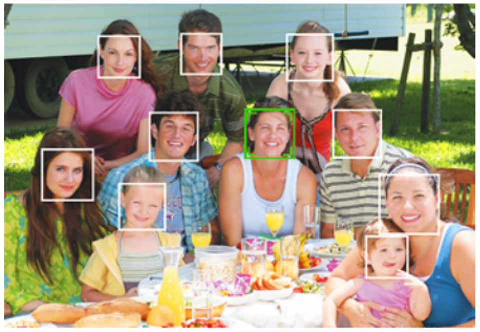

COMP 135

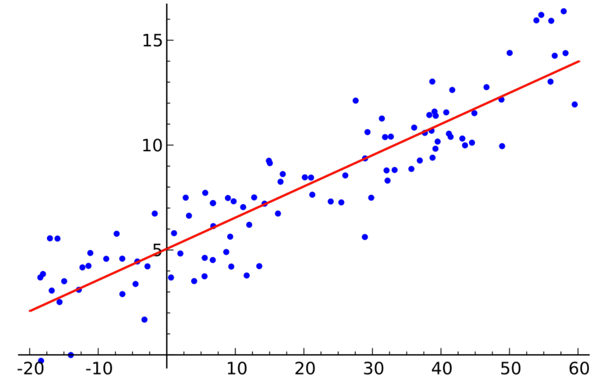

Introduction to Machine Learning

Prof. Michael C. Hughes ("Mike")

Fall 2020, First day of class

As you join, please check out:

* **Website:** https://www.cs.tufts.edu/comp/135/2020f/
  Read syllabus, skim schedule, waitlist info, etc.
* **Piazza** forum: https://www.piazza.com/tufts/fall2020/comp135
  Access code (today only): **validation2020**
  Ask LIVE questions throughout today's class (and every class)

Many slides attributable to: Emily Fox (UW), Finale Doshi-Velez (Harvard), Erik Sudderth (UCI), & Liping Liu (Tufts)

Today's Agenda

- Why take this course?
- What is Machine Learning?
- What skills/concepts will we learn?
- Who is teaching?
- How will we spend our time?

Q: Why should you take this course?

A: Machine Learning is everywhere! Those who know how to wield it effectively and *responsibly* can change the world.

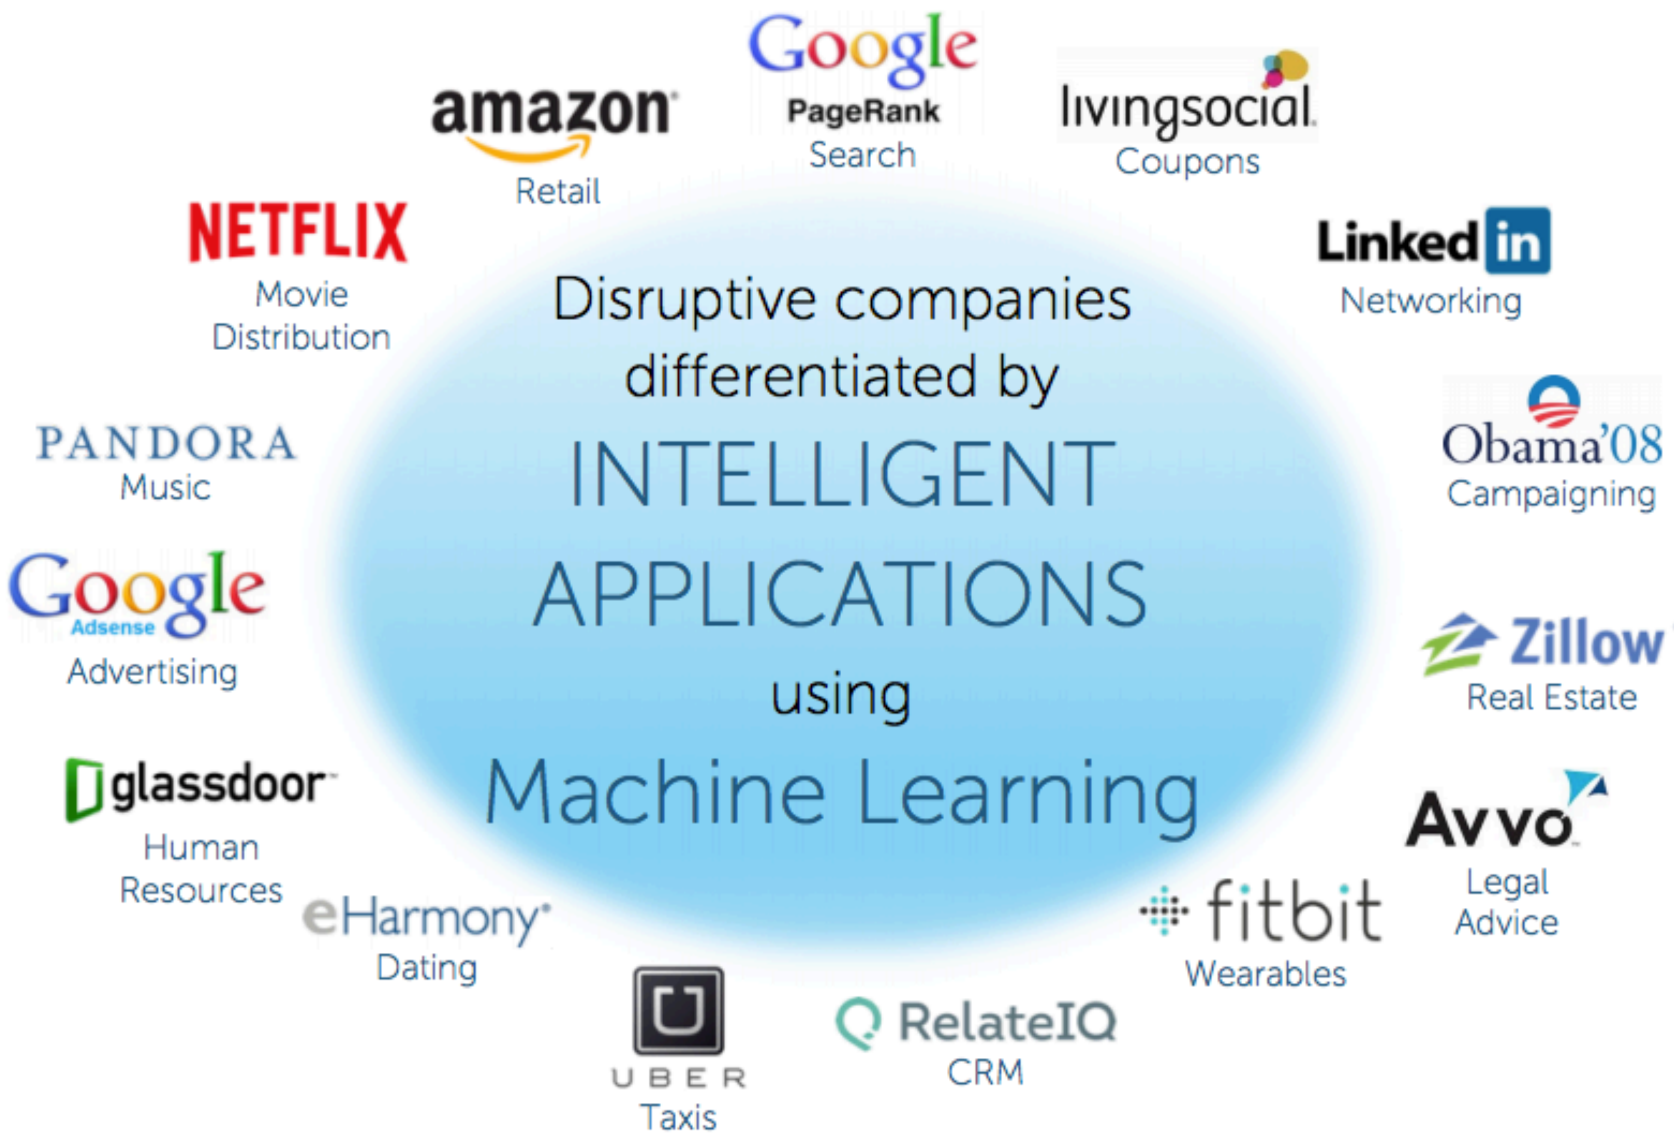

Image Credit: Emily Fox

Goals of this course

Our goal is to prepare students to effectively apply machine learning methods to problems that might arise in "the real world" -- in industry, medicine, government, education, and beyond.

