CMSC 478
Intro. to Machine Learning
Spring 2024

KMA Solaiman

ksolaima@umbc.edu

Instructor: KMA Solaiman (Salvi)

ITE 201C/Remote

ksolaima@umbc.edu

Wed 5:45 - 6:30 pm,

Thu 3 - 3:45 pm

by appointment

- Multimodal Information Retrieval
- Vision & language processing
- Learning with low-to-no supervision
- Novelties in Learning Models

Administrivia

Course Website

WWW

Schedule, slides, assignments, readings, materials, syllabus here

https://umbc-cmsc478.github.io/spring2024/

campuswire

- **Course announcements,** Q&A, discussion board here
- No public code, follow posted rules and etiquette

gradescope by Turnitin

- **Assignment Submission**
- **Rubrics Grading**
- **Peer Grading**

Text

- No specific text
- Hal Duame, CIML
- Tom Mitechell
- Lecture Notes
- Website

Academic Integrity

- Super important: I take it ***very*** seriously
- ***You*** are responsible for your (& your group's) own work: if in doubt, ask!
- Penalties could include 0 on the assignment, course failure, suspension, or expulsion (not exhaustive)

Final Grades

≥	Letter
90	A
80	B
70	C
60	D
0	F

Programming Languages for Assignments

Python, though individual assignments could vary

Remember: programming languages are *tools*. Don't get too caught up in not "knowing" a language. This course will not be grading software engineering prowess.

Libraries: Assignment dependent. Generally OK, as long as you don't use their implementation of what you need to implement

If in doubt, ask first

Courtesy: Frank Ferraro

Late Policy

Everyone has a budget of 10 *late days, maximum 3 per assignment*

If you have them left: assignments turned in after the deadline will be graded and recorded, no questions asked

Late Policy

Everyone has a budget of 10 *late days, maximum 3 per assignment*

If you have them left: assignments turned in after the deadline will be graded and recorded, no questions asked

If you don’t have any left: still turn assignments in. They could count in your favor in borderline cases

Courtesy: Frank Ferraro

Late Policy

Everyone has a budget of 10 *late days, maximum 3 per assignment*

Use them as needed throughout the course

They're meant for personal reasons and **emergencies**

Do not procrastinate

Courtesy: Frank Ferraro

Late Policy

Everyone has a budget of 10 *late days, maximum 3 per assignment*

Contact me privately if an extended absence will occur

You must know how many you’ve used

Courtesy: Frank Ferraro

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

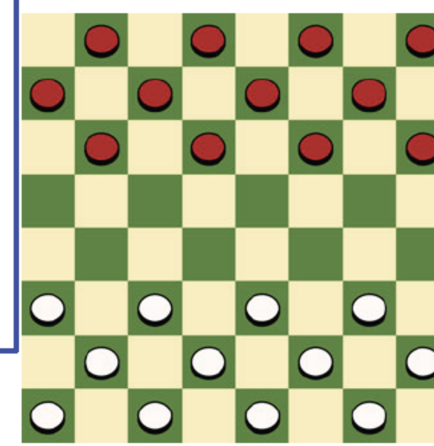

Photos from Wikipedia

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

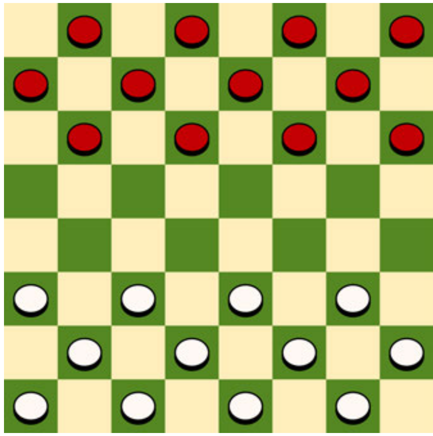

Image from Tom Mitchell's homepage

Taxonomy of Machine Learning (A Simplistic View Based on Tasks)

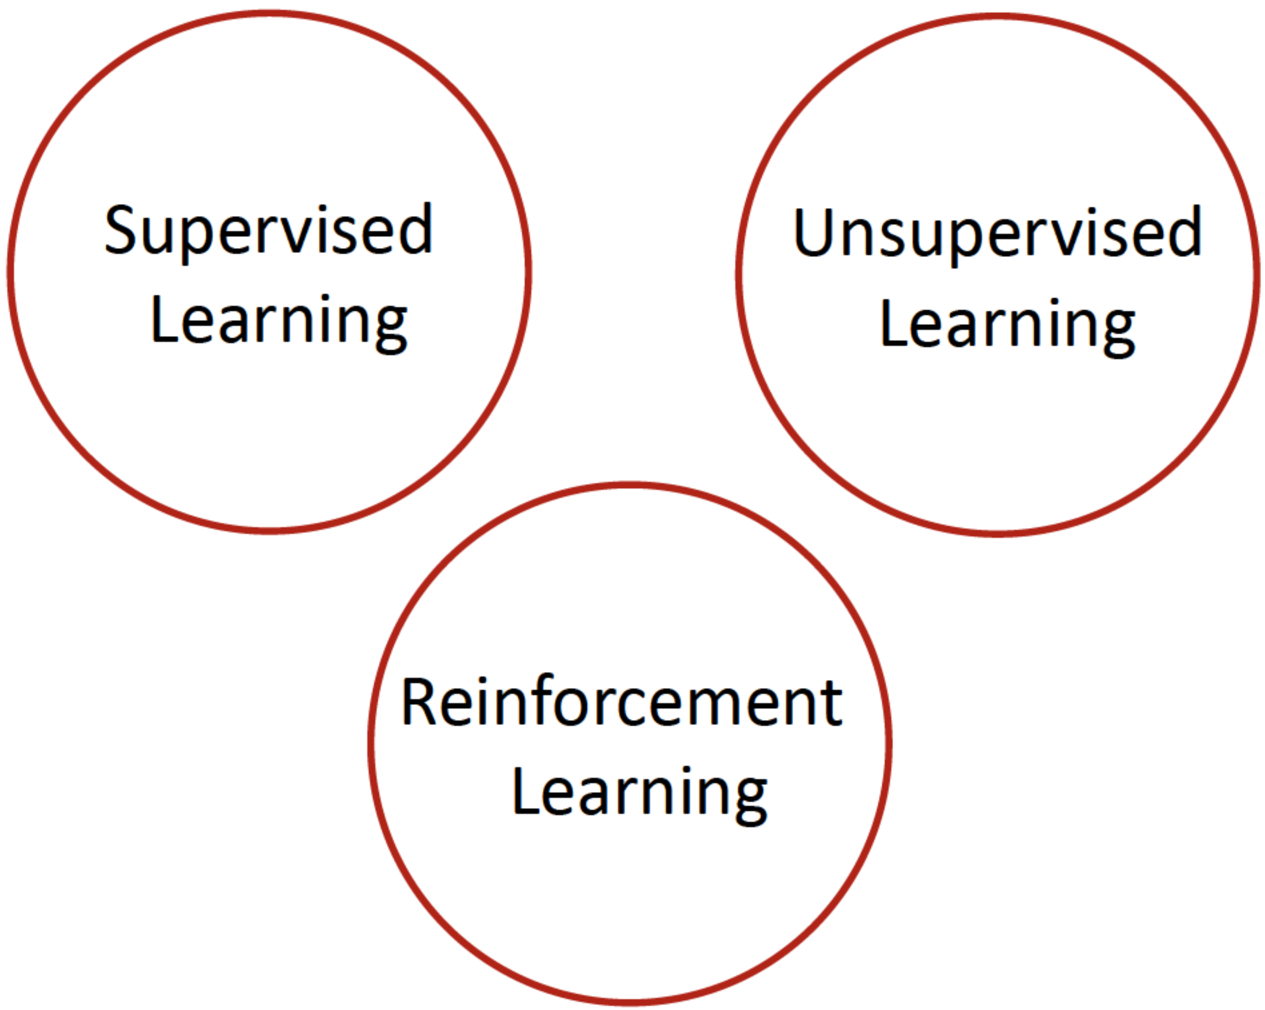

Taxonomy of Machine Learning (A Simplistic View Based on Tasks)

