

# CS181: Introduction to Machine Learning

## Spring 2021

Finale Doshi-Velez and David C. Parkes  
Harvard Computer Science

Staff: Nari Johnson (head TF), Prasad Chhabria, Jonathan Chu, Andrew Kim, Dylan Li, Yash Nair, Sanjana Narayanan, Ife Omidiran, Karthik Rao, Rylan Schaeffer, Katherine Tian, and Bill Zhang

# What this course is about



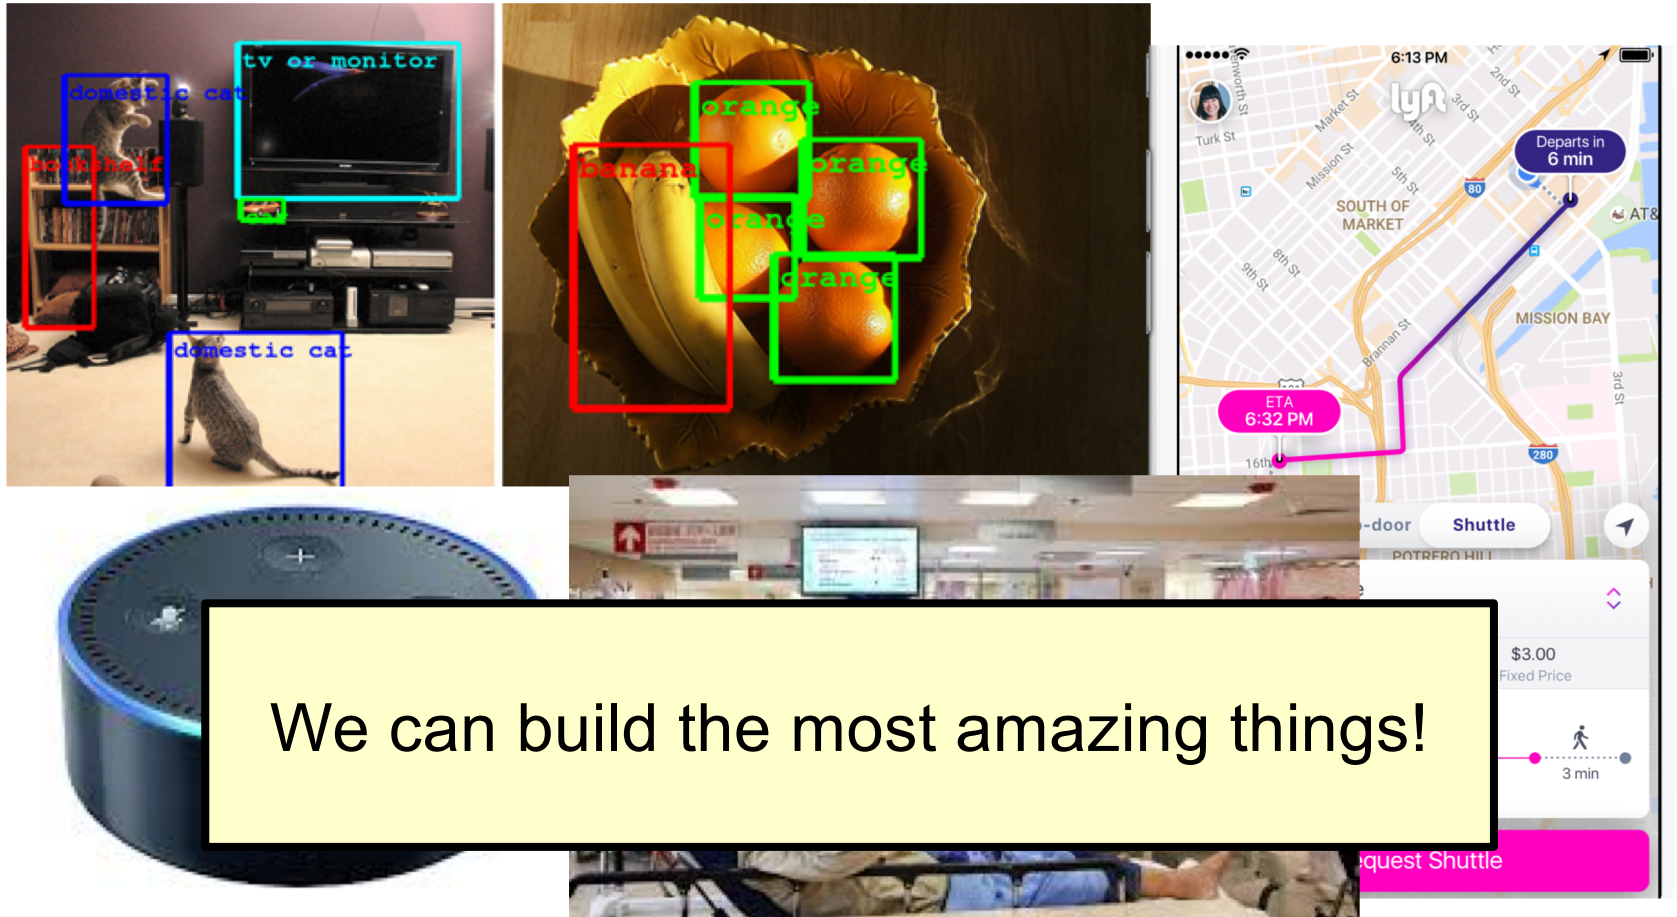
<https://s.aolcdn.com/hss/storage/midas/cb433fb6c234b014394ad56714ceee8b/200707597/google-image-recognition-tech.jpg>