Gain skills and *understanding* for a future as:

- Developer using ML “out-of-the-box”
- ML methods researcher

After taking this course, you will be able to:

- *Think systematically and ethically*
  - Compare/contrast each method's strengths & limitations
  - "Can ML solve this problem?"
  - "***Should*** ML solve this problem?"

- *Deploy and debug rapidly on real problems*
  - Hands-on experience with open-source libraries
  - Address issues in "real-world" data analysis
    - Numerical issues, convergence issues, class imbalance, missing values, etc.

- *Evaluate carefully and honestly*
  - Design experiments to assess generalization to never-before-seen data
  - Select task-appropriate performance metrics
  - Report confidence or uncertainty in performance numbers

- *Communicate insightfully and reproducibly*
  - Surface key insights via figures, tables, and text in a written report
  - Provide details for a peer to repeat your analysis and draw same conclusions

Q: What is Machine Learning?

Venn Diagram of Knowledge

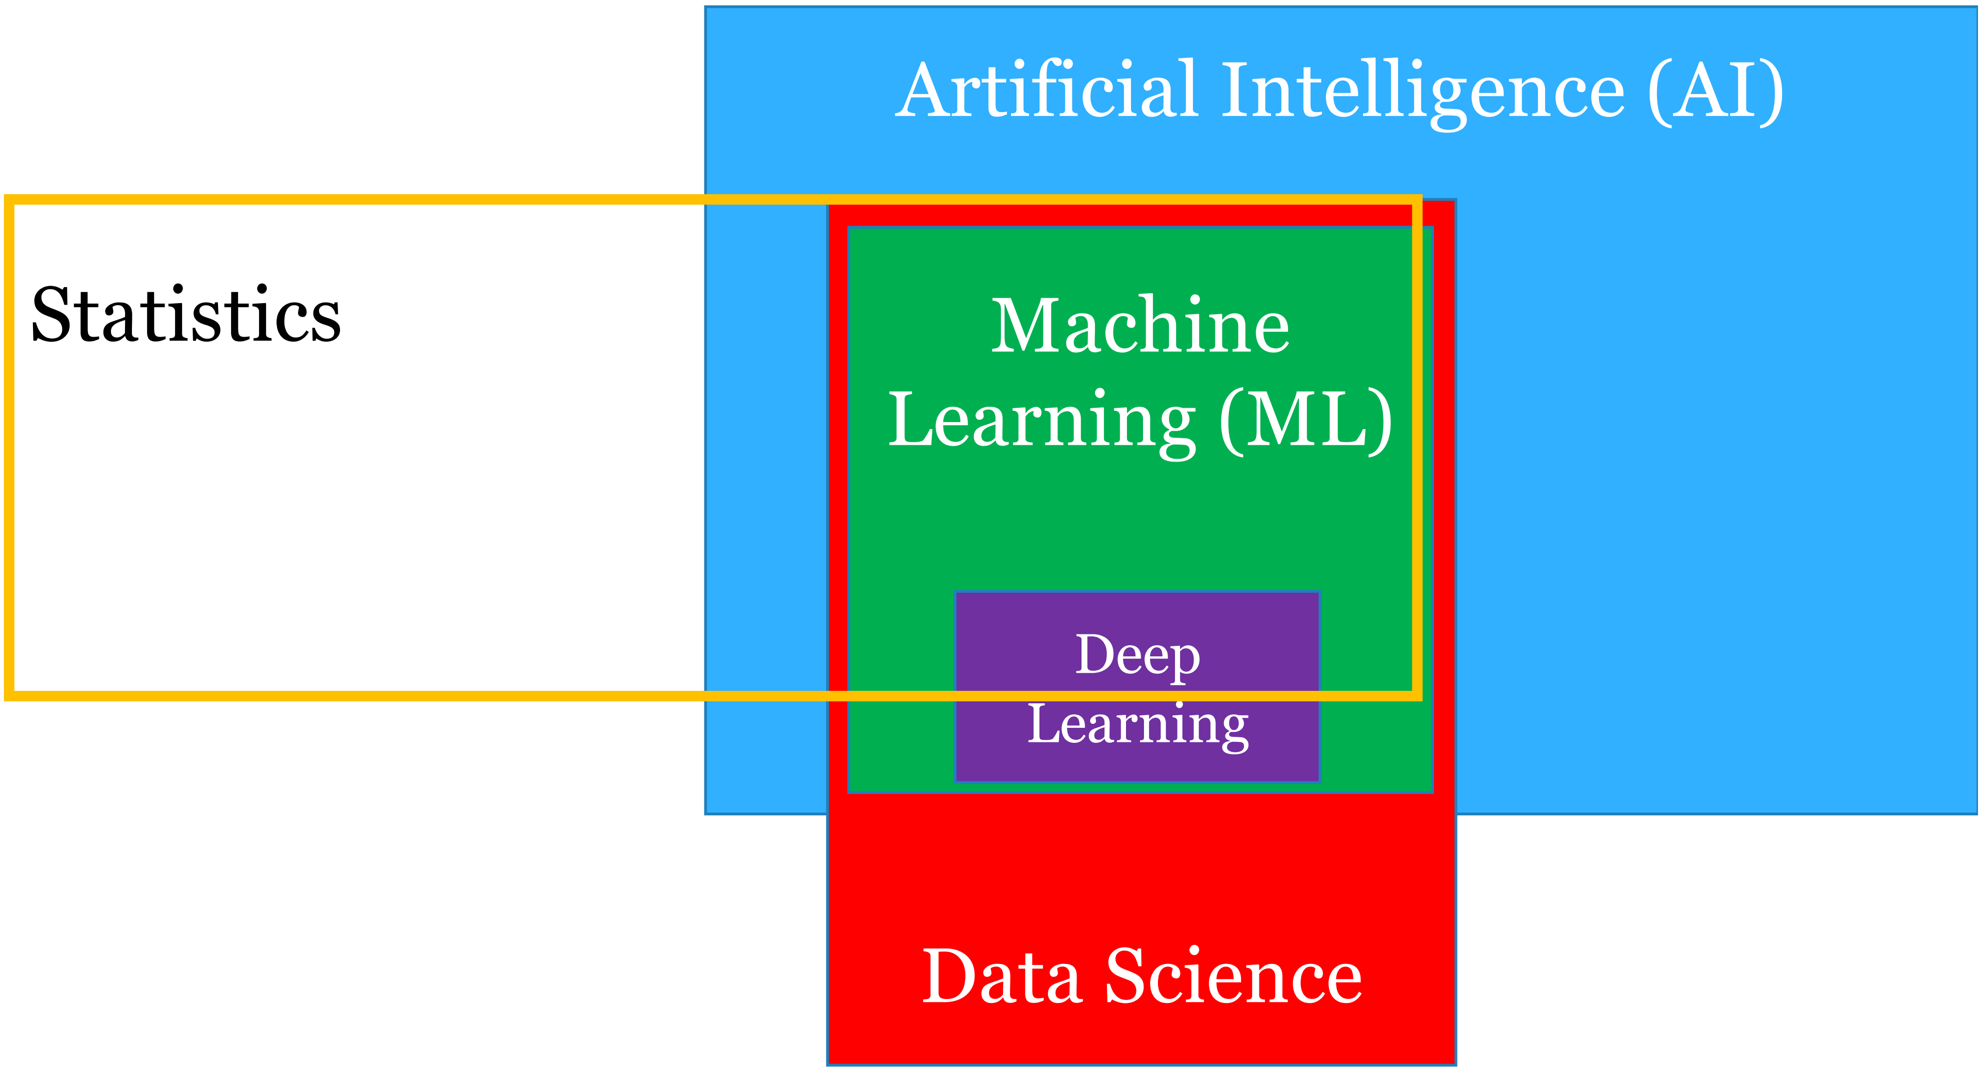

Artificial Intelligence (AI)

Study of “intelligent” systems, with many parts:

logic, planning, search, probabilistic reasoning, **learning from experience**, interacting with other agents, etc.

