

# Lecture 23:

## AI, ML, NLP, ...

**But first ...**  
**What's in the news?**

# Update to Top500 (Nov 16)

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	<b>Supercomputer Fugaku</b> - Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D, Fujitsu RIKEN Center for Computational Science Japan	7,630,848	442,010.0	537,212.0	29,899
2	<b>Summit</b> - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592	148,600.0	200,794.9	10,096
3	<b>Sierra</b> - IBM Power System AC922, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM / NVIDIA / Mellanox DOE/NNSA/LLNL United States	1,572,480	94,640.0	125,712.0	7,438
4	<b>Sunway TaihuLight</b> - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway, NRCPC National Supercomputing Center in Wuxi China	10,649,600	93,014.6	125,435.9	15,371

What could possibly go wrong?

# Amazon launches Amazon Pharmacy for prescription medicine delivery

*Prime members will get perks like free delivery and discounts on drugs*

By [James Vincent](#) | Nov 17, 2020, 6:54am EST

Amazon is making its biggest push into the healthcare industry yet with the launch today of [Amazon Pharmacy](#): a new service offering home delivery for prescription medication.

Customers can sign up to the new store by creating a "secure pharmacy profile," with the option of adding information about their health insurance, any outstanding medical issues like allergies, and any regular prescriptions. The store will offer a range of "generic and brand-name drugs," [reports CNBC](#), including "commonly prescribed drugs like insulin, triamcinolone steroid creams, metformin for controlling blood sugar, and sumatriptan for migraines." Notably, the pharmacy will not sell [Schedule II medications](#), which includes many common opioids like Oxycontin.

# Big brother is watching you



MENU

US

## Employee surveillance software demand increased as workers transitioned to home working

As people hunkered down to work from home during COVID-19, companies turned to employee surveillance software to track their staff.



By [Eileen Brown](#) for [Social Business](#) | November 16, 2020 -- 17:11 GMT (09:11 PST) | Topic: [Security](#)

What does the rise of intrusive tools such as employee surveillance software mean for workers at home?

A new study shows that the demand for employee surveillance software was up 55% in June 2020 compared to the [pre-pandemic average](#). From webcam access to random screenshot monitoring, these surveillance software products can record almost everything an employee does on their computer.

# ET, call home ?

ars TECHNICA

BIZ & IT TECH SCIENCE POLICY CARS GAMING & CULTURE  
STORE FORUMS

OCSP PRIVACY WORRIES —

## Mac certificate check stokes fears that Apple logs every app you run

Amid concern that macOS logs app usage in real time, Apple issues assurances.

DAN GOODIN - 11/16/2020, 9:25 PM

Getty Images

Digital Composite Image Of Businessman Using Laptop With Icons At Desk

Enlarge

102

Last Thursday afternoon, Mac users everywhere began complaining of a **crippling slowdown** when opening apps. The cause: online certificate checks Apple performs each time a user opens an app not downloaded from the App Store. The mass upgrade to Big Sur, it seems, caused the Apple servers responsible for these checks to slow to a crawl.

Apple quickly fixed the slowdown, but concerns about paralyzed Macs were soon replaced by an even bigger worry—the vast amount of personal data Apple, and possibly others, can glean from Macs performing certificate checks each time a user opens an app that didn't come from the App Store.

# Apple's response

## Privacy protections

macOS has been designed to keep users and their data safe while respecting their privacy.

Gatekeeper performs online checks to verify if an app contains known malware and whether the developer's signing certificate is revoked. We have never combined data from these checks with information about Apple users or their devices. We do not use data from these checks to learn what individual users are launching or running on their devices.

Notarization checks if the app contains known malware using an encrypted connection that is resilient to server failures.

These security checks have never included the user's Apple ID or the identity of their device. To further protect privacy, we have stopped logging IP addresses associated with Developer ID certificate checks, and we will ensure that any collected IP addresses are removed from logs.

In addition, over the the next year we will introduce several changes to our security checks:

- A new encrypted protocol for Developer ID certificate revocation checks
- Strong protections against server failure
- A new preference for users to opt out of these security protections

(August)

# *Airbnb, a 'Sharing Economy' Pioneer, Files to Go Public*

The home rental company, which was privately valued at \$31 billion, is trying to go public after its business was crushed by the pandemic.

The Airbnb headquarters in San Francisco. Its debut will most likely be helped by a strong stock market.

The Airbnb headquarters in San Francisco. Its debut will most likely be helped by a strong stock market. Jason Henry for The New York Times



By Erin Griffith

Aug. 19, 2020





# ***Airbnb Reveals Falling Revenue, With Travel Hit by Pandemic***

(November)

The drop was \$1.2 billion for the first nine months of 2020, in the first comprehensive look at the company's finances as it moves to go public.