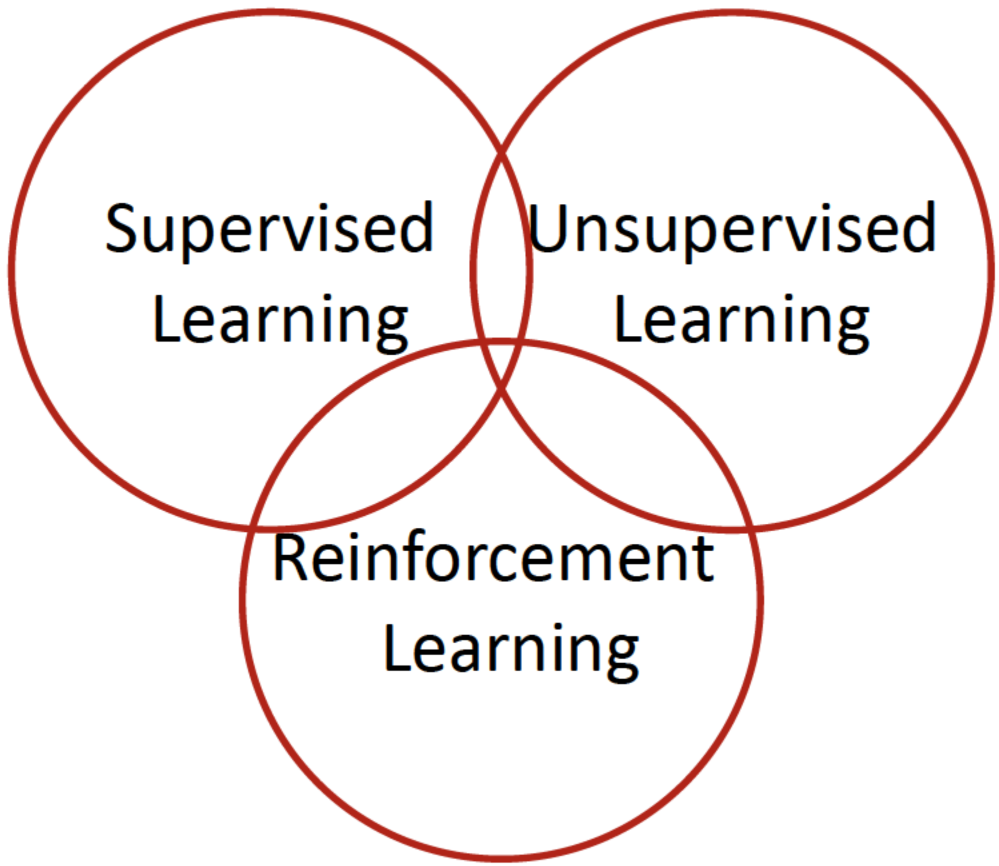

can also be viewed as tools/methods

Supervised Learning

Housing Price Prediction

- Given: a dataset that contains n samples

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

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

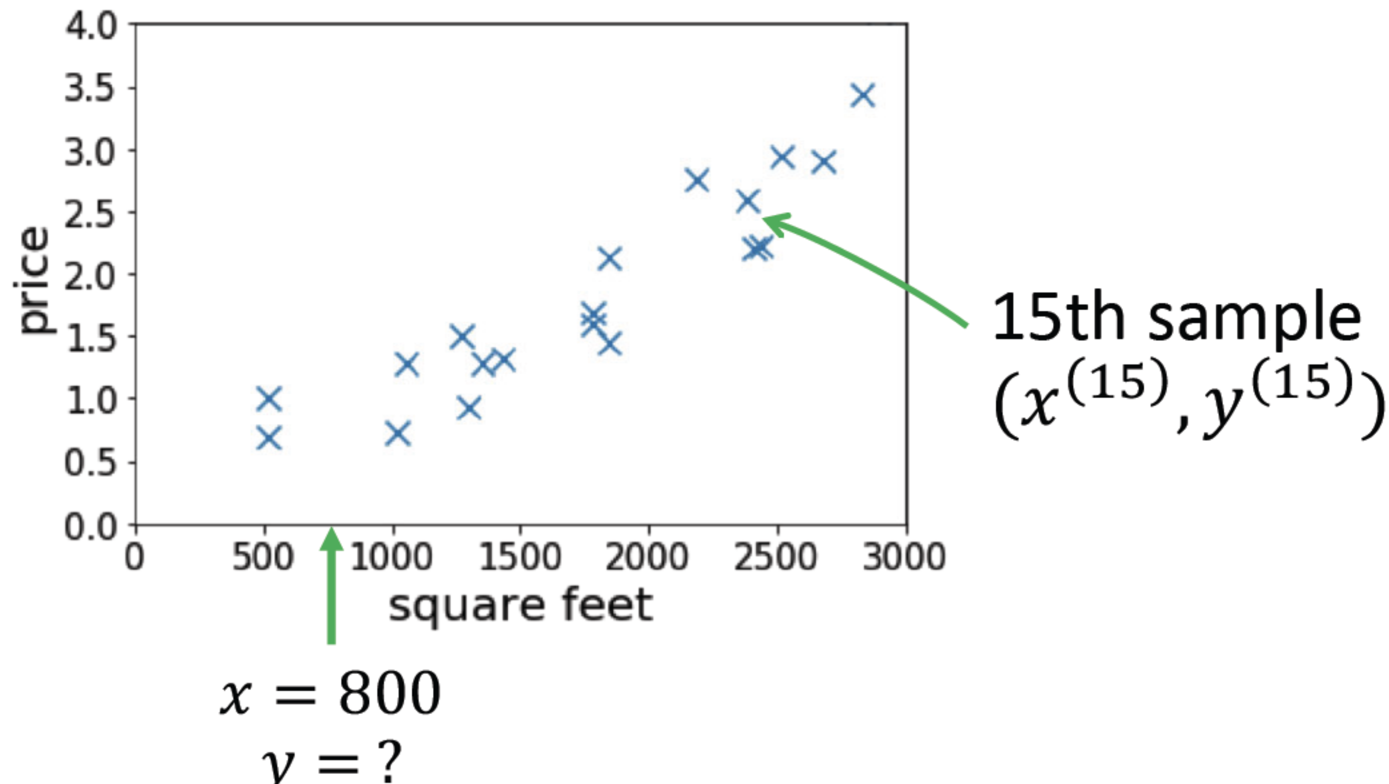

Housing Price Prediction

- Given: a dataset that contains n samples

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

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

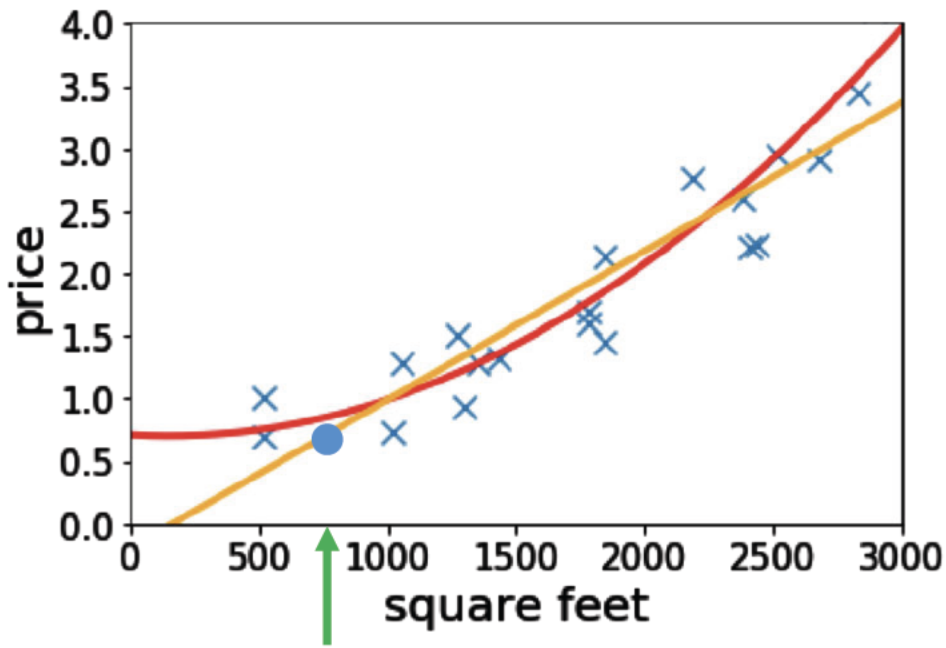

$\mathbf{x} = 800$

$y = ?$

- Lecture 2&3: fitting linear/ qaudratic functions to the dataset

More Features