Alpha Go
Computer that can beat best human players of the game of “Go” (harder than chess)

DARPA Grand Challenge
Autonomous vehicles can navigate a real-world course without humans at the wheel

Machine Learning (ML)

Study of computer programs that **learn from experience/data** to perform a task

- Output: a *prediction, decision, or summary*

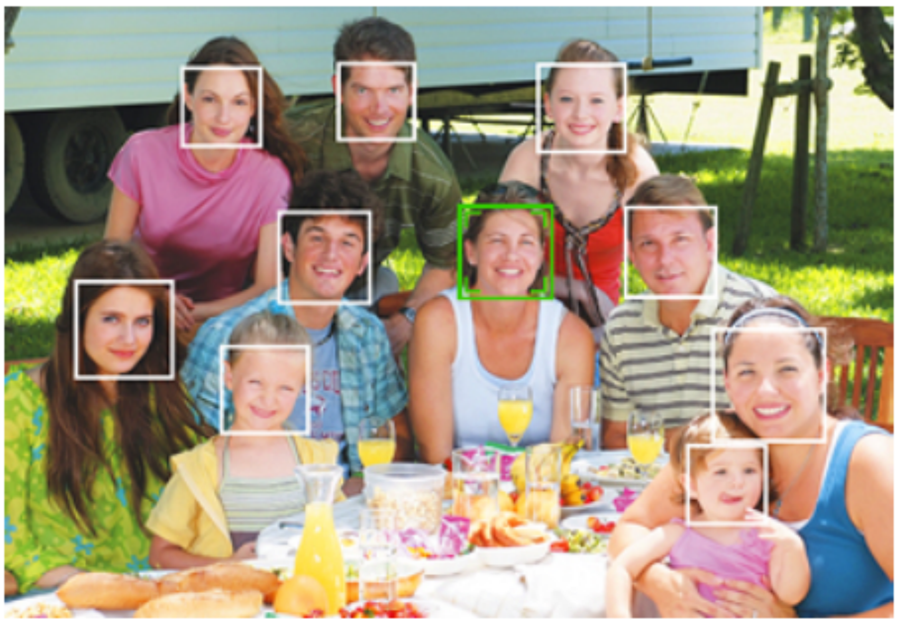

Face Detection
Predict location of human faces in natural images

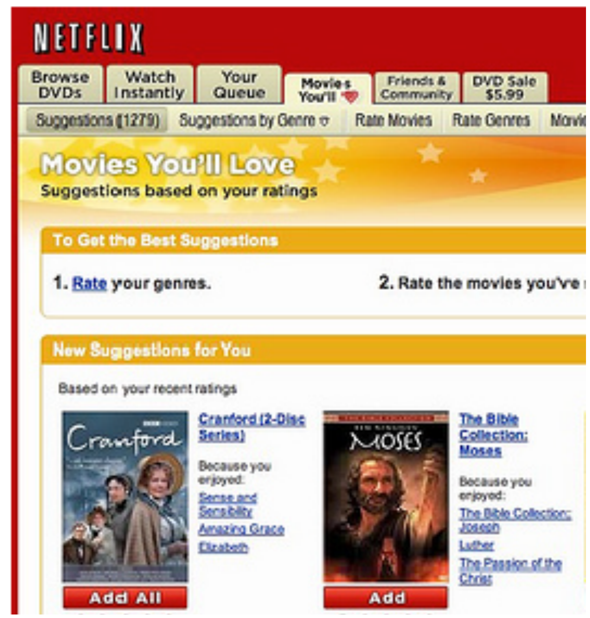

Movie Recommendation
Predict what to watch next

The Machine Learning Process

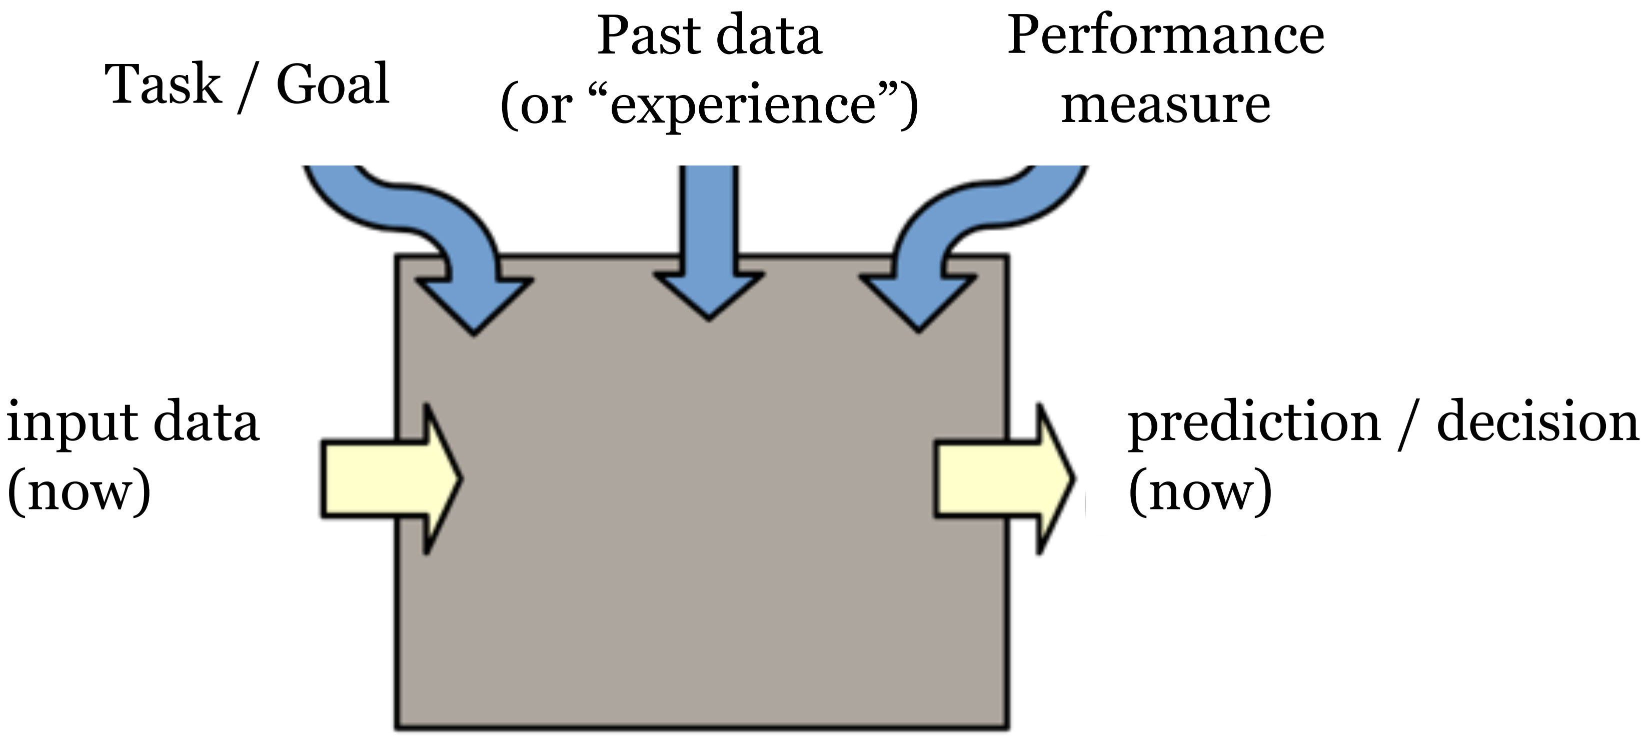

Q: What concepts will we learn?

