# CSE 417T: Introduction to Machine Learning

Prof. Sanmay Das
Spring 2019

# Plan for today

Welcome & introductions

What is this class about?

Class logistics

# What is machine learning?

“Enabling computers to learn from data”

Supervised Learning: Generalizing from seen data to unseen data

Unsupervised Learning: Finding patterns in input data

Reinforcement Learning: Learning how to act, based on rewards for actions

# Supervised Learning

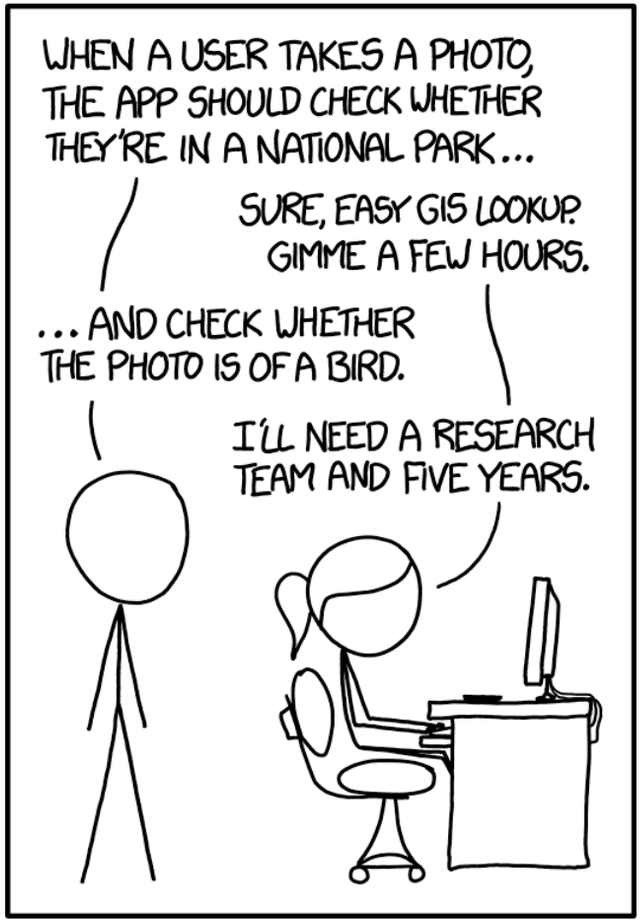


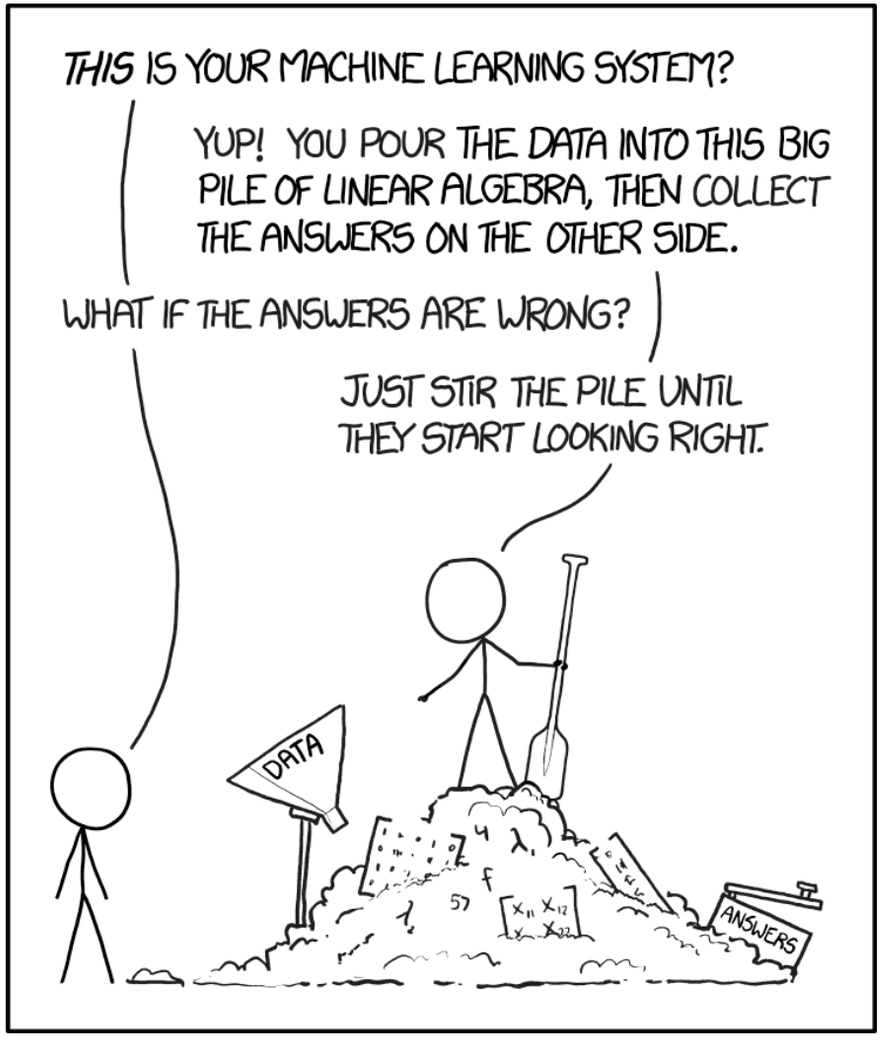

# Supervised Learning

## Can you define a tree?

A brown trunk coming up from the ground, with branches extending out?

# Are these trees?

Hard to define

“Know it when I see it”

I’ve **learned** it from data!

# The general supervised learning problem

Unknown target function: $f : \mathcal{X} \to \mathcal{Y}$
Training data: $(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \ldots, (\mathbf{x}_N, y_N)$
where $y_i = f(\mathbf{x}_i)$
Want to learn $g$ "close to" $f$

Two central questions:

- How do we learn $g$?
- What can we say about how close $g$ is to $f$?

# A concrete example: Credit approval

You apply for a credit card

Bank decides whether to approve or deny

What is the form of each *(x, y)* example?

What are we trying to learn, and from what?

- We have past data on customers, and want to learn the “ideal credit approval function”
- *x*’s can consist of demographic, income, personal data
- *y*’s are some relevant outcome

# How can I learn *f?*

- Pick a *hypothesis set* $\mathcal{H} = \{h_1, h_2, \ldots, \}$
- Use a *learning algorithm* to select a hypothesis $g$ from $\mathcal{H}$ based on the training data

The choice of $\mathcal{H}$ and the learning algorithm are deeply tied to each other

# A linear hypothesis space

Suppose we have data on annual income ($x_1$), debt ($x_2$), average income in ZIP code ($x_3$). A possible hypothesis space can be expressed as $(w_1, w_2, w_3, t)$ where the credit approval function is:

$$\text{Approve if } \sum_{i=1}^{3} w_i x_i \geq t$$

$$\text{Deny if } \sum_{i=1}^{3} w_i x_i < t$$

Note that the hypothesis space is infinite! How can we learn, and what can we say about what we've learned? That's what this class is all about!