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

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

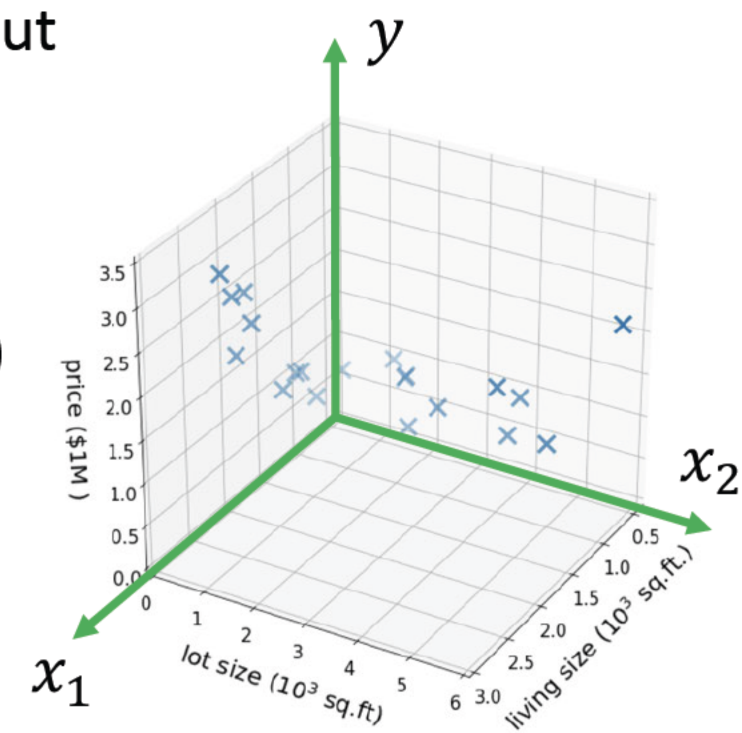

- Dataset: $\left(x^{(1)}, y^{(1)}\right), \ldots, (x^{(n)}, y^{(n)})$

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

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

High-dimensional Features

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

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

- Lecture 6-7: infinite dimensional features
- Lecture 10: 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?

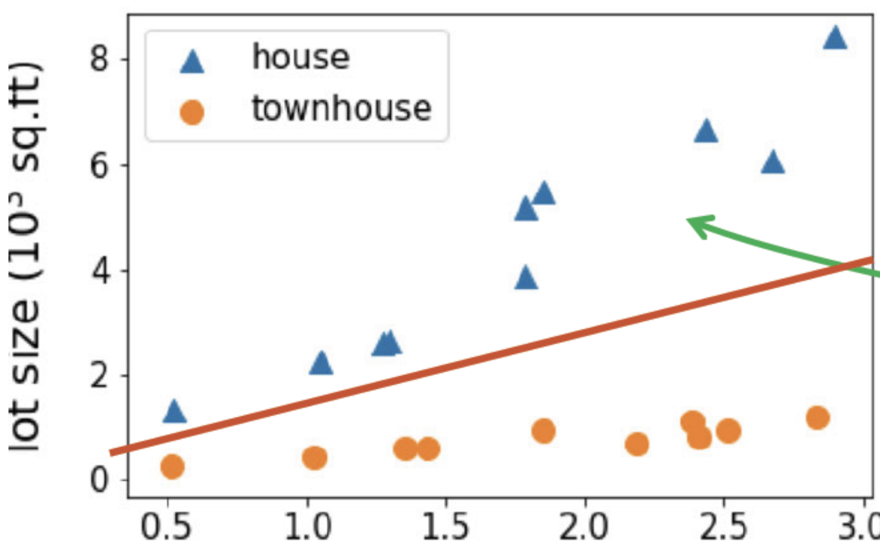

$y =$ house or townhouse?

