MAT8034: Machine Learning

Introduction

Fang Kong

https://fangkongx.github.io/Teaching/MAT8034/Spring2025/index.html

Part of slide credits: Stanford CS229

Self introduction

- Position
  - Assistant professor at the department of statistics and data science
- Education
  - Shanghai Jiao Tong University, PhD in computer science
  - Shandong University, bachelor's degree in software engineering
- Research experience
  - CUHK, Tencent, MSRA, Alibaba Damo Academy…
- Office hour
  - 10:00-12:00, every Wednesday @ Room 315, Business School Building

How am I related to ML?

- Research Interests
  - Interactive machine learning
    - Bandit algorithms
    - Online learning
    - Reinforcement learning
- Have published 10+ papers on top ML conferences ICML/NeurIPS/ICLR/COLT…
- Also the reviewer for ICML/NeurIPS/ICLR…

- Was recipient of Baidu Scholarship 2023
  - 10 graduate students worldwide in the field of Artificial Intelligence

Course staff

- Teaching assistants
  - Yunshan Li (李芸珊)
  - Email: 12331105@mail.sustech.edu.cn
  - 2nd year PhD student
  - Research on Low-rank Tensor Modeling
  - Office hour: Thursday 14:00-16:00 @ Room 342, Business School Building

  - Qiqian Li (李其谦)
  - Email: 12431478@mail.sustech.edu.cn
  - 1st year PhD student
  - Research on Financial Machine Learning
  - Office hour: Thursday 14:00-16:00
  - Room 337, Business School Building

  - Haoran Tang (唐浩然)
  - Email: 12432698@mail.sustech.edu.cn
  - 1st year master's student
  - Research on Expensive Evolutionary Algorithm
  - Office hour: Wednesday 14:30-16:30
  - Room 648, College of Engineering, South Tower

Grading

- Assignments (20%)
  - About 3 written assignments
- Project (40%)
- Final Exam (40%)

- Discount for late submissions
  - Final score = original score * (1 – 0.1*number of days late)
  - Submissions more than 5 days late will not be accepted

Course project

- We encourage you to form a group of 1-3 people
  - Same criterion for 1-3 people
- All topics related to ML are ok
- List of potential topics
  - Natural Language
  - Computer Vision
  - Machine Learning Theory
  - Reinforcement Learning
  - Finance & Commerce……
- More submission requirements would be released later

References

- Main reference:
  - Stanford CS229 Lecture Notes
  - https://cs229.stanford.edu/main_notes.pdf by Andrew Ng and Tengyu Ma

- Others
  - Reinforcement Learning: An Introduction by Richard S. Sutton and Andrew G. Barto
  - Pattern Recognition and Machine Learning by Christopher M. Bishop
  - Elements of Statistical Learning by Travor Hastie, Robert Tibshirani, Jerome Friedman

Resources

- Blackboard:
  - All resources (announcements, slides, assignments, grades, etc.) posted here
  - Also the course webpage (schedule and slides, announcements)
- WeChat Group

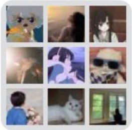

群聊: Machine Learning
Spring 2025

Collaboration and academic dishonesty

- Discussions are encouraged
- Independently finish your assignment
- Same solutions will be reported

Definition of Machine Learning

Arthur Samuel (1959): Machine Learning is the field of study that gives the computer the ability to learn without being explicitly programmed.

A. L. Samuel*

Some Studies in Machine Learning Using the Game of Checkers. II—Recent Progress

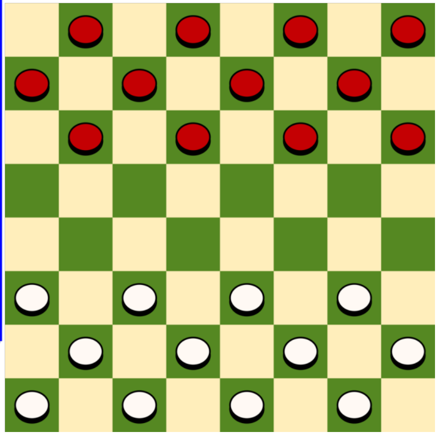

Photos from Wikipedia

Definition of Machine Learning

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.

Experience (data): games played by the program (with itself)

Performance measure: winning rate

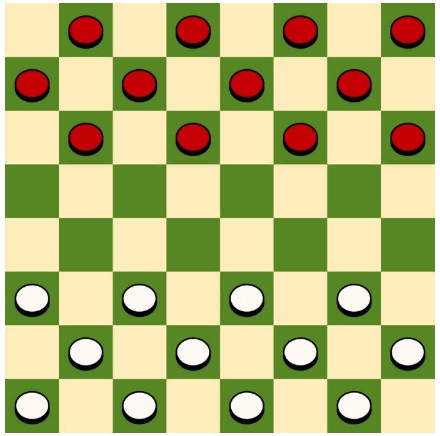

Image from Tom Mitchell's homepage

AI definition by John McCarthy

- What is artificial intelligence
  - It is the science and engineering of making intelligent machines, especially intelligent computer programs
- What is intelligence
  - Intelligence is the computational part of the ability to achieve goals in the world

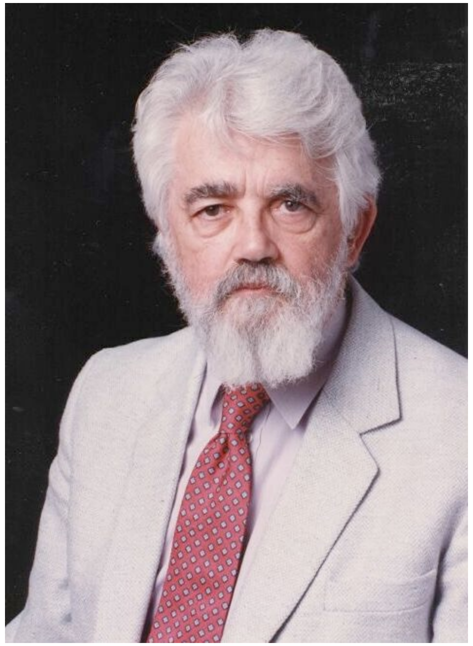

http://www-formal.stanford.edu/jmc/whatisai/whatisai.html

Difference between AI and ML

- AI is a **bigger** concept to create intelligent machines to achieve objectives
- Machine learning is an application or **subset** of AI that allows machines to learn from data without being programmed explicitly

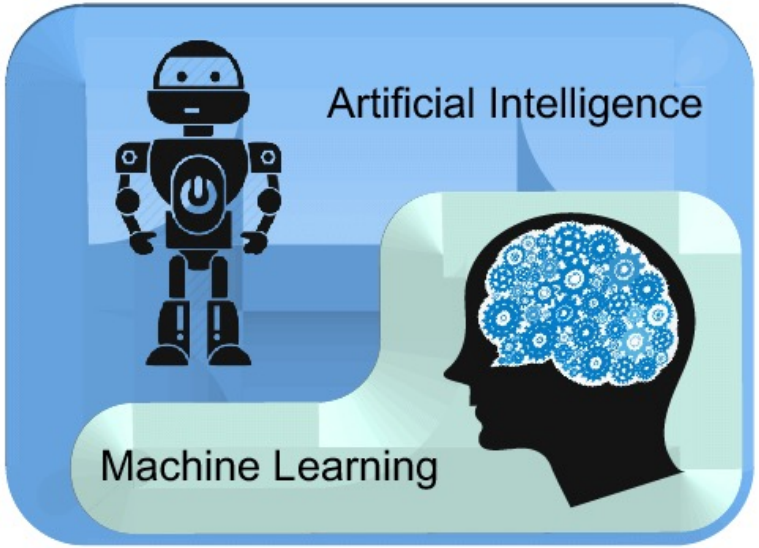

- Example of AI but is not ML
  - A* search algorithm

https://www.javatpoint.com/difference-between-artificial-intelligence-and-machine-learning

