Machine learning, government, and policy analysis

SUSAN ATHEY

THE ECONOMICS OF TECHNOLOGY PROFESSOR, STANFORD GSB

Artificial Intelligence and Machine Learning

1. What does it do, how is it changing
2. Implications for industry structure
3. Societal implications
4. Direct benefits to the public sector

Machine Learning Basics

Unsupervised

- I discovered cats on YouTube!
- Dimension reduction for text, images, etc.

Supervised

- Predict an outcome (y) or classify an object on the basis of its characteristics (x)
- Personalized policies, targeting, etc.

Robots

Algorithms for taking actions as a function of the state of the world now and in the past

- Many approaches
- Often a collection of algorithms working together

Recent Advances in ML: Causal Inference

Adapt ML methods to answer questions such as:

- How can we combine "big data" on individuals to learn about the **effectiveness of policies** in the absence of experiments (by controlling for individual differences)?
- **For which groups** does a policy work well?
- How can we estimate and evaluate (personalized) optimal assignments of individuals to policies?
- How can we **experiment** with personalized treatment allocation with minimum loss to users?
- When can an **algorithm improve upon a human**?

The ML "Production Function"

In industry, ML algorithms are constantly improved

- Learning by doing
- Randomized experiments with live user traffic
- Decentralized innovation

ML as a service, cloud, open source

- Allow smaller firms, governments to benefit from others' innovations
- Cloud computing dramatically reduces costs
- Extent to which it works also varies by application

Economies of scale

- Fixed costs of R&D
- Faster, more efficient experimentation with more users
- Rate of diminishing returns to users/data varies widely by application
- "Tree" algorithms keep splitting the data
  - By user location, by device/browser type, by user history, by context

Implications for Society

Massive changes in the types of jobs available, not just for manufacturing workers, but for financial services, all sectors

- Managing and complementing machines
- Perhaps enabling workers with differing physical abilities (e.g. the elderly) to help machines do physical labor, e.g. caregiving

Automation will lower the costs of goods, services, transport, and increase the effective supply of land

- If Price = Marginal Cost, cost of living goes down, offsetting some of the losses workers experience
- If economics of automation are “Google-like,” we may see Price > MC, and workers may see lower wages without offsetting decreases in price

Big Data and the Public Sector I: Using Algorithms to Make Decisions

Local governments allocate inspectors

Police determine who to stop/search

The health care system determines guidelines for allocating patients to treatments

Judges decide pre-trial release, sentences, etc.

Hiring systems determine who to interview

Predictive Policing and Firefighting

New York, United States

- Collect data on 60 risk factors for buildings
- FireCast predicts which are most likely to catch fire
- Prioritize risky buildings for inspections
- Iterative innovation
  - Firecast 2.0 is 10 times as good as predecessor
  - Firecast 3.0, in development, uses 7,500 factors

Fine Collection

Department of Justice, United Kingdom

Collect 1 million fines per year, with high staff costs

Use adaptive experimentation to learn what the most effective SMS text message reminders are

Increase fine payment by 200%

Most effective: personalized texts with name and amount due

- Signals effective government with appropriate information to collect

Table 2: The Riskiest People Receiving Joint Replacements

Predicted Mortality Percentile	Observed Mortality Rate	Total Number Annually
1	0.562	4905
	(.027)	
2	0.530	9810
	(.02)	
5	0.456	24525
	(.012)	
10	0.345	49045
	(.008)	
20	0.228	98090
	(.005)	
30	0.165	147135
	(.004)	
100	0.057	490450
	(.001)	

Approach: use ML methods to predict mortality as a function of covariates

A large number of joint replacements going to people who die within the year

Could we just eliminate the ones above a certain risk?

Who to Release?