Lecture 3&4: classification

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

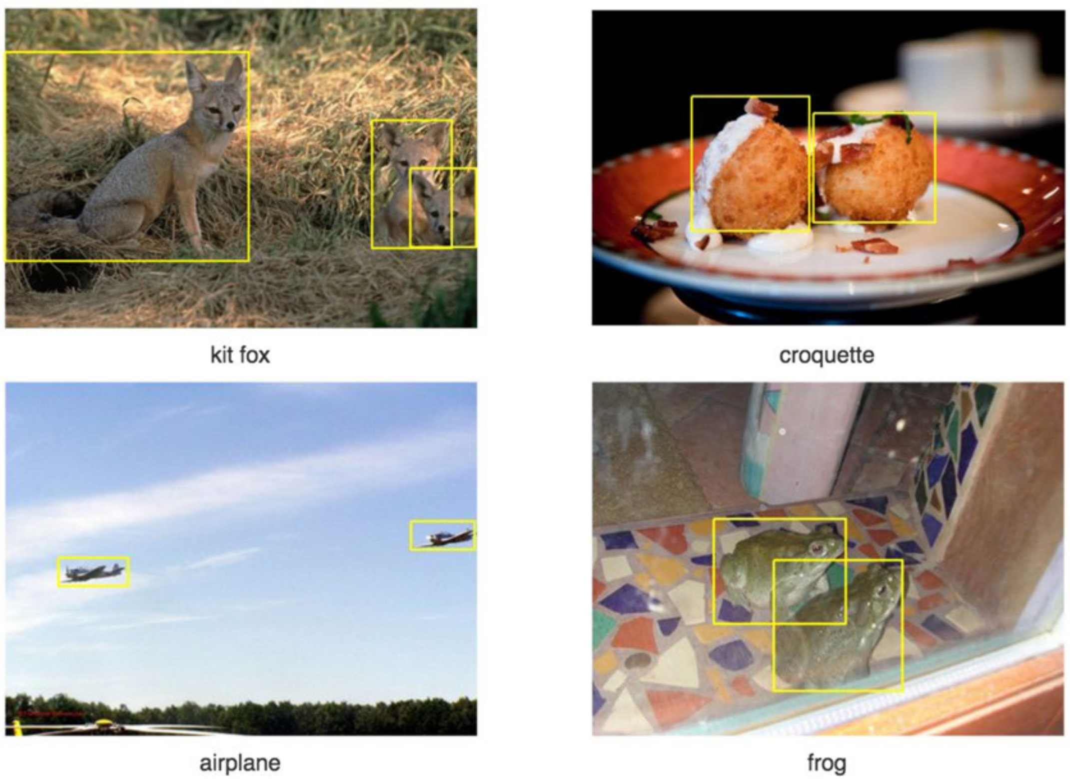

ImageNet Large Scale Visual Recognition Challenge. Russakovsky et al.'2015

Supervised Learning in Natural Language Processing

- Machine translation

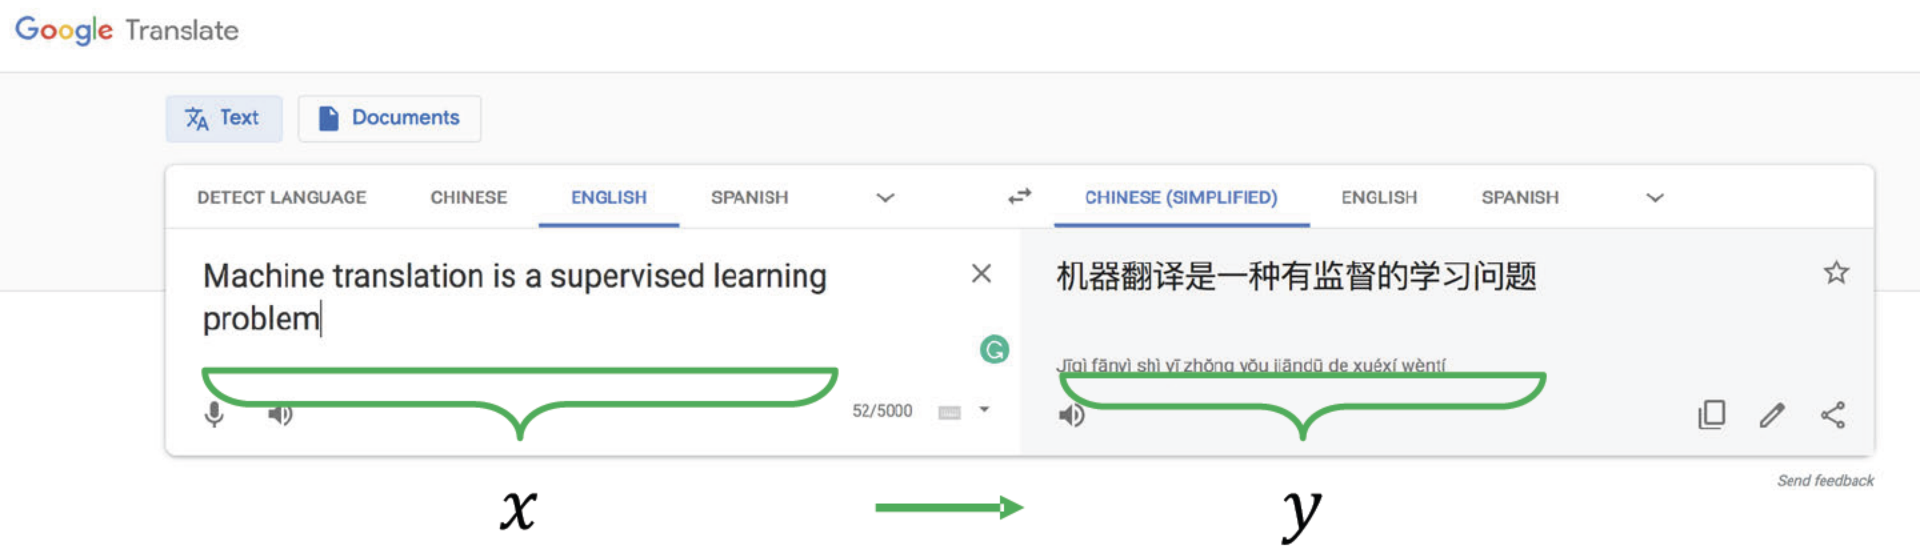

- Note: this course only covers the basic and fundamental techniques of supervised learning