Taxonomy of ML

- A simplistic view based on tasks

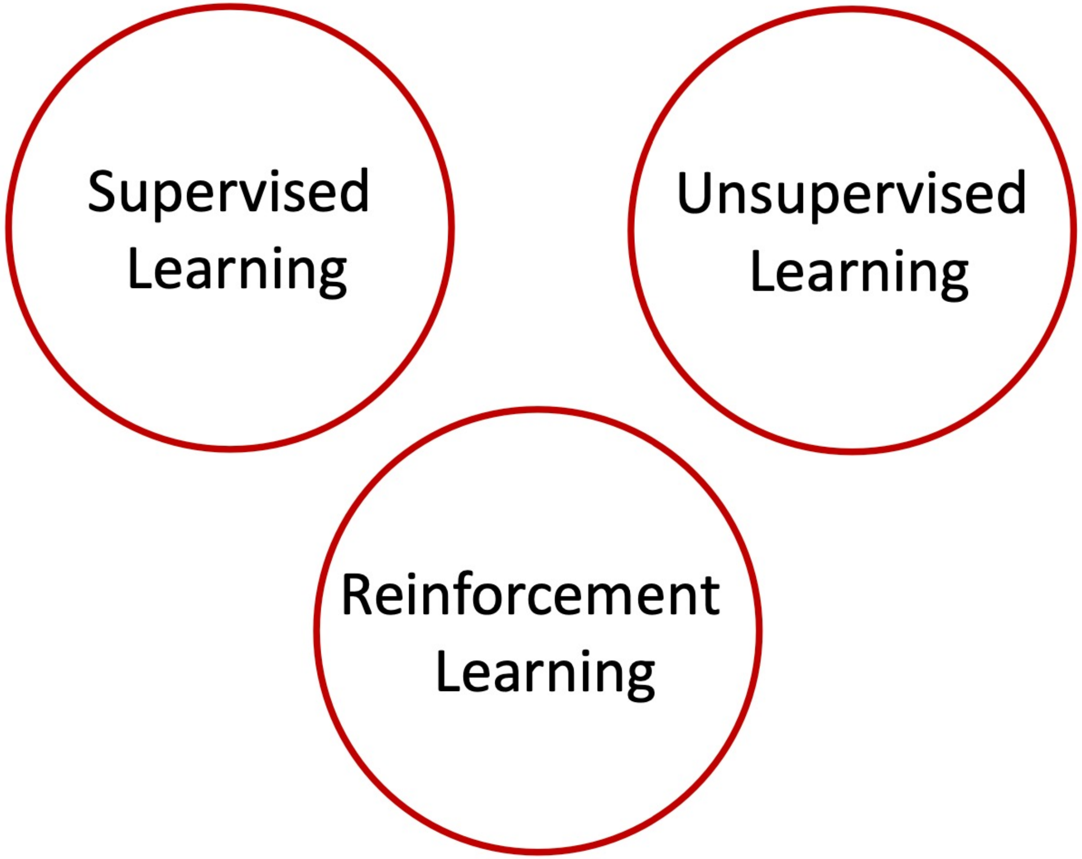

Taxonomy of ML

- A simplistic view based on tasks

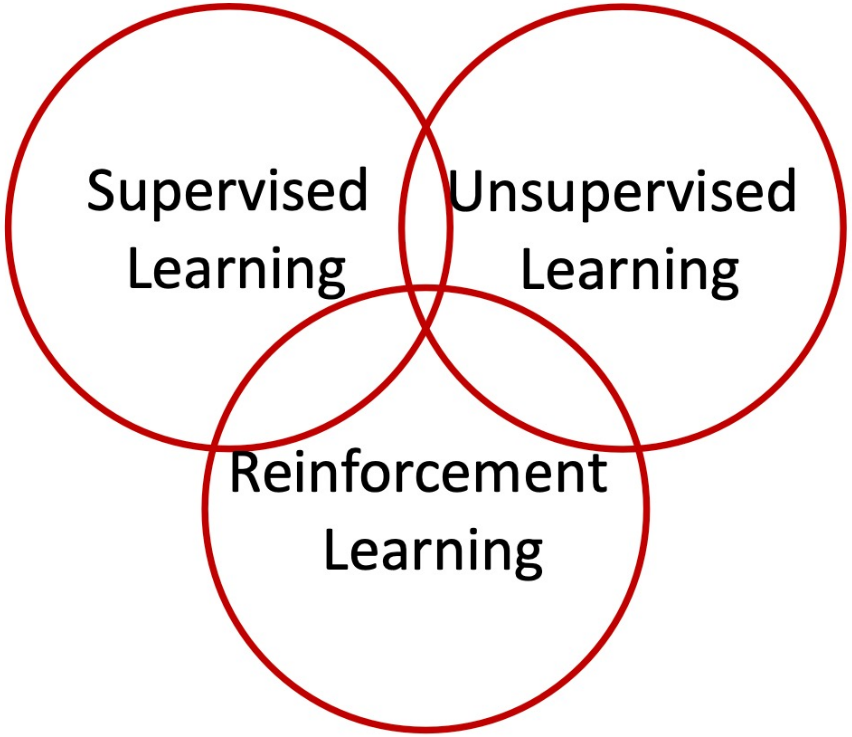

can also be viewed as tools/methods

Supervised Learning

House Price Prediction

- Given: a dataset that contains n samples

$$(x^{(1)}, y^{(1)}), \dots (x^{(n)}, y^{(n)})$$

- Task: if a residence has x square feet, predict its price?

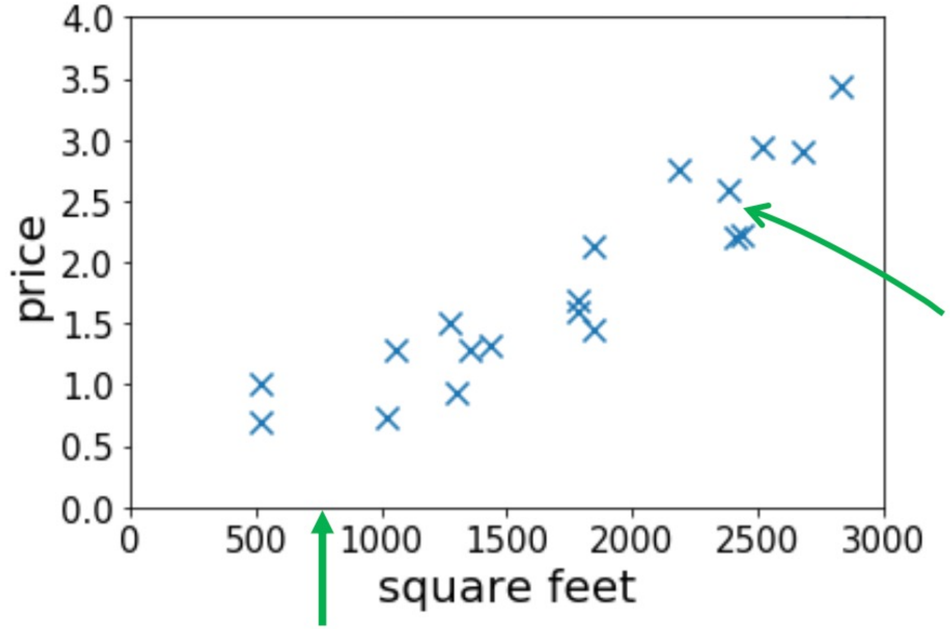

15th sample
$(x^{(15)}, y^{(15)})$

$x = 800$
$y = ?$

House Price Prediction

- Given: a dataset that contains n samples

$$(x^{(1)}, y^{(1)}), \dots (x^{(n)}, y^{(n)})$$

- Task: if a residence has x square feet, predict its price?

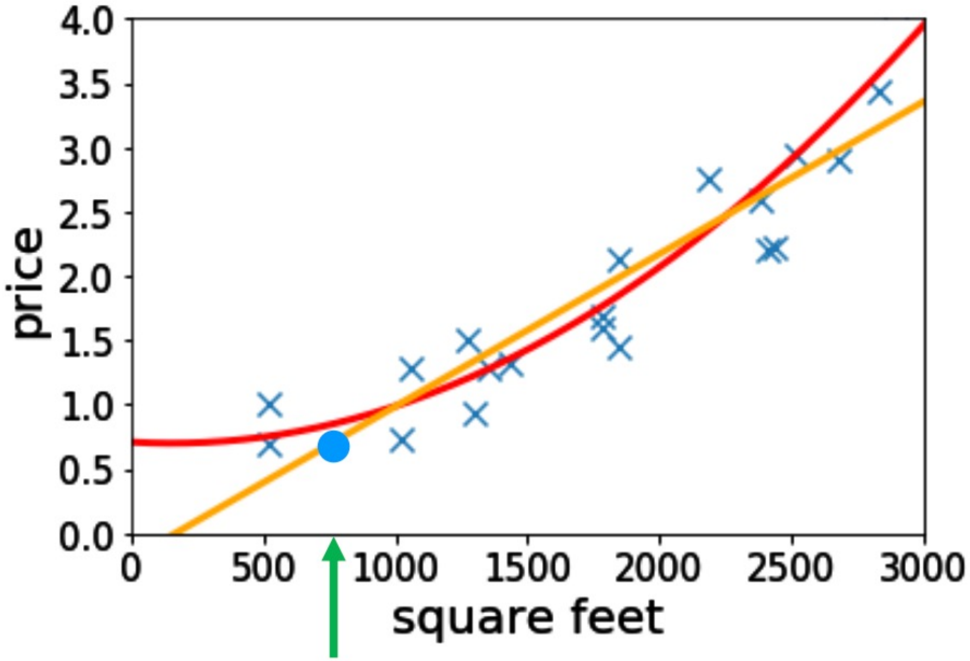

Some lectures : Fitting linear/quadratic/neural functions to the dataset

More Features

- Suppose we also know the lot size
- Task: find a function that maps

$$\underbrace{\text{(size, lot size)}}_{\substack{\text{features/input}\\ x \in \mathbb{R}^2}} \quad \rightarrow \quad \underbrace{\text{price}}_{\substack{\text{label/output}\\ y \in \mathbb{R}}}$$