[https://cdn.shopify.com/s/files/1/0684/3445/products/Partner\\_AmazonEcho\\_small\\_031932b3-c31d-4977-a922-5ebbe052bad6.jpg?v=1497504913](https://cdn.shopify.com/s/files/1/0684/3445/products/Partner_AmazonEcho_small_031932b3-c31d-4977-a922-5ebbe052bad6.jpg?v=1497504913)

<https://cdn.lyft.com/brochure/shuttle-screenshot.1e79cd66.png>

[https://cdn3.i-scmp.com/sites/default/files/styles/980x551/public/images/methode/2016/05/29/b0dc6826-2577-11e6-80b1-a87df553e801\\_1280x720.JPG?itok=QVZxRMvy](https://cdn3.i-scmp.com/sites/default/files/styles/980x551/public/images/methode/2016/05/29/b0dc6826-2577-11e6-80b1-a87df553e801_1280x720.JPG?itok=QVZxRMvy)

# What this course is about



We can build the most amazing things!

<https://s.aolcdn.com/hss/storage/midas/cb433fb6c234b014394ad56714ceee8b/200707597/google-image-recognition-tech.jpg>

[https://cdn.shopify.com/s/files/1/0684/3445/products/Partner\\_AmazonEcho\\_small\\_031932b3-c31d-4977-a922-5ebbe052bad6.jpg?v=1497504913](https://cdn.shopify.com/s/files/1/0684/3445/products/Partner_AmazonEcho_small_031932b3-c31d-4977-a922-5ebbe052bad6.jpg?v=1497504913)

<https://cdn.lyft.com/brochure/shuttle-screenshot.1e79cd66.png>

[https://cdn3.i-scmp.com/sites/default/files/styles/980x551/public/images/methode/2016/05/29/b0dc6826-2577-11e6-80b1-a87df553e801\\_1280x720.JPG?itok=QVZxRMvY](https://cdn3.i-scmp.com/sites/default/files/styles/980x551/public/images/methode/2016/05/29/b0dc6826-2577-11e6-80b1-a87df553e801_1280x720.JPG?itok=QVZxRMvY)

# What this course is about

BUSINESS NEWS    OCTOBER 9, 2018 / 11:12 PM / 2 MONTHS AGO

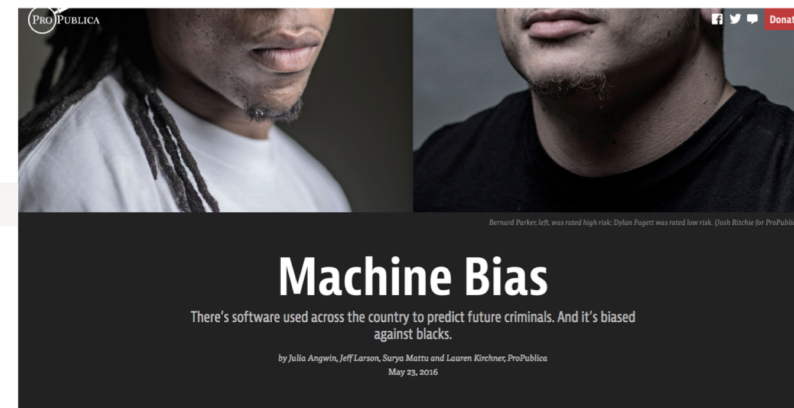
## Amazon scraps secret AI recruiting tool that showed bias against women



HEALTH TECH

## 3 lessons from Stanford's Covid-19 vaccine algorithm debacle

By CASEY ROSS @caseymross and ERIN BRODWIN @erbrod / DECEMBER 21, 2020



<https://twitter.com/geraldmellor/status/712880710328139776>

<https://www.statnews.com/2020/12/21/stanford-covid19-vaccine-algorithm/>

<https://www.theguardian.com/education/2020/aug/20/england-exams-row-timeline-was-ofqual-warned-of-algorithm-bias>

<https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>

<https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>

# What this course is about

BUSINESS NEWS OCTOBER 9, 2018 / 11:12 PM / 2 MONTHS AGO

## Amazon scraps secret AI recruiting tool that showed bias against women



Rigor, professionalism, ethics, and inclusion are essential!



- <https://twitter.com/geraldmellor/status/712880710328139776>
- <https://www.statnews.com/2020/12/21/stanford-covid19-vaccine-algorithm/>
- <https://www.theguardian.com/education/2020/aug/20/england-exams-row-timeline-was-ofqual-warned-of-algorithm-bias>
- <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scrap-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>

David and Finale chat  
about machine learning  
in real life, in our work

# Your Turn: Detecting Fake Videos



Obama: <https://www.youtube.com/watch?v=AmUC4m6w1wo&feature=youtu.be>

Nixon: [https://www.youtube.com/watch?time\\_continue=2&v=yaq4sWFvnAY&feature=emb\\_logo](https://www.youtube.com/watch?time_continue=2&v=yaq4sWFvnAY&feature=emb_logo)



Ethics Question: Why might we care about deepfakes, when Impersonation is nothing new?