Airbnb's prospectus painted an optimistic picture, advertising its brand's association with unique travel experiences. "We have helped millions of people satisfy a fundamental human need for connection," the company said. "And it is through this connection that people can experience a greater sense of belonging."

In total, Airbnb brought in \$2.5 billion in revenue in the first nine months of the year, down from \$3.7 billion a year earlier. Its net loss more than doubled during that period to \$697 million.

**AI, ML, NLP, ...**

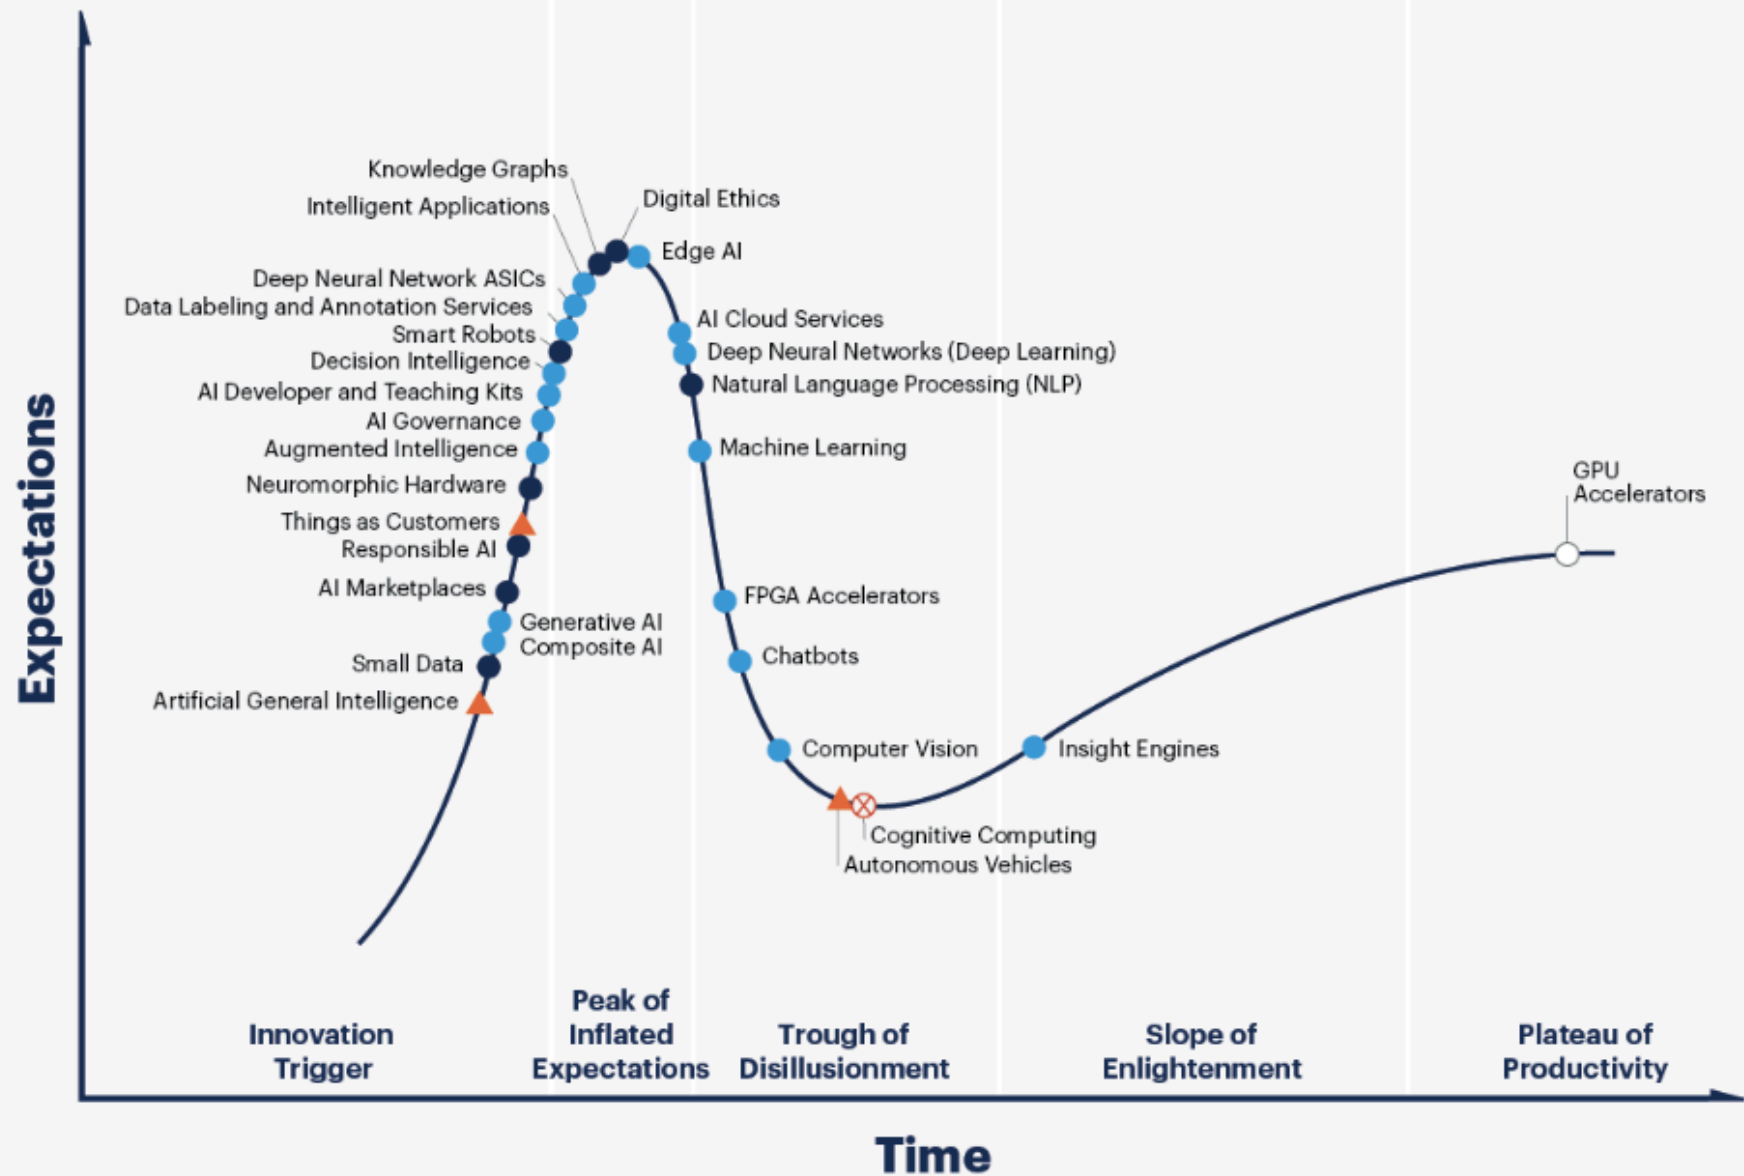
# Artificial intelligence, machine learning, machine intelligence, natural language processing, ...