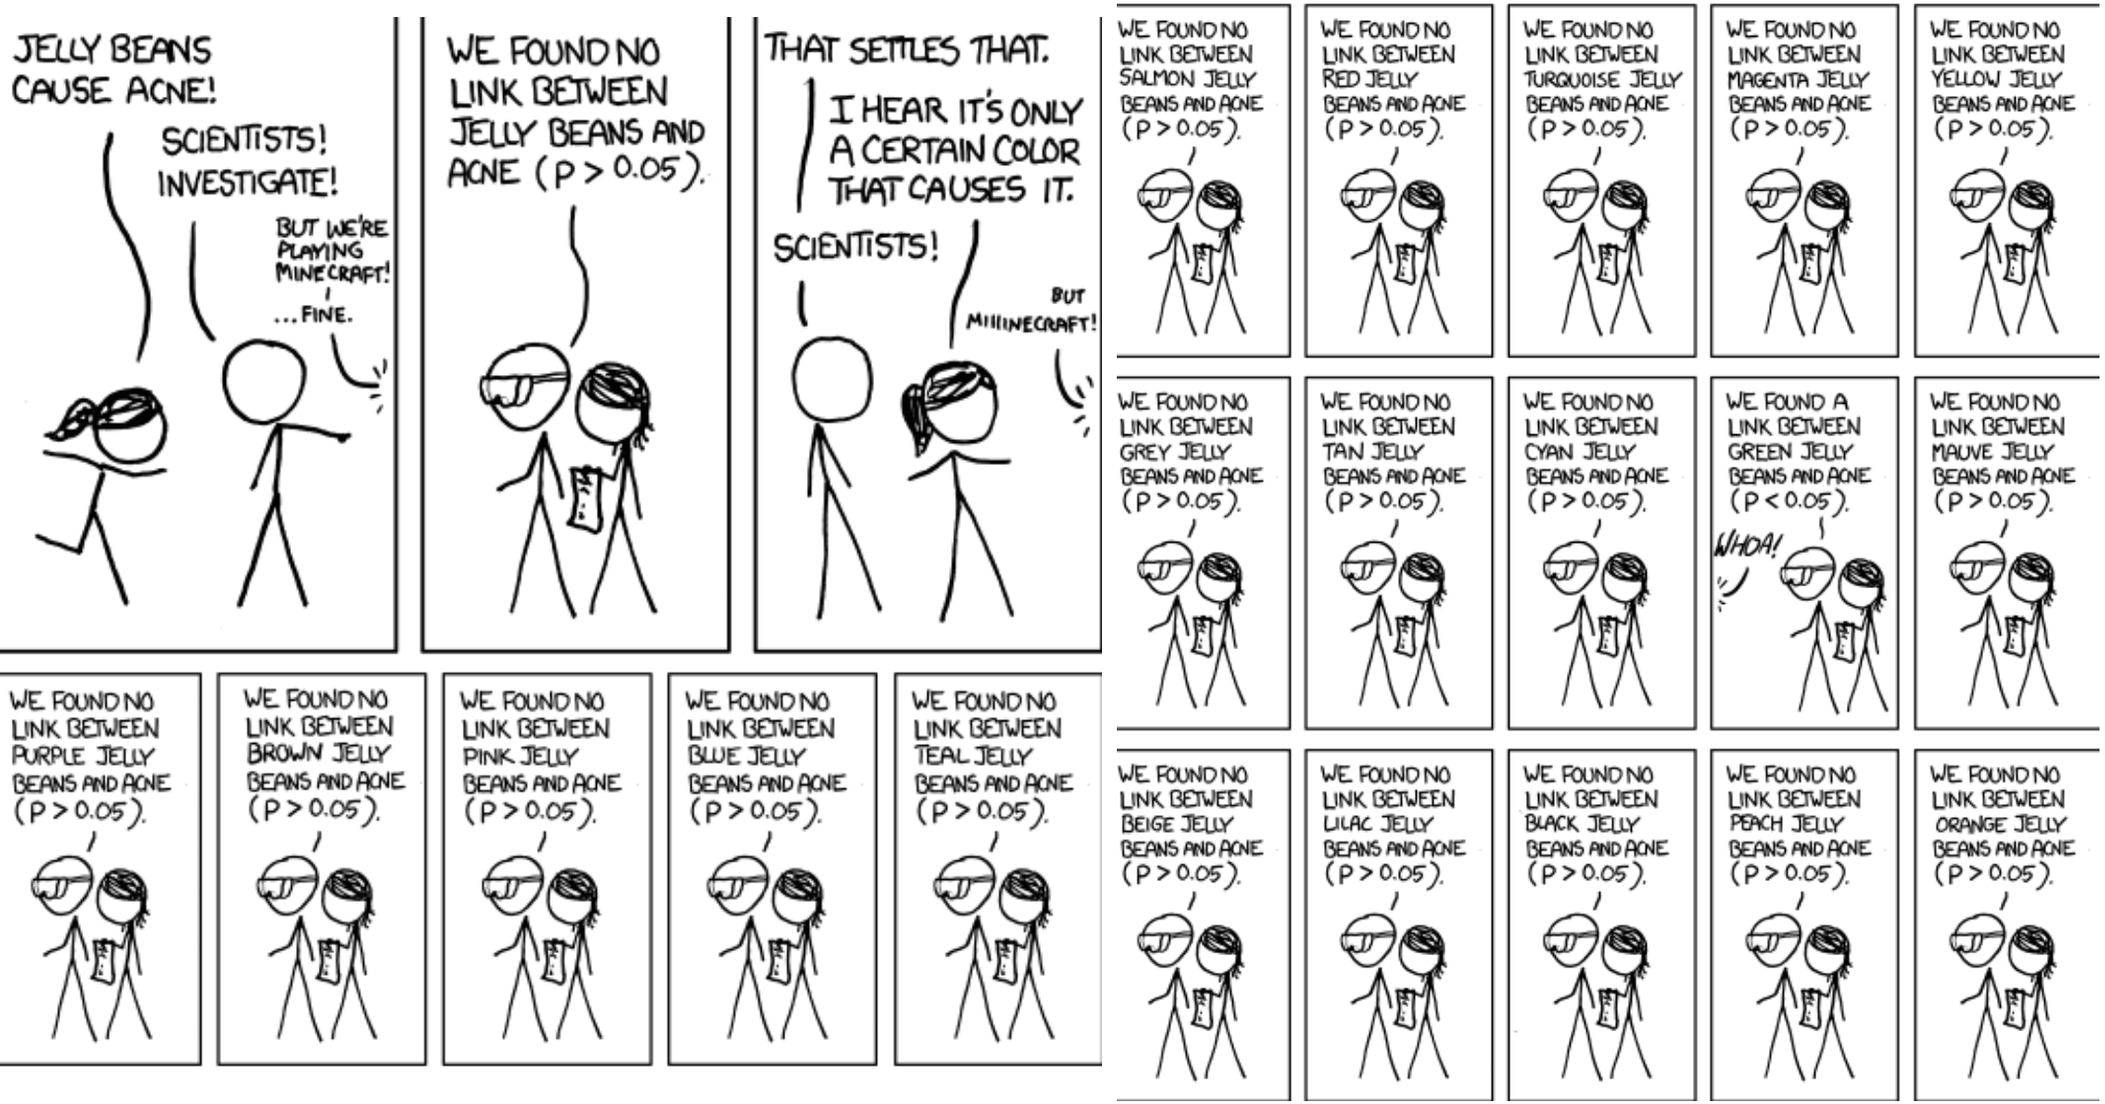


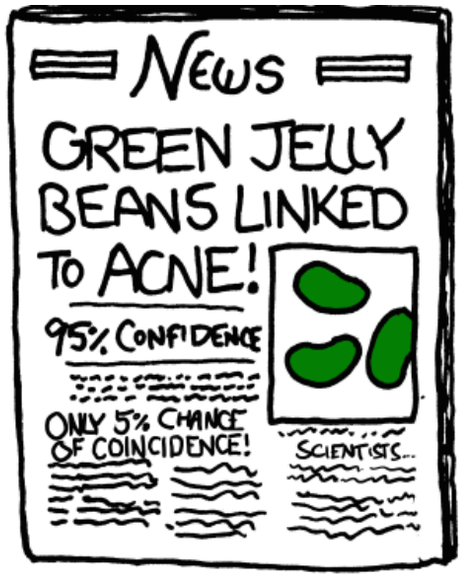


From xkcd, by Randall Munroe: http://xkcd.com/882/

http://fivethirtyeight.com/features/science-isnt-broken/

*Bernard Parker, left, was rated high risk; Dylan Fugett was rated low risk. (Josh Ritchie for ProPublica)*

# Machine Bias

There's software used across the country to predict future criminals. And it's biased against blacks.

*by Julia Angwin, Jeff Larson, Surya Mattu and Lauren Kirchner, ProPublica*

May 23, 2016

Monkey Cage

# A computer program used for bail and sentencing decisions was labeled biased against blacks. It's actually not that clear.

