



Doing Some Good with Machine Learning

LESTER MACKEY

MICROSOFT RESEARCH

Act I: Grad School

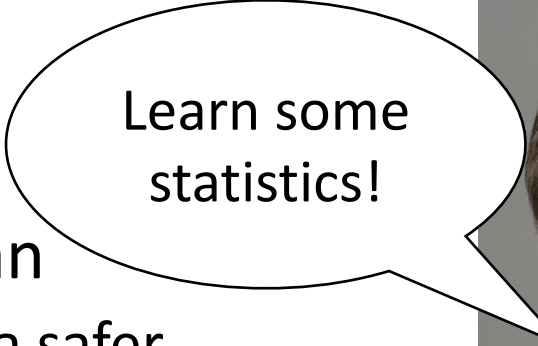
- **Fall 2007:** Started CS grad school at UC Berkeley

Quixotic though it may sound, I hope to use computer science **and statistics** to change the world for the better. If you have thoughts on how to do this, feel free to contact me.

- **2007-2008:** Learned some statistics

- **Dec. 2008:** Fortuitous email from Mike Jordan

Any of you interested in trying to make the world a safer place with machine learning?




Learn some statistics!



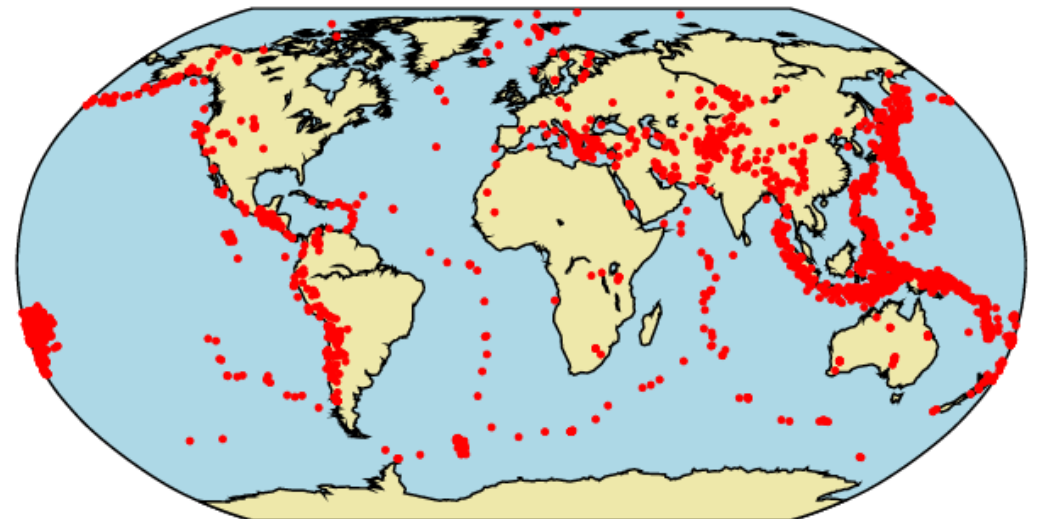
Combating Nuclear Proliferation

- **United Nations' 1996 Comprehensive Test Ban Treaty (CTBT)**
 - Bans all nuclear weapons testing on Earth
 - Enforced by monitoring the globe's seismic activity to identify violations
- **Problem:** Automated seismic bulletin yields *many* false positives, hand-corrected by data analysts
- **Opportunity:** Over 10 years of daily hand-corrected seismic bulletins

SEL3 Event Locations



LEB Event Locations



Combating Nuclear Proliferation

- **United Nations' 1996 Comprehensive Test Ban Treaty (CTBT)**
 - Bans all nuclear weapons testing on Earth
 - Enforced by monitoring the globe's seismic activity to identify violations
- **Problem:** Automated seismic bulletin yields *many* false positives, hand-corrected by data analysts
- **Opportunity:** Over 10 years of daily hand-corrected seismic bulletins
- **Idea:** Learn to correct automated bulletin and prioritize seismic events for analyst review (Mackey, Kleiner, & Jordan, AGU, 2009)
 - **Event features:** time, location, error ellipse, waveform profile, ...
 - 81% accuracy, 89% AUC for flagging invalid seismic events
 - Deployed at the CTBT International Data Centre



Ariel Kleiner

Act I: Grad School


- **Fall 2007:** Started CS grad school at UC Berkeley
 - Quixotic though it may sound, I hope to use computer science and statistics to change the world for the better. If you have thoughts on how to do this, feel free to contact me.
- **2007-2008:** Learned some statistics
- **Dec. 2008:** Fortuitous email from Mike Jordan
- **2008-2012:** Recommender systems, algorithms, theorems
- **May 2012:** Graduated!
- **July 2012:** Lilly finished her bar exam
 - Decided to work together on the [ALS Prediction Prize](#)

I'm bored!