- Dataset: $\left(x^{(1)}, y^{(1)}\right), \ldots, \left(x^{(n)}, y^{(n)}\right)$

where $x^{(i)} = \left(x_1^{(i)}, x_2^{(i)}\right)$

- "Supervision" refers to $y^{(1)}, \ldots, y^{(n)}$

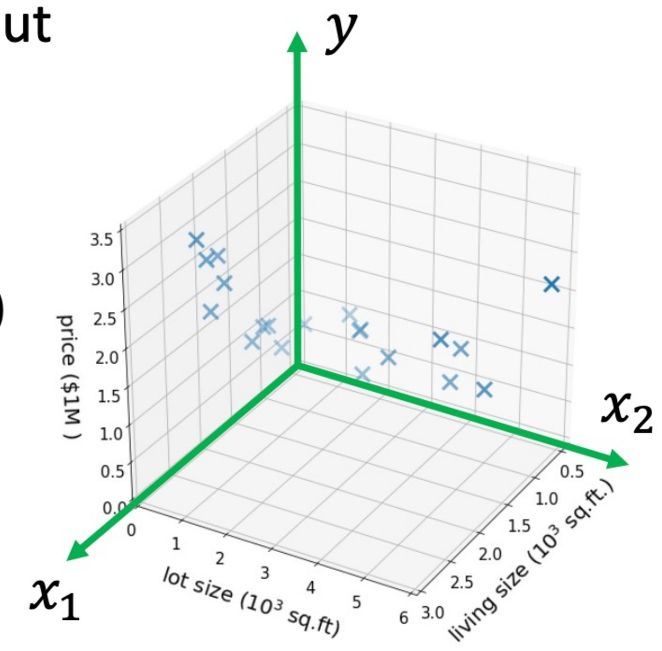

High-dimensional Features

- $x \in \mathbb{R}^d$ for large d
- E.g.,

$$x = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \\ \vdots \\ \vdots \\ x_d \end{bmatrix} \begin{array}{l} \text{--- living size} \\ \text{--- lot size} \\ \text{--- \# floors} \\ \text{--- condition} \\ \text{--- zip code} \\ \quad\quad \vdots \\ \\ \end{array} \longrightarrow \quad y \text{ --- price}$$

Some lectures: Select features based on the data

Regression vs Classification

- regression: if $y \in \mathbb{R}$ is a continuous variable
  - e.g., price prediction
- classification: the label is a discrete variable
  - e.g., the task of predicting the types of residence

(size, lot size) $\rightarrow$ house or townhouse?

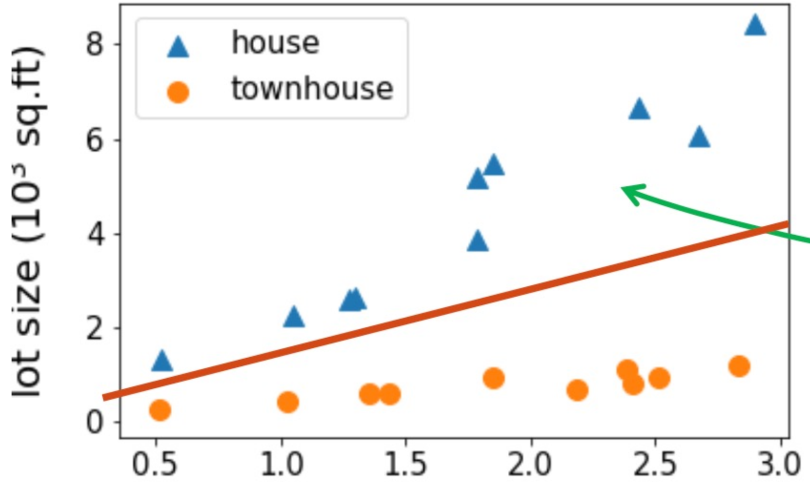

$y =$ house or townhouse?

Some lectures: Classfication

Supervised Learning in Computer Vision

- Image Classification
  - x = raw pixels of the image, y = the main object

ImageNet Large Scale Visual Recognition Challenge. Russakovsky et al.'2015

Supervised Learning in Computer Vision

- Object localization and detection
  - x = raw pixels of the image, y = the bounding boxes

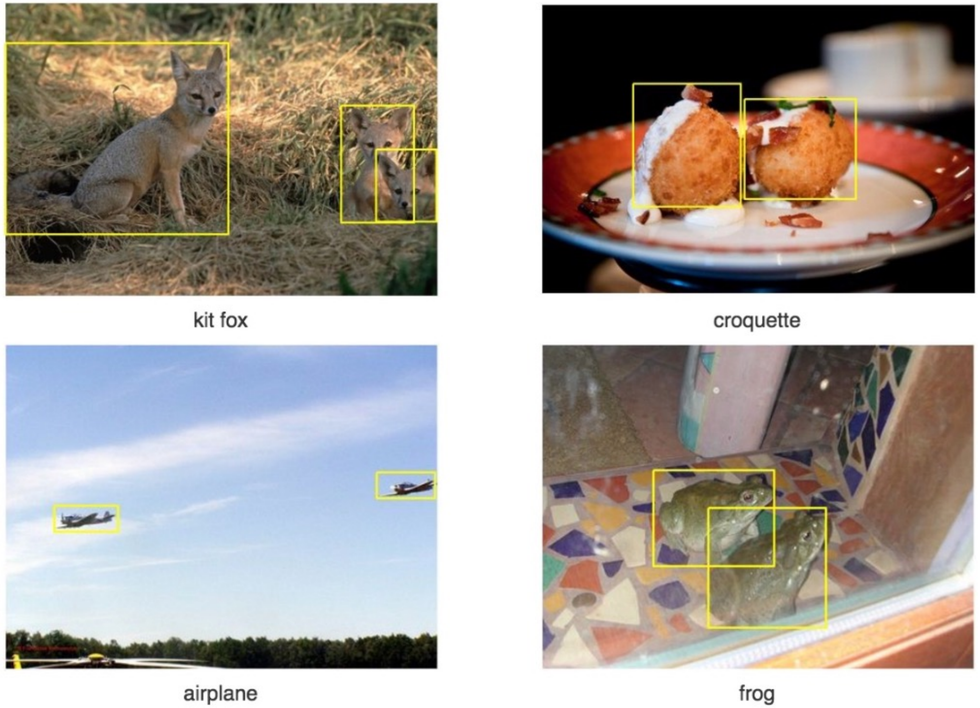

ImageNet Large Scale Visual Recognition Challenge. Russakovsky et al.'2015

Supervised Learning in Natural Language Processing

- Machine translation

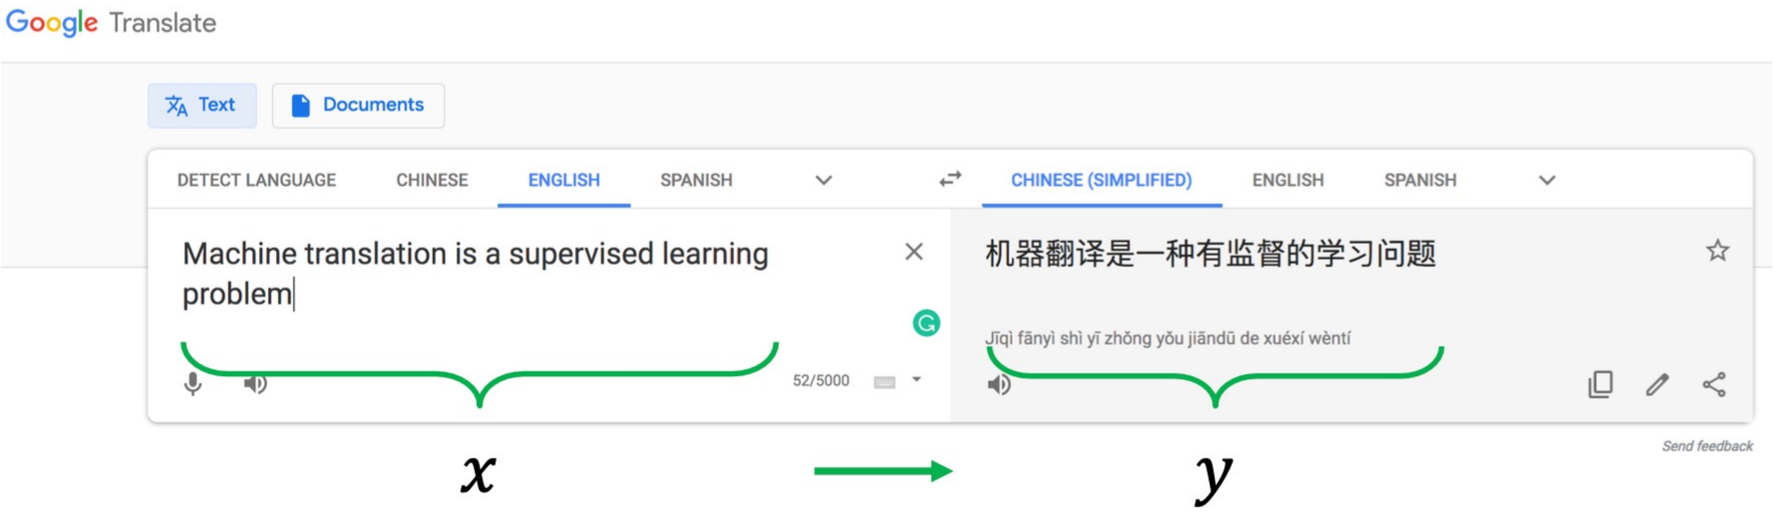

- Note: this course only covers the basic and fundamental techniques of supervised learning