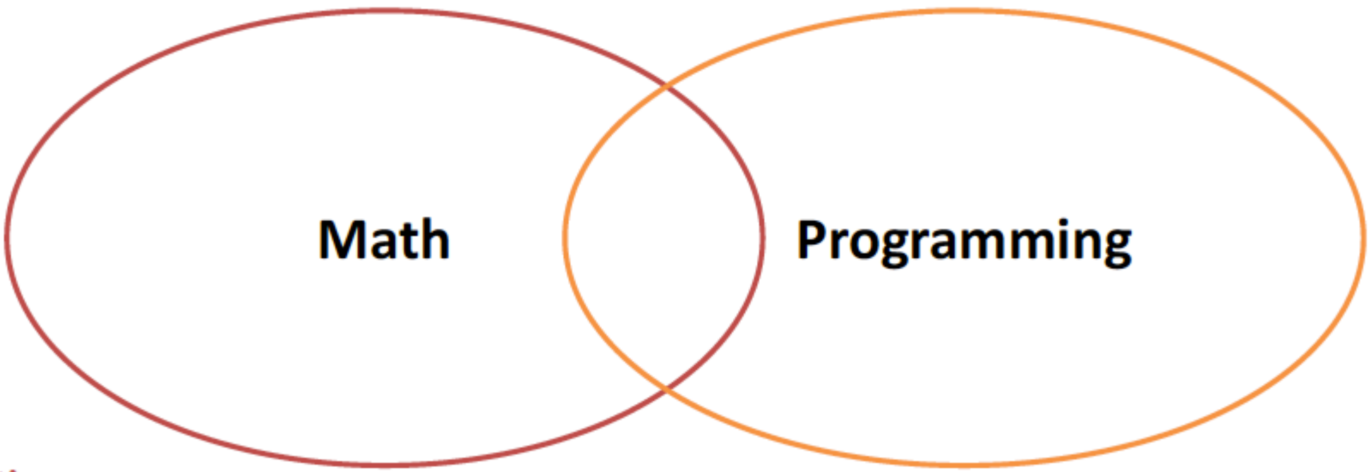

Statistics,
Probability,
Linear Algebra,
Optimization

Data Structures,
Algorithms,
Computational Complexity,
Data Management

What will we learn?

Supervised Learning

Unsupervised Learning

Reinforcement Learning

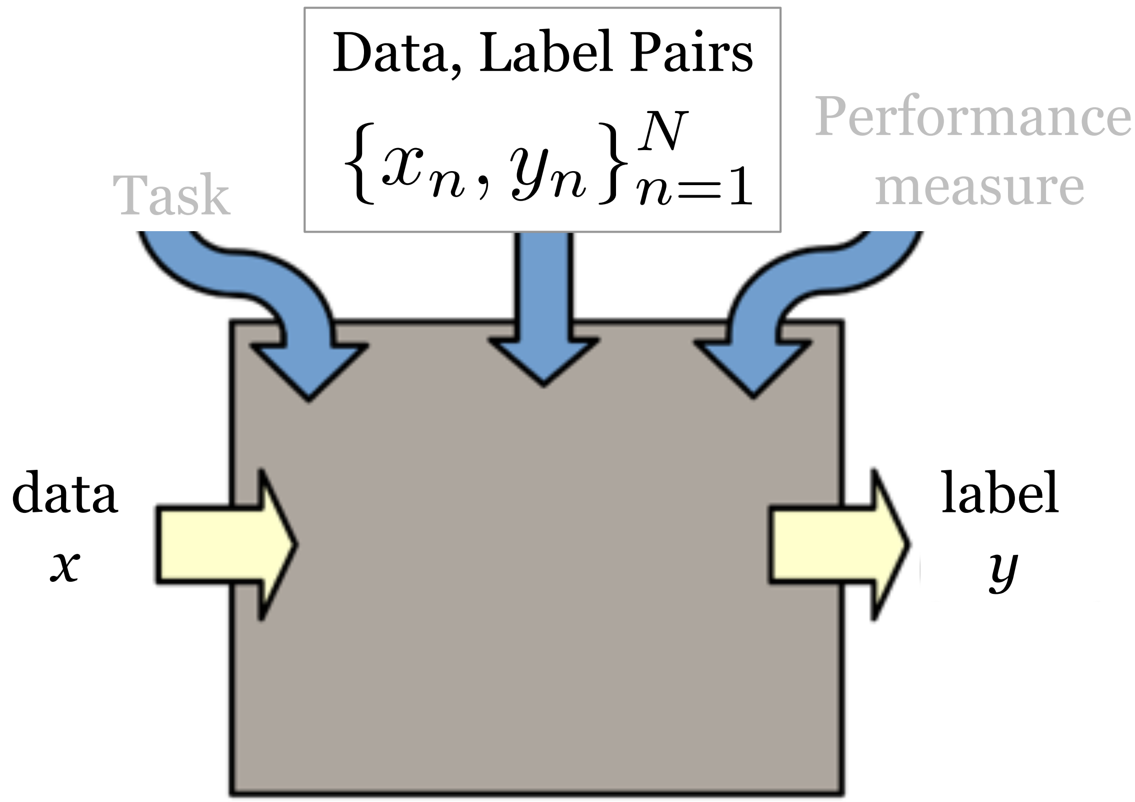

Task: Regression

Supervised Learning

regression

Unsupervised Learning

Reinforcement Learning

y is a continuous variable
e.g. sales in $$

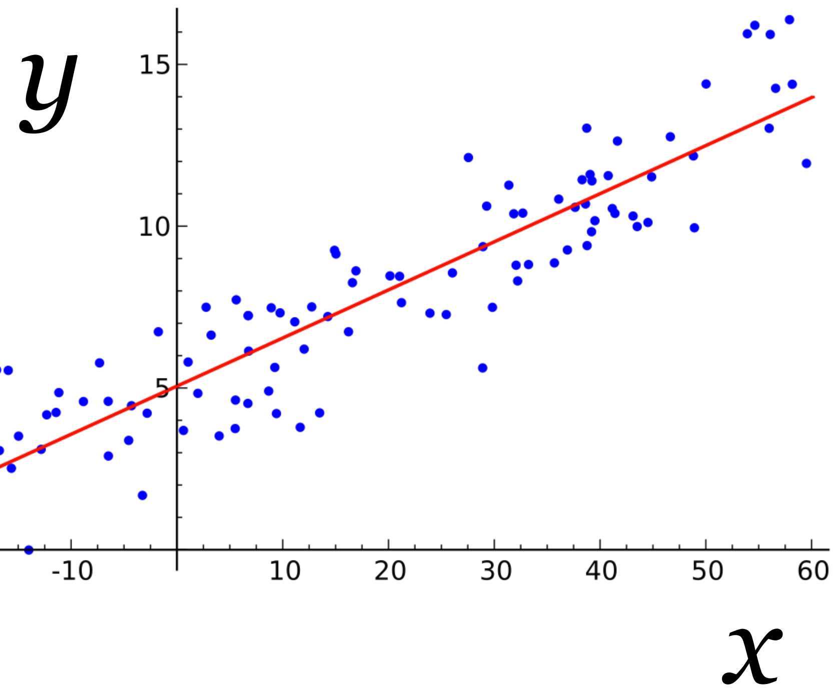

Regression Example: Uber

Supervised Learning

regression

Unsupervised Learning

Reinforcement Learning

Predictions of travel time, price, supply, demand

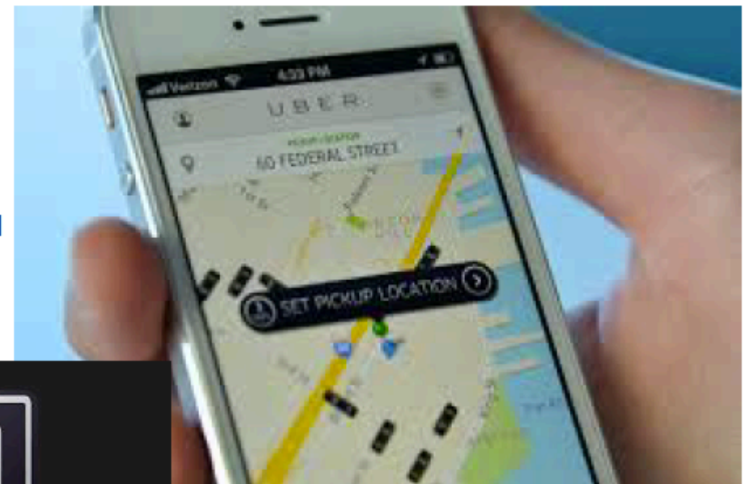

UBER

Regression Example: Uber

(Keith Chen)

