Autonomous vehicles
- Intelligent assistants: Activity recognition, recommender systems
- Computational finance: Stock trading, portfolio optimization

This course

- **Supervised** and **unsupervised** machine learning

- But not reinforcement learning
- See CS 486/686/885 for Reinforcement Learning
  - https://cs.uwaterloo.ca/~ppoupart/teaching/cs885-fall22/