What is ALS?

- **Amyotrophic lateral sclerosis or Lou Gehrig's Disease**
 - A neurodegenerative disease that targets motor neurons
 - Leads to muscle atrophy, paralysis, and ultimately death



What is ALS?

- **Amyotrophic lateral sclerosis or Lou Gehrig's Disease**
 - A neurodegenerative disease that targets motor neurons
 - Leads to muscle atrophy, paralysis, and ultimately death
 - 100% fatal, typically within 3-5 years, but not always

**Fast
progressor**



Lou Gehrig
(died within 2 years of diagnosis)



**Slow
progressor**

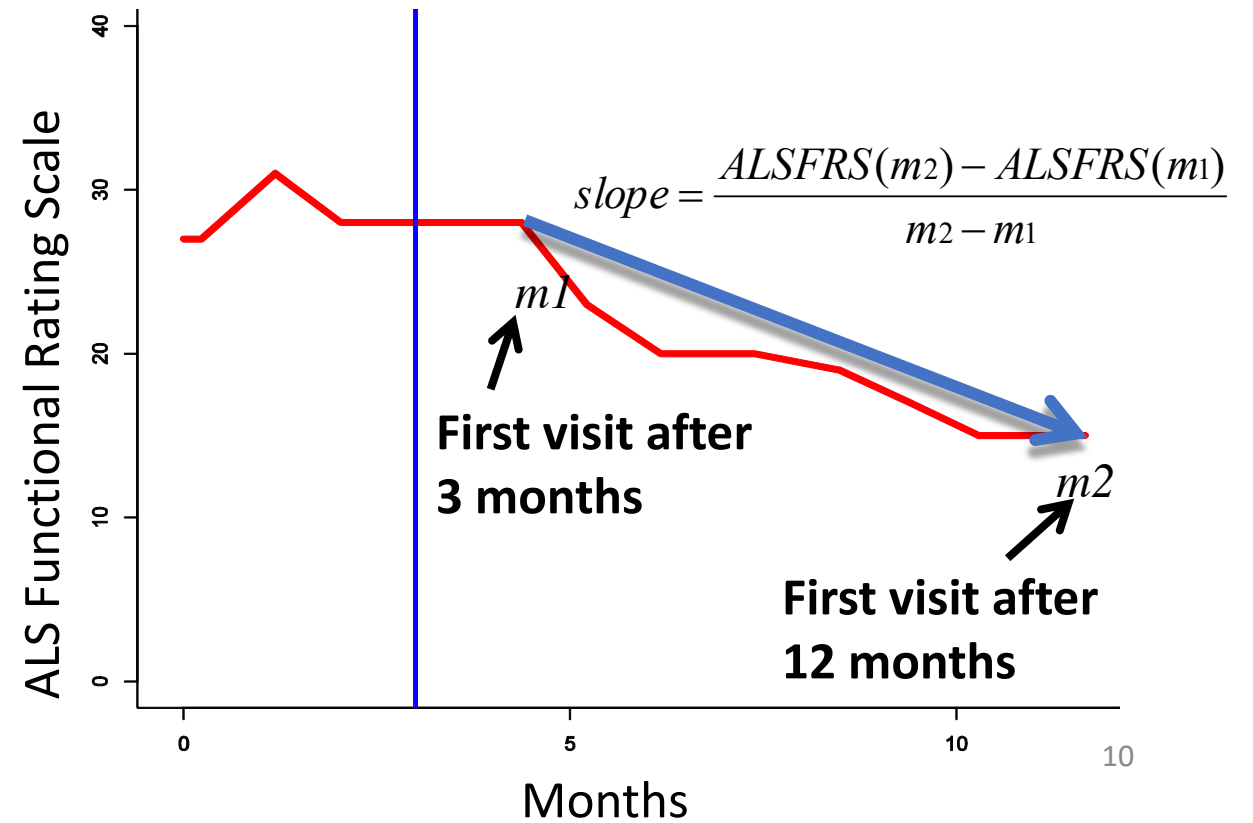
Stephen Hawking
(lived with the disease for over 50 years)