SNL: <https://www.youtube.com/watch?v=BG-ZDrypec>

Russian Robot: <https://www.youtube.com/watch?v=om5z3Uck9IY>



**How we will get there:  
Zen and Hard Work**

# Zen



- Make appropriate model choices
- Have sufficient understanding to learn and apply the new techniques
- Identify sources of error
- Evaluate carefully

# Zen and Hard Work

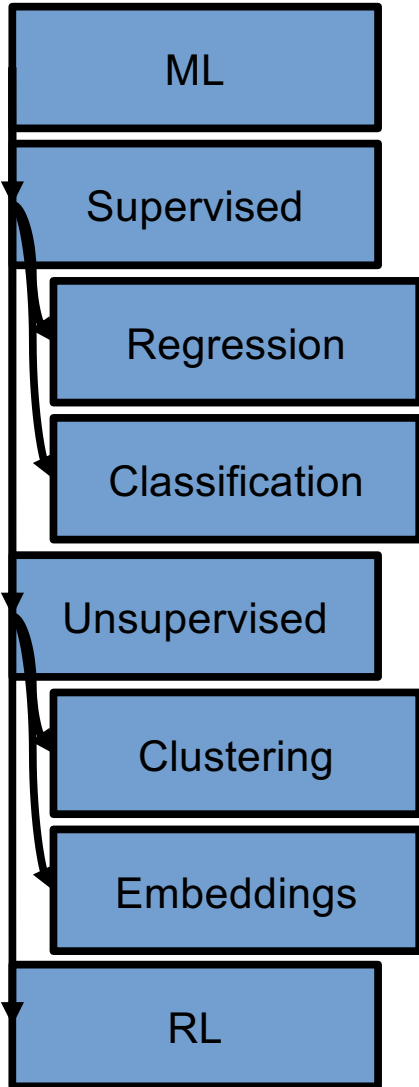


- Make appropriate model choices
- Have sufficient understanding to learn and apply the new techniques
- Identify sources of error
- Evaluate carefully