Unsupervised Learning

Unsupervised Learning

- Dataset contains no labels: $x^{(1)}, \dots x^{(n)}$
- Goal (vaguely-posed): to find interesting structures in the data

supervised

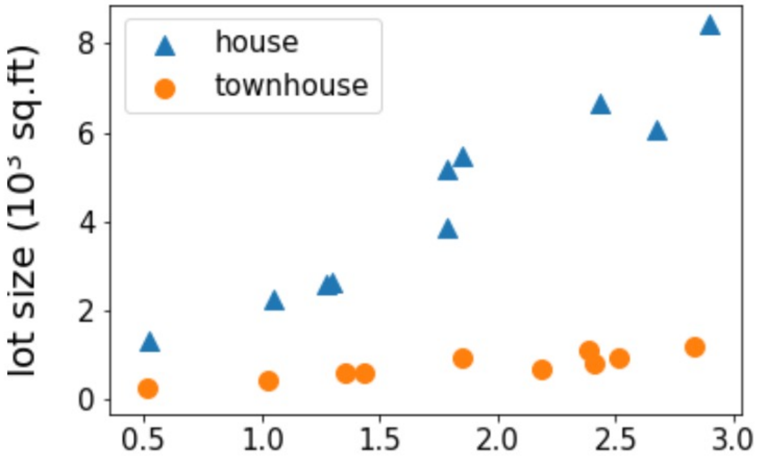

unsupervised

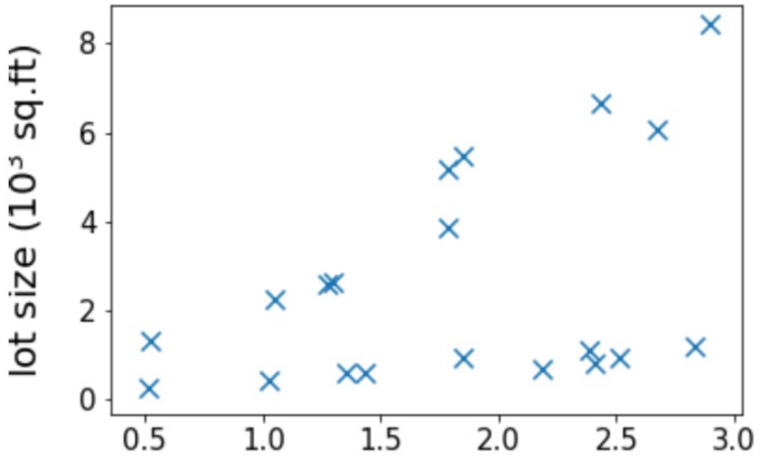

Clustering

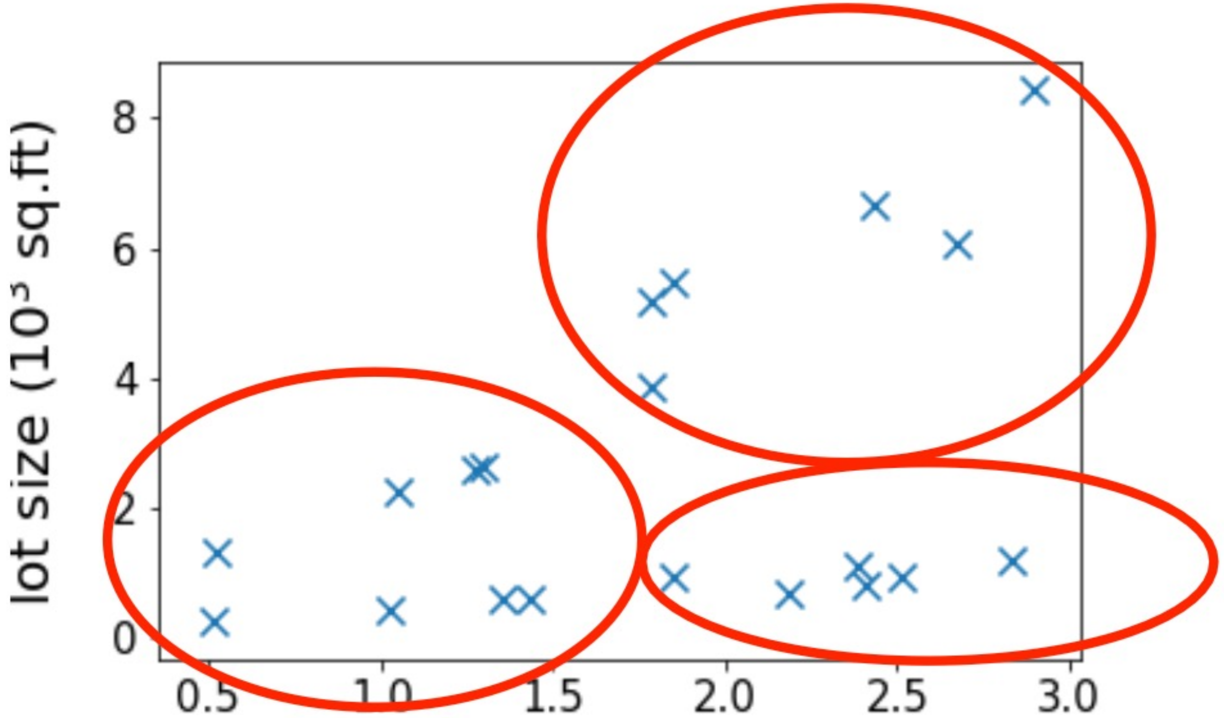

Clustering

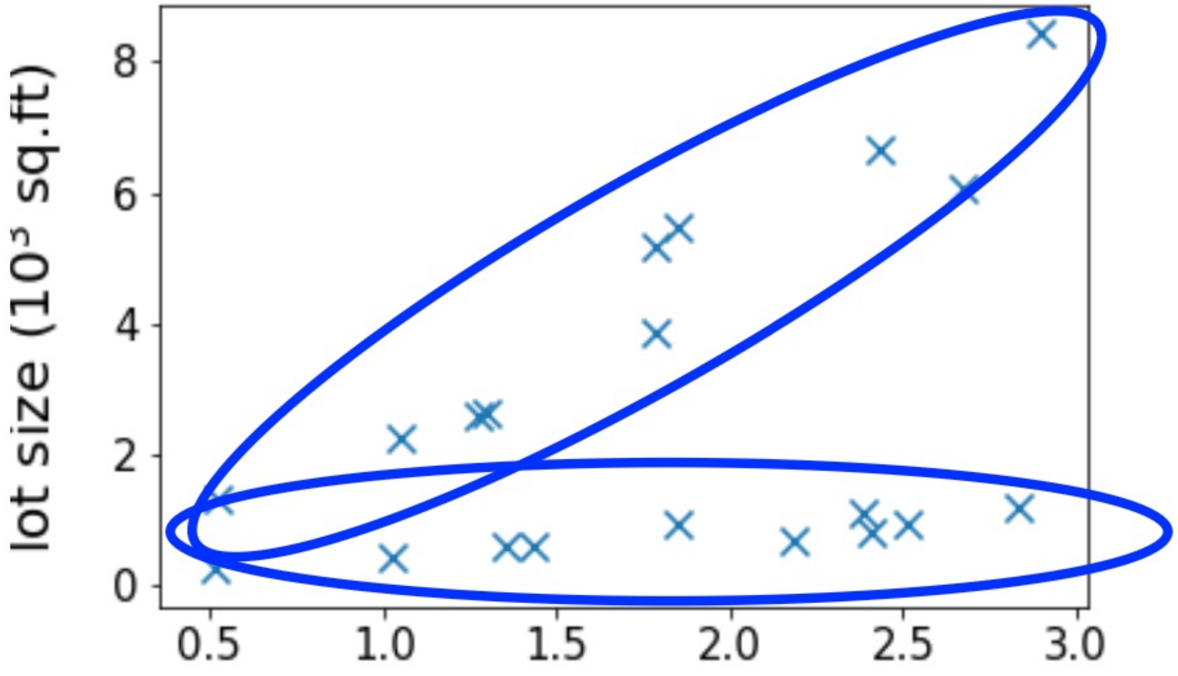

Some lectures: k-means clustering, mixture of Gaussians

Clustering Genes

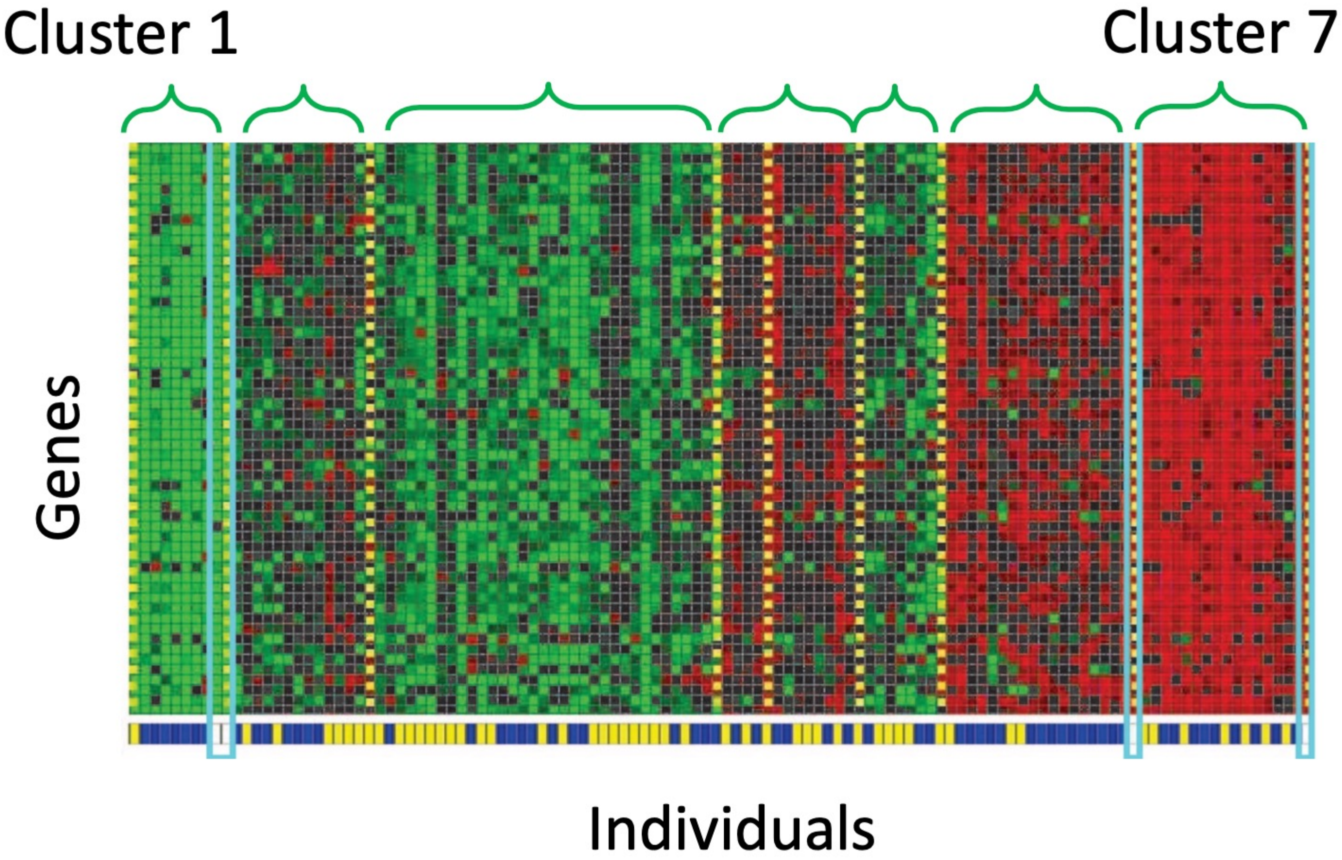

Identifying Regulatory Mechanisms using Individual Variation Reveals Key Role for Chromatin Modification. [Su-In Lee, Dana Pe'er, Aimee M. Dudley, George M. Church and Daphne Koller. ’06]

Latent Semantic Analysis (LSA)

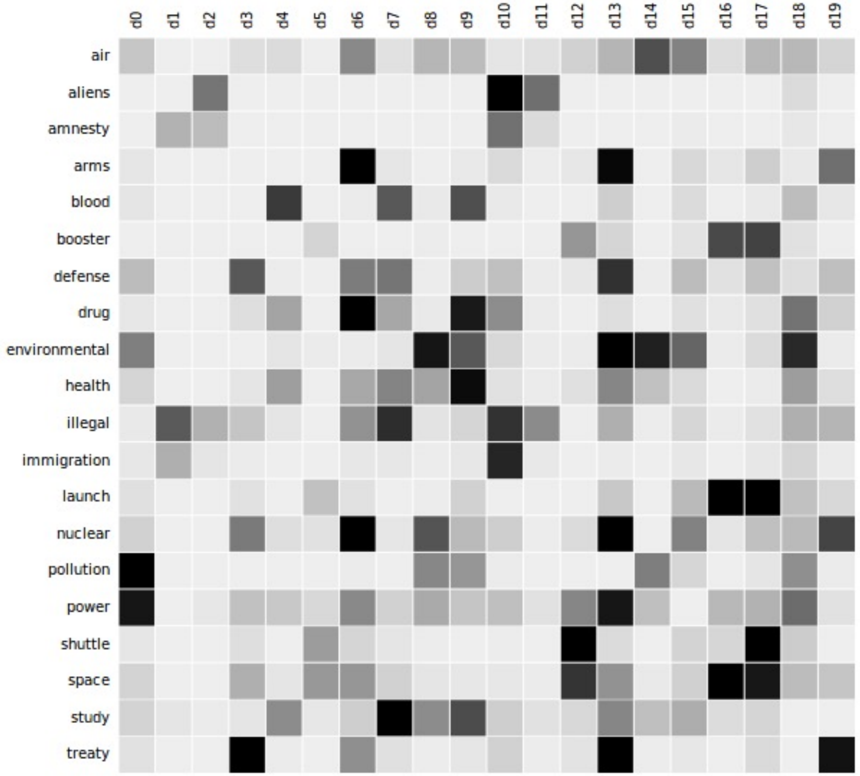