Unsupervised Learning

Unsupervised Learning

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

supervised

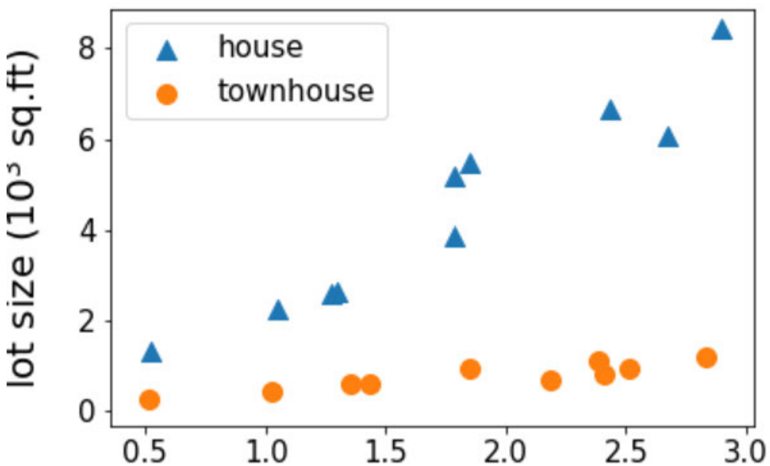

unsupervised

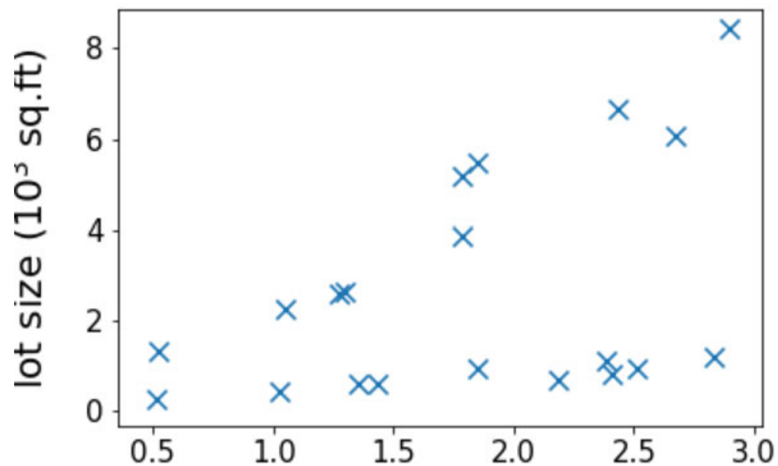

Clustering

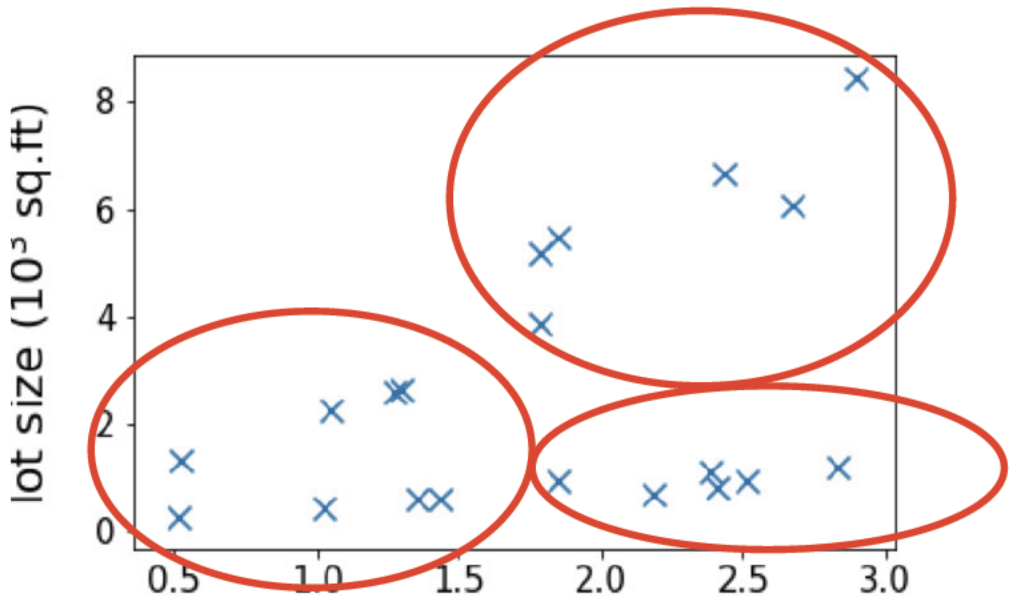

Clustering

➢ Lecture 12&13: k-mean clustering, mixture of Gaussians

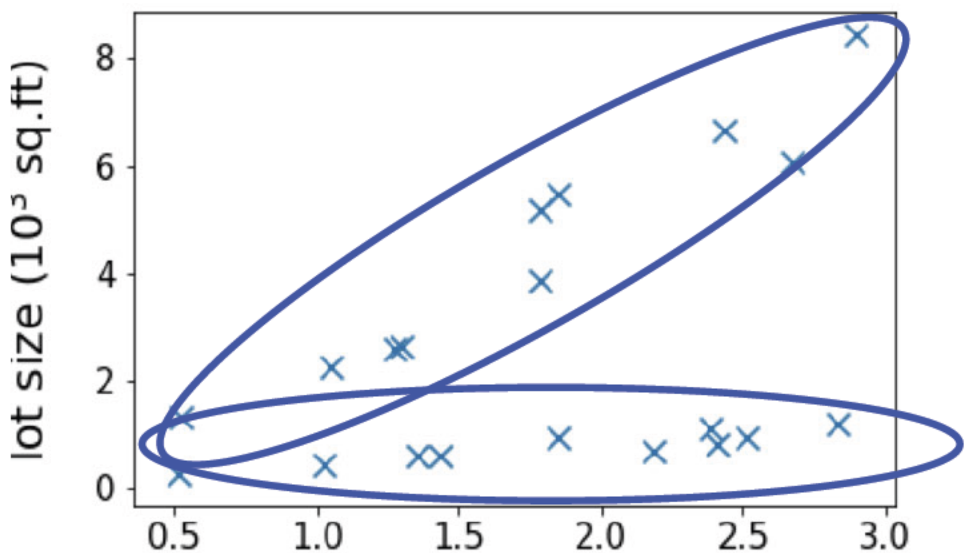

Clustering Genes

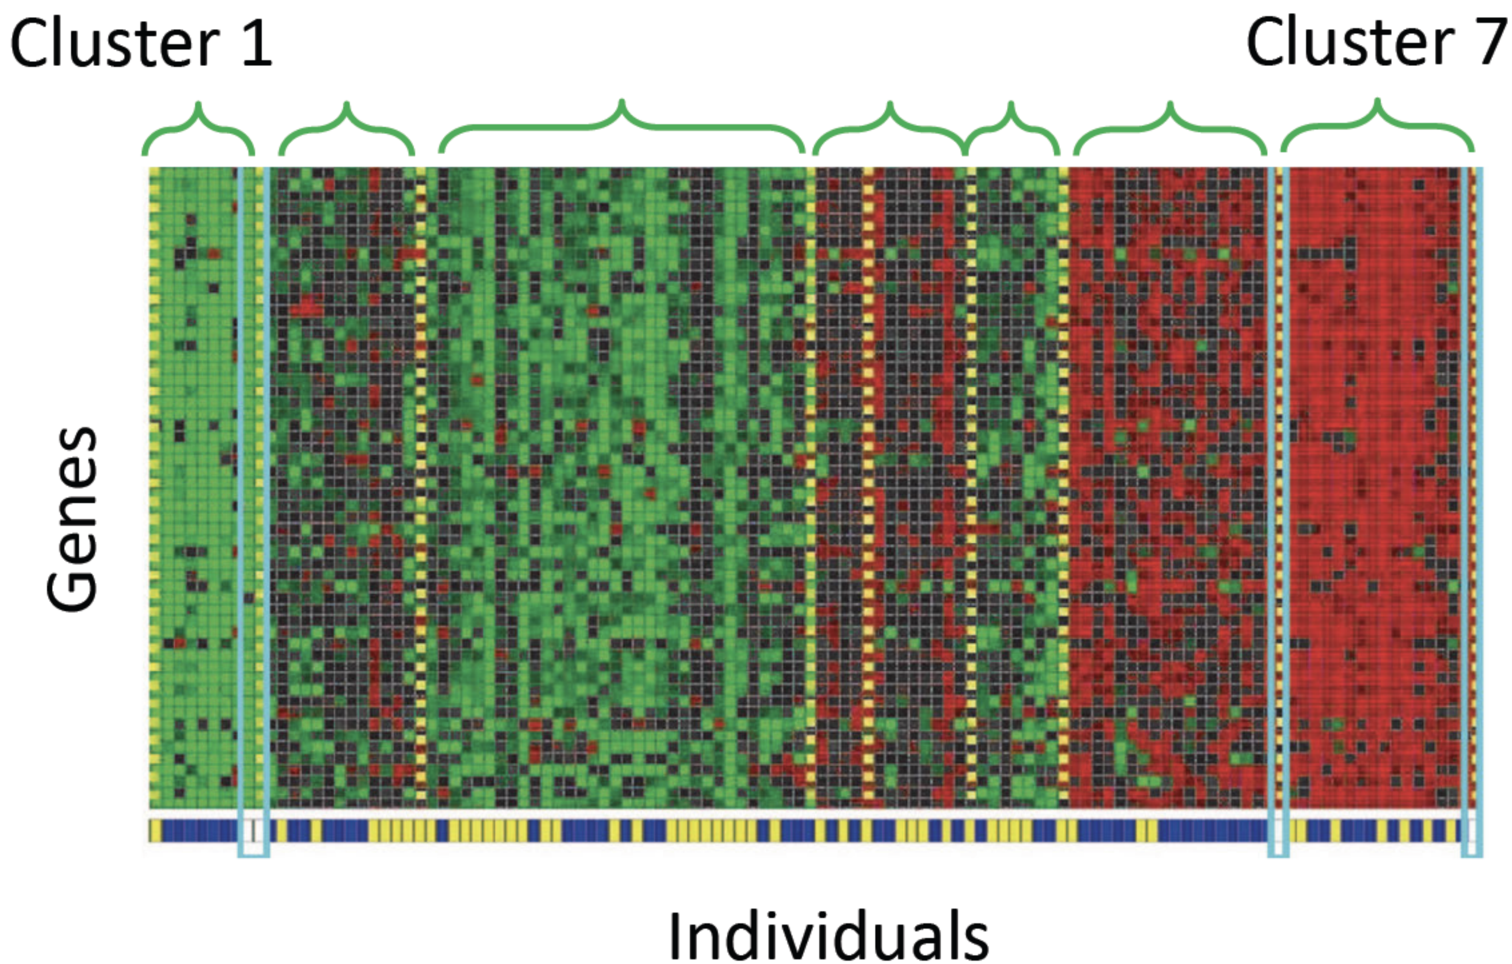

Identifying Regulatory Mechanisms using Individual Variation Reveals Key Role for Chromatin Modification. [Su-In Lee, Dana Pe'er, Aimee M. Dudley, George M. Church and Daphne Koller. '06]

