

CS109: Probability for Computer Scientists

Jerry Cain
January 8th, 2024

Ed Discussion: <https://edstem.org/us/courses/51412/discussion/4097193>

Live Lecture!

- Gates B01
- MWF, 3:00 – 4:20pm
- Perfectly acceptable if you need to watch lecture videos later or simply prefer to watch from the comfort of your dorm room
- Ask your questions in class and on [Ed](#)
 - Students in lecture are encouraged to interrupt me, ask questions, or even request I explain something a second time. Don't be shy.
 - The Ed form can be used for questions, too (e.g., questions that arise while reviewing lecture slides)
 - Teaching staff and I can answer questions after lecture.

If you were enrolled in the course as of this morning, you're already in the Ed forum.

Today's discussion thread: <https://edstem.org/us/courses/51412/discussion/4097193>

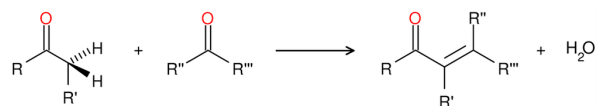


Welcome to
CS109!

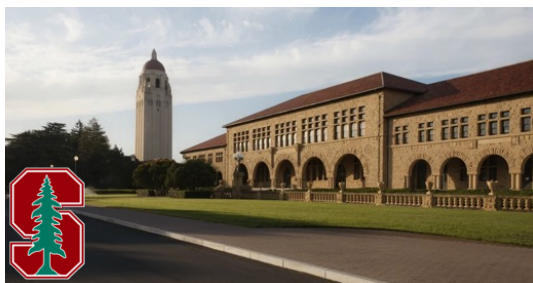
Jerry Cain



I went here from 1987 through 1991 and majored in chemistry.



Then I came here for a PhD in chem, switched to CS



Received MSCS 1998
Lecturer: nearly 28 years

My interests over time

Chemistry
and Physics



Computer
Science



STEM
Education



Why Jerry likes probability

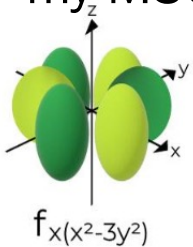
- I majored in chemistry and focused on physical chemistry, and my undergraduate research was rooted in surface science and **statistical** mechanics.
- When I switched to CS as a grad student here, I focused on CS theory and all the beautiful mathematics that comes with it.
- Probability has revived parts of AI and information theory that were thought to be borderline dead when I was getting my MSCS degree here during the 90's.



1974



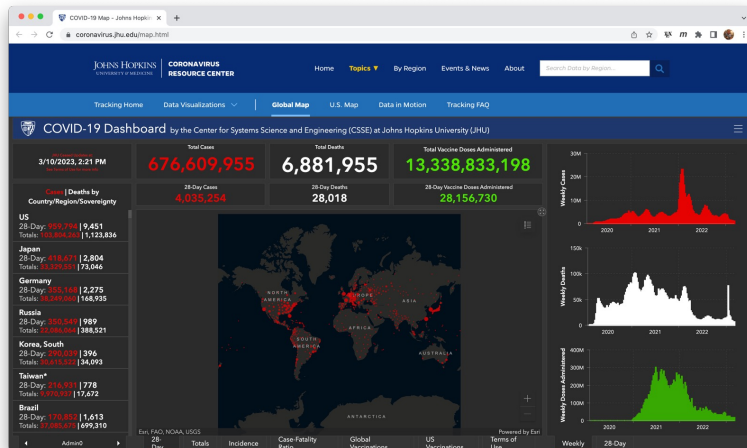
1996



$$PV = \frac{1}{3} N m v_{\text{rms}}^2. \quad f(v) = 4\pi \left(\frac{m}{2\pi kT} \right)^{\frac{3}{2}} v^2 e^{-\frac{mv^2}{2kT}} \quad v_{\text{rms}}^2 = \int_0^\infty v^2 f(v) dv = 4\pi \left(\frac{m}{2\pi kT} \right)^{\frac{3}{2}} \int_0^\infty v^4 e^{-\frac{mv^2}{2kT}} dv$$

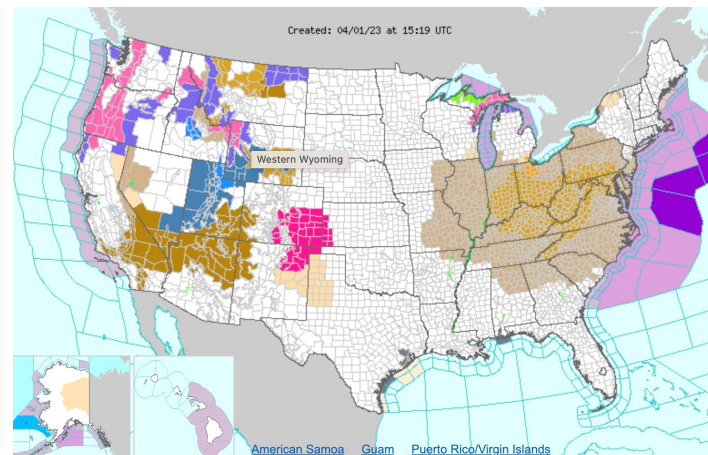
What makes this quarter important

We are seeing a huge surge in **statistics, predictions, and probabilistic models** shared through global news, governing bodies, and social media.



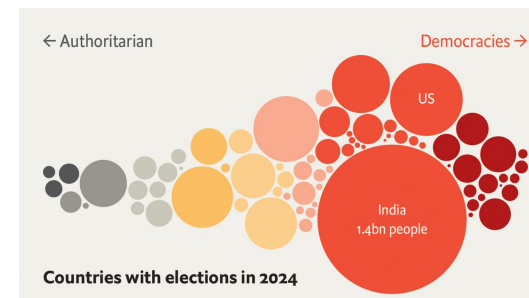
Global cases of COVID-19
as of March 10, 2023 (JHU)

<https://coronavirus.jhu.edu/map.html>



National Weather Service Alerts

<https://www.weather.gov/>



World Politics

<https://abcnews.go.com/538>

<https://www.nytimes.com/>

<https://www.economist.com/>

What makes this quarter important

We are seeing a huge surge in **statistics, predictions, and probabilistic models** shared through global news, governing bodies, and social media.

The **technological and social innovation** we develop during this time will strongly influence how we solve interesting problems impacting the **lives of countless people across the globe**.

Instructor goals:

To teach how probability applies to real life problems that truly matter

To foster and maintain human connection throughout the course



Course Mechanics

Prerequisites

CS106B

Programming
Recursion
Hash tables
Binary trees



MATH 51

Multivariate differentiation
Multivariate integration
Working knowledge of linear
algebra (e.g., vectors)



Important!

CS103

(co-requisite OK)

Proofs (induction)
Set theory
Mathematical
maturity

Companion class: CS109A

- CS109A is an extra 1-unit "ACE" section with additional support, practice, and instruction
- Meets for an additional weekly section and has additional review sessions, office hours, and practice problems
- Admission is via [application](#). You can ignore the published deadline in the form, as our CS109A application is due this Friday, January 12th at 5:00pm.
- CS109A meets on Tuesdays and Thursdays from 8:30 – 9:20am in STLC 118 and starts on January 16th.
- Feel free to email Michelle Qin at **`mdqin@stanford.edu`** with any questions.