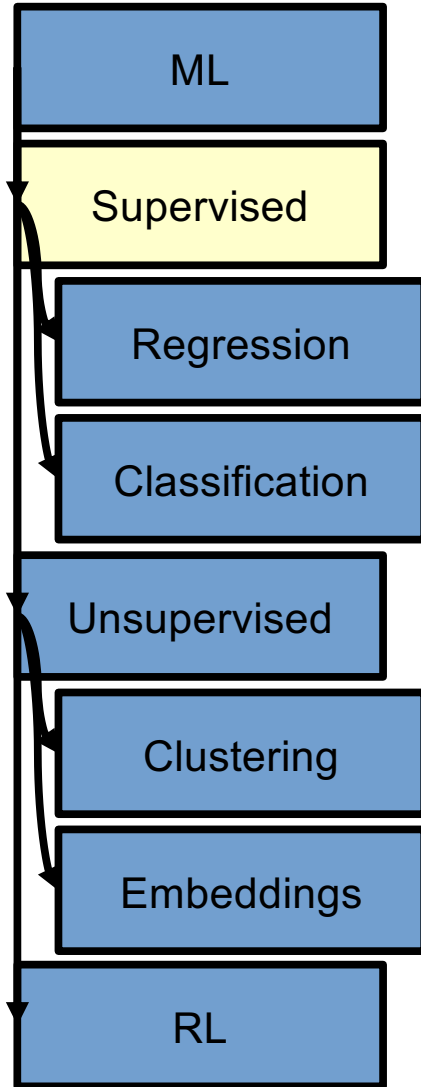


# **How we will get there: Taxonomy of Problems**


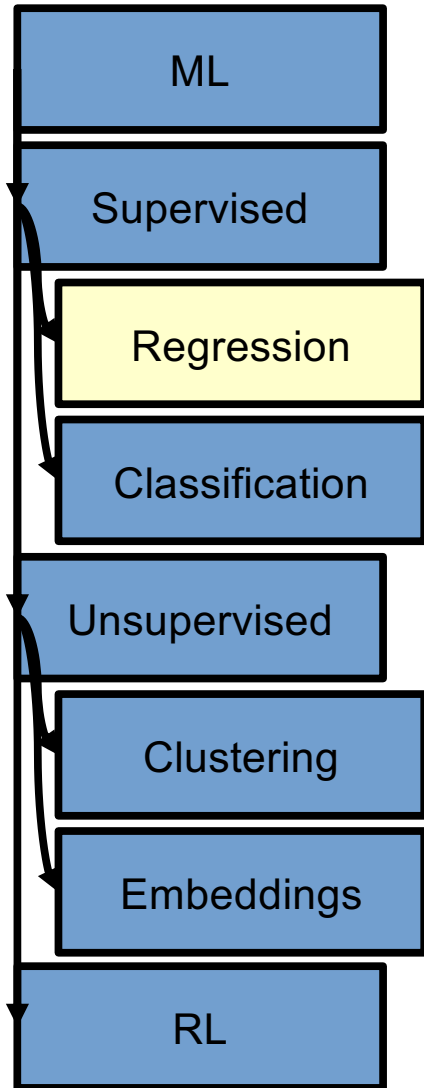
# Machine Learning Taxonomy



# Machine Learning Taxonomy



# Terminology: Regression



# Example: Virtu Financial


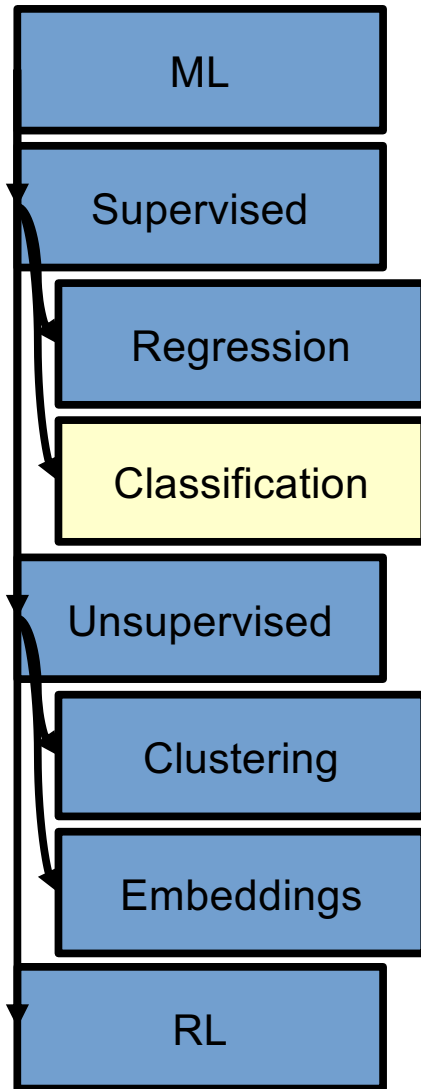


Core technology:  
regression model to predict  
stock price, often in the  
very short-term

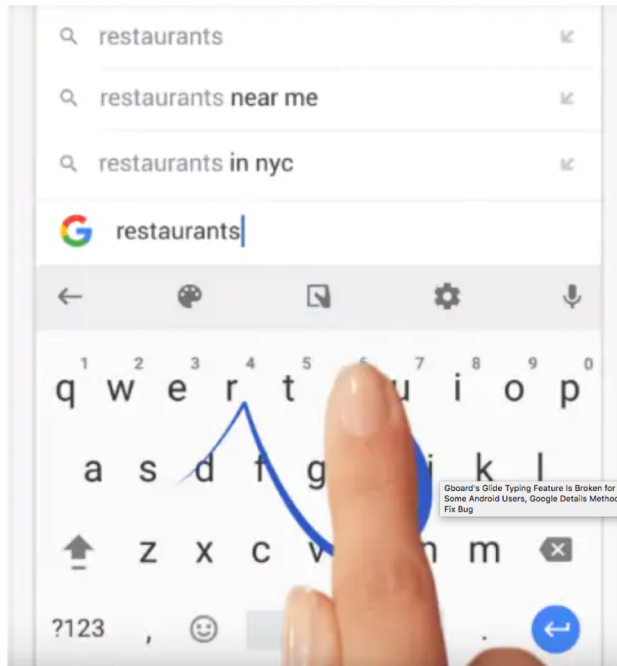
<https://fortune.com/2015/04/16/how-virtu-financial-went-from-pr-pariah-to-ipo-idol/>  
<https://www.virtu.com/>