Latent Semantic Analysis (LSA)

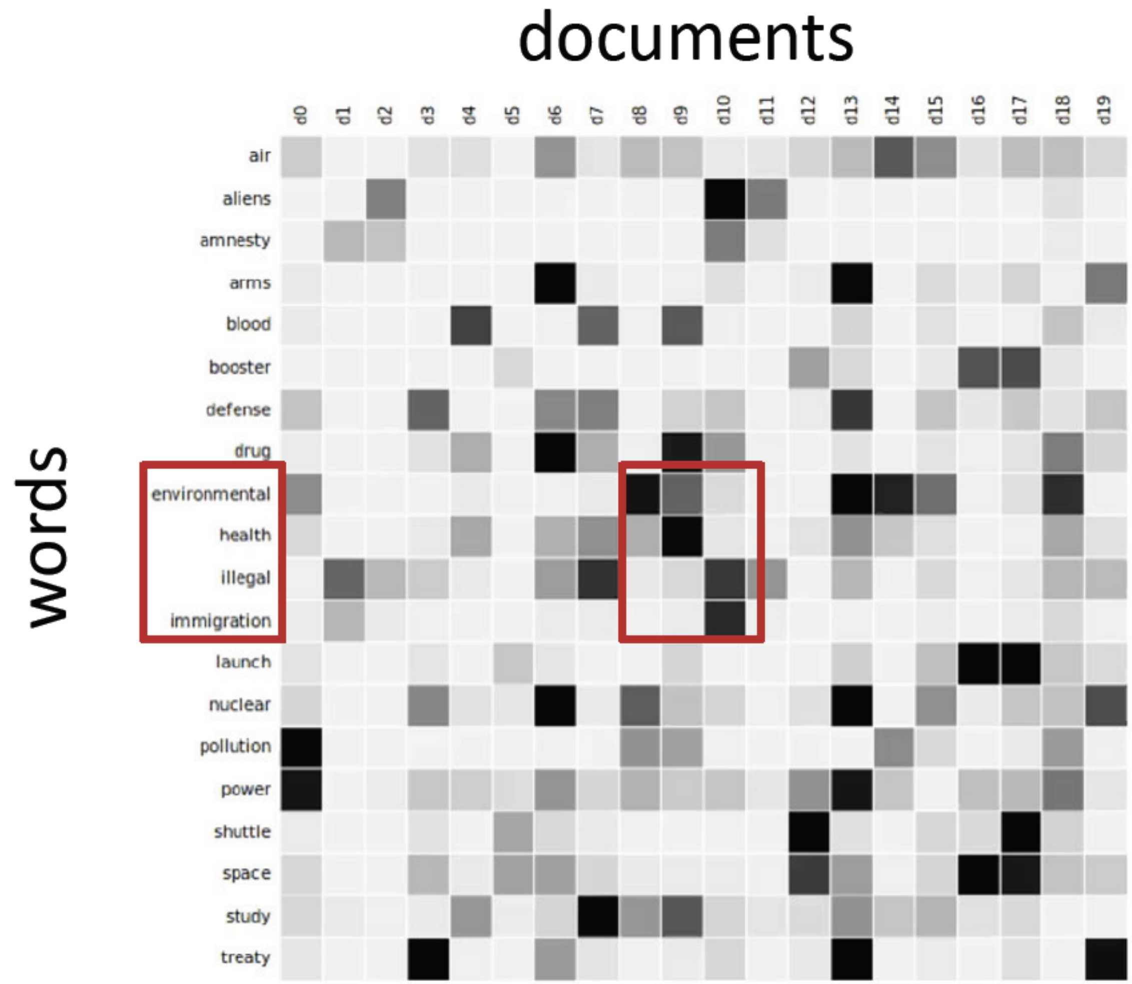

- Lecture 14: principal component analysis (tools used in LSA)

Image credit: https://commons.wikimedia.org/wiki/File:Topic_detection_in_a_document-word_matrix.gif

Word Embeddings

WIKIPEDIA
The Free ~~Encyclopedia~~
CORPUS

Unlabeled dataset

Represent words by vectors

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

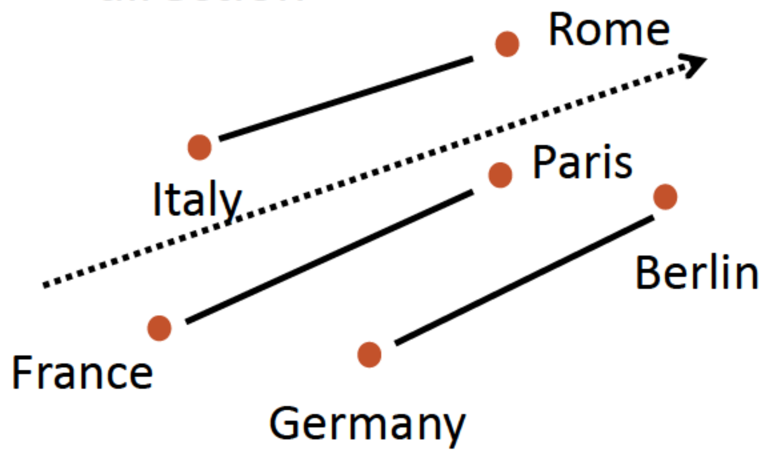

Word2vec [Mikolov et al'13]
GloVe [Pennington et al'14]

Clustering Words with Similar Meanings (Hierarchically)