- **buzzwords, hype, real accomplishments, wishful thinking**
  - big data, deep learning, neural networks, ...
- **brief history**
- **examples**
  - classification (spam detection)
  - prediction (future prices)
  - recommendation systems (Netflix, Amazon, Goodreads, ...)
  - natural language processing (sentiment analysis, translation, generation)
  - games (chess, Go)
- **issues and concerns**
- **Beware: on this topic,  
I am even less of an expert than normal.**

Good idea? Pure hype?



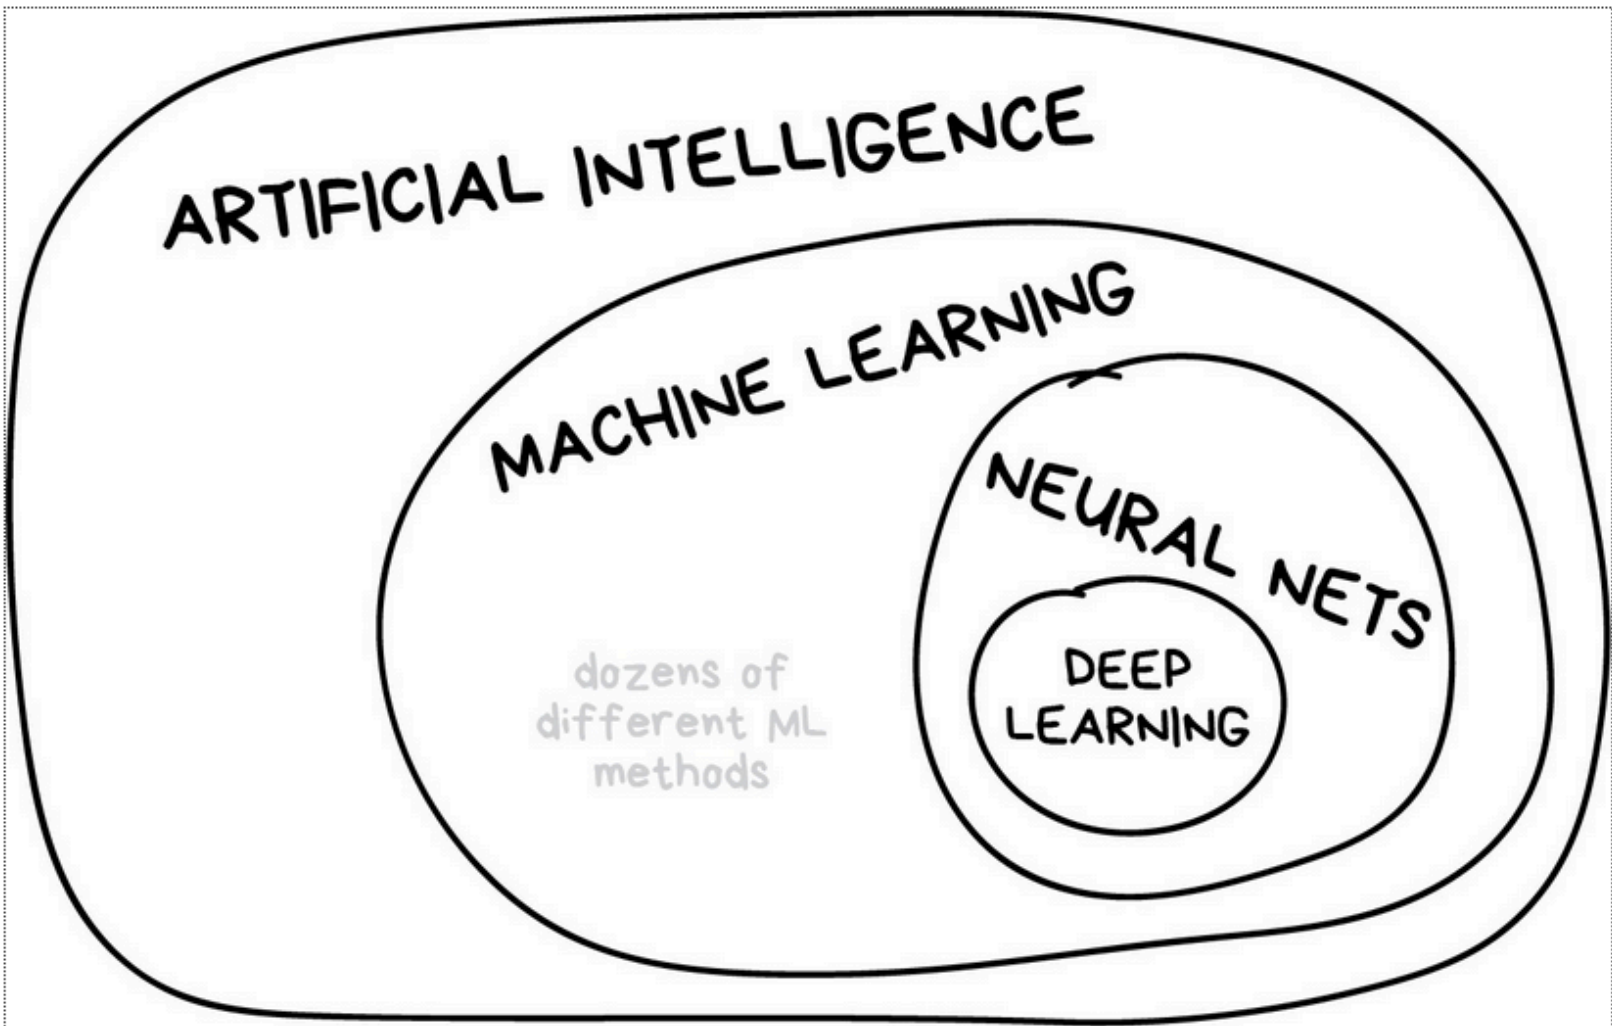
# Hype Cycle for Artificial Intelligence, 2020



# Revisionist history (non-expert perspective)

- **1950s, 1960s: naive optimism about artificial intelligence**
  - checkers, chess, machine translation, theorem proving, speech recognition, image recognition, vision, ...
  - almost everything proved to be much harder than was thought
- **1980s, 1990s: expert or rule-based systems**
  - domain experts write down lots of rules, computers apply them to make decisions
  - it's too hard to collect the rules, and there are too many exceptions
  - doesn't scale to large datasets or new problem domains
- **2010s: machine learning, big data, ...**
  - provide a "training set" with lots of examples correctly characterized
  - define "features" that might be relevant
  - write a program that "learns" from its successes and failures on the training data (basically by figuring out how to combine feature values)
  - turn it loose on new data

# The big picture [\(vas3k.com/blog/machine\\_learning\)](http://vas3k.com/blog/machine_learning)



# Examples of ML applications (tiny subset)

- **classification**
  - spam detection, digit recognition, optical character recognition, authorship, ...
  - image recognition, face recognition, ...
- **prediction**
  - house prices, stock prices, credit scoring, resume screening, ...
  - tumor probabilities, intensive care outcomes, ...
- **recommendation systems**
  - e.g., Netflix, Amazon, Goodreads, ...
- **natural language processing (NLP)**
  - language translation
  - text to speech; speech to text
  - sentiment analysis
  - text generation
- **games**
  - checkers, chess, Go