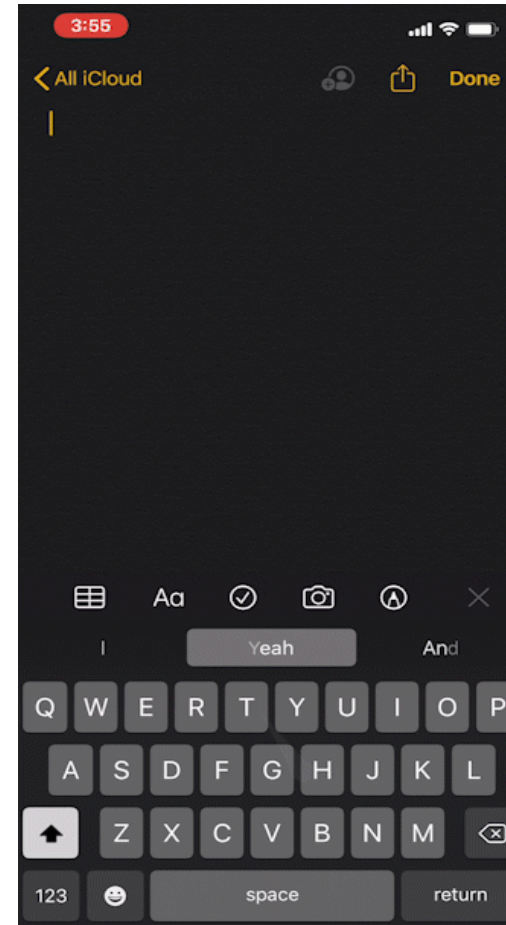
# Terminology: Classification



# Example: “Swipe typing”




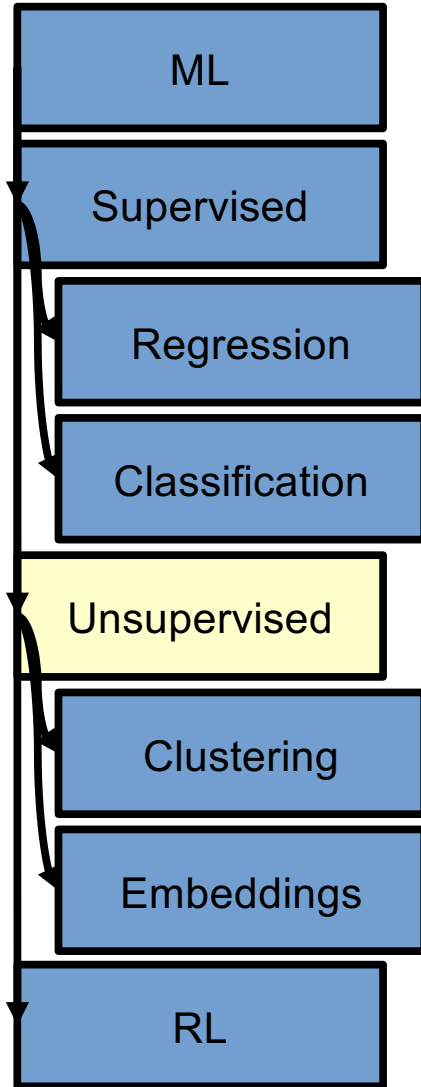
Android Glide Typing




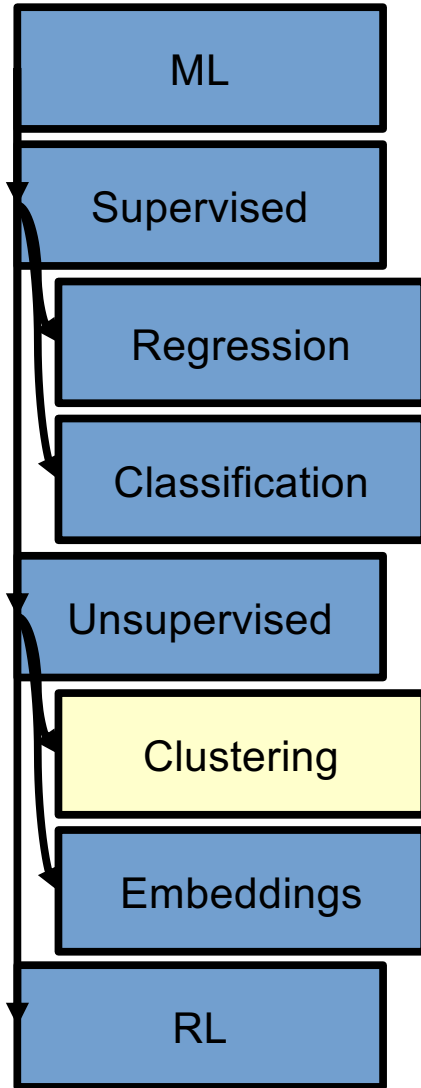
iOS QuickPath

Core technology:  
language model  
to predict which  
word is intended

# Machine Learning Taxonomy



# Terminology: Clustering



# Example: Google News

## Protesters Across Russia Call For Release Of Alexei Navalny

NPR · 1 hour ago

- **Protests erupt in Russia after arrest of opposition leader Alexey Navalny**  
CBS This Morning · 7 hours ago
- **Hundreds, including Navalny's wife Yulia, detained as protests in his support sweep across Russia**  
CNN · 5 hours ago
- **Biden must act to save Navalny's life — and hopes of freedom in Russia**  
The Washington Post · 1 hour ago · Opinion
- **Poison, Arrest and a Secret Palace Fire Up Russians**  
Bloomberg · Yesterday · Opinion

 [View Full Coverage](#)




## State of play: Impeachment trial date will give more time to confirm Biden Cabinet

CNN · 8 hours ago

- **Trump's second impeachment trial to begin week of February 8**  
CBS This Morning · 8 hours ago

 [View Full Coverage](#)