Some lectures: principal component analysis (tools used in LSA)

https://upload.wikimedia.org/wikipedia/commons/6/6e/Topic_detection_in_a_document-word_matrix.gif

Word Embeddings

Represent words by vectors

- word $\xrightarrow{\text{encode}}$ vector
- relation $\xrightarrow{\text{encode}}$ direction

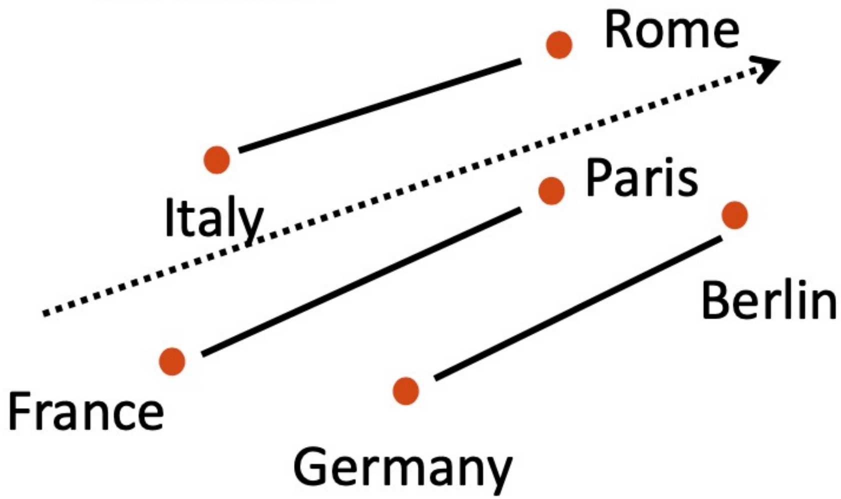

Unlabeled dataset

Word2vec [Mikolov et al'13]
GloVe [Pennington et al'14]

Clustering Words with Similar Meanings

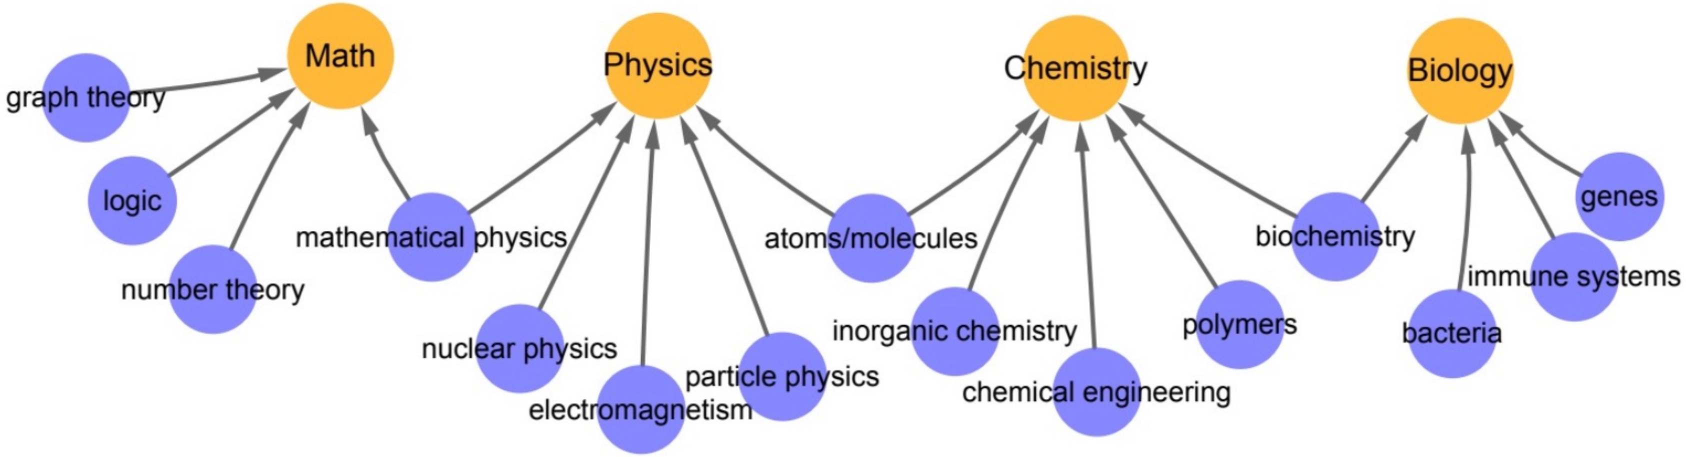

	logic deductive propositional semantics	graph subgraph bipartite vertex	boson massless particle higgs	polyester polypropylene resins epoxy	acids amino biosynthesis peptide
tag	*logic*	*graph theory*	*particle physics*	*polymer*	*biochemistry*

[Arora-Ge-Liang-M.-Risteski, TACL'17,18]