ALS Prediction Prize

- Data science competition run by PRIZE4LIFE
 - Non-profit focused on accelerating ALS treatment through contests
- **Goal:** Predict rate of disease progression in ALS patients
- **Helping clinicians**
 - More accurate prognosis
 - Identifying predictive patient characteristics: Which lab tests worthwhile?
- **Stratifying clinical trial patients**
 - Recent 1000 patient trial cost over \$100 million
 - Less variability \Rightarrow fewer subjects \Rightarrow less expensive, more interpretable clinical trials

ALS Prediction Prize

- **Observe:** 3 months of longitudinal clinical trial measurements
 - Functional measures, vital signs, lab tests, demographics, medical / family history
 - From 11000-patient **PRO-ACT** database (<https://nctu.partners.org/ProACT>)
- **Goal:** Predict rate of disease progression over next 9 months



Our Approach

- **Challenges and solutions**

- **Irregular time series:** roughly monthly measurements, unequally spaced

- Featurize time series as **fixed-dimensional vectors of summary statistics**

- **Outliers + substantial missingness**

- **Anomaly detection + median imputation**

- **Most features** believed to be **irrelevant**

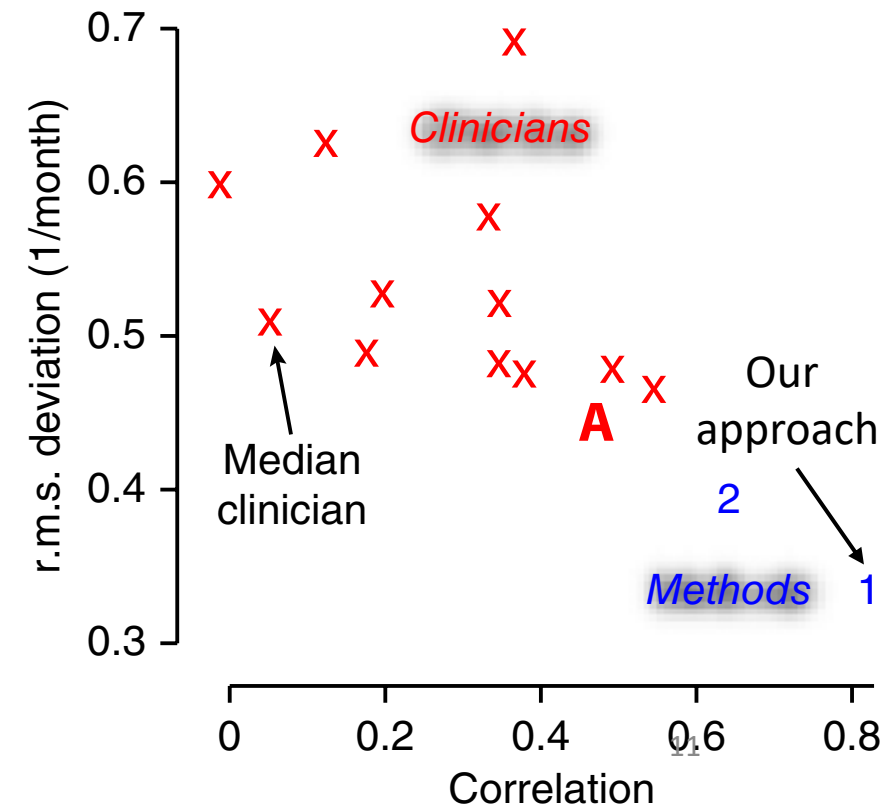
- **Bayesian Additive Regression Trees**

(Chipman, George, & McCulloch, Ann. Appl. Stat., 2010)

- Placed first & outperformed 12 clinicians

- Prize4Life estimates predictions would reduce drug trial sizes by 20%

Küffner et al. Crowdsourced analysis of clinical trial data to predict amyotrophic lateral sclerosis progression, Nature Biotechnology, 2015, doi.org/10.1038/nbt.3051



Act II: Postdoc

- **Fall 2012:** Started postdoc at Stanford
 - Series of impactful conversations with a social good theme

How exactly are
you changing the
world for the
better?