By **Sam Corbett-Davies**, **Emma Pierson**, **Avi Feller** and **Sharad Goel**
October 17, 2016

**Most Read Politics**

1 House Republican leaders move to strip Rep. Steve King of his committee assignments over comments about white nationalism

2 Old land deal quietly haunts Mick Mulvaney as he serves as Trump's chief of staff

3 Trump and lawmakers paralyzed over shutdown as both sides remain dug in

4 The shutdown is giving some Trump advisers what they've long wanted: A

# Unsupervised Learning (More in 517A)

Suppose you only have the feature vectors ($x$) and no labels.

Still want to describe the data in some useful way

Example from the book:


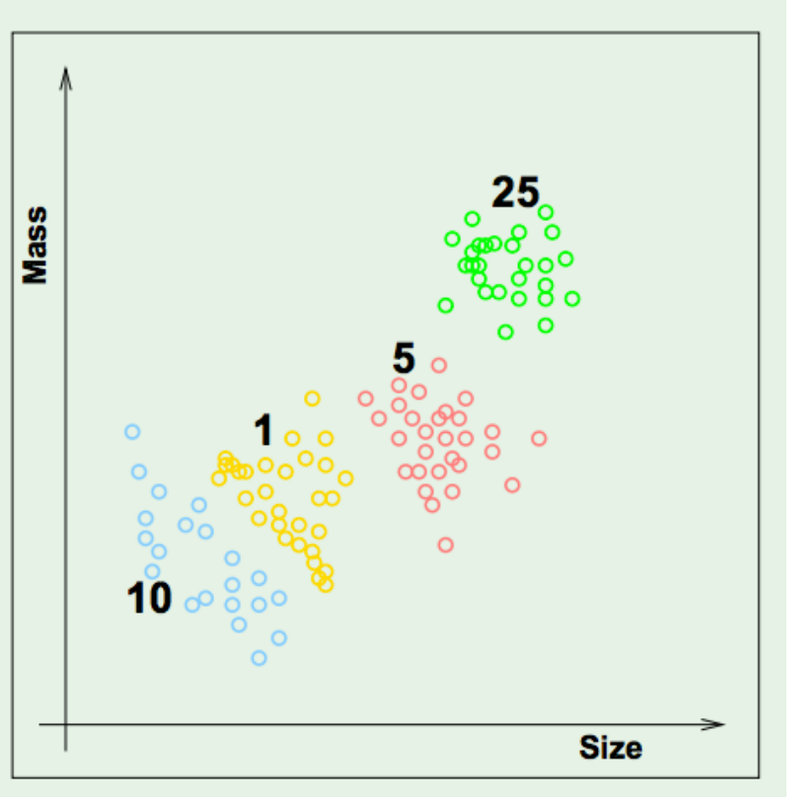

# Figure 1 Estimated population structure.

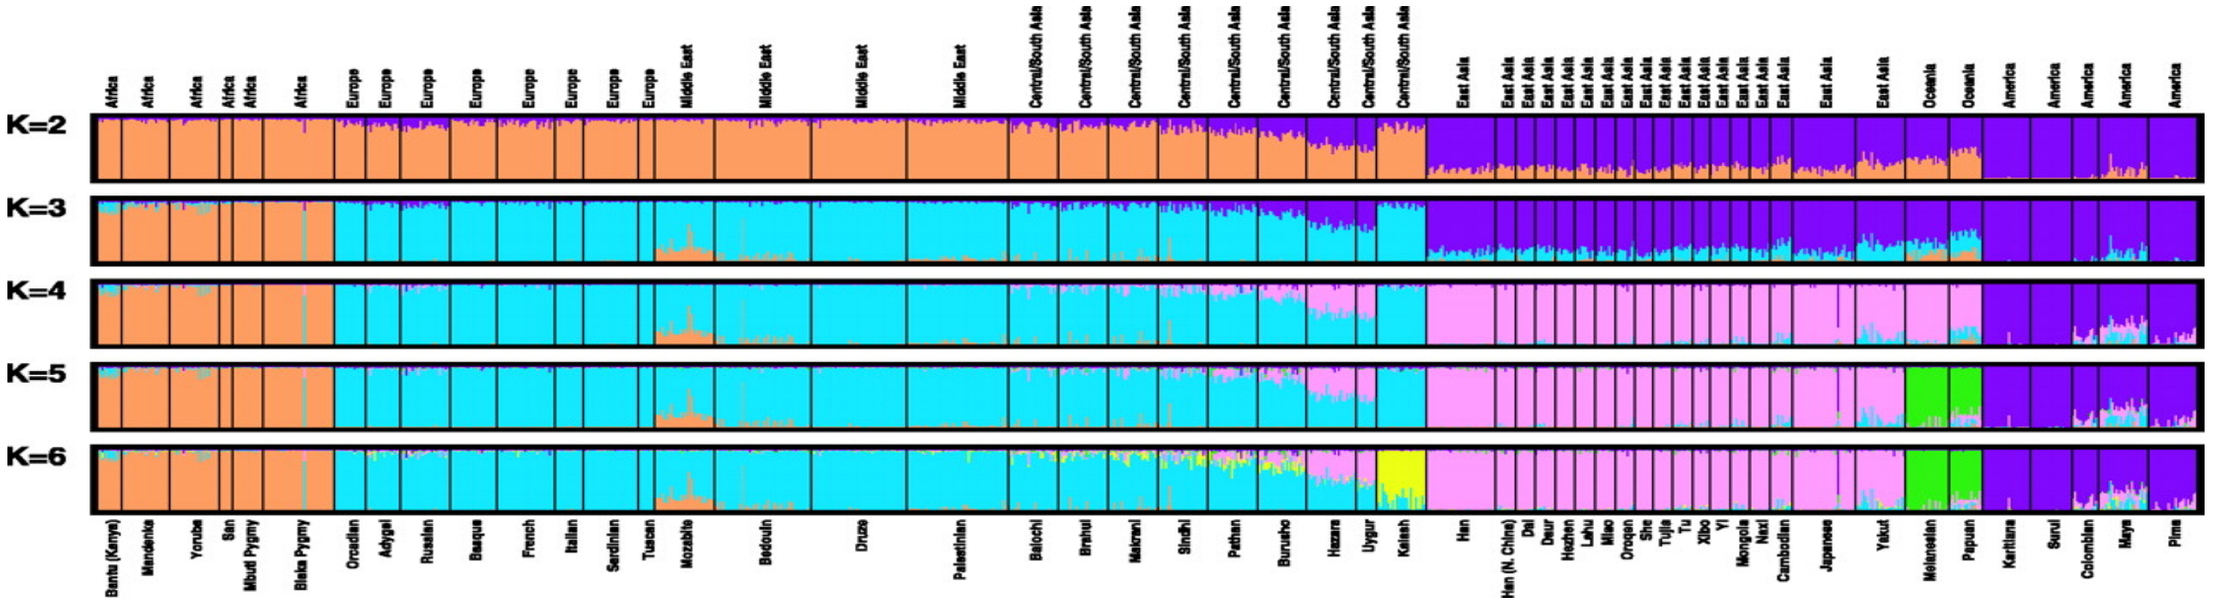


N A Rosenberg et al. Science 2002;298:2381-2385

Published by AAAS

# Reinforcement Learning (More in 511A)

Agent interacts with the world by taking *actions*

Feedback is in the form of *rewards* (or *costs*)

Agent must learn a *policy*, which maps from the state of the world to an action

Major issues:

- Delayed reward / credit assignment
- Exploration / exploitation

# Syllabus and course logistics

- Role of lectures
- Programming and math
  - Matlab is fun!
  - In a city of 1 million, 100 people are radicalized. The government installs machine-learning-based profiling software in a popular mall to detect “radicalized” behavior, which has 99% accuracy (if it sees a radicalized person, it will trigger an alarm 99% of the time, and if it sees a non-radicalized person, it will not trigger an alarm 99% of the time). The alarm goes off! What’s the probability the person who set it off is radicalized?
  - (a) 99% (b) 95% (c) 50% (d) 2% (e) <1%
- Let’s look over the website and syllabus