Michelle Qin

Course components

42%	6 Problem Sets
22%	Two Midterms
21%	Final Exam
5%	Section Participation
10%	Concept Checks

Course components

42% **6 Problem Sets**

22% Two Midterms

21% Final Exam

5% Section Participation

10% Concept Checks

L^AT_EX

Written portion

- LaTeX for powerful typesetting
- Tutorial on CS109 website



python

Coding portion in Python

- Review session #1 on Thursday 01/11 at noon, in Gates B01

Late policy

- Submit by deadline and you're set!
- Need a short extension? No need to ask! Take an extra class period.
- Need a longer extension? Just ask us and we'll probably be okay with it.
- Extensions can be at most a week.

Course components

42% 6 Problem Sets

22% **Two Midterms**

21% Final Exam

5% Section Participation

10% Concept Checks

- In person! But held outside of class so we can let you work *sans* time pressure.
- Closed-book, mostly-closed-notes, closed-computer, no calculators.
- You can bring **two** 8.5" x 11" pages of notes—using both sides—and refer to them during the exams.
- Held on Wednesdays.
 - Week 4: Wed, 01/31, 7:00 – 9:00pm
 - Week 7: Wed, 02/21, 7:00 – 9:00pm
- Irreconcilable Conflict? Let Jerry know and we'll work something out.

Course components

42% 6 Problem Sets

22% Two Midterms

21% **Final Exam**

5% Section Participation

10% Concept Checks

- Scheduled for Tuesday, March 19th from 8:30 until 11:30am (our official time).
- Closed-book, mostly-closed-notes, closed computer, no calculators.
- You may prepare **four** 8.5" x 11" pages of notes—using both sides—and refer to them and a reference sheet I provide during the exams. Content must be visible to naked eye. (Otherwise, no notes.)
- Conflict with another final? Final exam can then be taken later that same day from 12:15 -3:15pm. Let Jerry know.

Course components

42% 6 Problem Sets

22% Two Midterms

21% Final Exam

5% **Section Participation**

10% Concept Checks

- Sections meet on Thursdays and Fridays (times to be released today at 5:00pm)
- Sections start Week 2
- Your section grade is 100%, but each absence reduces the weight and increases the weight of the final exam
- Go to section!

Course components

42% 6 Problem Sets

22% Two Midterms

21% Final Exam

5% Section Participation

10% **Concept Checks**

- Short set of questions released after each lecture.
- Questions are straightforward and there to ensure you've absorbed the key points and formulas from class.
- All of Week n 's concept checks are due the Tuesday of Week $n + 1$ at noon.
- No late submissions accepted unless truly extenuating circumstances make it truly impossible to meet deadline.

CS109 Contest

- Announced mid-quarter, genuinely optional
- Boost final course grades after letter grade buckets have been determined



Your baseline is CS109, and the sky is the limit.

Previous winning submissions:

- Recidivism Risk: Algorithmic Prediction and Racial Bias
- A Better Way to Reform the Electoral College
- Monte Carlo Tree Search for Tic Tac Toe
- COVID's impact on Student Interest in Post-Secondary School Education

Stanford Honor Code

Permitted

- Talking to the course staff
- Talking with classmates (though cite collaboration)
- Looking up general material online

Cite all references aside from instructors, staff, lecture slides/notes, course reader, and the optional Ross textbook.

NOT permitted:

- Copying answers:
 - from classmates
 - from former students
 - from previous quarters
- Copying answers from the internet
- Ask for answers on the internet
- Relying on generative AI (e.g., ChatGPT, Bard, others) to answer problem set questions.

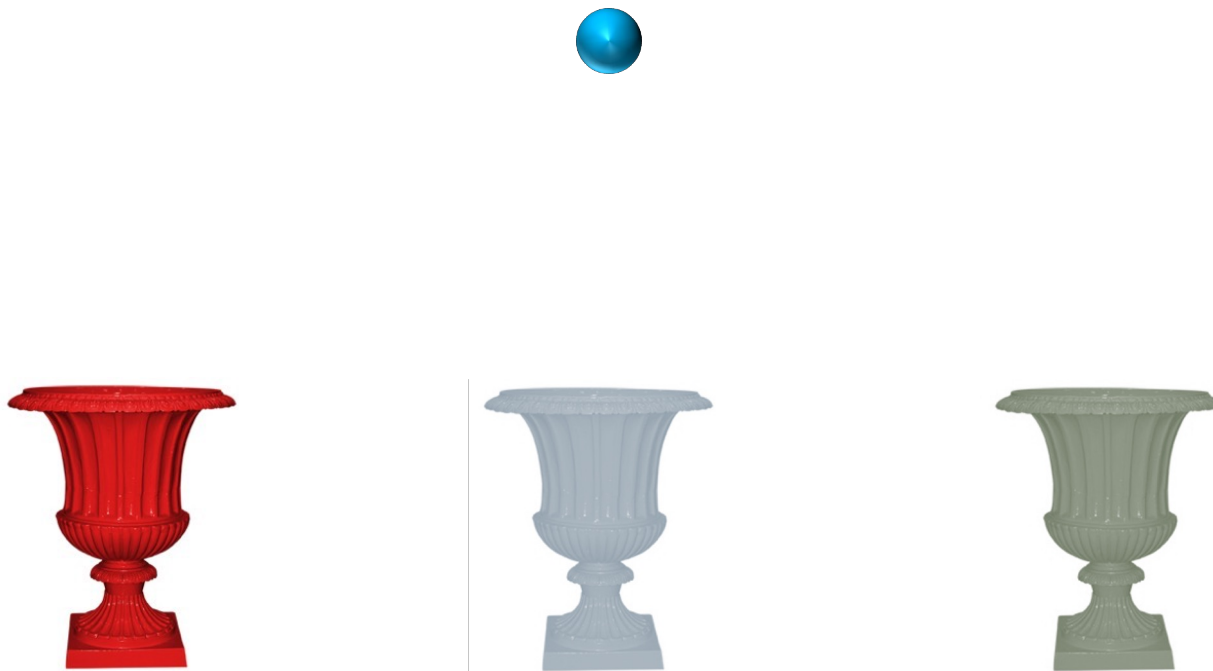
Dramatic
Pause

Questions?