# Types of learning algorithms

- **supervised learning (labeled data)**
  - teach the computer how to do something with training examples
  - then let it use its new-found knowledge to do it on new examples
- **unsupervised learning (unlabeled data)**
  - let the computer learn how to do something without training data
  - use this to determine structure and patterns in data
- **reinforcement learning**
  - some kind of "real world" system to interact with
  - feedback on success or failure guides/teaches future behavior
- **recommender systems**
  - look for similarities in likes and dislikes / behaviors / ...
  - use that to predict future behaviors

# Classification example: spam detection

How do we know these are spam?


- Not everyone likes being cold. Get a heated vest
- Well-rounded investors should invest in biotechnology
- Trusted biotech stock advice
- Thermal vest with 5 built in heaters
- Recharge your phone with the power of the sun
- Last Day - Half-Off Portable Mini Heaters
- You have received a Hundred dollar CVS reward
- Steep laser engraver discount - in time for xmas
- Early alerts on promising biotech investments
- Is sitting all day killing you?
- Drop 35-lbs by spring
- One more day left to unlock your \$100 bonus

# Classification example: spam detection

- **rule-based: look for odd words & spellings, known bad sources, etc.**
  - V1^6R^, M0NE`/, spamRus.com, ...
- **machine learning: choose a set of features like**
  - odd spelling, weird characters, language and grammar, origin, length, ...
- **provide a training set of messages that are marked "spam" or "not spam"**
- **ML algorithm figures out parameter settings that let it do the best job of separating spam from not spam in the training set**
- **then apply that to real data**
- **potential problems:**
  - training set isn't good enough or big enough
  - creating it is probably done manually
  - "over-fitting": does a great job on training set but little else
  - spammers keep adapting so we always need new training material

# Classification example: hand-written digit recognition


- NIST test suite: 60,000 training images and 10,000 testing images
- best error rates  $\sim 0.25\%$



# Classification example: image recognition



# Classification example: identifying cancerous tumors

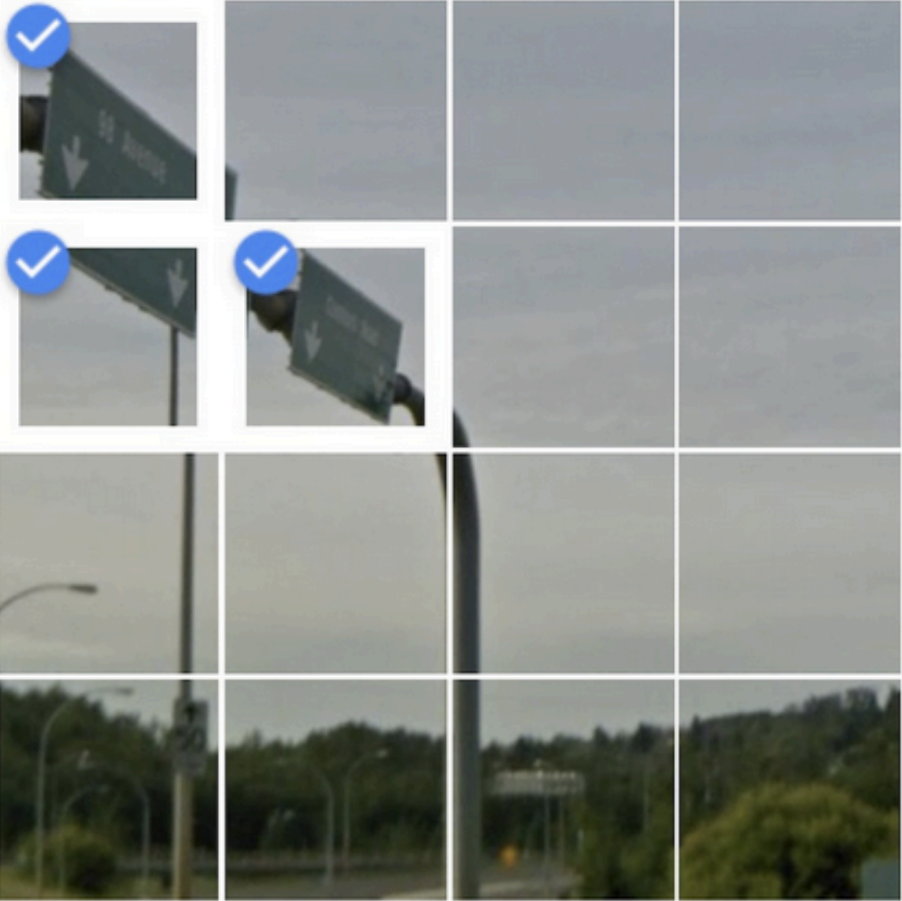


# Classification example: medical diagnosis

- **input: symptoms**
  - fever, cough, rapid breathing, shaking, nausea, ...
- **output: diagnosis**
  - pneumonia, flu, common cold, bronchitis, ...
- **multiclass classification**
  - one of several distinct possibilities
- **answer is a set of probabilities**
  - 70% pneumonia
  - 20% flu
  - 5% bronchitis
- **how would this change in Covid era?**