Large Language Models

- Machine learning models for language learnt on large-scale language datasets
- Can be used for many purposes

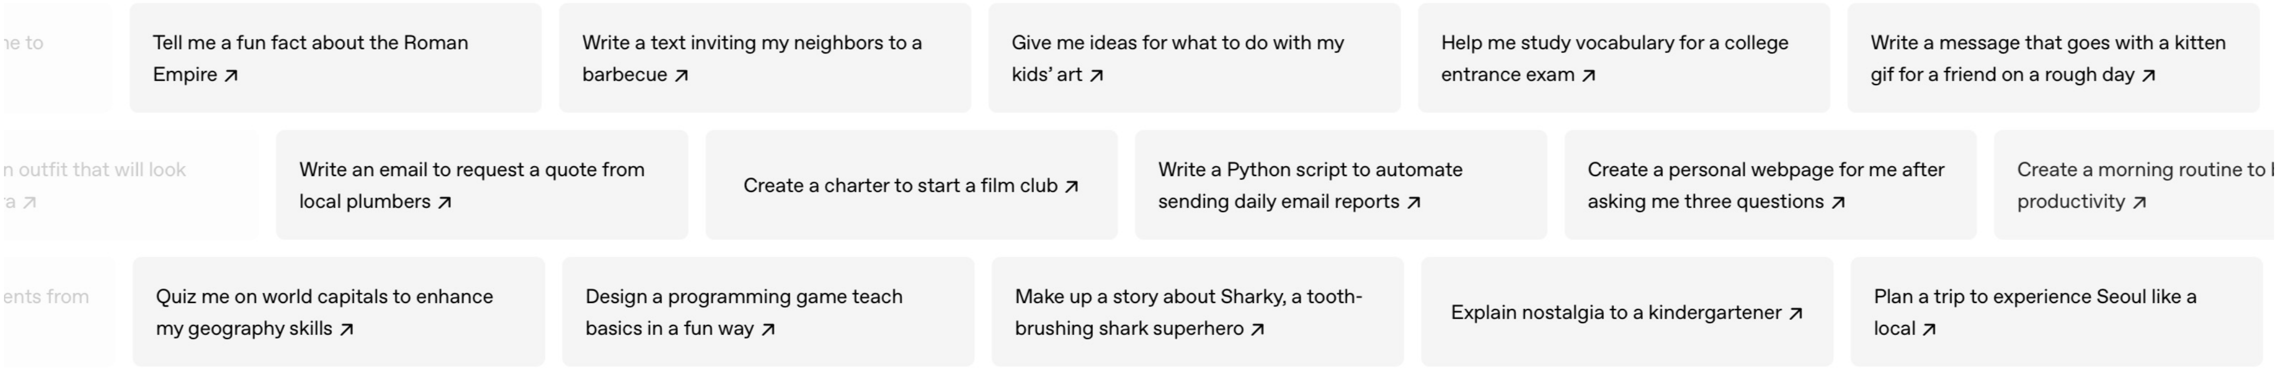

https://openai.com/chatgpt/overview/

Reinforcement Learning

Reinforcement Learning

- Learning to make sequential **decisions**
  - decisions vs predictions

AlphaGo

Reasoning Capability in LLMs

DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning

DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z.F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J.L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R.J. Chen, R.L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S.S. Li et al. (100 additional authors not shown)

We introduce our first-generation reasoning models, DeepSeek-R1-Zero and DeepSeek-R1. DeepSeek-R1-Zero, a model trained via large-scale reinforcement learning (RL) without supervised fine-tuning (SFT) as a preliminary step, demonstrates remarkable reasoning capabilities. Through RL, DeepSeek-R1-Zero naturally emerges with numerous powerful and intriguing reasoning behaviors. However, it encounters challenges such as poor readability, and language mixing. To address these issues and further enhance reasoning performance, we introduce DeepSeek-R1, which incorporates multi-stage training and cold-start data before RL. DeepSeek-R1 achieves performance comparable to OpenAI-o1-1217 on reasoning tasks. To support the research community, we open-source DeepSeek-R1-Zero, DeepSeek-R1, and six dense models (1.5B, 7B, 8B, 14B, 32B, 70B) distilled from DeepSeek-R1 based on Qwen and Llama.

Subjects: **Computation and Language (cs.CL)**; Artificial Intelligence (cs.AI); Machine Learning (cs.LG)
Cite as: arXiv:2501.12948 **[cs.CL]**
(or arXiv:2501.12948v1 **[cs.CL]** for this version)
https://doi.org/10.48550/arXiv.2501.12948

September 12, 2024 Release

Learning to reason with LLMs

We are introducing OpenAI o1, a new large language model trained with reinforcement learning to perform complex reasoning. o1 thinks before it answers—it can produce a long internal chain of thought before responding to the user.

Example: Learning to Walk

Initial

[Kohl and Stone, ICRA 2004]

[Video: AIBO WALK – initial]

Example: Learning to Walk

Finished

[Kohl and Stone, ICRA 2004]

[Video: AIBO WALK – finished]

The Crawler!

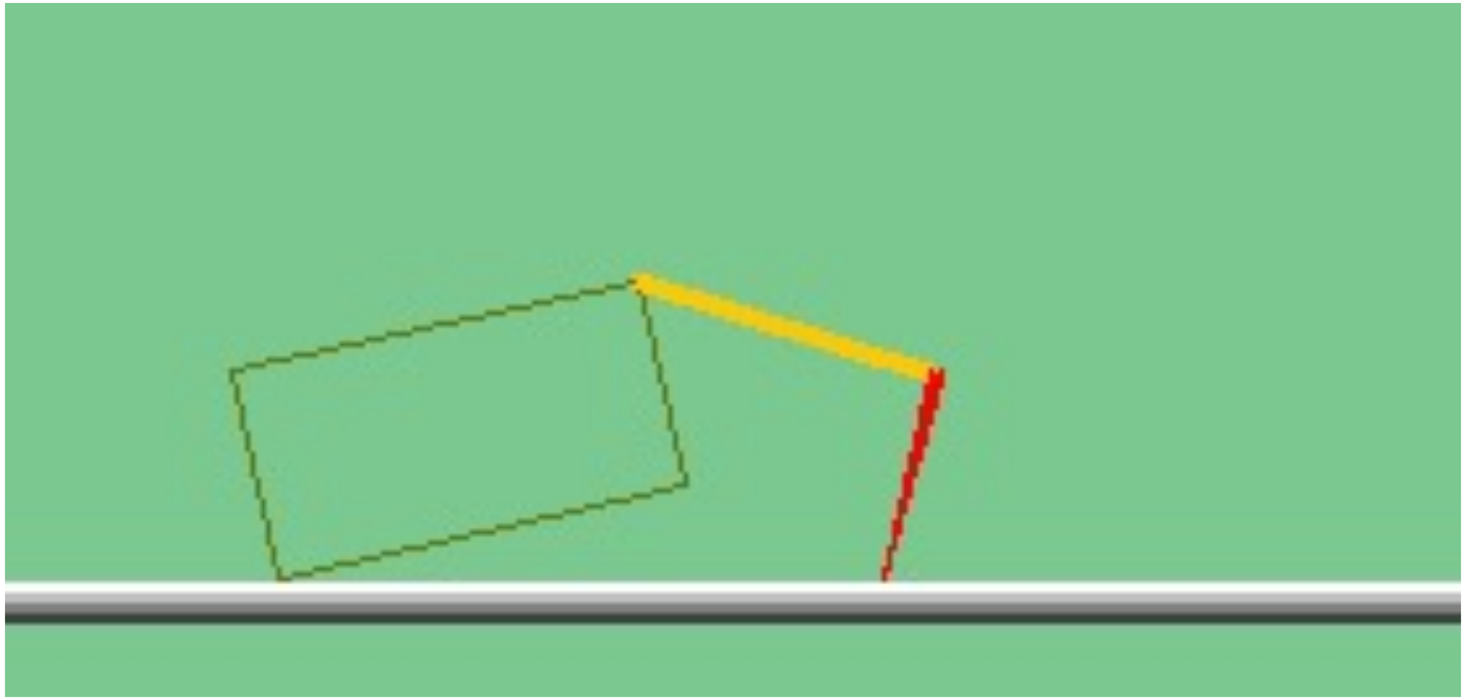

[Demo: Crawler Bot (L10D1)]