Why you
should take
CS109

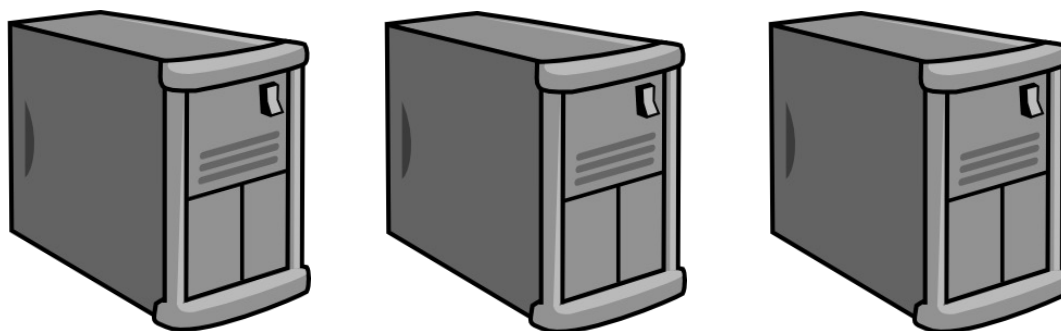
Traditional View of Probability



CS view of probability

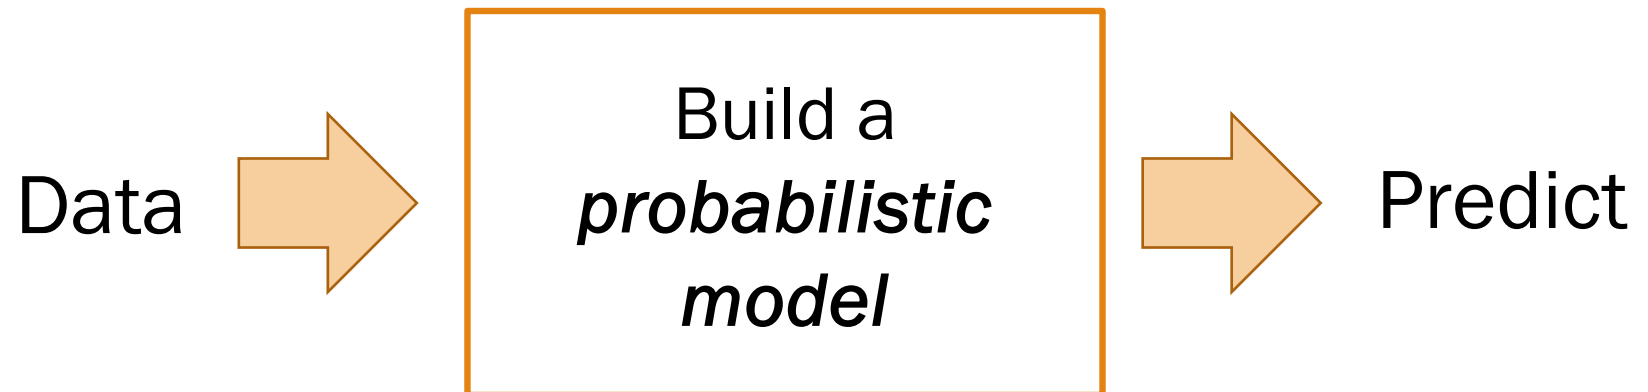
<http://www.site.com>

But wait...
There's MORE!!



Machine Learning
= Machine (compute power)
+ Probability
+ Data

Machine Learning Algorithm



Binary Classification Silliness

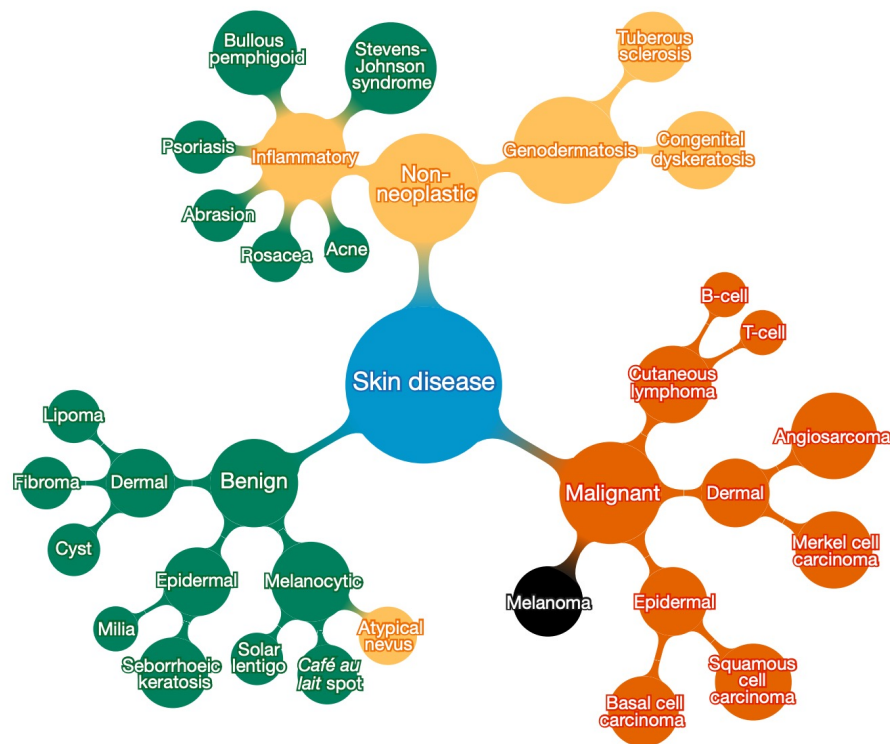


chihuahua or muffin?



poodle or fried chicken?

Classification: Where is this useful?



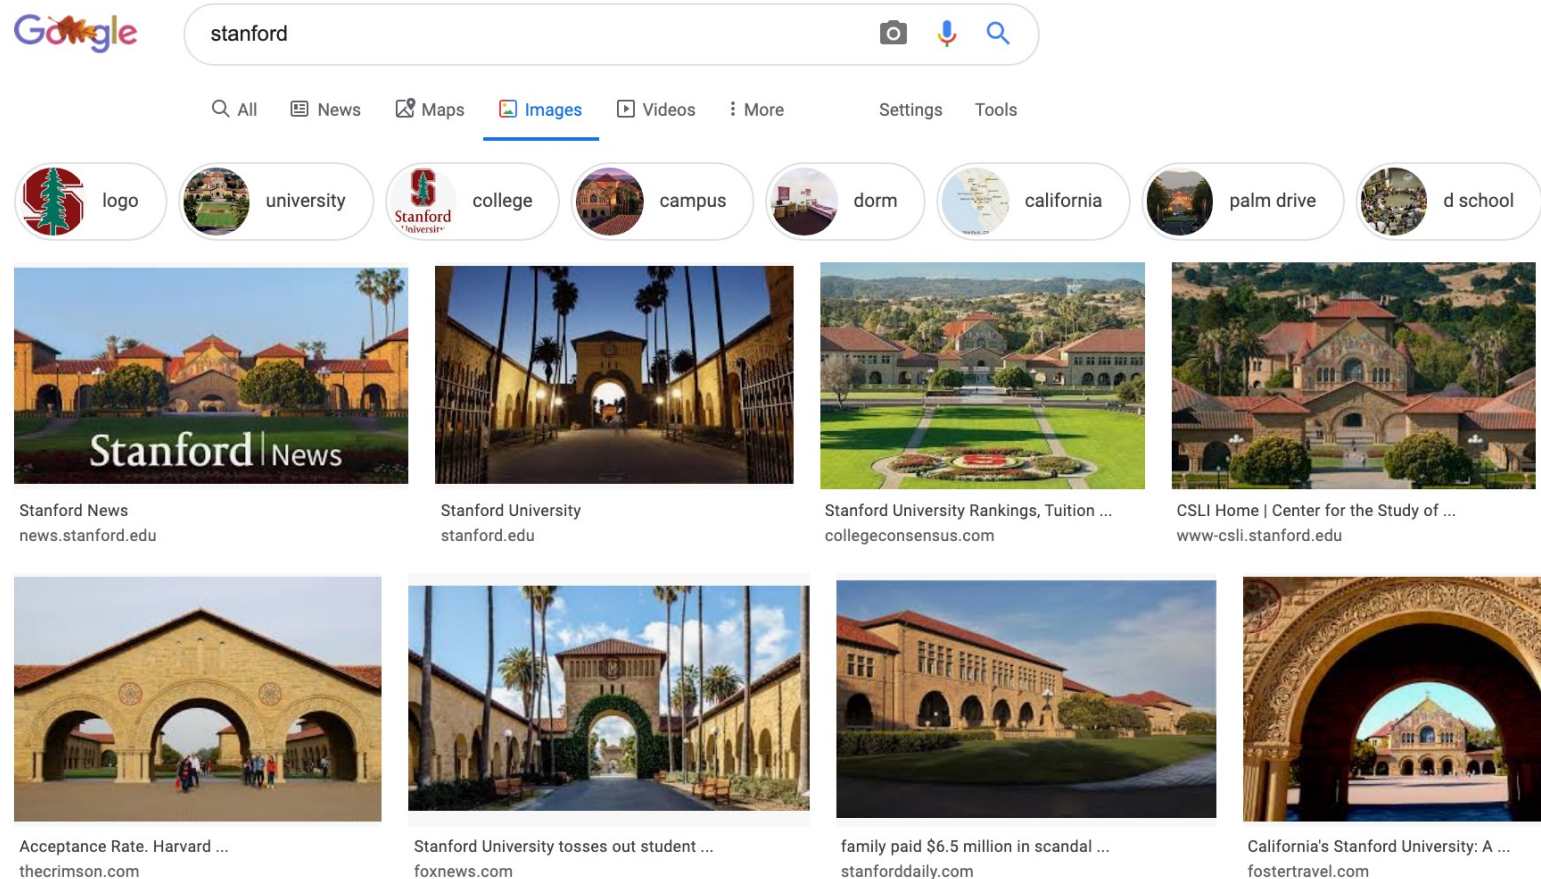
A machine learning algorithm performs **better** than the best dermatologists.