How do we get  
labeled data for  
training?  
(for example,  
image recognition)

Select all squares with  
**street signs**  
If there are none, click skip



⌂

ⓘ


VERIFY




# Prediction example: house prices

- only one feature here: square footage
- straight line? ("linear regression")
- some kind of curve?


## Housing price prediction.




# Over- and under-fitting



High bias  
(underfit)




“Just right”



High variance  
(overfit)


# Predicting the price of art

(Pablo Gutierrez '18)




# Predicting the price of art

- what features might you use?
  - without specifying how to weight them



# Predicting the price of art

- what features might you use?
  - without specifying how to weight them
- artist
- previous sale prices
- time period
- provenance, previous owners, story
- topic, style, theme
- medium, size, colors, ...
- ...



How much?



Claude Monet Claude Monet, La route de la Ferme Saint-Siméon en hiver

**Sotheby's Estimate:**


**\$5,924,170 - \$7,898,894**

**Model Estimate:**

**\$7,703,839 - \$10,929,222**


**Selling Price (2008):**

**\$8,444,905**




*Claude Monet* Claude Monet, La route de la Ferme Saint-Siméon en hiver

How much?







How much? \$179M 2015



How much?




How much?




# Clustering (learning from unlabeled data)

- **contrast with supervised learning**
  - supervised learning
    - given a set of labels, fit a hypothesis to it
  - unsupervised learning
    - try and determine structure in the data
    - clustering algorithm groups data together based on data features
- **good for**
  - market segmentation - group customers into different market segments
  - social network analysis - Facebook "smartlists"
  - topic analysis
  - authorship



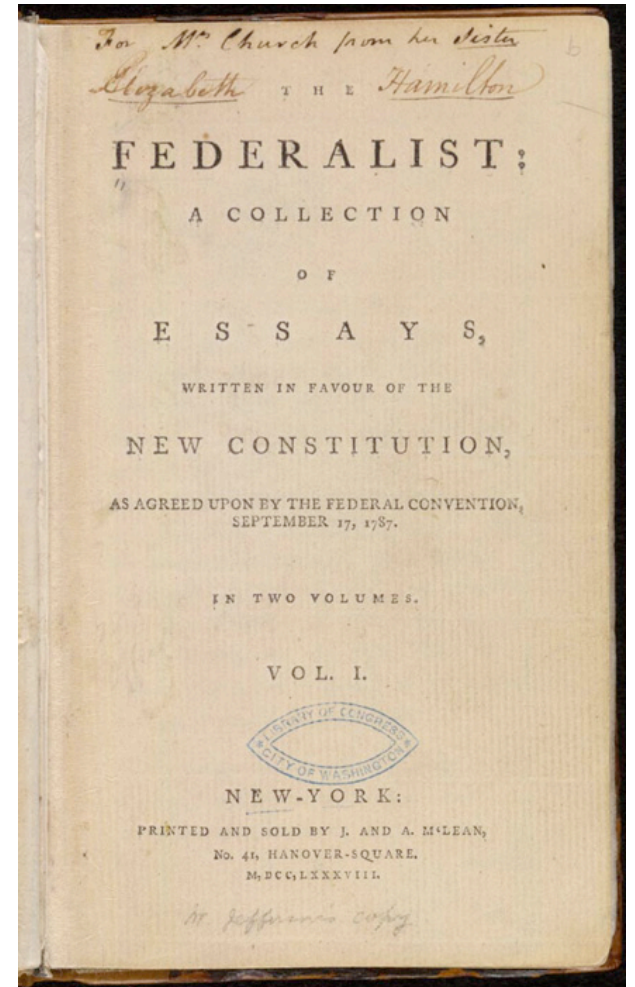
# Who wrote the Federalist Papers?

- 85 articles / essays in 1787-8 by Hamilton, Madison, Jay as "Publius"
- who wrote which ones?
- a classic example of author attribution
- what features might we use to classify?



# Who wrote the Federalist Papers?

- 85 articles / essays in 1787-8 by Hamilton, Madison, Jay as "Publius"
- who wrote which ones?
- a classic example of author attribution
- what features?
- sentence lengths
- words used
- distribution of parts of speech
- syntactic structures
- punctuation
- ...