Regression Example: Uber

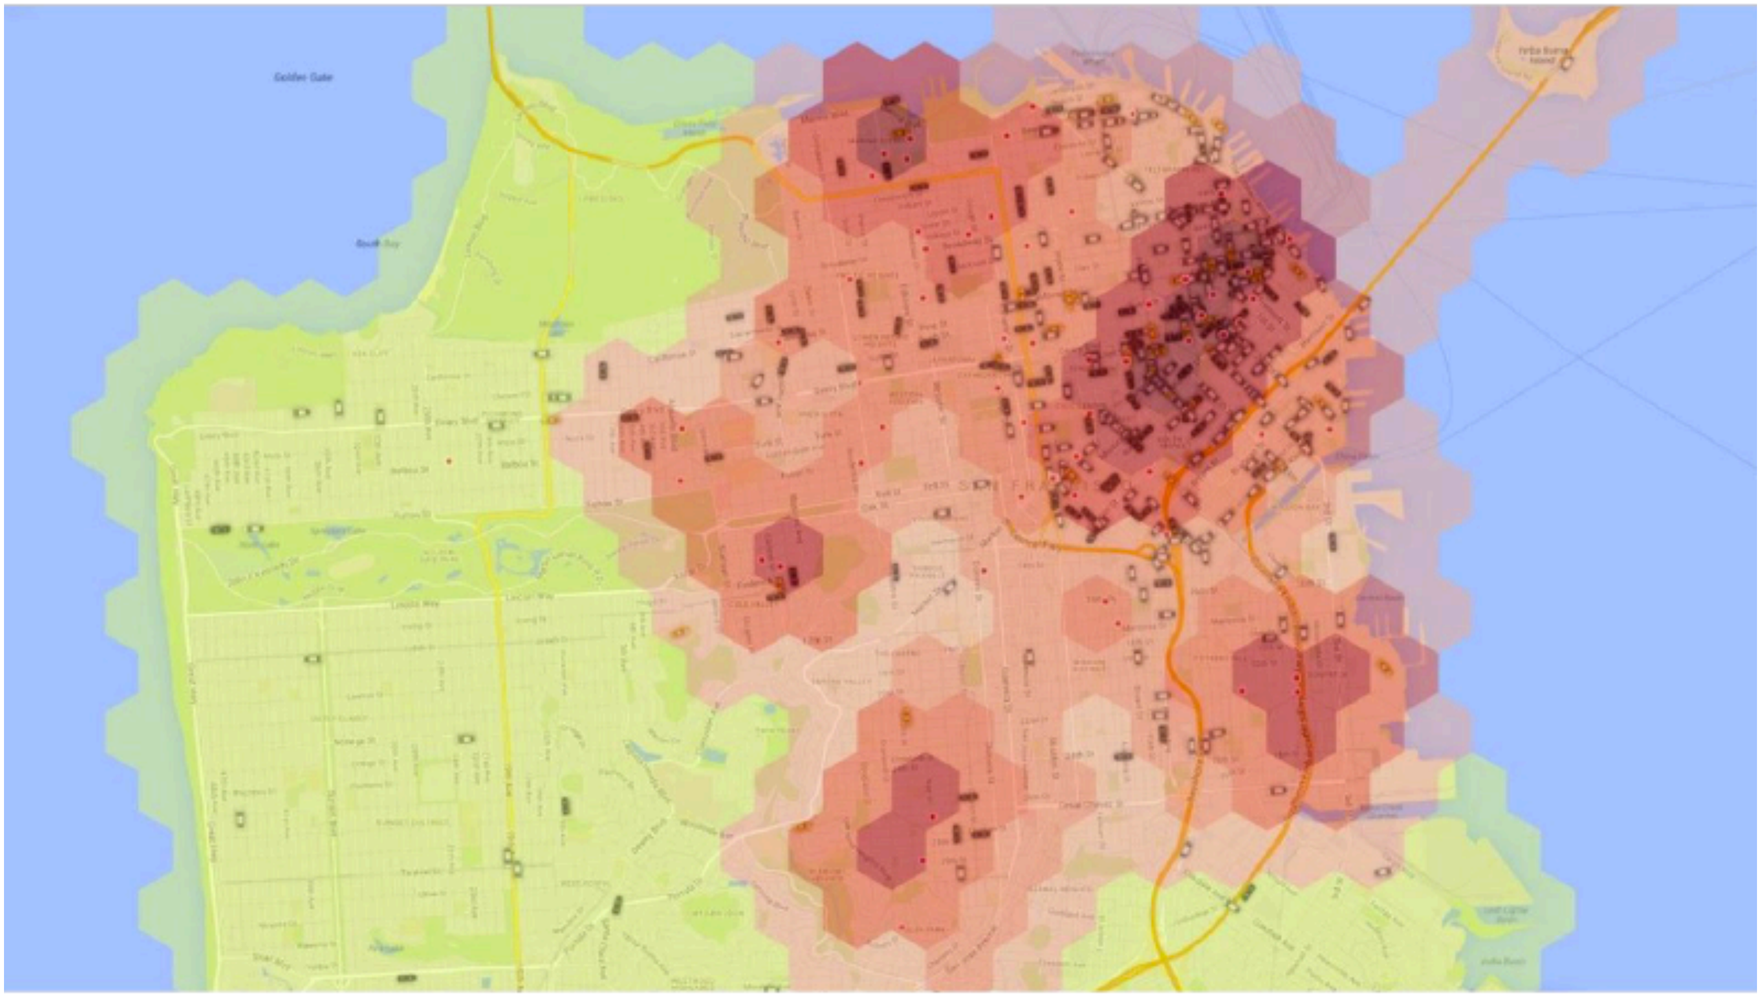

(Keith Chen)

Task: Classification

Supervised Learning

classification

Unsupervised Learning

Reinforcement Learning

y is a discrete variable (red or blue)