Video of Demo Crawler Bot

Reinforcement Learning

- The algorithm can collect data interactively

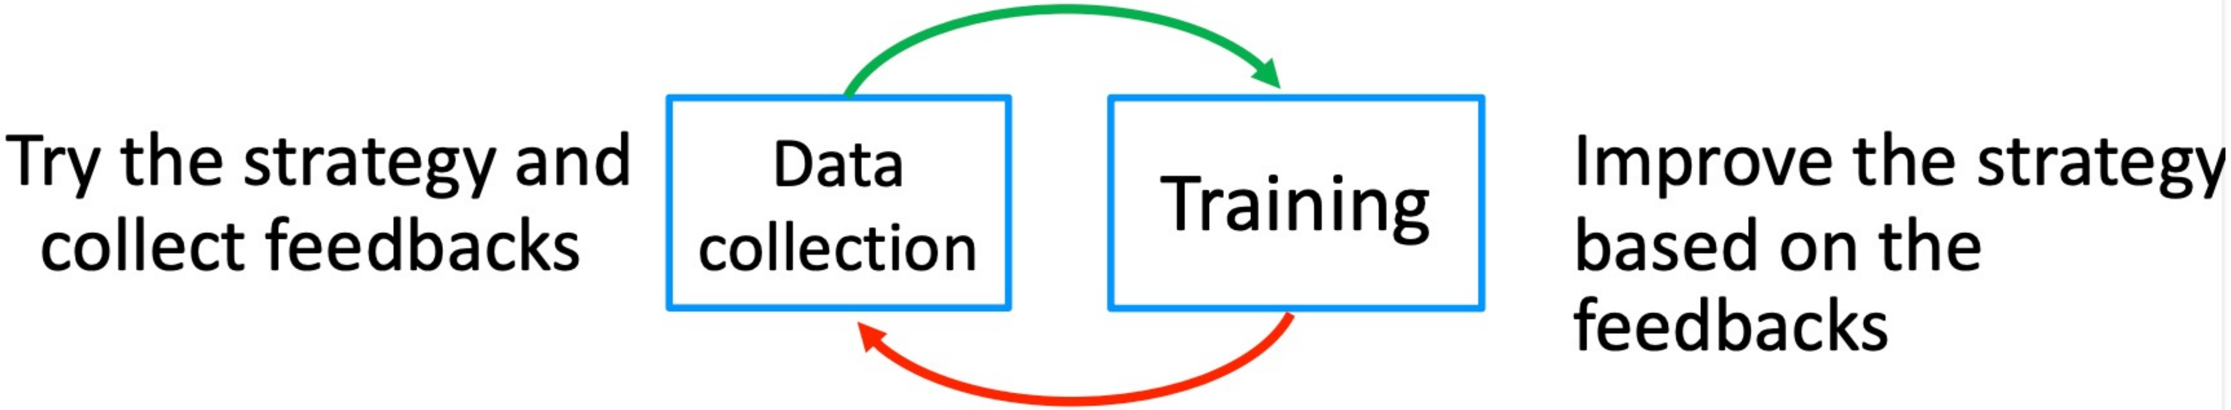

Taxonomy of ML

- A simplistic view based on tasks

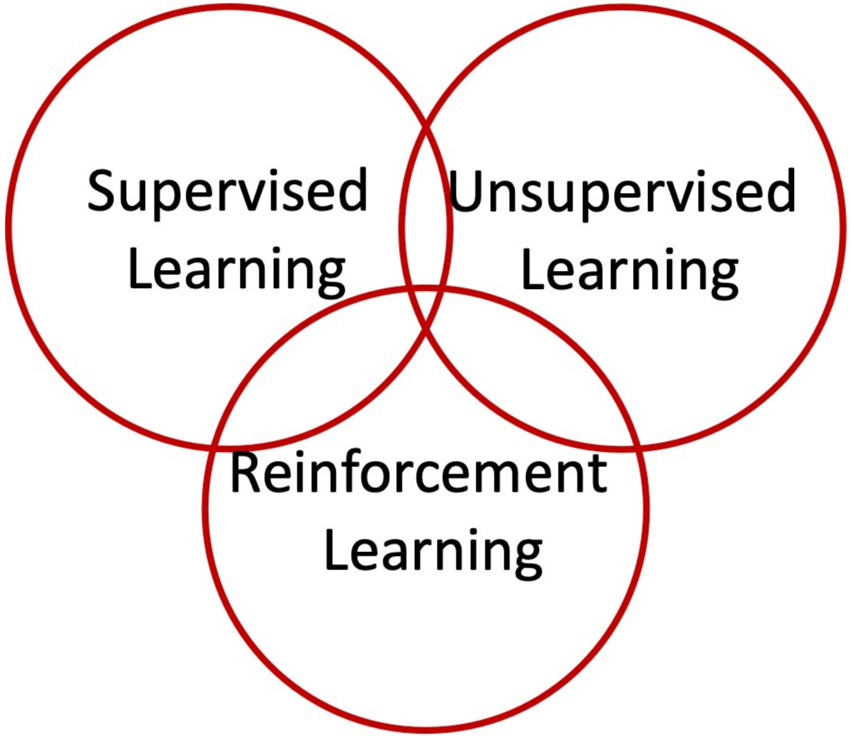

can also be viewed as tools/methods

Other Learning Settings

Other Learning Settings

- Semi-supervised
- Weakly supervised
- Multi-task
- Transfer
- Few-shot
- Zero-shot
- Federated
- ……

Semi-supervised Learning

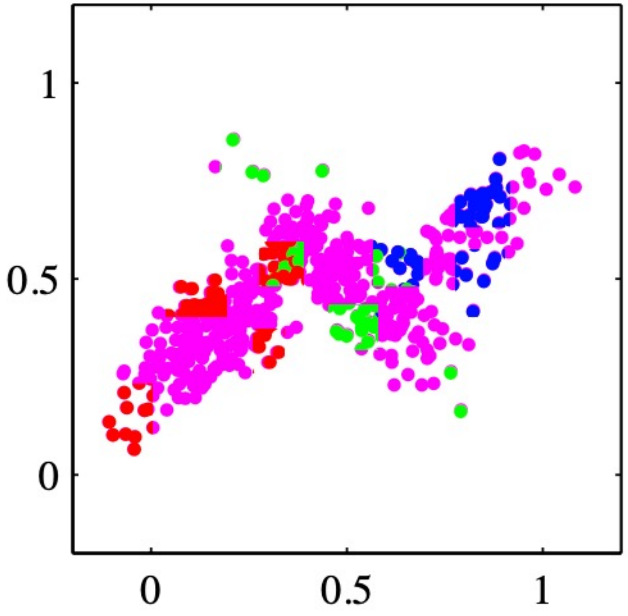

- Observe:
  - Features **x** for all data points
  - Labels y only for some data points
- Learning goal:
  - Model to predict y from **x**

Weakly Supervised Learning

Label y: Perfect bounding box

Imprecise label

Inaccurate label

- Decrease cost or complexity of labeling by using “surrogate” labels
- Observe:
  - Features **x**
  - Some signal z related to true label y:
    - Imprecise labels – simpler, high-level labels
    - Inaccurate labels – inexpensive, lower-quality labels
    - Existing resources – knowledge bases or heuristics to generate labels
- Learning goal:
  - Model to predict y from **x**

Multi-task Learning

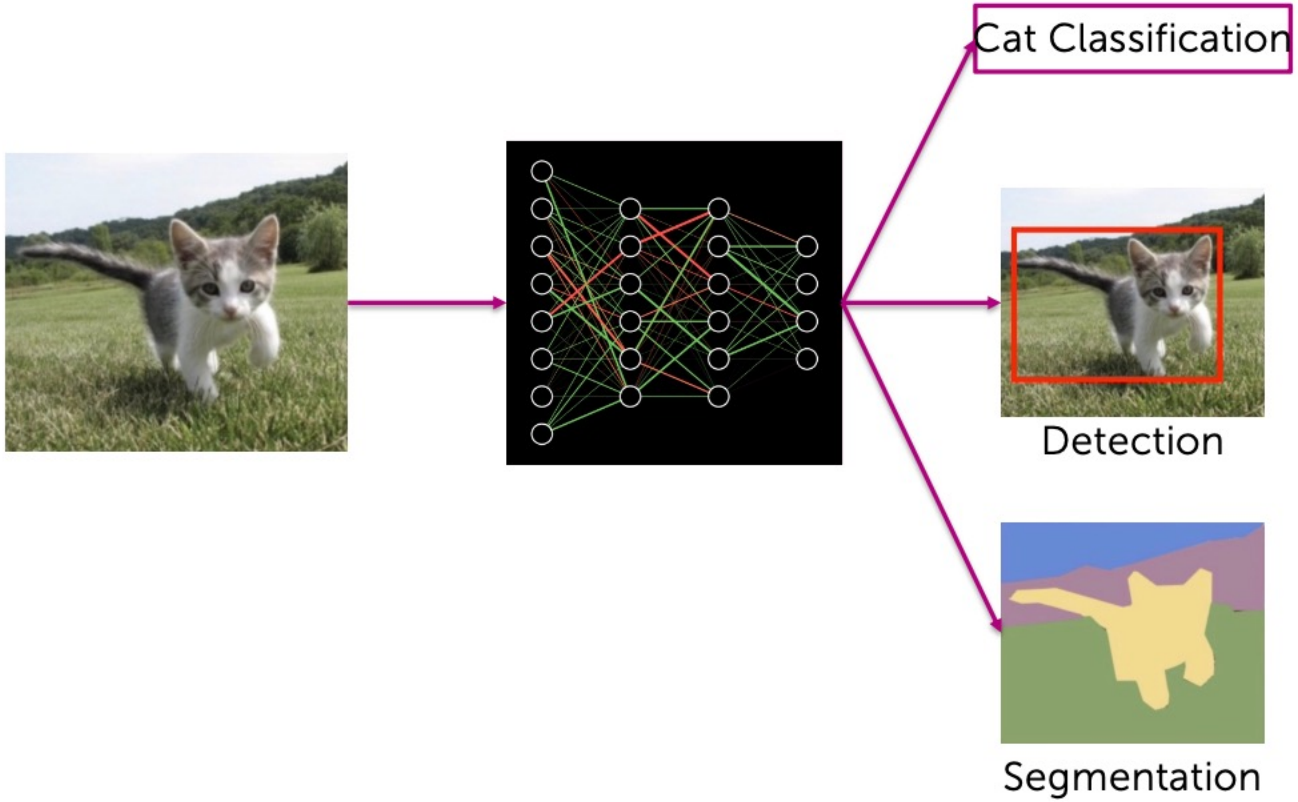

- **Observe:**
  - k tasks
  - Each data point:
    - Features **x**
    - Labels y_j for task j
      - Potentially labels for multiple tasks
- **Learning goal:**
  - Model to predict $y_1,\ldots,y_k$ from **x**