H 51

M 29

J 5

# Recommendation systems




# Recommendation systems


"Customers who viewed this item also viewed"




**Understanding the Digital World: What You Need to Know about Computers,...**  
Brian Kernighan  
★★★★☆ 10  
Kindle Edition  
**\$14.46**



**Millions, Billions, Zillions: Defending Yourself in a World of Too Many...**  
Brian Kernighan  
★★★★☆ 7  
Kindle Edition  
**\$13.49**



**The Practice of Programming (Addison-Wesley Professional...**  
> Brian W. Kernighan  
★★★★☆ 62  
Kindle Edition  
**\$31.19**



**Adventures of a Computational Explorer**  
> Stephen Wolfram  
★★★★☆ 13  
Kindle Edition  
**\$9.99**



**The Go Programming Language (Addison-Wesley Professional Computing...**  
Alan A. A. Donovan  
★★★★☆ 171  
Kindle Edition  
**\$22.33**



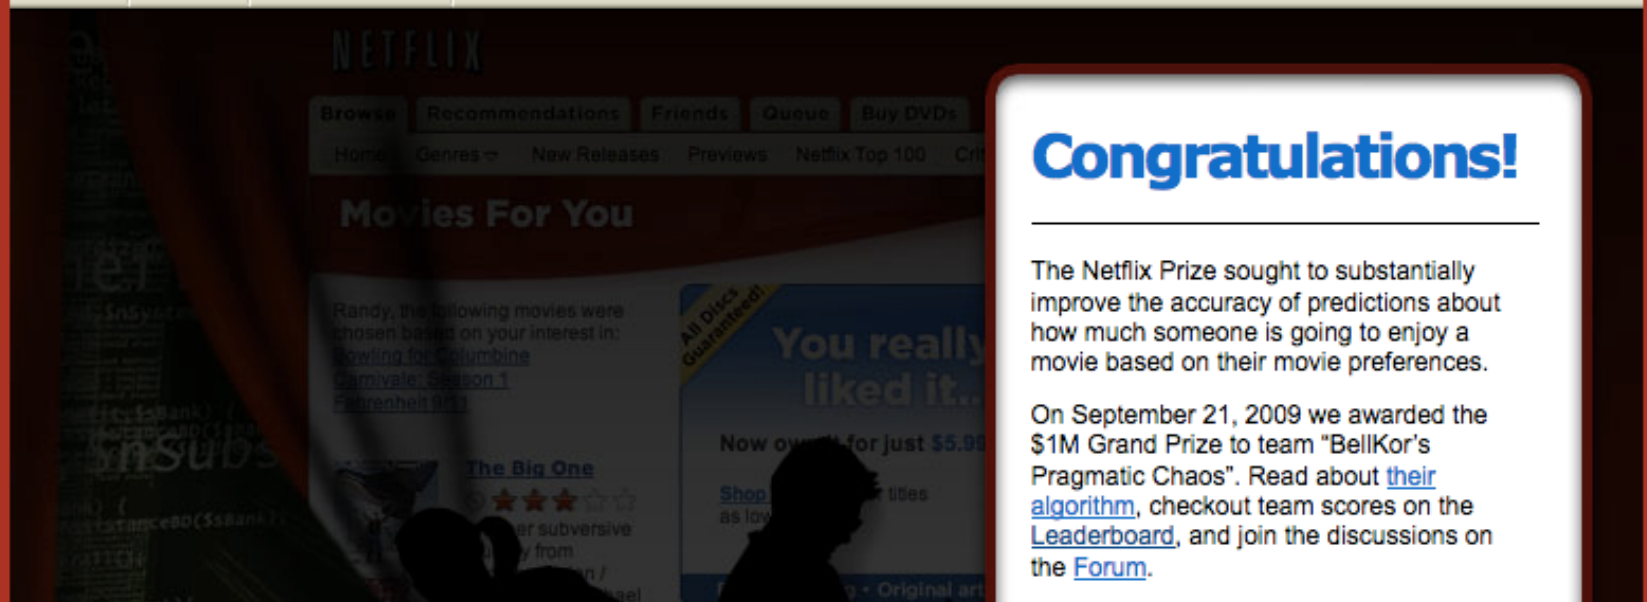
**Sex Drugs and Unix**  
Alex Morton  
★★★★☆ 4  
Kindle Edition  
**\$4.99**



## Netflix Prize

COMPLETED

Home Rules Leaderboard Update



## Congratulations!

The Netflix Prize sought to substantially improve the accuracy of predictions about how much someone is going to enjoy a movie based on their movie preferences.