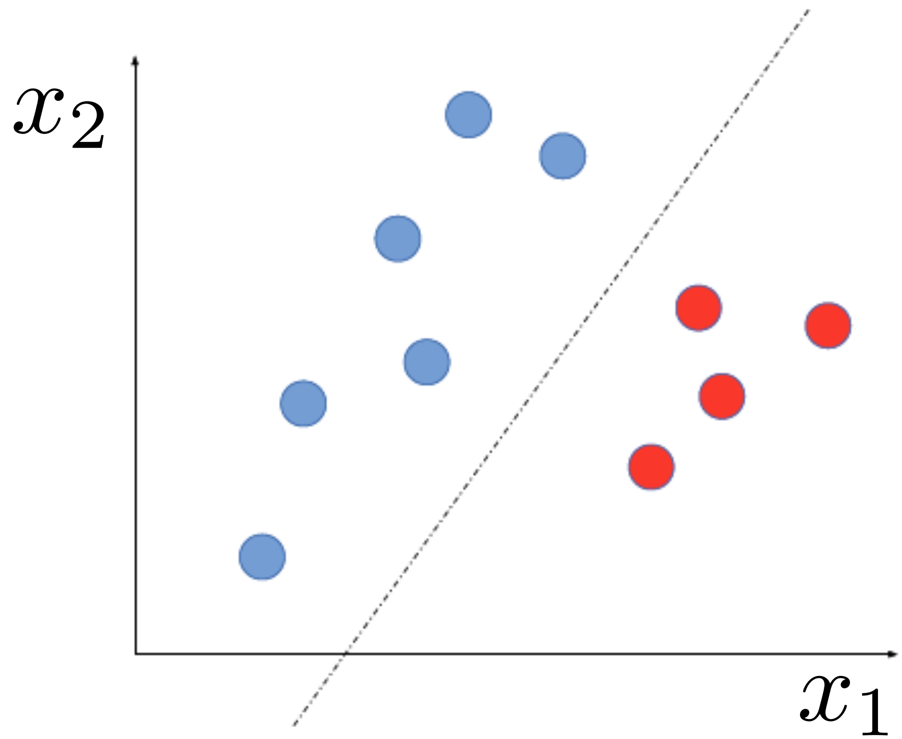

Classification Example: Swype

Predict words from keyboard trajectories

Novel Product:
An easier way
to input text on
mobile devices

What will we learn?

Supervised Learning

Unsupervised Learning

Reinforcement Learning

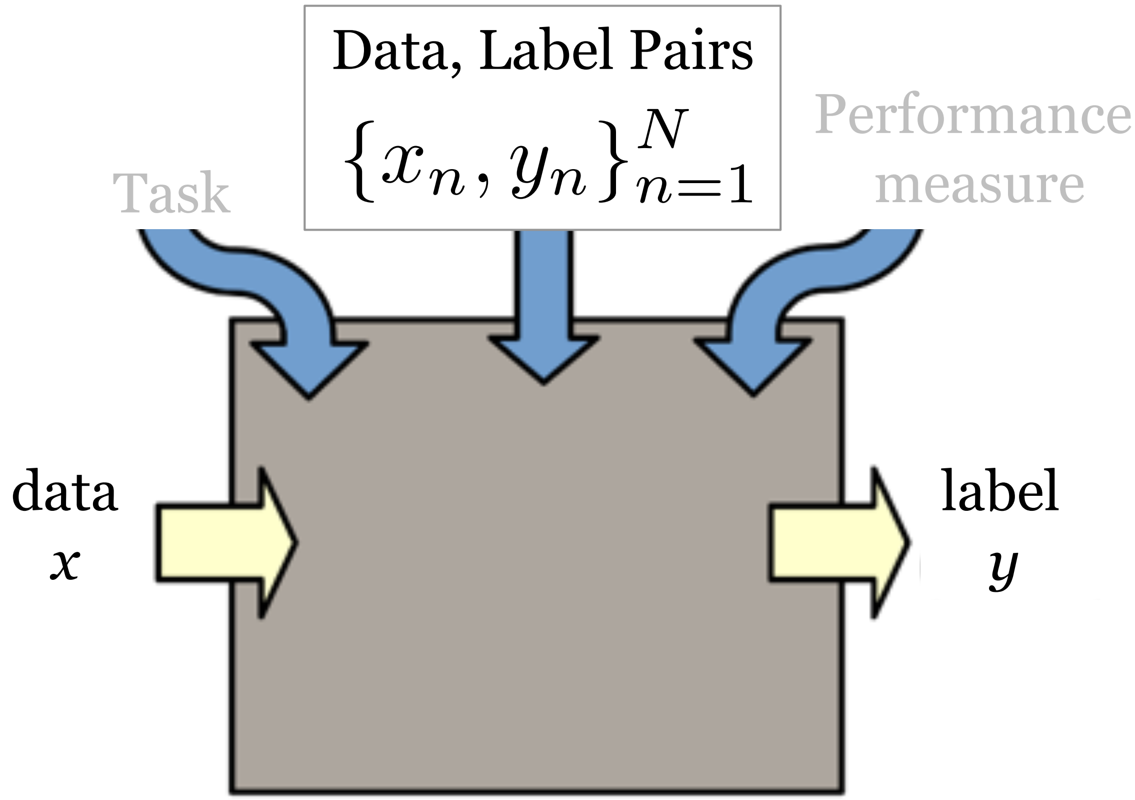

What will we learn?