Transfer Learning

- Observe:
  - Model M for previous task
    - Maps **x** → z
  - New task
    - Features **x**
    - Labels y
- Learning goal:
  - Model to predict y from **x**

Transfer Learning

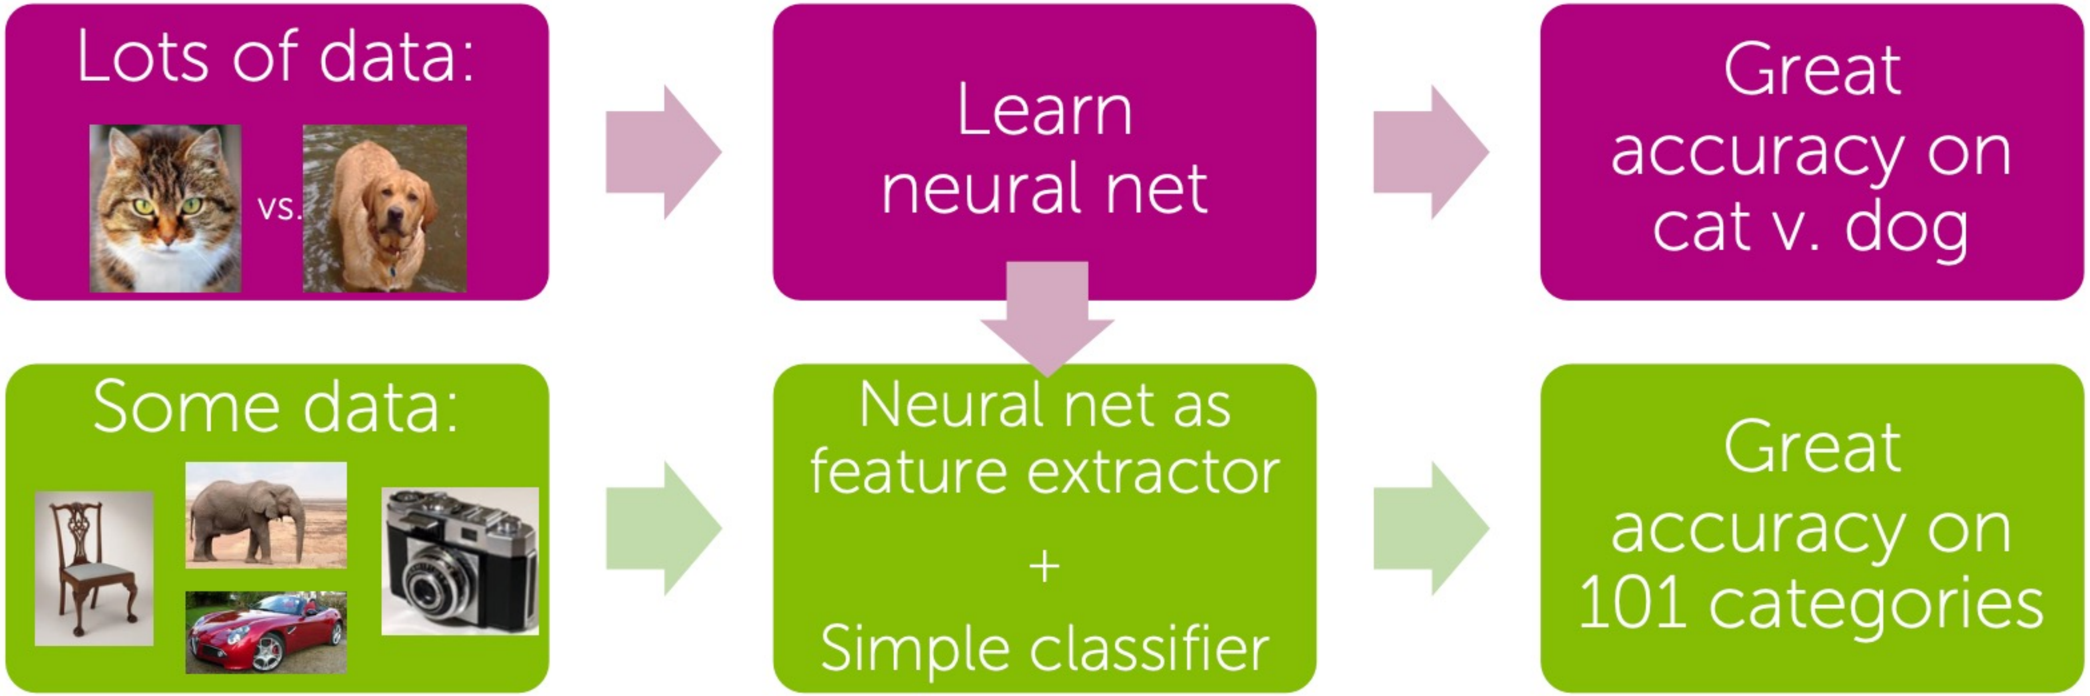

Few-Shot Learning

- Observe:
  - Very few data points: (1 – 100)
    - Features **x**
    - Labels y
- Learning goal:
  - Model to predict y from **x**

Zero-Shot Learning

Zebra???

- Observe:
  - Features **x**
  - Labels y
- Learning goal:
  - Model to predict y’ from **x**
    - For a new class y’ not seen in training data?????

Federated Learning

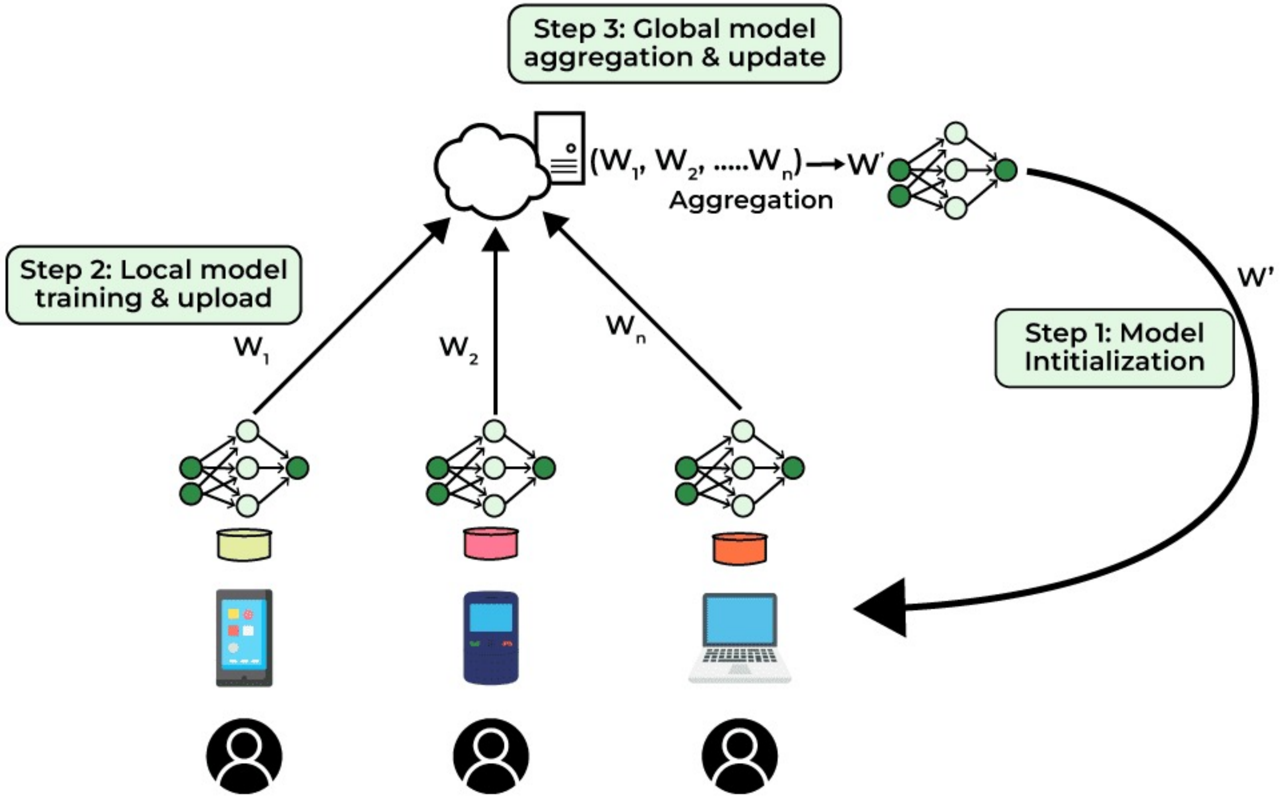

https://www.geeksforgeeks.org/collaborative-learning-federated-learning/

Topics in this course

- Supervised learning
- Unsupervised learning
- Reinforcement learning

- Introduction to learning theory
  - Bias variance tradeoff
  - Feature selection
- Broader aspects of ML
  - Fairness/robustness