On September 21, 2009 we awarded the \$1M Grand Prize to team "BellKor's Pragmatic Chaos". Read about [their algorithm](#), checkout team scores on the [Leaderboard](#), and join the discussions on the [Forum](#).

Rank	Team Name	Best Test Score	% Improvement	Best Submit Time
<b>Grand Prize - RMSE = 0.8567 - Winning Team: BellKor's Pragmatic Chaos</b>				
1	<a href="#">BellKor's Pragmatic Chaos</a>	0.8567	10.06	2009-07-26 18:18:28
2	<a href="#">The Ensemble</a>	0.8567	10.06	2009-07-26 18:38:22
3	<a href="#">Grand Prize Team</a>	0.8582	9.90	2009-07-10 21:24:40
4	<a href="#">Opera Solutions and Vandelay United</a>	0.8588	9.84	2009-07-10 01:12:31
5	<a href="#">Vandelay Industries !</a>	0.8591	9.81	2009-07-10 00:32:20
6	<a href="#">PragmaticTheory</a>	0.8594	9.77	2009-06-24 12:06:56
7	<a href="#">BellKor in BigChaos</a>	0.8601	9.70	2009-05-13 08:14:09

# How To Break Anonymity of the Netflix Prize Dataset

Arvind Narayanan, Vitaly Shmatikov




*(Submitted on 18 Oct 2006 (v1), last revised 22 Nov 2007 (this version, v2))*

We present a new class of statistical de-anonymization attacks against high-dimensional micro-data, such as individual preferences, recommendations, transaction records and so on. Our techniques are robust to perturbation in the data and tolerate some mistakes in the adversary's background knowledge.


We apply our de-anonymization methodology to the Netflix Prize dataset, which contains anonymous movie ratings of 500,000 subscribers of Netflix, the world's largest online movie rental service. We demonstrate that an adversary who knows only a little bit about an individual subscriber can easily identify this subscriber's record in the dataset. Using the Internet Movie Database as the source of background knowledge, we successfully identified the Netflix records of known users, uncovering their apparent political preferences and other potentially sensitive information.

# Neural networks, deep learning

- simulate human brain structure with artificial neurons in simple connection patterns



# Neural networks (from [vas3k.com/blog/machine\\_learning](http://vas3k.com/blog/machine_learning))



# Reinforcement learning

- learns by feedback from "real world"
- very successful for things like games: Go, chess
  - AlphaGo, AlphaGo Zero
  - AlphaZero

# Google's AlphaZero Destroys Stockfish In 100-Game Match



FM MikeKlein  

Dec 6, 2017, 12:50 PM |  352 | Chess Event Coverage

M English <

Chess changed forever today. And maybe the rest of the world did, too.

A little more than a year after **AlphaGo** sensationally won against the top Go player, the artificial-intelligence program **AlphaZero** has obliterated the [highest-rated chess engine](#).

**Stockfish**, which for most top players is their go-to preparation tool, and which won the **2016 TCEC Championship** and the **2017 Chess.com Computer Chess Championship**, didn't stand a chance. AlphaZero won the closed-door, 100-game match with 28 wins, 72 draws, and zero losses.

Oh, and it took AlphaZero only four hours to "learn" chess. Sorry humans, you had a good run.

That's right -- the programmers of AlphaZero, housed within the **DeepMind** division of Google, had it use a type of "[machine learning](#)," specifically reinforcement learning. Put more plainly, AlphaZero was not "taught" the game in the traditional sense. That means no opening book, no endgame tables, and apparently no complicated algorithms dissecting minute differences between center pawns and side pawns.

# Natural language processing (NLP)

- **understanding text**
  - parsing, syntactic structure
  - topic modeling
  - sentiment analysis
  - text generation
- **text to speech**
- **speech to text**
- **translation**

# ML / AI issues

- **algorithmic fairness**
  - results can't be better than training data
  - if that has implicit or explicit biases, results are biased
  - can we detect and eliminate bias?
- **accountability and explainability**
  - what is the algorithm really doing?
  - can its results be explained
- **appropriate uses?**
  - prison sentencing
  - drone strikes
  - weapon systems
  - resume evaluation
  - medical decisions
  - ...
- **to learn more:**
  - <https://fairmlbook.org>



# Amazon scraps secret AI recruiting tool that showed bias against women (10/10/18)

- SAN FRANCISCO (Reuters) - Amazon.com Inc's (AMZN.O) machine-learning specialists uncovered a big problem: their new recruiting engine did not like women.
- [...] Amazon's computer models were trained to vet applicants by observing patterns in resumes submitted to the company over a 10-year period. Most came from men, a reflection of male dominance across the tech industry.
- In effect, Amazon's system taught itself that male candidates were preferable. It penalized resumes that included the word "women's," as in "women's chess club captain." And it downgraded graduates of two all-women's colleges, according to people familiar with the matter.