Jacob Steinhardt

DataKind

- Brings together volunteer data scientists and social organizations in need of data aid
- DataDives: weekend events to tackle social data problems
 - **Example:** Collecting food pricing and consumption data for the World Bank to manage poverty and thwart food crises
- DataCorps: longer-term volunteer commitments

Let's use data to change the world. Here's how.

The Eric & Wendy Schmidt

Data Science for Social Good

Summer Fellowship






We're training data scientists to tackle problems that really matter.

- Three-month summer fellowship for students and post-docs to learn data science skills
- Partner with non-profits and governments to tackle real problems in education, health, energy, transportation, ...
- Now at Carnegie Mellon University

Act III: Professor

- **9/1/13:** Joined Stanford statistics faculty
- **9/4/13:** Inquired about data for social good initiatives on Stanford's campus
 - Response: "This is a great idea! I am not aware of anything similar here."
- **9/27/13:** Emailed Stanford machine learning and statistics communities with a proposal for a **statistics for social good** working group

Statistics for Social Good: The Proposal

- I'd like to organize a **reading group** this quarter on 'statistics for social good,' with a particular focus on economic inequities like poverty, hunger, access to education, and human trafficking
(Later expanded to environment, transportation, energy, health care, ...)
 Wrong: **working group**
- Much of our effort will be **exploratory** as we ferret out specific social problems ..., **track down data sources and data partners**, and **determine which statistical and computational tools** will be needed to tackle these tasks.
 Wrong: **problem partners**
- This **will require work each week** from everyone involved
 **That was true**
- P.S. This idea was inspired in no small way by U. Chicago's Data Science for Social Good summer program.

Statistics for Social Good


- **Divide-and-conquer approach:** every member tasked with
 - **Contacting** potential **problem partners** with expertise in a social domain and knowledge of how data analysis could help
 - **Distilling** pressing inferential questions and relevant data sources
 - **Documenting** findings (<http://stats-for-change.github.io/>)
 - **Carving out projects** surrounding the most promising leads



Suzanne Tamang



Shea Shelton, Co-chair



Toren Fronsdal, Co-chair



Jack Collison, Co-Chair



Lester Mackey




Paul Switzer




AJ Alvero



Jialu Liu Streeter



Abubakar Abid



Hoda Magid




Alex Kim




Shuo Xie



Lucas Janson



Jason Hirshman



Jason Huang




David Flatow



Jessica Hwang



Sven Schmit



Wen-wai Yim



Kris Sankaran



Jessica Hwang



Kris Sankaran

Linking service profiles with client progress

- Free financial coaching centers for low-income clients
- **Target outcomes:** Improved credit scores, debt levels, income, savings
- **Services:** Matched savings programs, resume improvement, balancing household budget, ...
- **Detailed records:** Baseline assessments and follow-up measurements
- **Key questions:** For a given target outcome, which services should SparkPoint recommend and in what order?

Identifying patients in need of health-care interventions

- Stanford Clinical Excellence Research Center develops new healthcare interventions for at-risk populations
- **First step:** Predict individuals who will (without intervention) incur the greatest future healthcare costs
- Detailed healthcare records of 2M Denmark residents, 2004-2011
- 30% improvement in healthcare cost capture over standard forecasts

Tamang et al. Predicting patient 'cost blooms' in Denmark: a longitudinal population-based study, *BMJ Open*, 2017.
doi.org/10.1136/bmjopen-2016-011580