Developed in 2017 at Stanford.

Esteva, Andre, et al. "Dermatologist-level classification of skin cancer with deep neural networks." *Nature* 542.7639 (2017): 115-118.

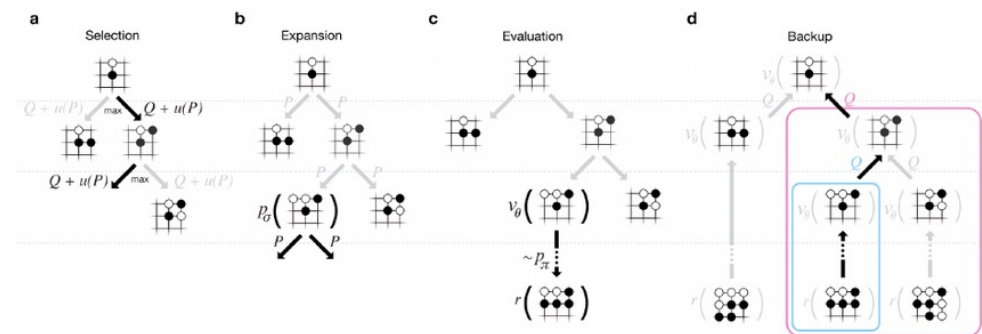
Lisa Yan, Chris Piech, Mehran Sahami, and Jerry Cain, CS109, Winter 2024

Classification: Image tagging



Lisa Yan, Chris Piech, Mehran Sahami, and Jerry Cain, CS109, Winter 2024

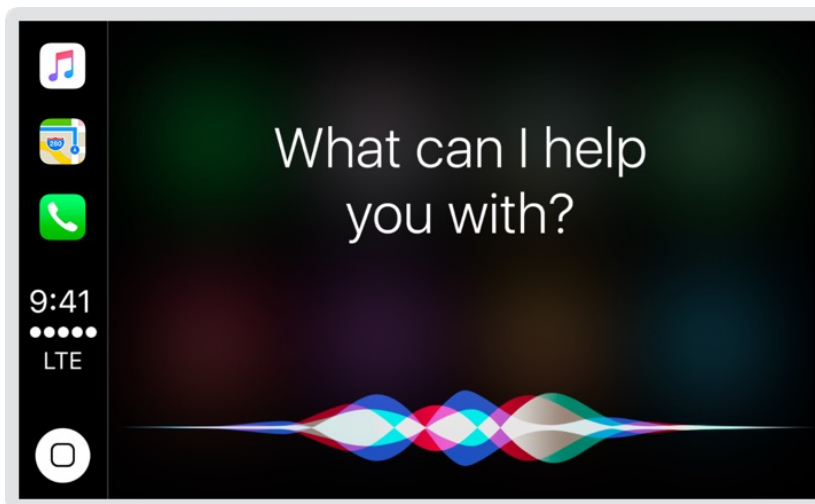
Decision Making: The last frontier in board games



Natural language and speech processing



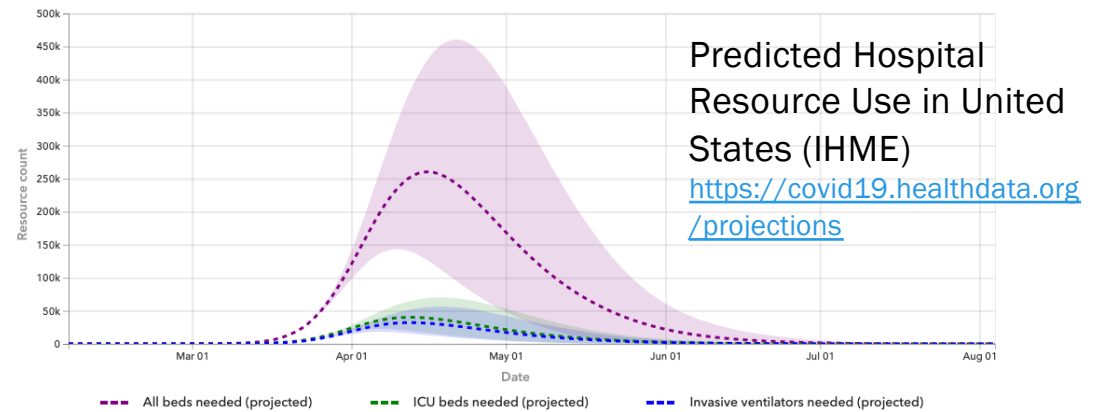
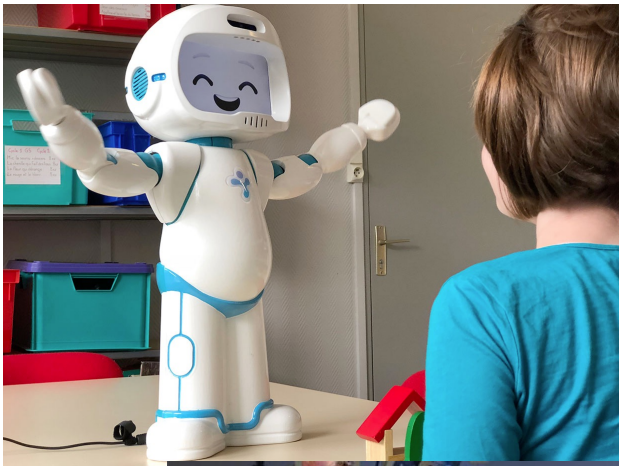
Augmented reality
machine translation
on Google Translate



Voice assistants:
voice to text to answer
(Siri allegedly getting [much better](#) in 2024)

Probability is *more* than
just machine learning.