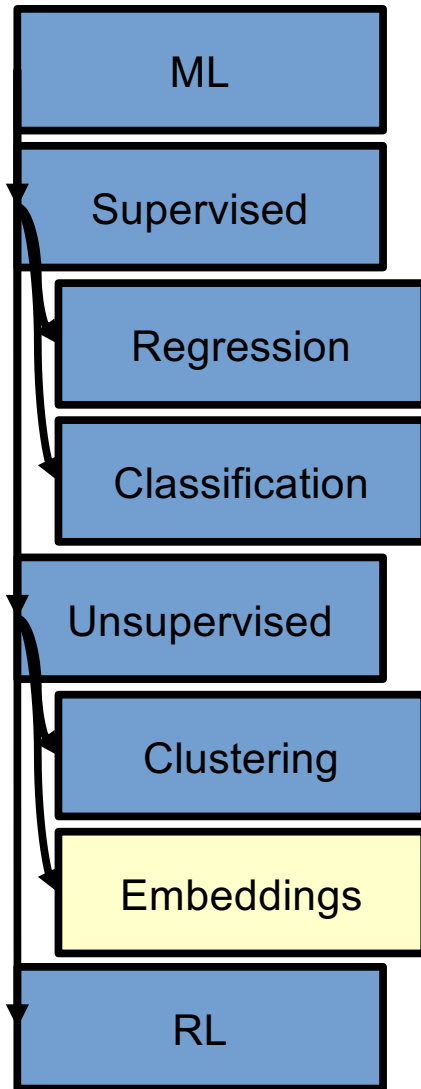


## Biden signs executive orders on stimulus checks, food stamps and



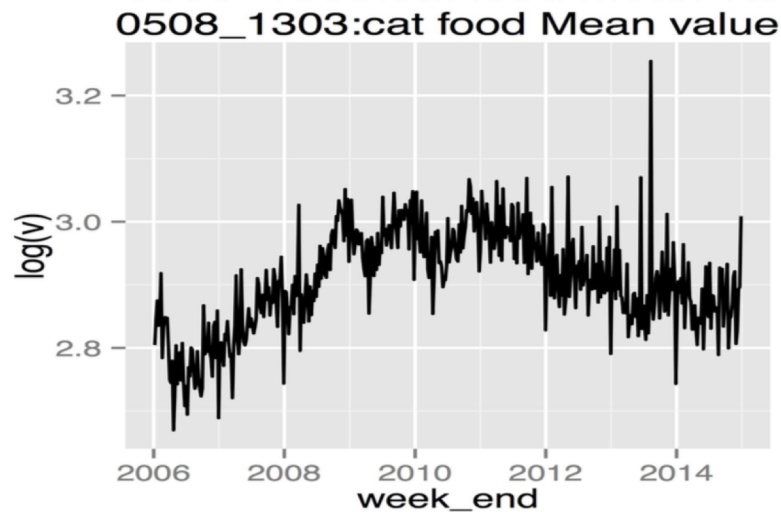
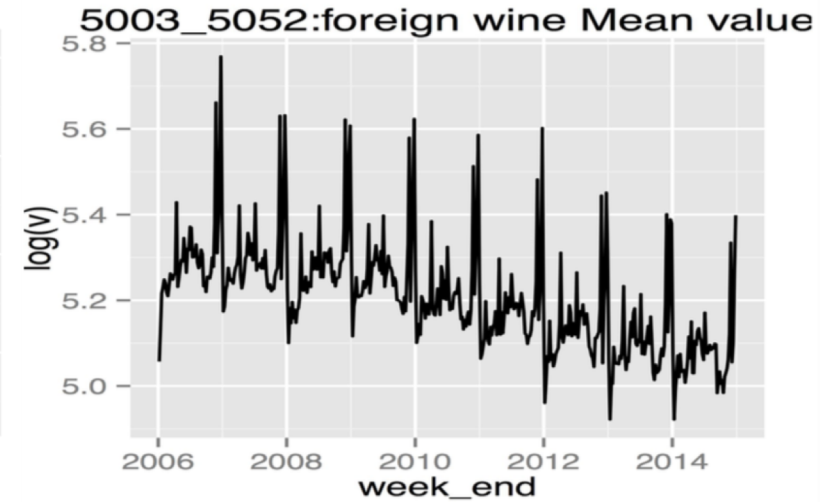
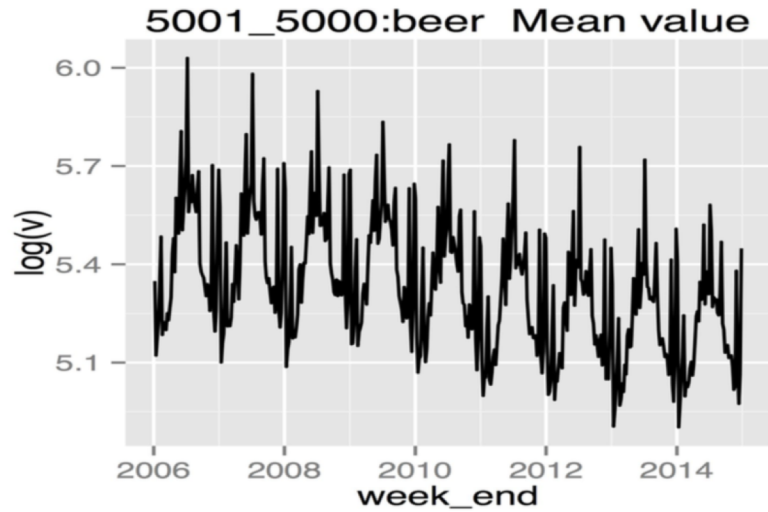
Core technology:  
clustering based  
on language and  
links to  
understand  
which stories are  
related

# Terminology: Embedding



# Example: Point-of-Sales Data

(S. Ng 2016)