Suzanne Tamang



Jean-Raymond Betterton



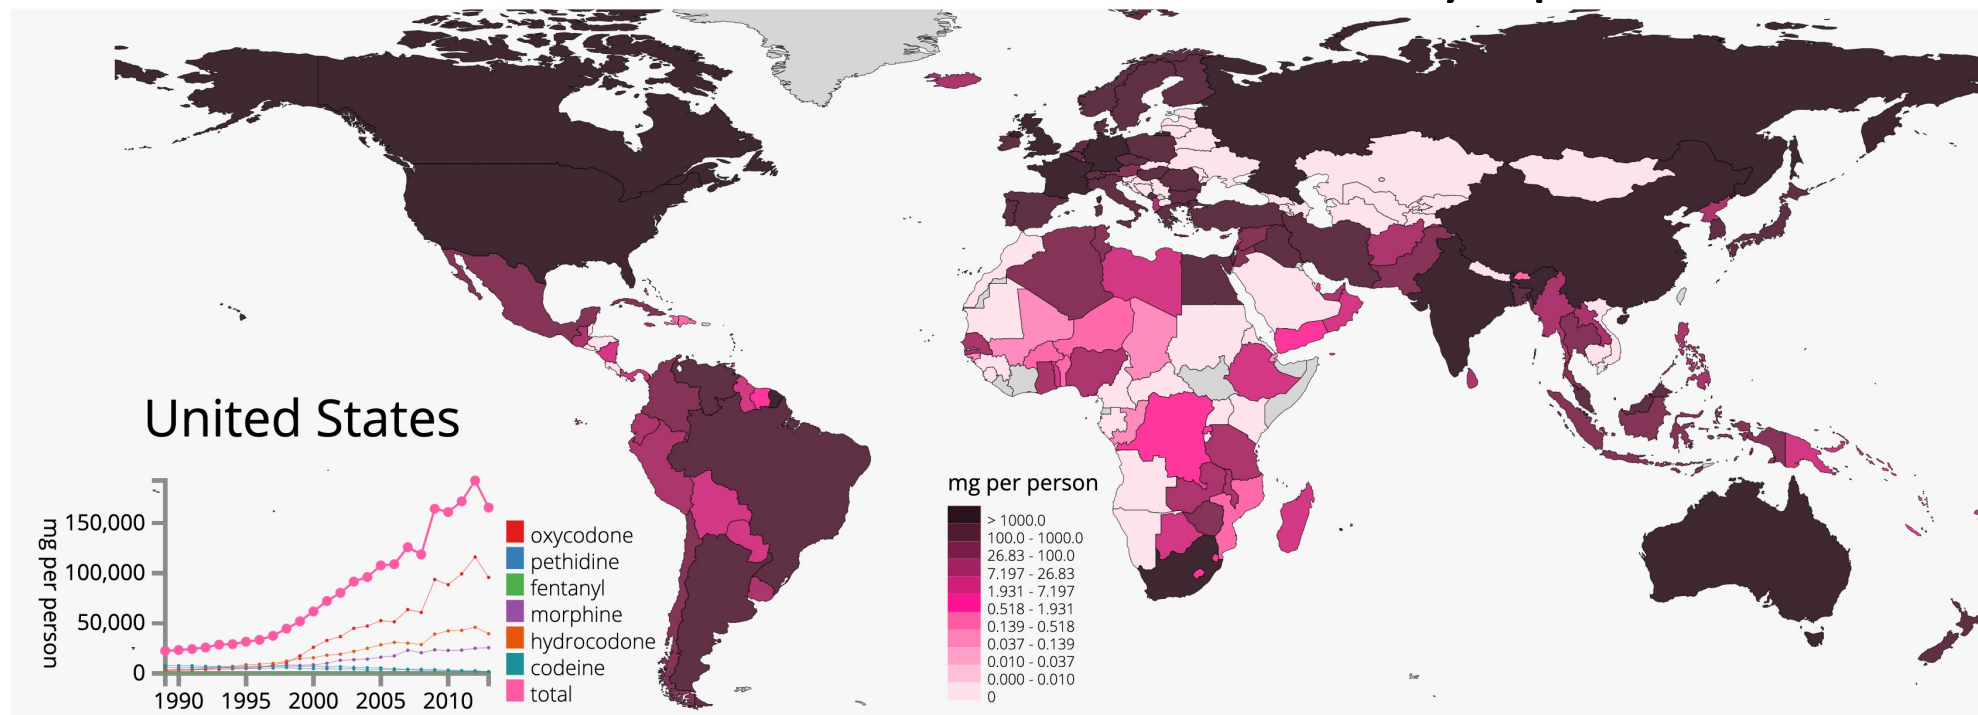
Lucas Janson



Global Oncology Initiative

Investigating inequities in palliative care

- Volunteer organization dedicated to improving cancer care worldwide
- Asked Stats for Good to find, consolidate, and visualize opiate consumption data in an [Opioid Atlas](#) for palliative care researchers
- Showcased at 2016 Global Cancer Research Symposium