Probability and medicine

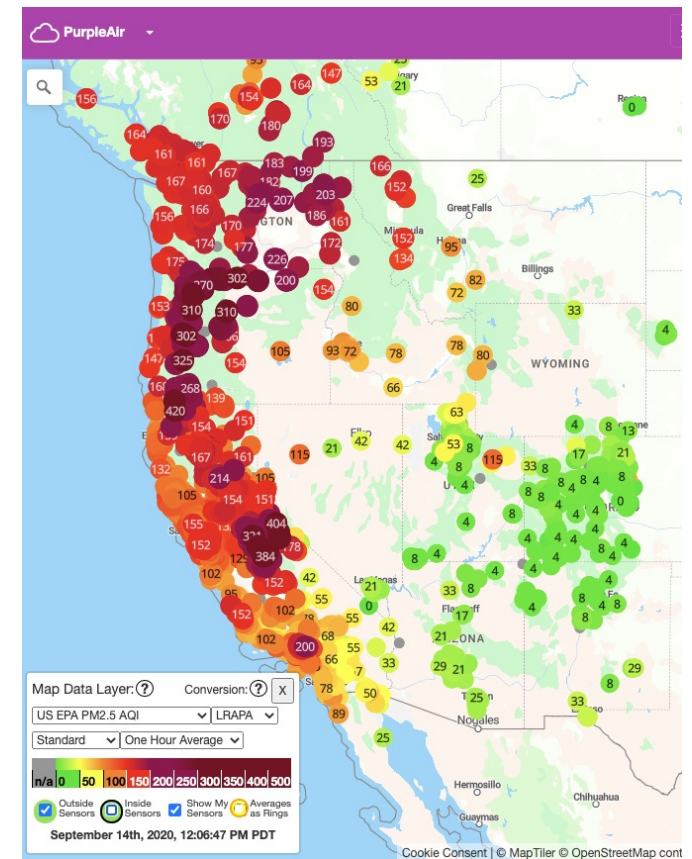
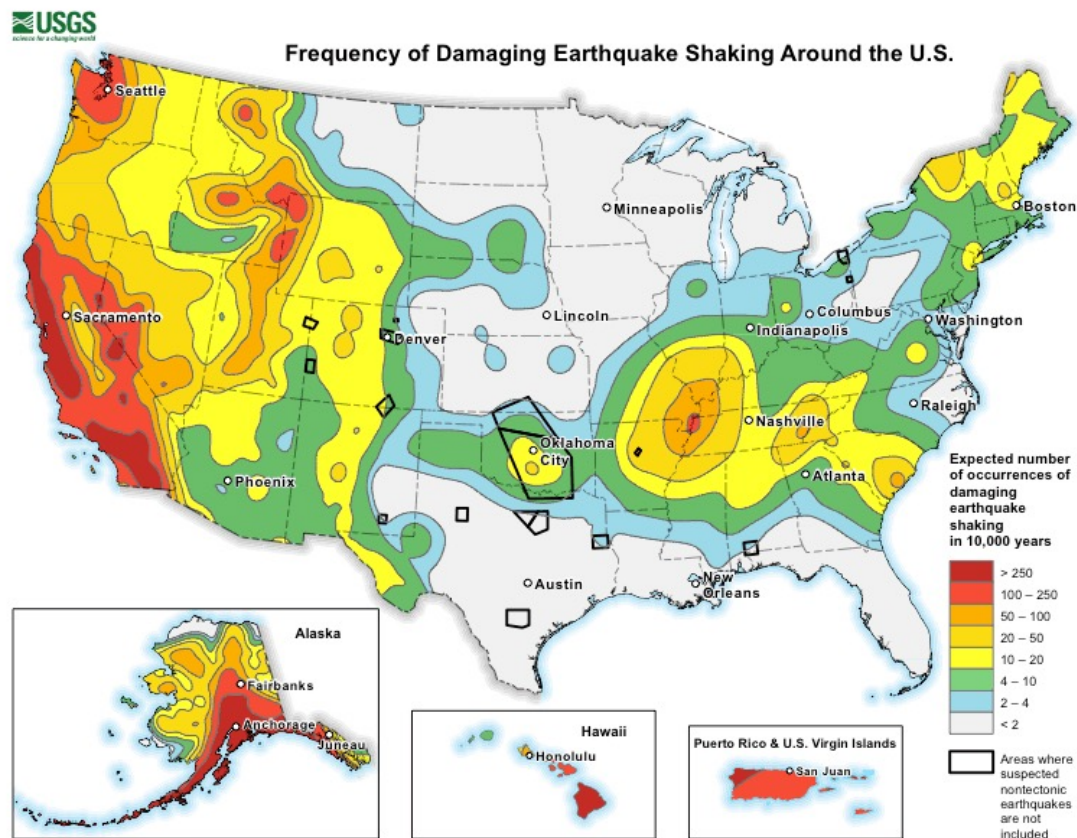


How do COVID-19, RSV, and monkeypox testing rates in a region correlate with the actual spread of the disease?

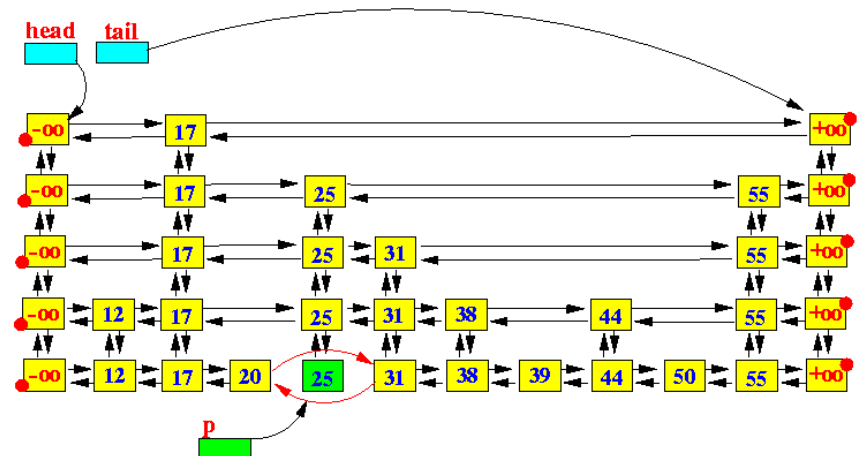
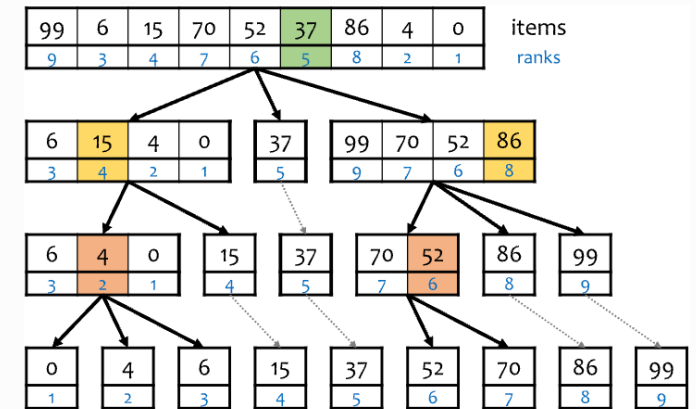
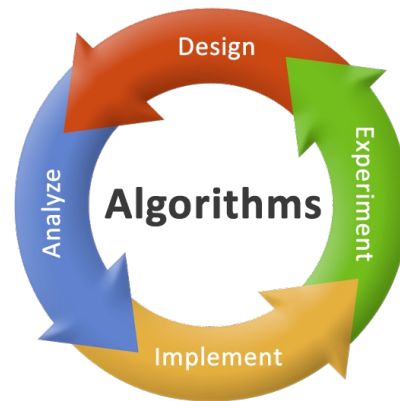
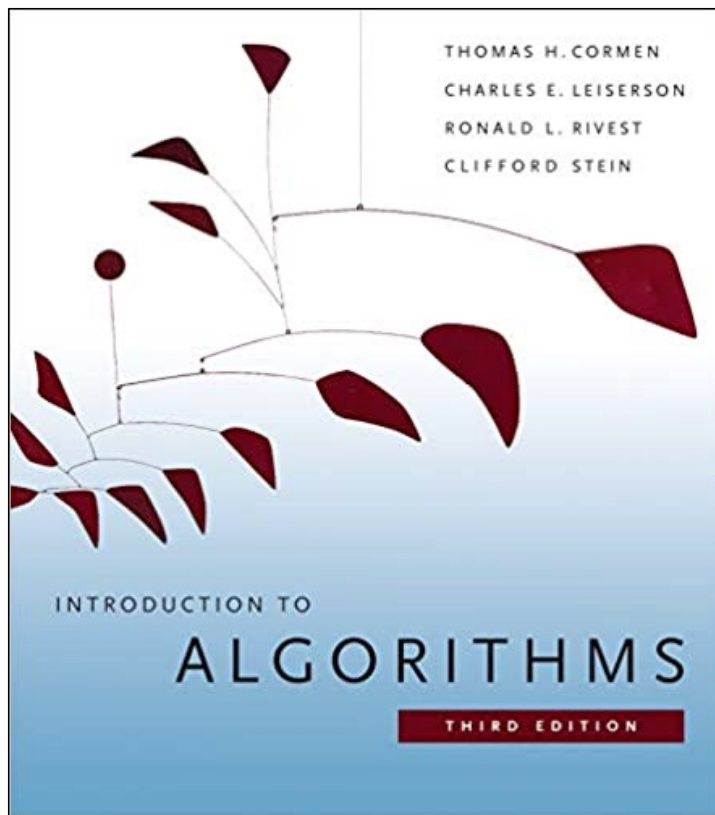
Probability and art