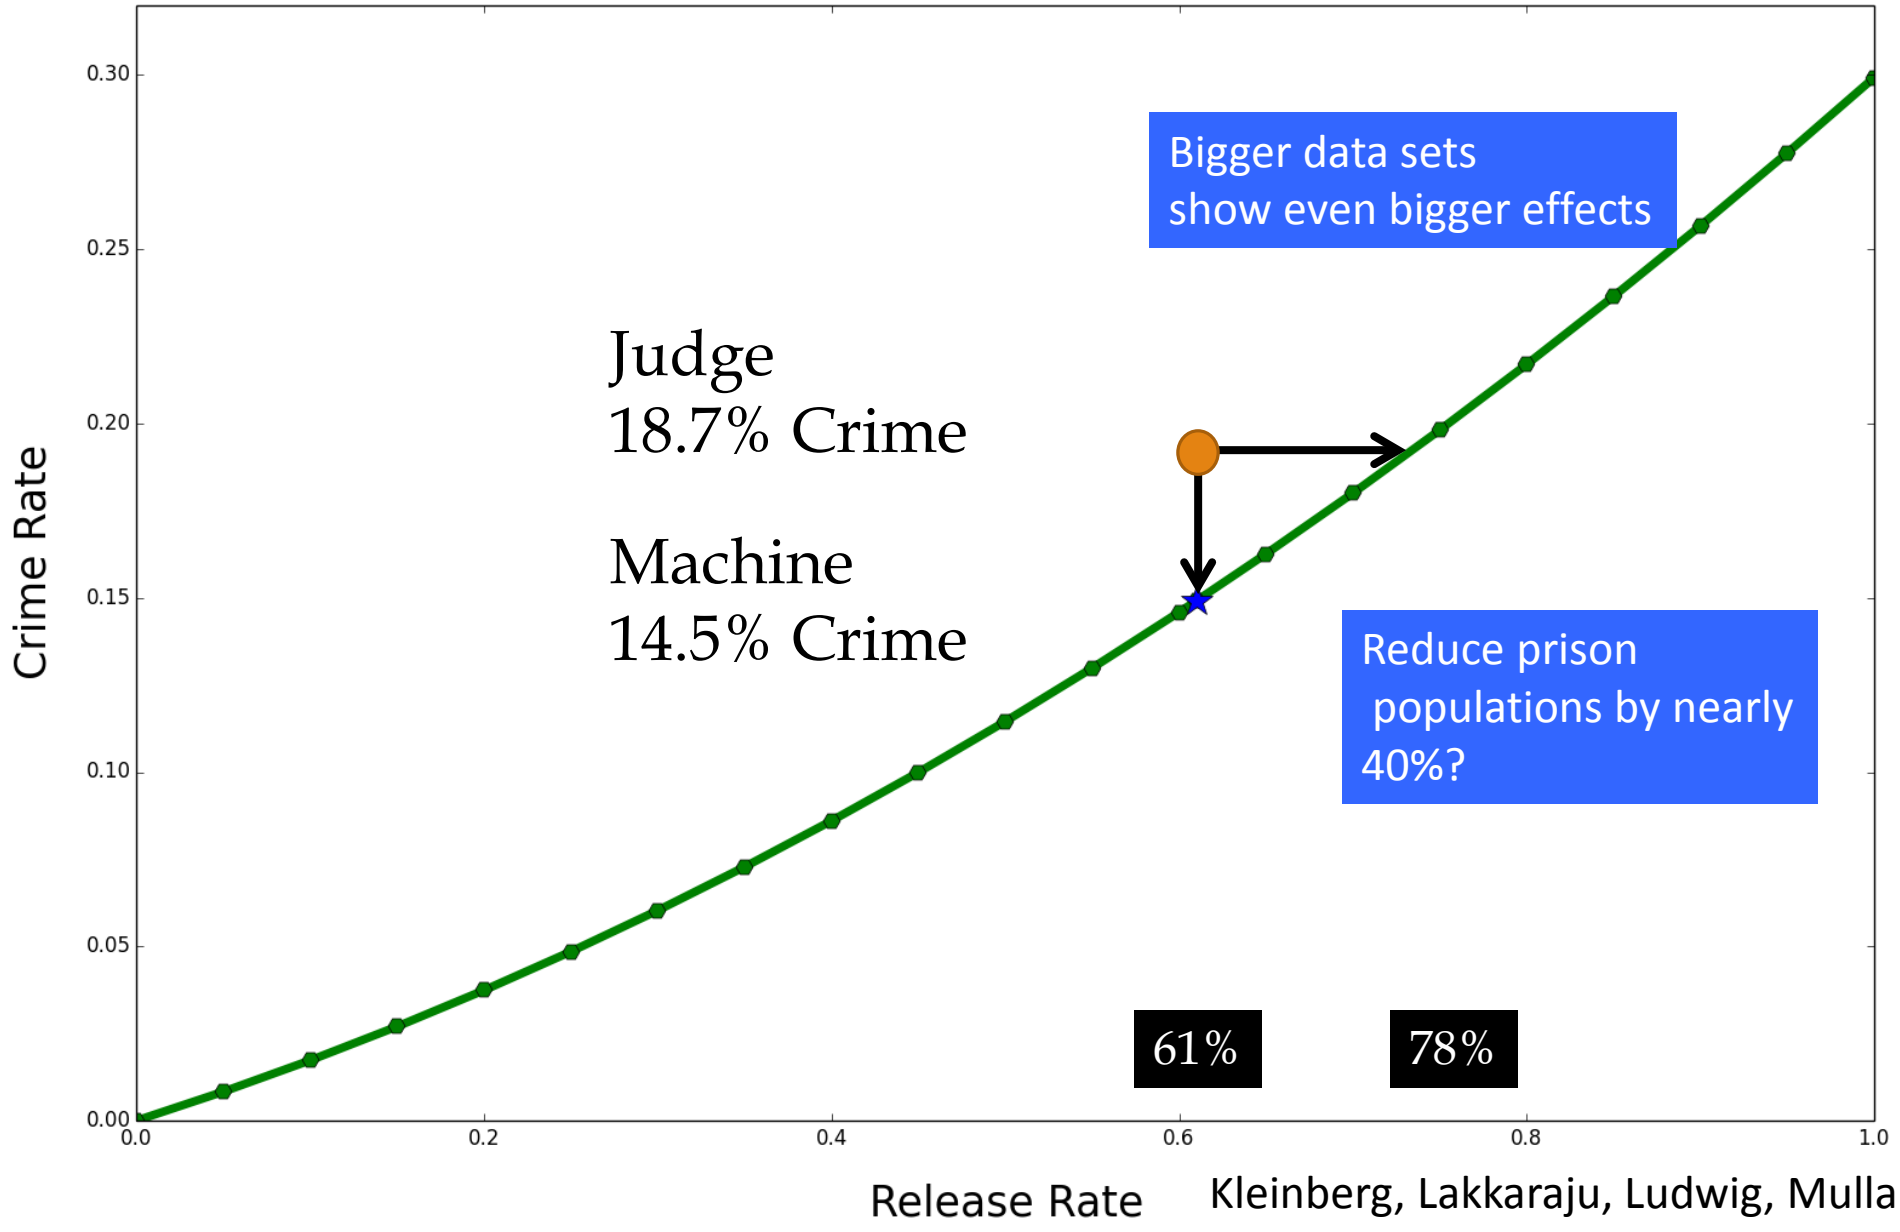

Kleinberg, Lakkaraju, Ludwig, Mullainathan (2015)

Why do Algorithms beat Humans?

Kleinberg et al fit algorithm that mimics judge release rule

- Interpret as what judges would do if they didn't use subjective information
- Finding: fitted rule beats what judges actually do
- Not difficult to preserve racial composition etc. achieved by judges

General result: humans over-react to perceived subjective factors

- Similar evidence from hiring—machines do a better job
- Avoid unproductive biases (e.g. went to my college)

Big Data and the Public Sector II: Releasing Information and Scores

Restaurant hygiene scores
- Publicizing them on Yelp improves hygiene

Hospital/doctor quality scores
- Doctors respond positively to information about own performance

Open data initiatives
- Allow firms to create apps that show transport delays, travel times