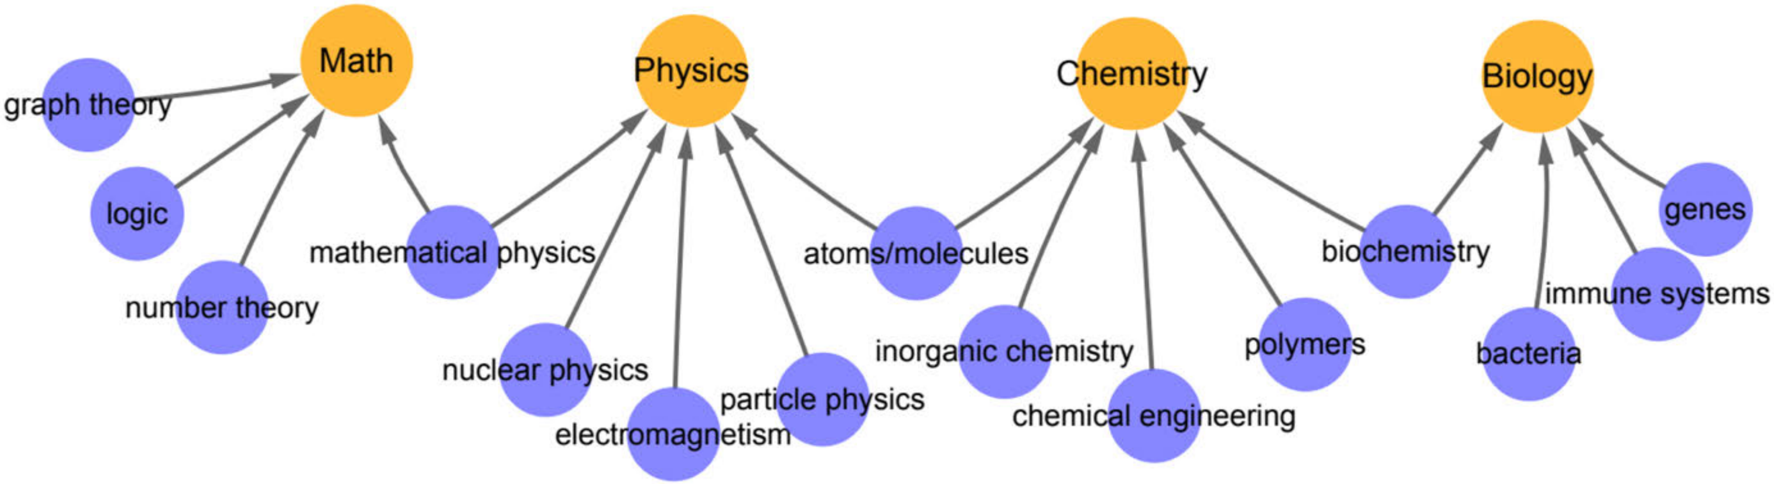

	logic deductive propositional semantics	graph subgraph bipartite vertex	boson massless particle higgs	polyester polypropylene resins epoxy	acids amino biosynthesis peptide
tag	*logic*	*graph theory*	*particle physics*	*polymer*	*biochemistry*

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

Large Language Models (Lecture 16)

- machine learning models for language learnt on large-scale language datasets
- can be used for many purposes

SYSTEM PROMPT (HUMAN-WRITTEN)

In a shocking finding, scientist discovered a herd of unicorns living in a remote, previously unexplored valley, in the Andes Mountains. Even more surprising to the researchers was the fact that the unicorns spoke perfect English.

MODEL COMPLETION (MACHINE-WRITTEN, 10 TRIES)

The scientist named the population, after their distinctive horn, Ovid's Unicorn. These four-horned, silver-white unicorns were previously unknown to science.

Now, after almost two centuries, the mystery of what sparked this odd phenomenon is finally solved.

Dr. Jorge Pérez, an evolutionary biologist from the University of La Paz, and several companions, were exploring the Andes Mountains when they found a small valley, with no other animals or humans. Pérez noticed that the valley had what appeared to be a natural fountain, surrounded by two peaks of rock and silver snow.

Pérez and the others then ventured further into the valley. "By the time we reached the top of one peak, the water looked blue, with some crystals on top," said Pérez.

Pérez and his friends were astonished to see the unicorn herd. These creatures could be seen from the air without having to move too much to see them – they were so close they could touch their horns.

Language Models are Few-Shot Learners [Brown et al.'20]
https://openai.com/blog/better-language-models/