Probability, Seismology, and Meteorology



Probabilistic analysis of algorithms



Probability in practice

amazon prime Deliver to Jerry San Franc... 94114 All harry potter

EN Hello, Jerry Account & Lists Returns & Orders Cart

All Medical Care Gift Cards Buy Again Customer Service Groceries Find a Gift Household, Health & Baby Care Handmade Reset with a grocery restock

Movies & TV New Releases Best Sellers Deals Blu-ray 4K Ultra HD TV Shows Kids & Family Anime All Genres Prime Video Your Video Library

Trends International Harry Potter 5 Dumbledore's Army Wall Poster, 22.375" x... ★★★★★ 47 \$19.00 prime

Back to results Sponsored

Harry Potter COMPLETE 8-FILM COLLECTION

Harry Potter: The Complete 8-Film Collection
Daniel Radcliffe (Actor), Rupert Grint (Actor), David Yates (Director), & 1 more | Rated: NR | Format: Blu-ray
4.6 ★★★★★ 2,284 ratings
600+ bought in past month

\$51.50
Or \$9.51 /mo (6 mo). Select from 1 plan
prime Two-Day
FREE Returns

Blu-ray \$51.50 prime DVD from \$50.62

Click image to open expanded view

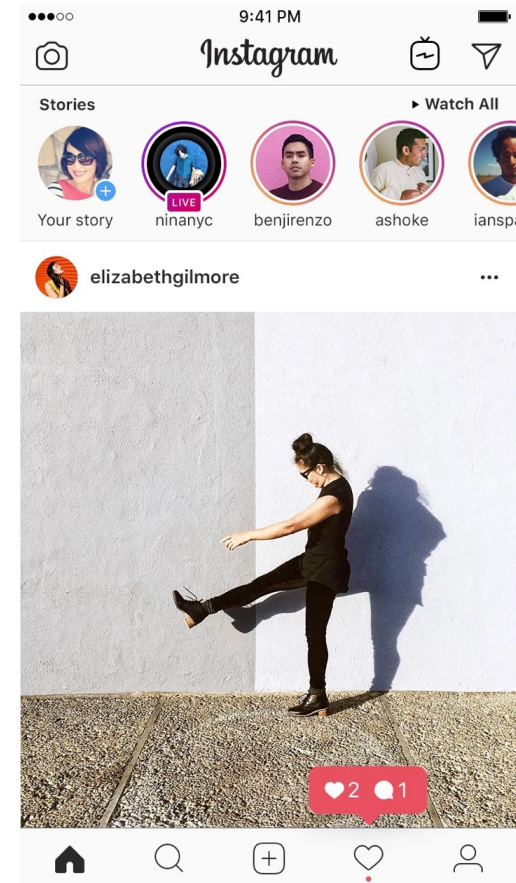
Frequently bought together

Harry Potter + The Prisoner of Azkaban + The Prisoner of Azkaban
Total price: \$117.03
Add all three to Cart
Add all three to List

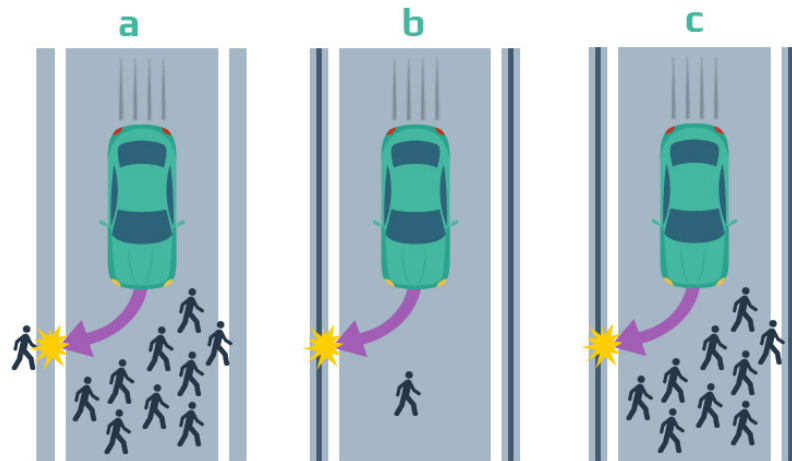
Delivery Pickup

Buy new: \$51.50
prime Two-Day
FREE Returns
FREE delivery Monday, January 8.
Order within 7 hrs 32 mins
Deliver to Jerry - San Francisco 94114
In Stock
Qty: 1
Add to Cart
Buy Now

Probability at your fingertips



Probability and ethics



The golden rule for autonomous car ethics doesn't exist



So far, there are no unified ethical standards ... for autonomous cars. The big [Moral Machine study](#) conducted by MIT showed that it's hard to identify universal ethical values. The moral choices that people made in the MIT survey were different and varied even at a local level. That's why it's hard to create a universal ethics of self-driving cars that won't be controversial. [\[source\]](#)

Probability is not always
intuitive.

Disease testing

A patient takes a virus test that returns positive.
What is the probability they really have the virus?