Kris Sankaran



Suzanne Tamang



Shuo Xie



Kris Sankaran

Helping nonprofits understand client feedback

- Nonprofit review platform: Helps donors, volunteers, and clients learn about nonprofits matching their interests
- Stats for Good studied the relationship between reviewer background and review content
- Debunked prevailing “courtesy bias” hypothesis that those depending on nonprofit services for basic needs avoid negative feedback

Xie and Sankaran. [The Surprising Truth Behind Beneficiary Feedback](#),
Stanford Social Innovation Review, 2014.

Kindred Spirits

- DS 4 Social Good (<https://groups.google.com/forum/#!forum/ds4-social-good>)
 - Google groups forum started by Patrick Meier, open to the public
- Stanford Data Lab (<https://datalab.stanford.edu/>)
 - Stanford courses designed to teach students data science skills in the context of tackling real social challenges, run by Bill Behrman
- Data-Pop Alliance (<https://datapopalliance.org/>)
 - Collaboration of the Harvard Humanitarian Initiative, MIT Connection Science, and Overseas Development Institute, cofounded by Emmanuel Letouzé
- Berkeley D-Lab (<https://dlab.berkeley.edu/>)
 - Intelligent research design for data intensive social science
 - Dav Clark created a collaborative website for anyone interested in statistics or data science for social good: <http://stats-for-change.github.io/>
 - We used it to document our findings

Statistics for Social Change

Effecting social change through effective data analysis

This is a collaborative portal intended to document ongoing social data analysis efforts, summarize relevant resources and knowledge, and connect those with technical skills and a passion for social change to organizations and individuals with pressing needs.

- [Potential Partners](#)
Relevant organizations, with information about projects, conversations, and contacts
- [Data Sources](#)
Public data sources related to social issues
- [Discussion Group](#)
Feel free to introduce yourself here!

If you are an organization, research group, or individual engaged in data analysis for social good, we want to hear from you!

Data Sources

Public data sources related to social issues

- **Development**

- [UNDP Human Development Index](#)

- Huge database of International Human Development Indicator(s), Gender Inequality Index and Multidimensional Poverty Index, insearchable by country or by indicator. Also contains great graphs and maps comparing trends over time and cross-national data.

- [AidData](#)

- Searchable database of nearly one million past and present aid activities around the world.

- [Africa Open Data](#)

- The Open Africa Platform initiative aims to be largest repository of data on the African Continent.

- [World Bank](#)

- Open datasets that are related to the financing and delivery of public goods, works, and services, including procurement and contracting. Also has World Development Indicators.

- **Education**

- [PISA](#)

Potential Partners

Relevant organizations, groups, and individuals

Search:

Name, description, and external link ▼	Contact	Tags
<p>United Nations Refugee Agency (UNHCR)</p> <p>The UNHCR is mandated to lead and co-ordinate international action to protect refugees and resolve refugee problems worldwide.</p> <p>http://www.unhcr.org</p>	<p>Kimberly Roberson, Chief of Field Information and Coordination Section</p>	<p>refugees, migration</p>
<p>United Nations Global Pulse</p> <p>Partner with experts from UN</p>	<p>Rene Clausen Nielsen, http://www.unglobalpulse.org/contact</p>	<p>United Nations,</p>

United Nations Refugee Agency (UNHCR)

The UNHCR is mandated to lead and co-ordinate international action to protect refugees and resolve refugee problems worldwide.

Details

- Contact: [Kimberly Roberson](#), Chief of Field Information and Coordination Section
- Tags: refugees, migration
- External: <http://www.unhcr.org>

Conversations

- [Meeting on November 20, 2013](#)

Act IV: Researcher

- **Spring 2016:** Noah was born
- **Fall 2016:** Moved to East Coast / MSR
- **Winter 2016:** Met Ernest Fraenkel and Judah Cohen



You should work
on climate
forecasting!