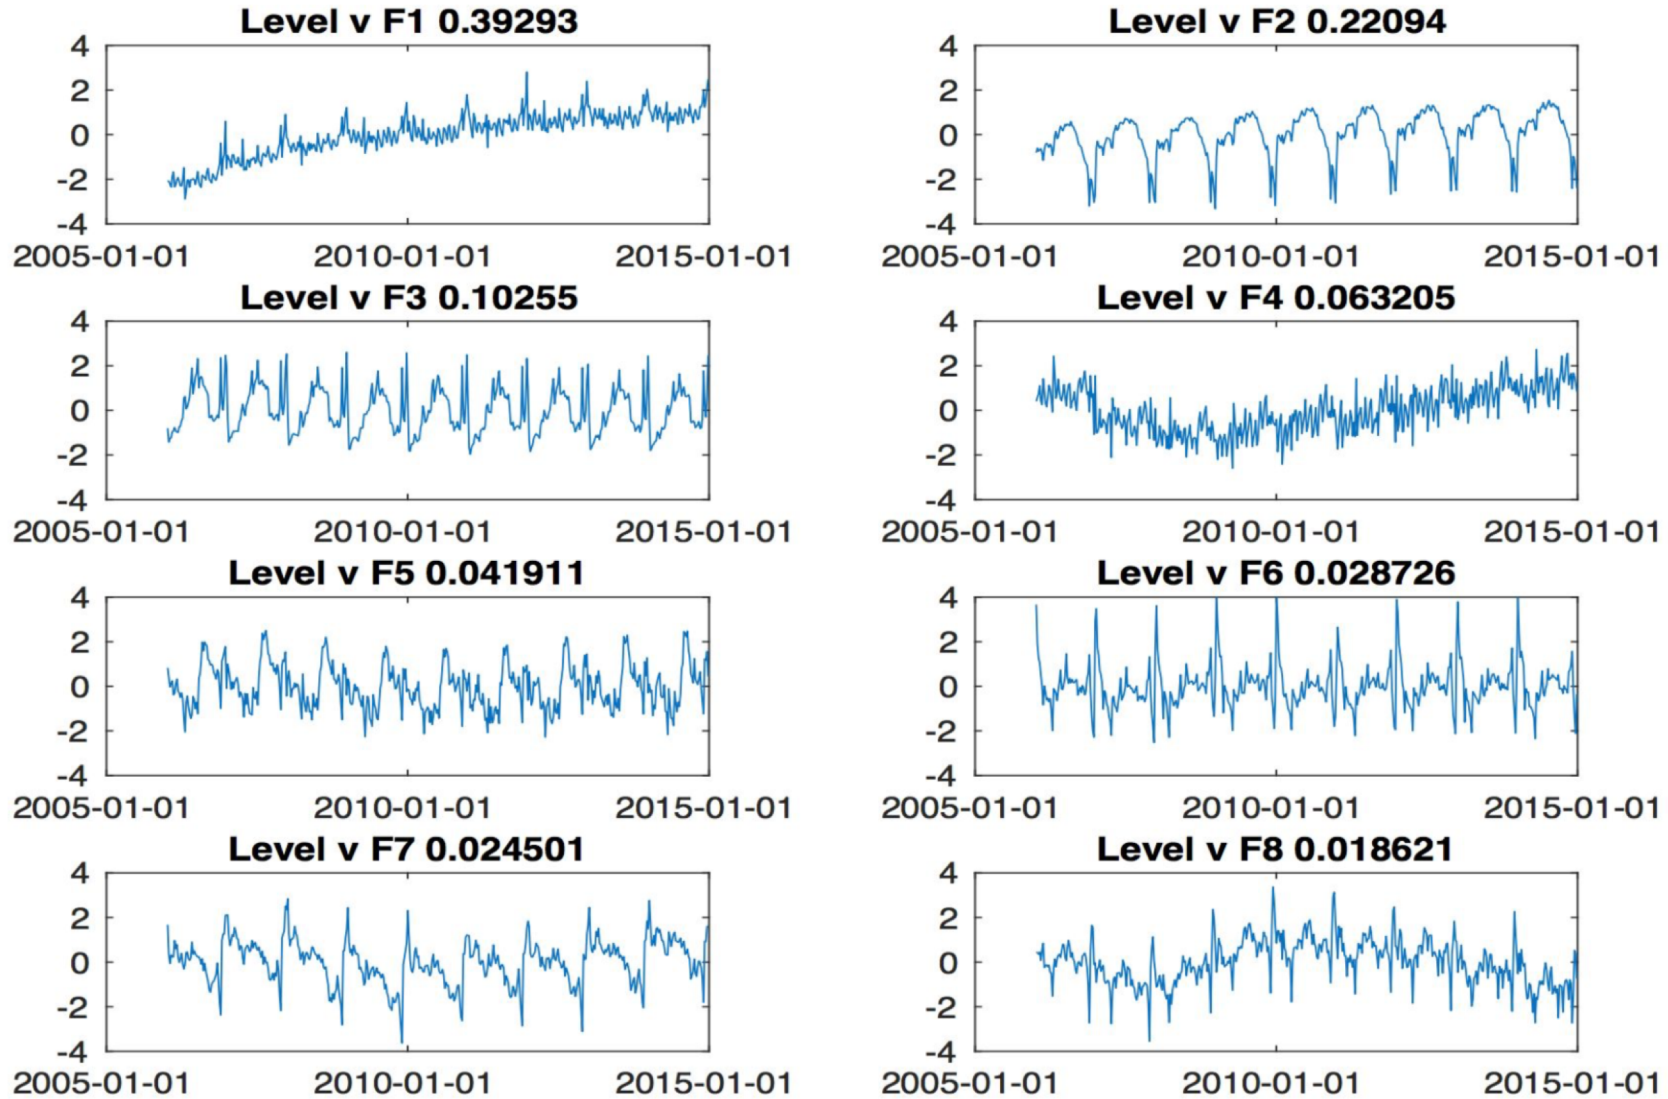


Nielsen scanner data,  
2006-2014 on 1073  
products in different  
locations ~6 TB

Can we see the effect on  
demand of the 2008-09  
financial crisis?


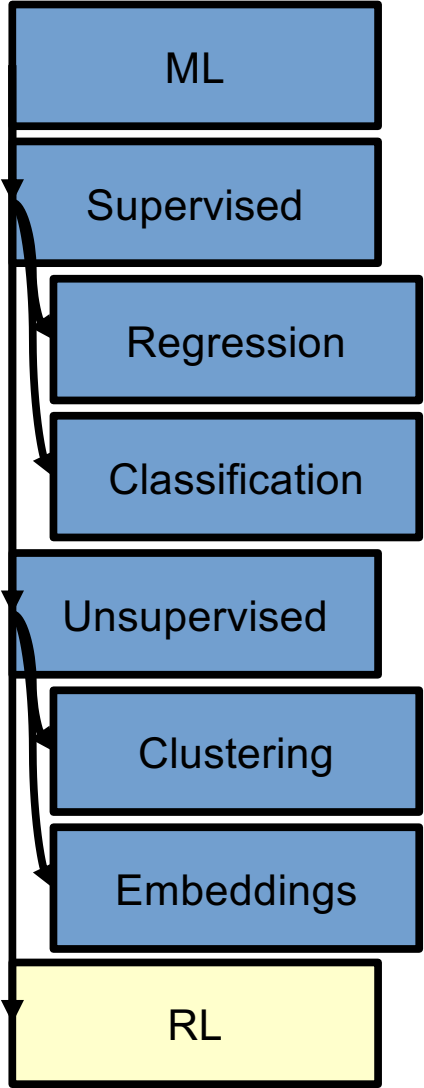
# Eight Components (PCA)

(S. Ng 2016)