- 0.03% of people have the virus
- Test has 99% positive rate for people with the virus
- Test has 7% positive rate for people without the virus

Answer: 0.42%

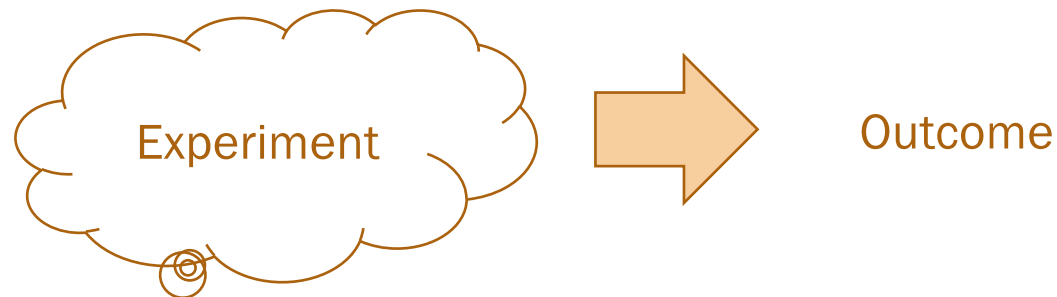




Counting I

What is Counting?

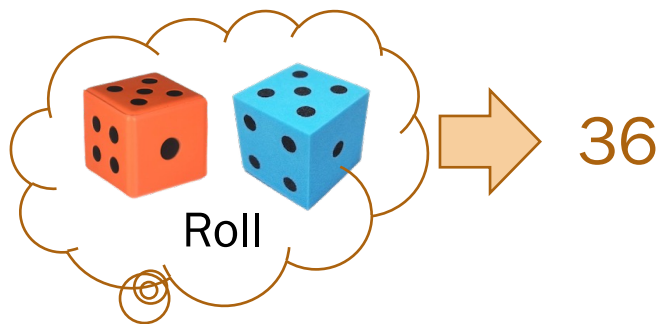
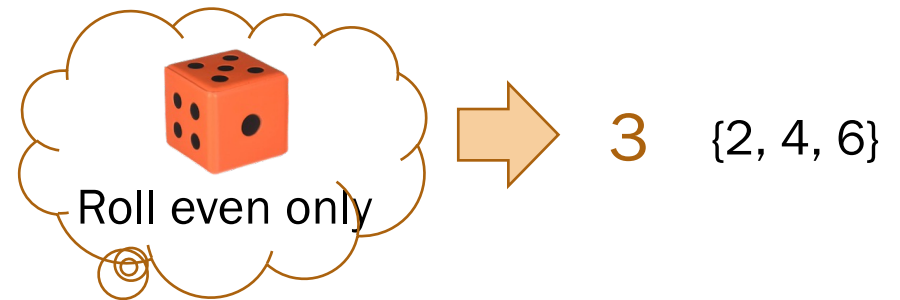
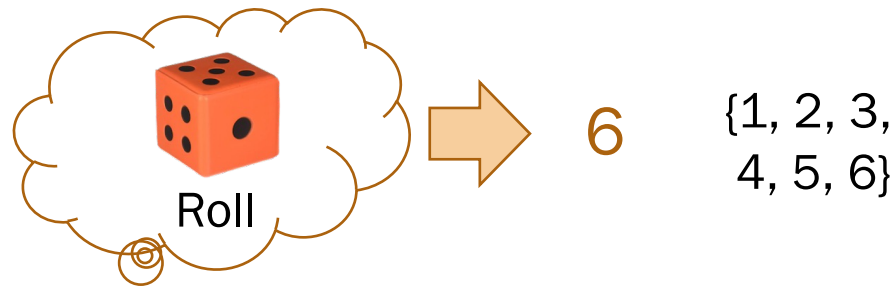
An experiment
in probability:



Counting:

How many possible **outcomes** can occur by performing this **experiment**?

What is Counting Combinatorial Analysis?



{(1, 1), (1, 2), (1, 3), (1, 4), (1, 5), (1, 6),
(2, 1), (2, 2), (2, 3), (2, 4), (2, 5), (2, 6),
(3, 1), (3, 2), (3, 3), (3, 4), (3, 5), (3, 6),
(4, 1), (4, 2), (4, 3), (4, 4), (4, 5), (4, 6),
(5, 1), (5, 2), (5, 3), (5, 4), (5, 5), (5, 6),
(6, 1), (6, 2), (6, 3), (6, 4), (6, 5), (6, 6)}

Sum Rule of Counting

If the outcome of an experiment can be either from

Set A , where $|A| = m$,

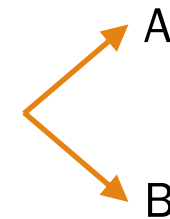
or Set B , where $|B| = n$,

where $A \cap B = \emptyset$,

Then the number of outcomes of the experiment is

$$|A| + |B| = m + n.$$

One experiment



Product Rule of Counting

If an experiment has two parts, where

the first part's outcomes are drawn from A , where $|A| = m$,
and the second part's outcomes are drawn from B , where $|B| = n$,

Then the number of outcomes of the experiment is

$$|A||B| = mn.$$

Two-step experiment



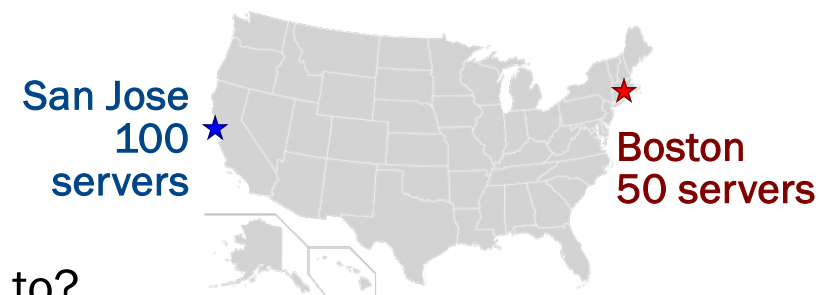
This generalizes to multistep experiments.

Let's try it out

Sum Rule, Product Rule, or something else? How many outcomes?

1. Video streaming application

- Your application has distributed servers in 2 locations (SJ: 100, Boston: 50).
- If a web request is routed to a server, how large is the set of servers it can get routed to?



2. Dice

- How many possible outcomes are there from rolling fifteen six-sided dice?



3. Strings

- How many **different** orderings of letters are possible for the string PEPPER?

PEPERP, EPPPRE,
PEERPP, REPPEP...



Let's try it out

Sum Rule, Product Rule, or something else? How many outcomes?

1. Video streaming application

- Your application has distributed servers in 2 locations (SJ: 100, Boston: 50).
- If a web request is routed to a server, how large is the set of servers it can get routed to?

2. Dice

- How many possible outcomes are there from rolling fifteen six-sided dice?

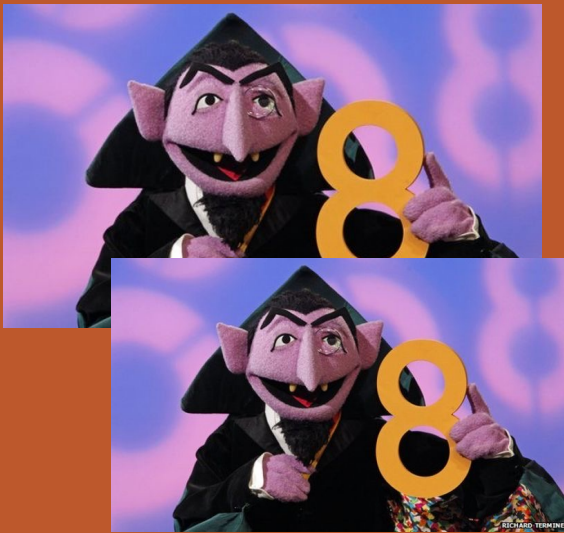
3. Strings

- How many **different** orderings of letters are possible for the string PEPPER?