Supervised Learning

Unsupervised Learning

Reinforcement Learning

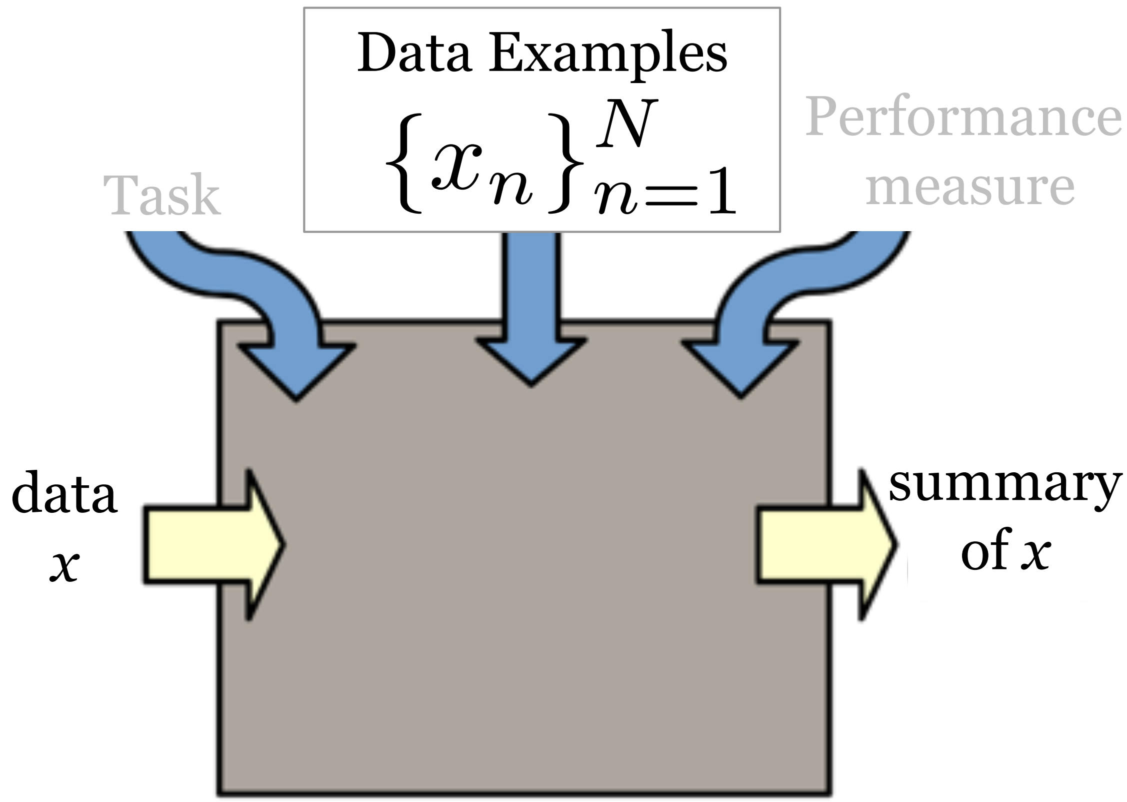

Task: Embedding

Supervised Learning

Unsupervised Learning

embedding

Reinforcement Learning

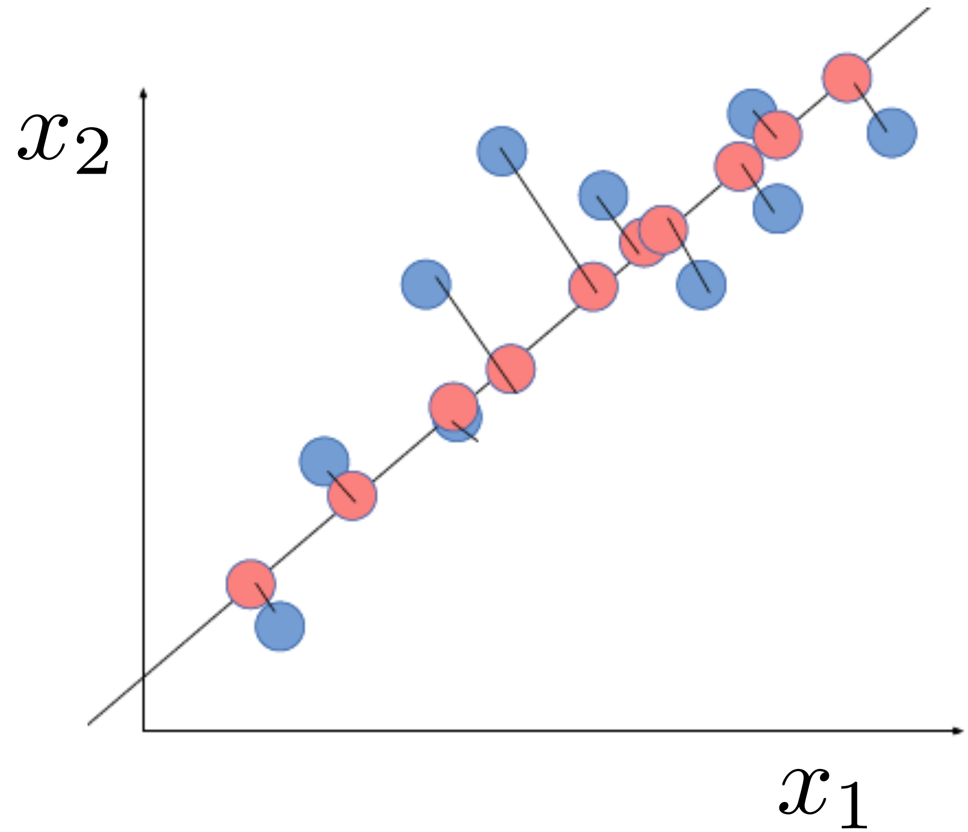

Example: Genes vs. geography

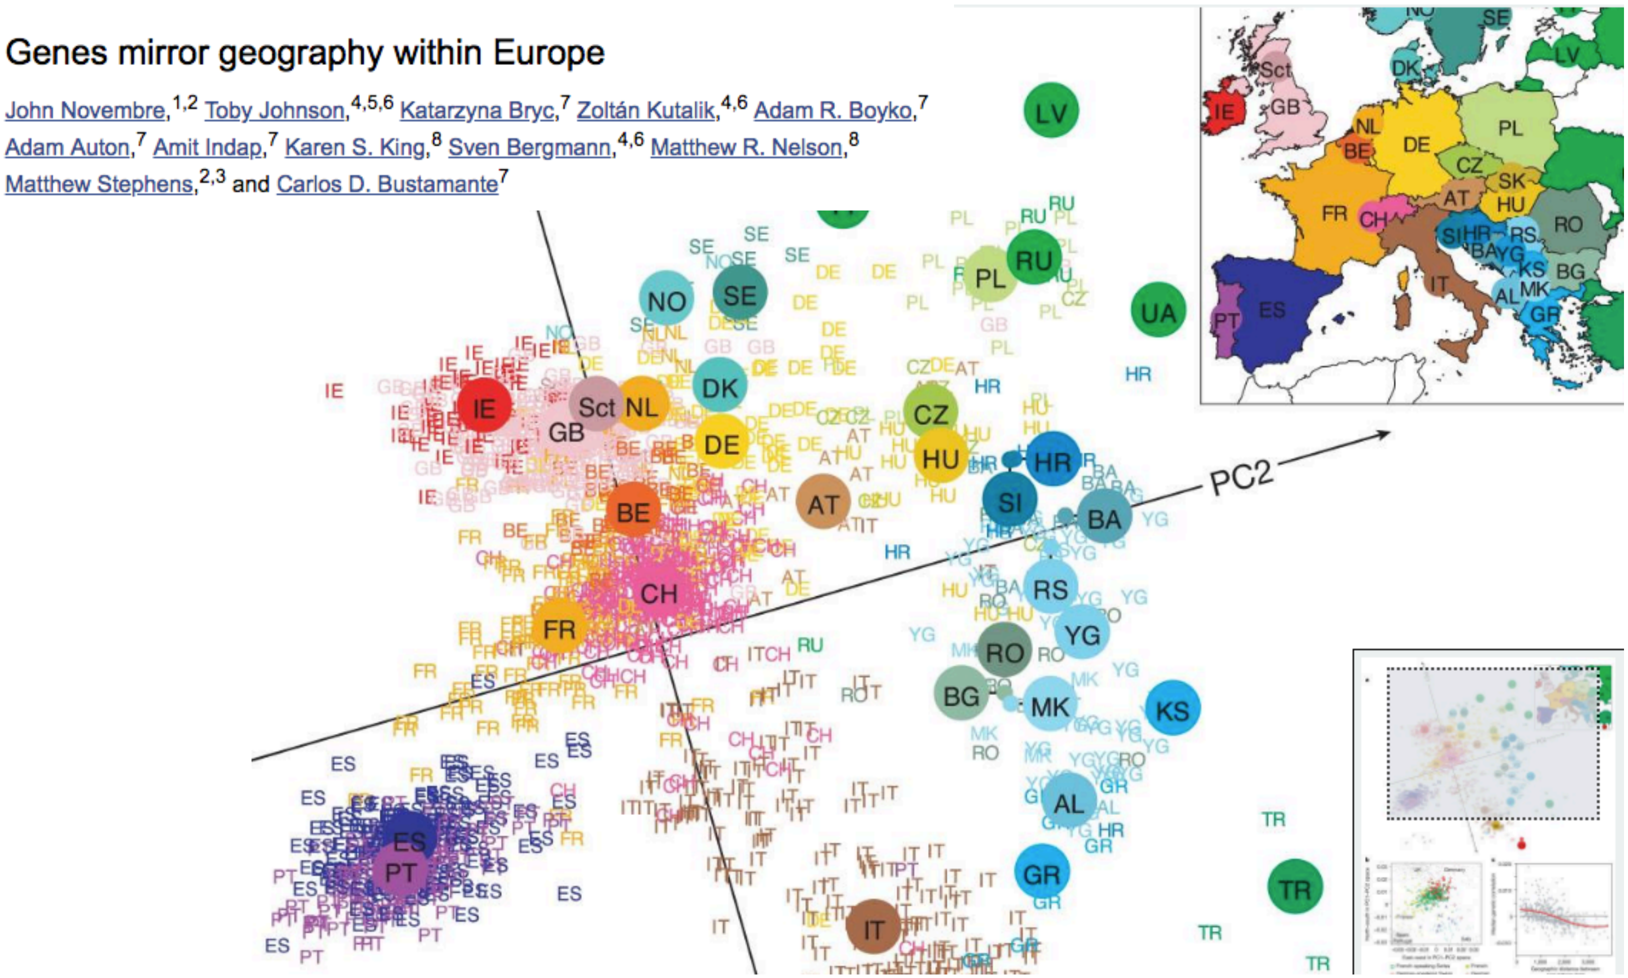

Task: Recommendation

Supervised Learning

recommendation

rvised Learning

Reinforcement Learning

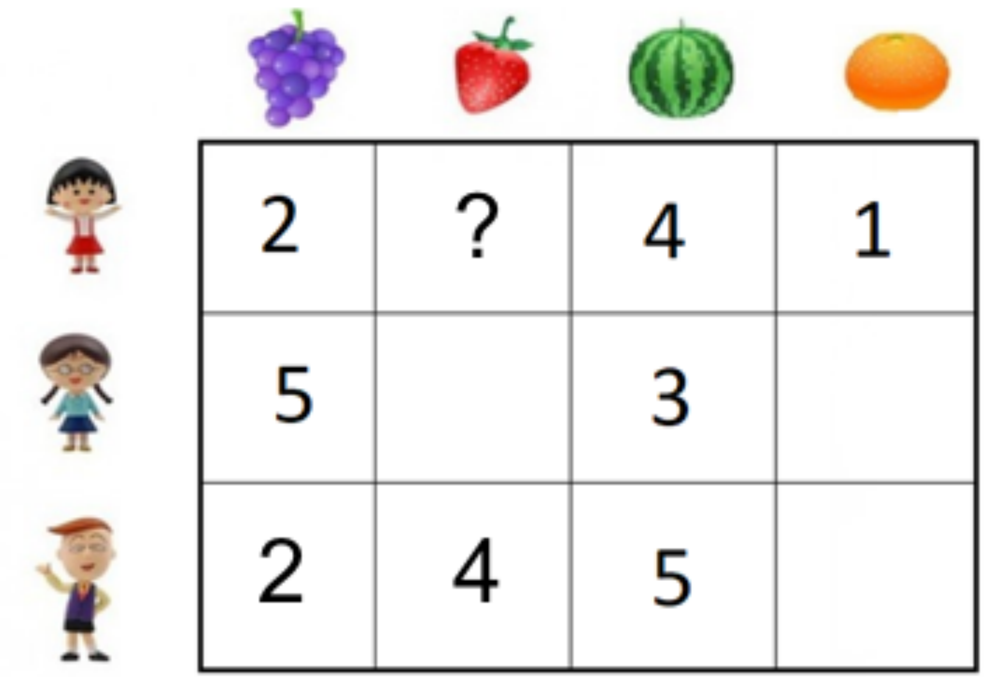