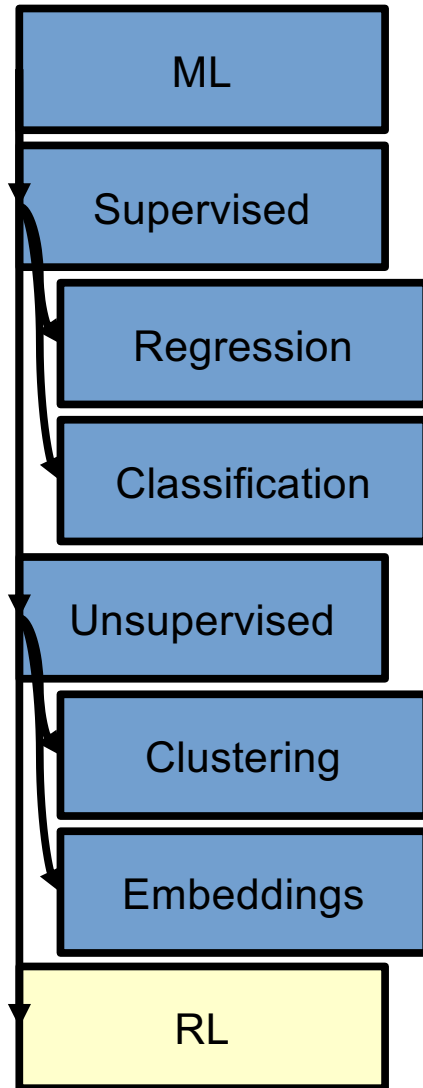




# Terminology: Reinforcement Learning



# Terminology: Reinforcement Learning



Kormushev et al., 2010  
[https://www.youtube.com/watch?v=W\\_gxLKSsSIE](https://www.youtube.com/watch?v=W_gxLKSsSIE)




2016



**When you're teaching Atlas to do new things, is it using any kind of machine learning? And if not, why not?**

As a company, we've explored a lot of things, but Atlas is not using a learning controller right now. I expect that a day will come when we will. [currently] reacting to forces, online and offline trajectory optimization, and model predictive control... We plan on using learning to extend and build on the foundation of software and hardware that we've developed...

<https://spectrum.ieee.org/automaton/robotics/humanoids/how-boston-dynamics-taught-its-robots-to-dance>



+ graphical models, reinforcement learning

# **How we will get there: Course Structure**

# Structure of the Course

- Follow the “cube”++
  - Supervised: regression, classification
  - Unsupervised: clustering, embeddings
  - Graphical Models, Reinforcement Learning
- Excursions to discuss model classes, model selection, evaluation...



# Lecture Structure

- Chit chat
- Session one
- Break
- Session two
- Concept checks

*Please join us!*

Lectures will include real world examples and concept exercises; we'll also discuss ethics

# Logistics

- Attend  $\geq 7$  live lectures (exemptions!), lectures recorded
- Regular homeworks: biweekly (8p Friday), done individually, theoretical and some programming elements (LaTeX, Python)
- Practical: groups of 2-3 people, 1 week, working on a real-world problem
- Six late days, at most two per problem set
- Two timed, closed-book midterms proctored over Zoom
- Sections: section zero with math preliminary material this week, then weekly sections on content; new “Beyond CS 181” sections
- Lots of office hours (almost all on *Collaborate*)
- Ed: for all discussion and questions on course content (content, clarification, private message)
- And more: concept checks, lecture recaps, CS 181 “textbook”

# This week only

## Office hours

All on **Congregate** except for Finale's

Finale Tuesday 3.30-4.30p (today!)

David Friday 2.30-4p

Nari + Sanjana Weds 8-9:30PM

Kat + Yash Tue 10:30pm-12am  
(today!)

Andrew + Ife Fri 6-7:30PM

Dylan + Rylan Sunday 1-2:30pm

Jonathan + Karthik Saturday 1-2:30pm

Prasidh Thurs 10:15am-11:45am

## Section 0

All on **Zoom** (will be recorded)

- Andrew + Jonathan. Thurs 7-8pm. Probability
- Kat + Yash Wed 8-9pm Lin Alg
- Dylan + Rylan Friday 1-2 pm Coding  
Numpy
- Sanjana + Prasidh. Wed 9-10 am Hw 0

# Collaboration Policy

- You may work with others, but your write-up must be entirely written by yourself in your own words. You may help each other debug code, but again, the code must be written by you.
- *Include the names of anyone you worked with in your write-up.*
- Do not copy parts of another person's assignment or jointly type up an assignment.
- You can make use of textbooks and online sources to help in answering questions, but you must cite your sources (and you should be ready to explain your answer to a member of the teaching staff)
- Do not look up solutions to specific questions from the internet or other sources (e.g. friends from previous years)

# FAQ: Is this the right course?

- CS181: Rigorous conceptual grounding of a broad range of machine learning ideas with math and code
- Goal: Reason about when ML algorithms work, and why
- Alternatives:
  - For mostly theory: Stat195, etc.
  - For mostly practice: CS109, etc.
  - For depth rather than breadth: CS282, etc.

FAQ: Can I sit in/audit/simultaneously  
enroll/take CS181?

**Yes!**

# FAQ: Do I need the Prereqs?


- What you need: programming, probability, calculus, and linear algebra
- See our section zero “Math review” resource (use this as a reference)
- See Math for ML resource <https://mml-book.com/>

Problem 1.1:  
Take the gradient of training loss  
with respect to the vector of  
weights  $w$


...

# The Thinking Machine (CBS, 1961)

<https://techtv.mit.edu/videos/10268-the-thinking-machine-1961---mit-centennial-film>







Arthur Samuel (1952): Learning to play checkers through self play

A Short Break  
(and then, ...  
nonparametric  
regression)