Dramatic
Pause

Questions?

Counting II



Inclusion-Exclusion Principle

If the outcome of an experiment can be either from

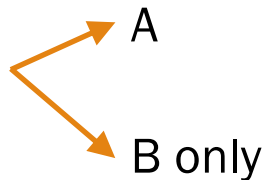
Set A or set B ,

where A and B may overlap,

Then the total number of outcomes of the experiment is

$$|A \cup B| = |A| + |B| - |A \cap B|.$$

One experiment



Sum Rule of Counting:
A special case

Transmitting bytes over a network

An 8-bit string is sent over a network.

- The receiver only accepts strings that either start with 01 or end with 10.

How many 8-bit strings will the receiver accept?

Define

A : 8-bit strings
starting with 01

B : 8-bit strings
ending with 10

01001100
byte (8 bits)



Transmitting bytes over a network

An 8-bit string is sent over a network.

- The receiver only accepts strings that either start with 01 or end with 10.

How many 8-bit strings will the receiver accept?

Define

A : 8-bit strings
starting with 01

B : 8-bit strings
ending with 10

01001100
byte (8 bits)

General Principle of Counting

If an experiment has r steps, such that

Step i has n_i outcomes for all $i = 1, \dots, r$,

Then the number of outcomes of the experiment is

$$n_1 \times n_2 \times \cdots \times n_r = \prod_{i=1}^r n_i .$$

Multi-step
experiment

Product Rule of Counting:
A special case

→ 1 → 2 → ...

License plates

How many CA license plates are possible if...



(pre-1982)



(present day)

Lisa Yan, Chris Piech, Mehran Sahami, and Jerry Cain, CS109, Winter 2024



License plates

How many CA license plates are possible if...



(pre-1982)

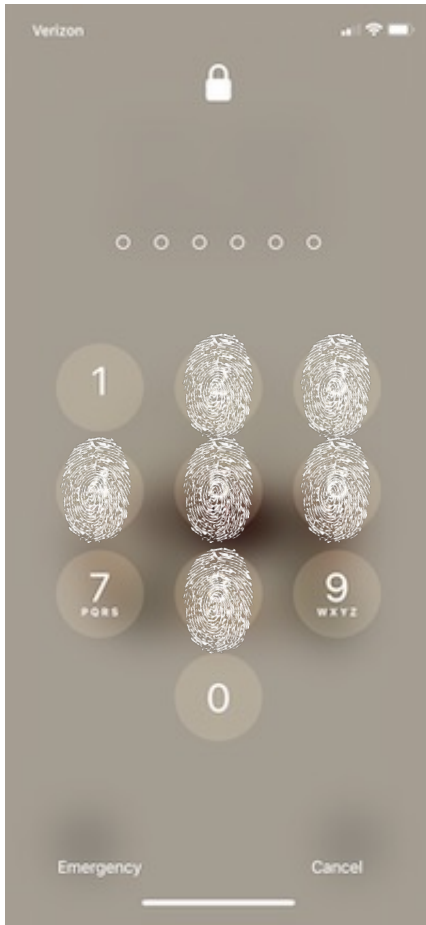


(present day)



Permutations I

Unique 6-digit passcodes with **six** smudges



How many unique 6-digit passcodes are possible if a phone password uses each of **six** distinct numbers?

Arrange n indistinct objects



Arrange n distinct objects



Michelle



Jacob



Groucho



Isabel



Kathleen

Arrange n distinct objects



Steps:

1. Choose 1st can 5 options
2. Choose 2nd can 4 options
- ...
5. Choose 5th can 1 option

$$\begin{aligned}\text{Total} &= 5 \times 4 \times 3 \times 2 \times 1 \\ &= 120\end{aligned}$$

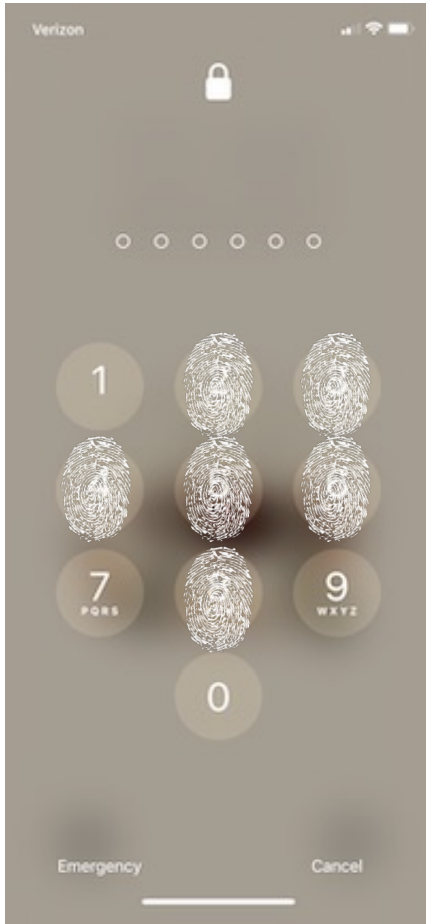
Permutations

A **permutation** is an ordered arrangement of objects.

The number of unique orderings (**permutations**) of n distinct objects is

$$n! = n \times (n - 1) \times (n - 2) \times \cdots \times 2 \times 1.$$

Unique 6-digit passcodes with **six** smudges

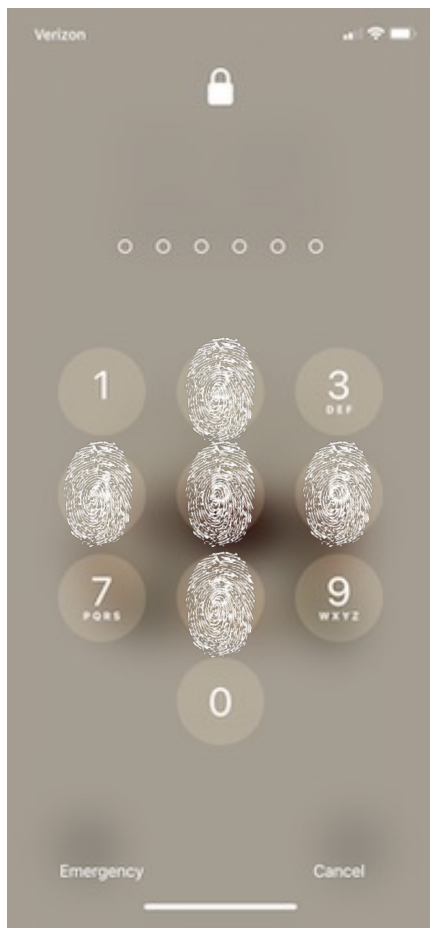


How many unique 6-digit passcodes are possible if a phone password uses each of **six** distinct numbers?

Total = $6!$
= 720 passcodes

```
>>> import math
>>> math.factorial(6)
720
```

Unique 6-digit passcodes with **five** smudges



How many unique 6-digit passcodes are possible if a phone password uses each of **five** distinct numbers?