Recommendation Example

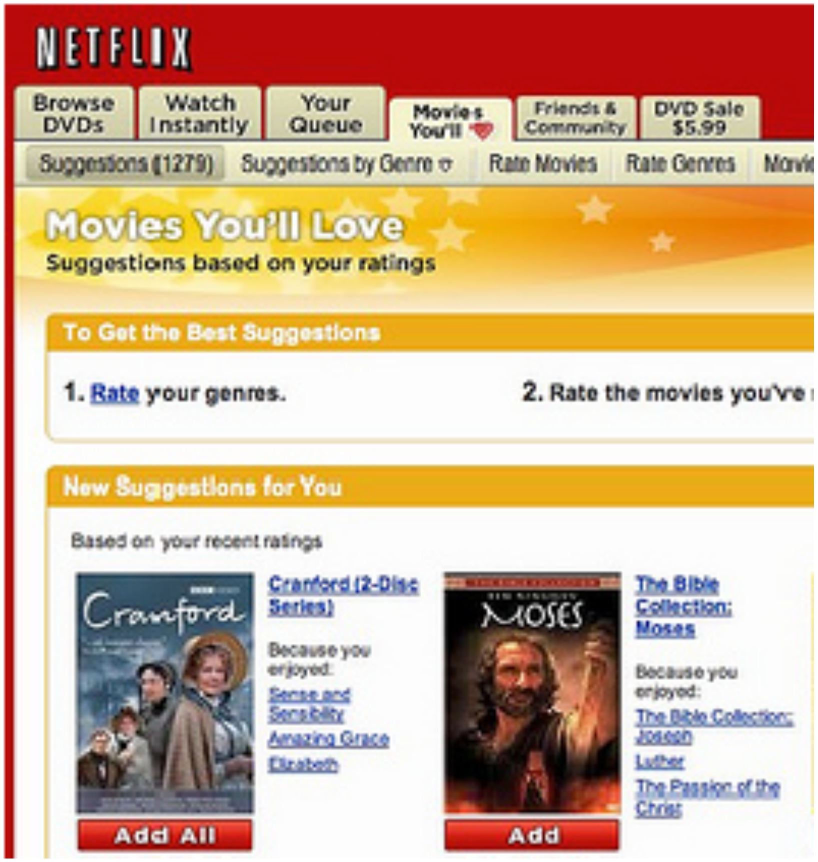

What will we learn?

Supervised Learning

Unsupervised Learning

Reinforcement Learning

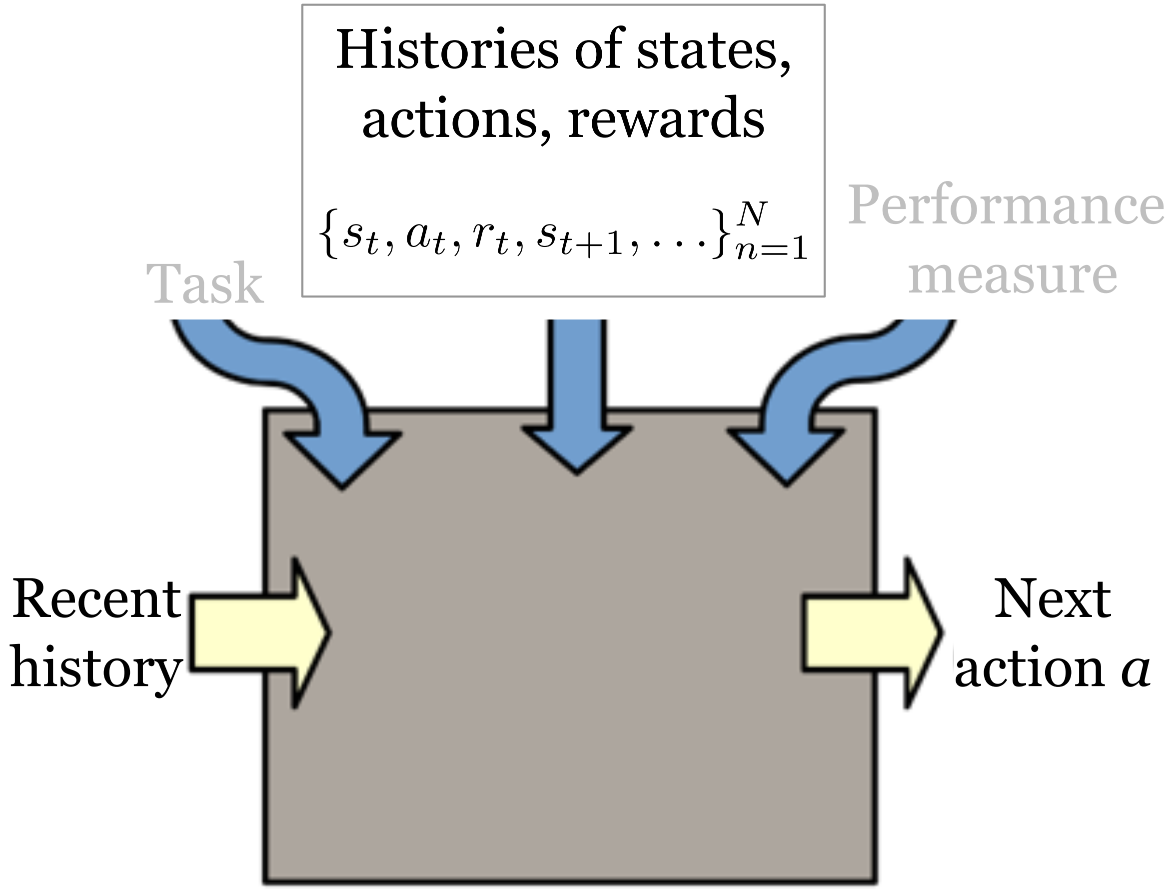

RL example: Pancake robot

Peter Kormushev, Imperial College

What **won’t** we cover?

- Clustering
- Probabilistic models
- Graphical models

- Active learning
- Transfer learning
- Semi-supervised learning
- Learning theory
- ….. lots more