Reinforcement Learning

- Learning to make sequential decisions


ALPHAGO

Albert learns to walk

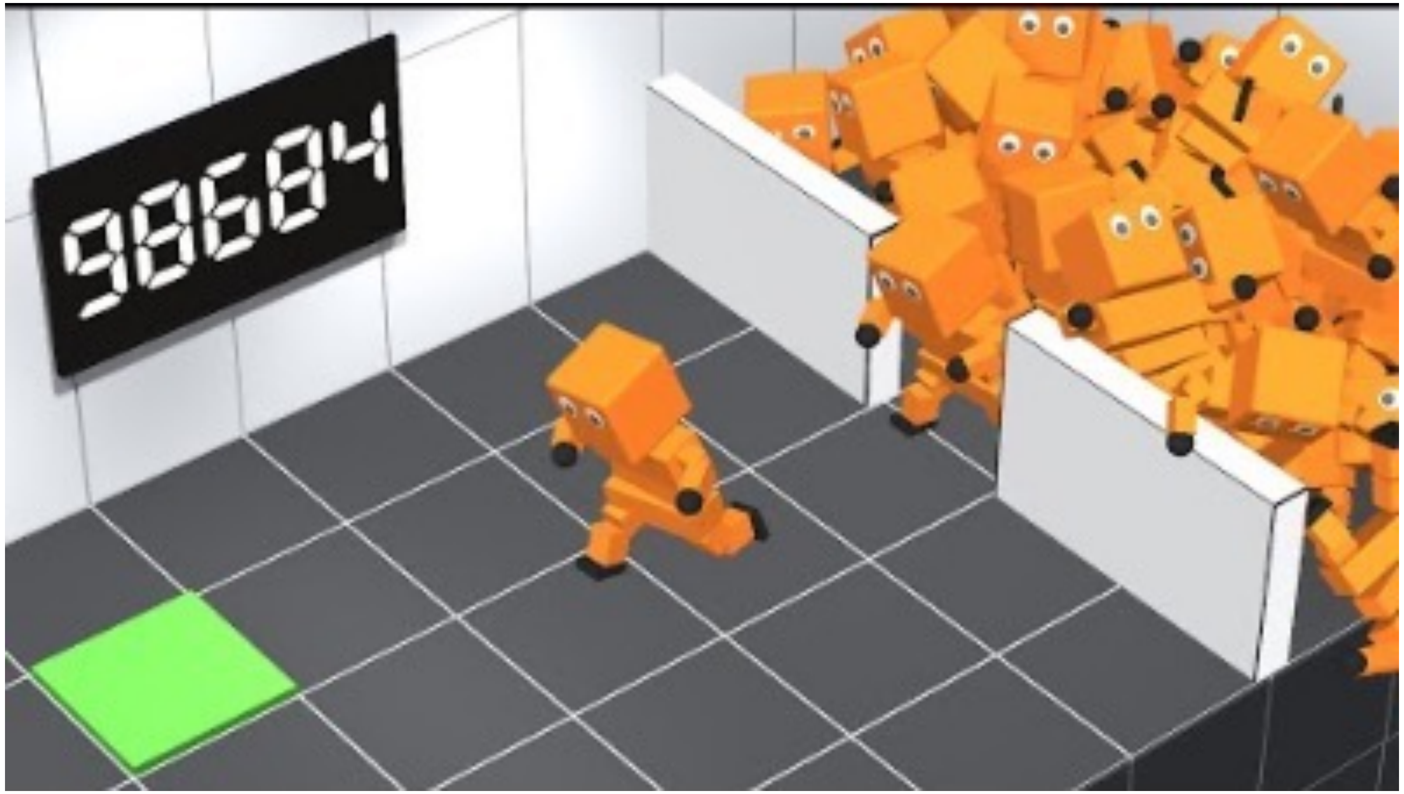

https://www.youtube.com/watch?v=L_4BPjLBF4E&t=95s

Albert learns to walk

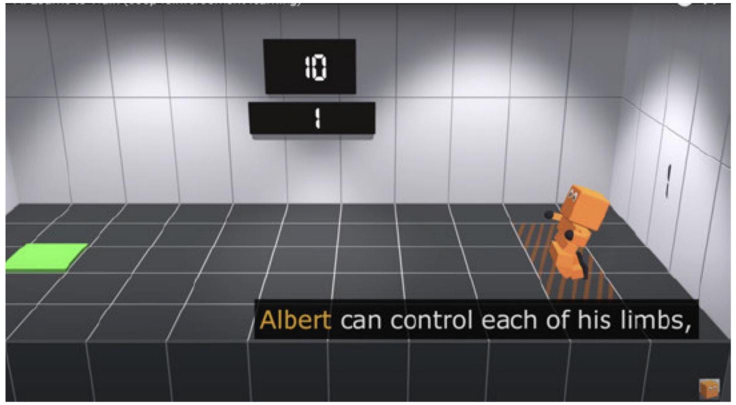

Iteration 1

Iteration 62

Iteration 163

Iteration 163

Albert learns to walk

Iteration 17

With new objective

Iteration 17

Iteration 163

Iteration 932

Reinforcement Learning

- The algorithm can collect data interactively

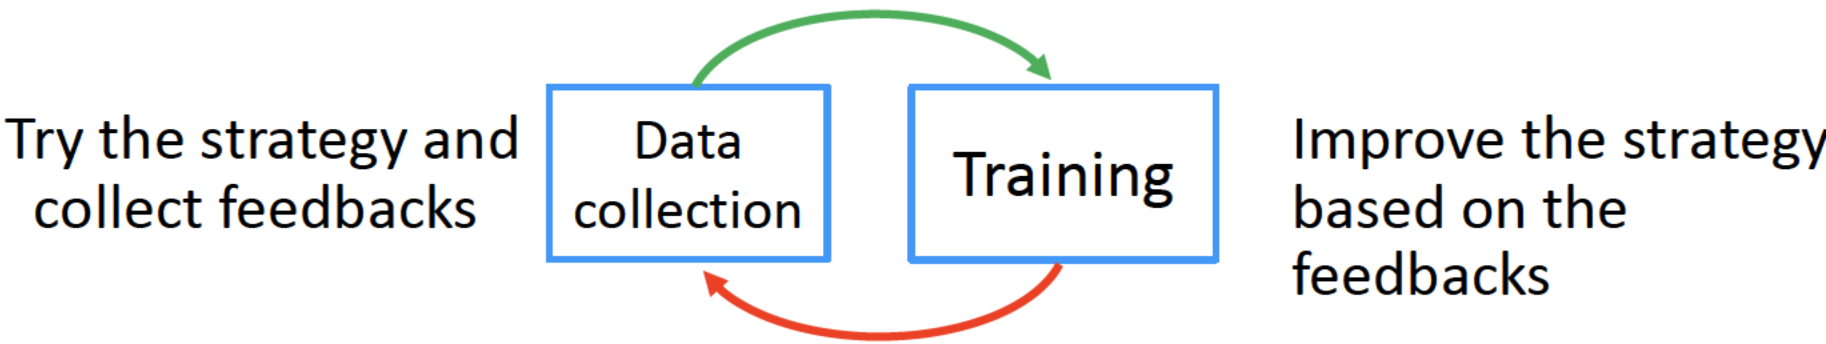

Taxonomy of Machine Learning (A Simplistic View Based on Tasks)

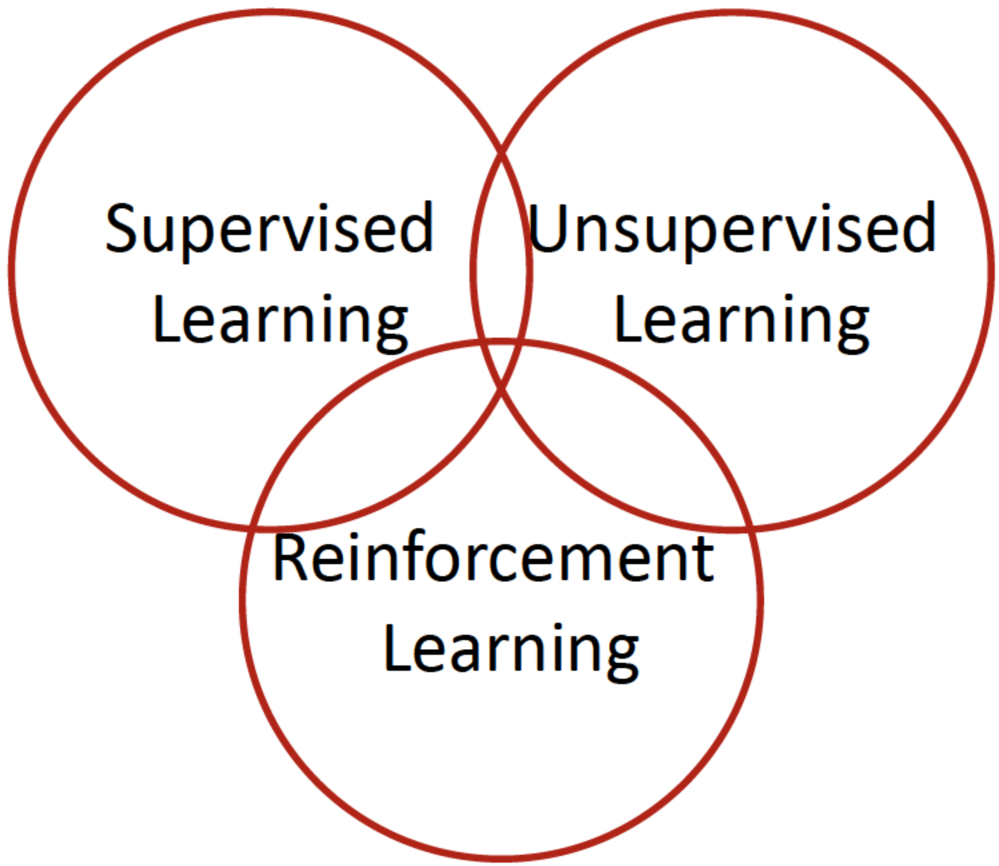

can also be viewed as tools/methods

Other Tools/Topics In This Course

- Deep learning basics

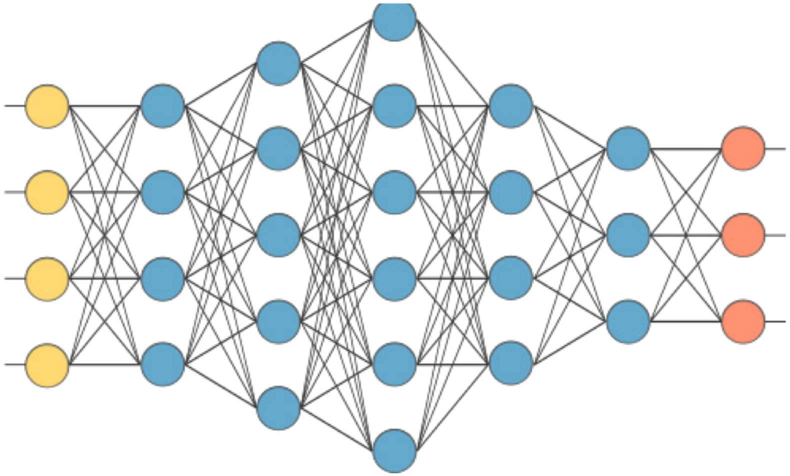

- Introduction to learning theory
  - Bias variance tradeoff
  - Feature selection
  - ML advice
- Broader aspects of ML
  - Robustness/fairness

Questions?

Thank you!