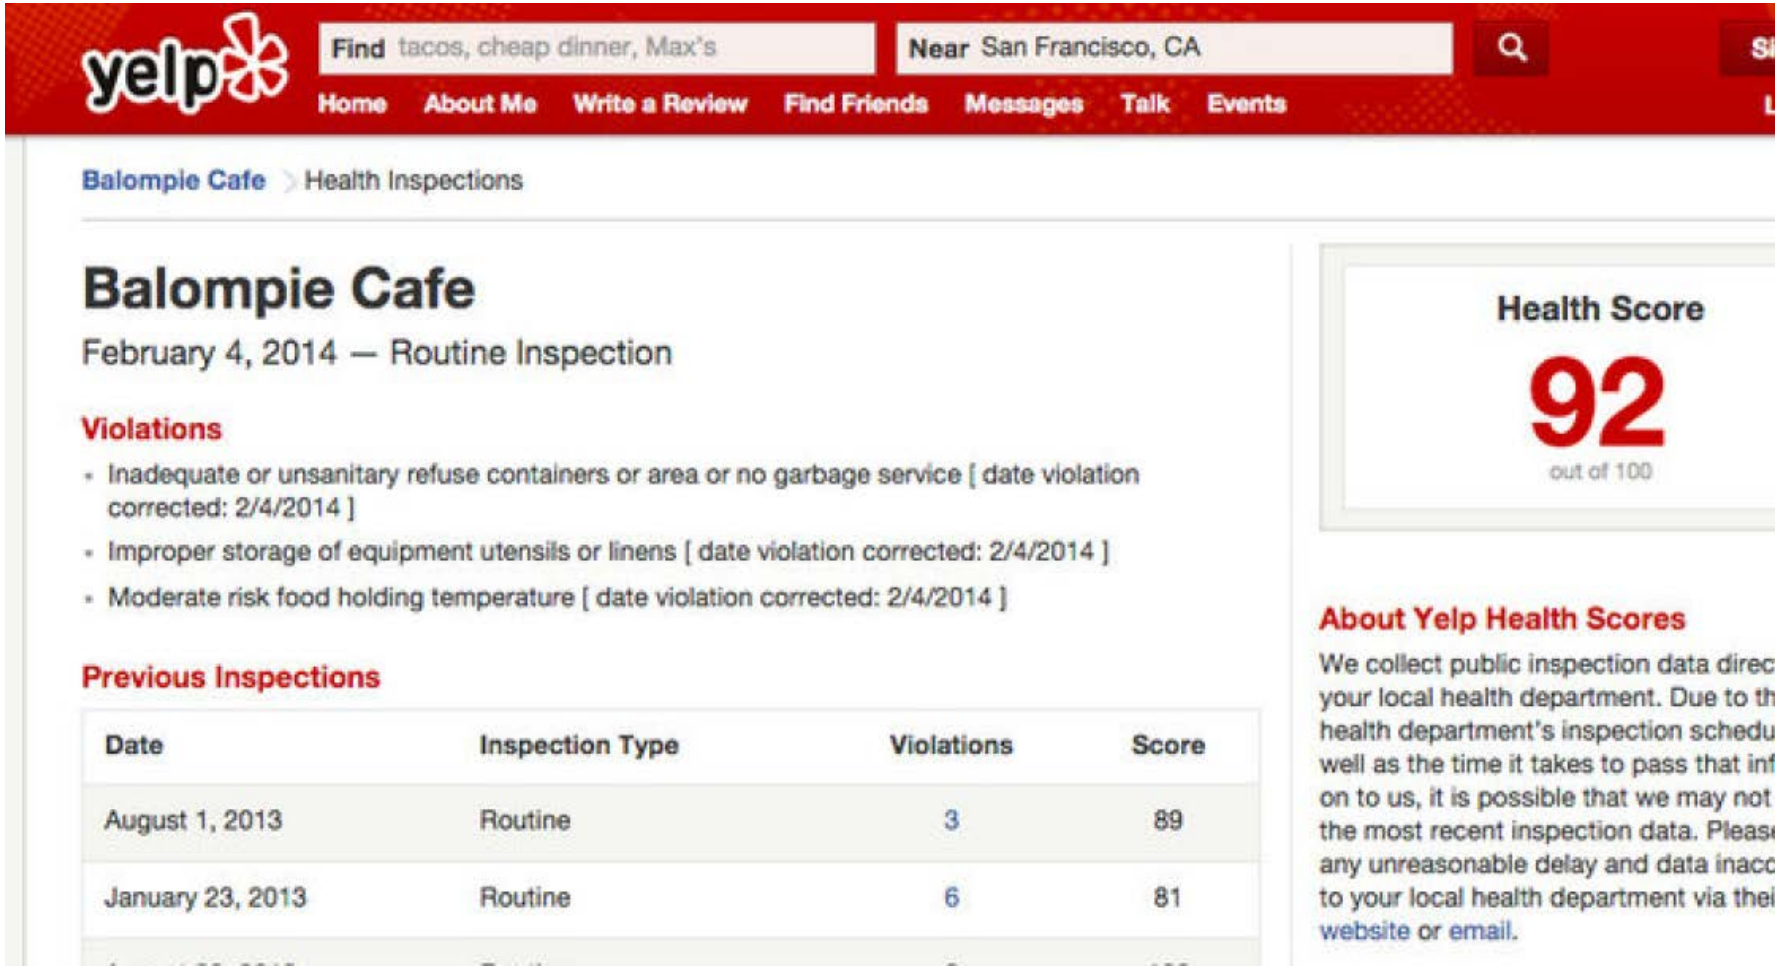

What might go wrong with this initiative?

Big Data and the Public Sector III: Measurement and Statistics

Big data/micro data for evaluating heterogeneous impact of public policies, health policies, etc.

Government releases numbers three months after economic changes that are “knowable” using credit card data, etc.

Large scale automated microdata collection

- Satellite images, streetview: neighborhood value changes, poverty
- Sensors: noise, air quality, energy consumption, drug use (sewers), gunshots, GPS

Big Data and Financial Services for Poor

Credit scoring using mobile, social, or other data

- ZestFinance: combines credit history with social and other publicly available data
- Cignifi (serving Mexico, Ghana, Brazil): mobile data, utility bills, auto insurance, remittance data; mobile data is behavioral, but not based on location data or identity of calls (prioritizes privacy)
- Mines.io (serving Africa): mobile, public data
- Lenndo (Philippines, Columbia, other emerging markets): community lending + social data

Credit scoring for small businesses

- Kabbage: Unique credit scoring provides answer in 7 minutes

Conclusions

ML/AI brings great benefits in terms of:

- Efficient resource utilization
- Savings for governments
- Fair treatment of citizens
- Measurement of policy effectiveness
- Personlization
- Speed

Risks for society

- Benefits of technology unequally distributed
- Large transition costs for big segments of workforce
- Cost of living may not fall as fast as wages, particularly if high concentration/market power
- Public policy should prepare now