U.S. Bureau of Reclamation

- “The mission of the [USBR] is to manage, develop, and protect water and related resources in an environmentally and economically sound manner in the interest of the American public.”
- **Manages water in 17 western states**
 - Provides 1 out of 5 Western farmers with irrigation water for 10 million farmland acres
 - Generates enough electricity to power 3.5M U.S. homes
- **“During the past eight years, every state in the Western United States has experienced drought** that has affected the economy both locally and nationally through impacts to agricultural production, water supply, and energy.”



Judah Cohen



- Climatologist, director of seasonal forecasting at Atmospheric and Environmental Research
- **Concern:** Community not making the best use of historical data in climate forecasting
 - Landscape dominated by **dynamical models**, purely physics-based models of the atmosphere and oceans
 - Accuracy limited by **chaos** of differential equations: errors in inputs rapidly amplified
- **Concern: Subseasonal forecasts** especially poor

Weather forecasts


predictability comes from initial atmospheric conditions

Sub-seasonal forecasts

predictability comes from monitoring the Madden-Julian Oscillation, land surface data, and other sources

Seasonal forecasts

predictability comes primarily from sea-surface temperature data accuracy dependent on ENSO state





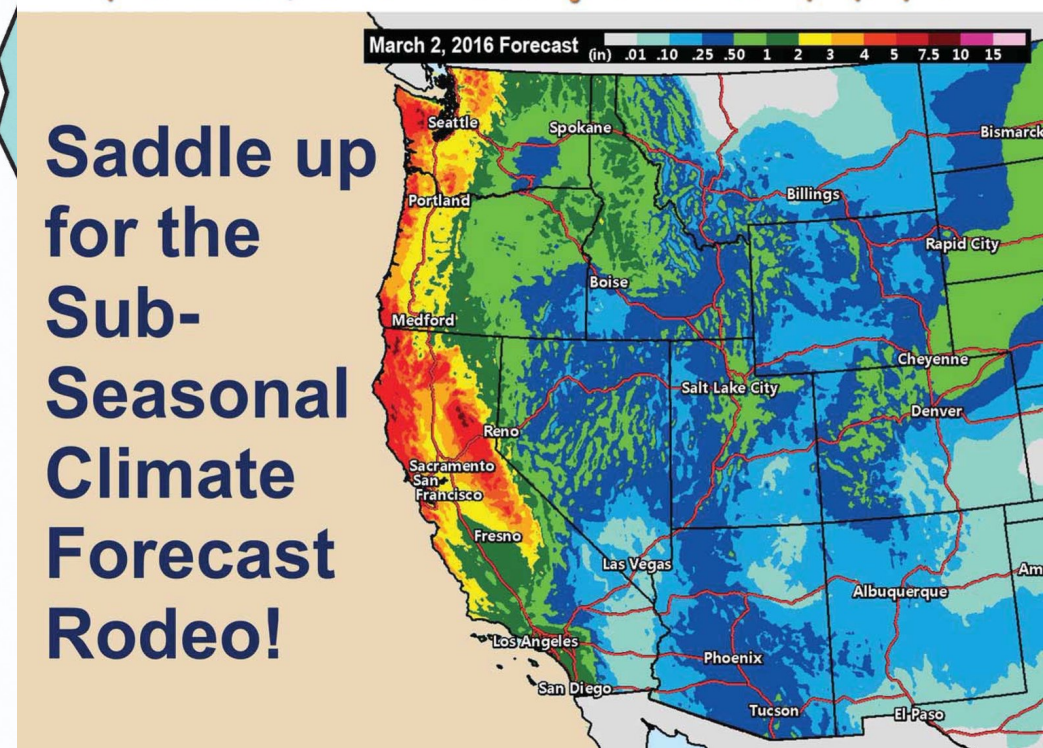
U.S. Bureau of Reclamation

- “The mission of the [USBR] is to manage, develop, and protect water and related resources in an environmentally and economically sound manner in the interest of the American public.”
- **Manages water in 17 western states**
 - Provides 1 out of 5 Western farmers with irrigation water for 10 million farmland acres
 - Generates enough electricity to power 3.5M U.S. homes
- “**During the past eight years, every state in the Western United States has experienced drought** that has affected the economy both locally and nationally through impacts to agricultural production, water supply, and energy.”

Credit: David Raff, USBR



\$800,000 in prize \$\$\$!

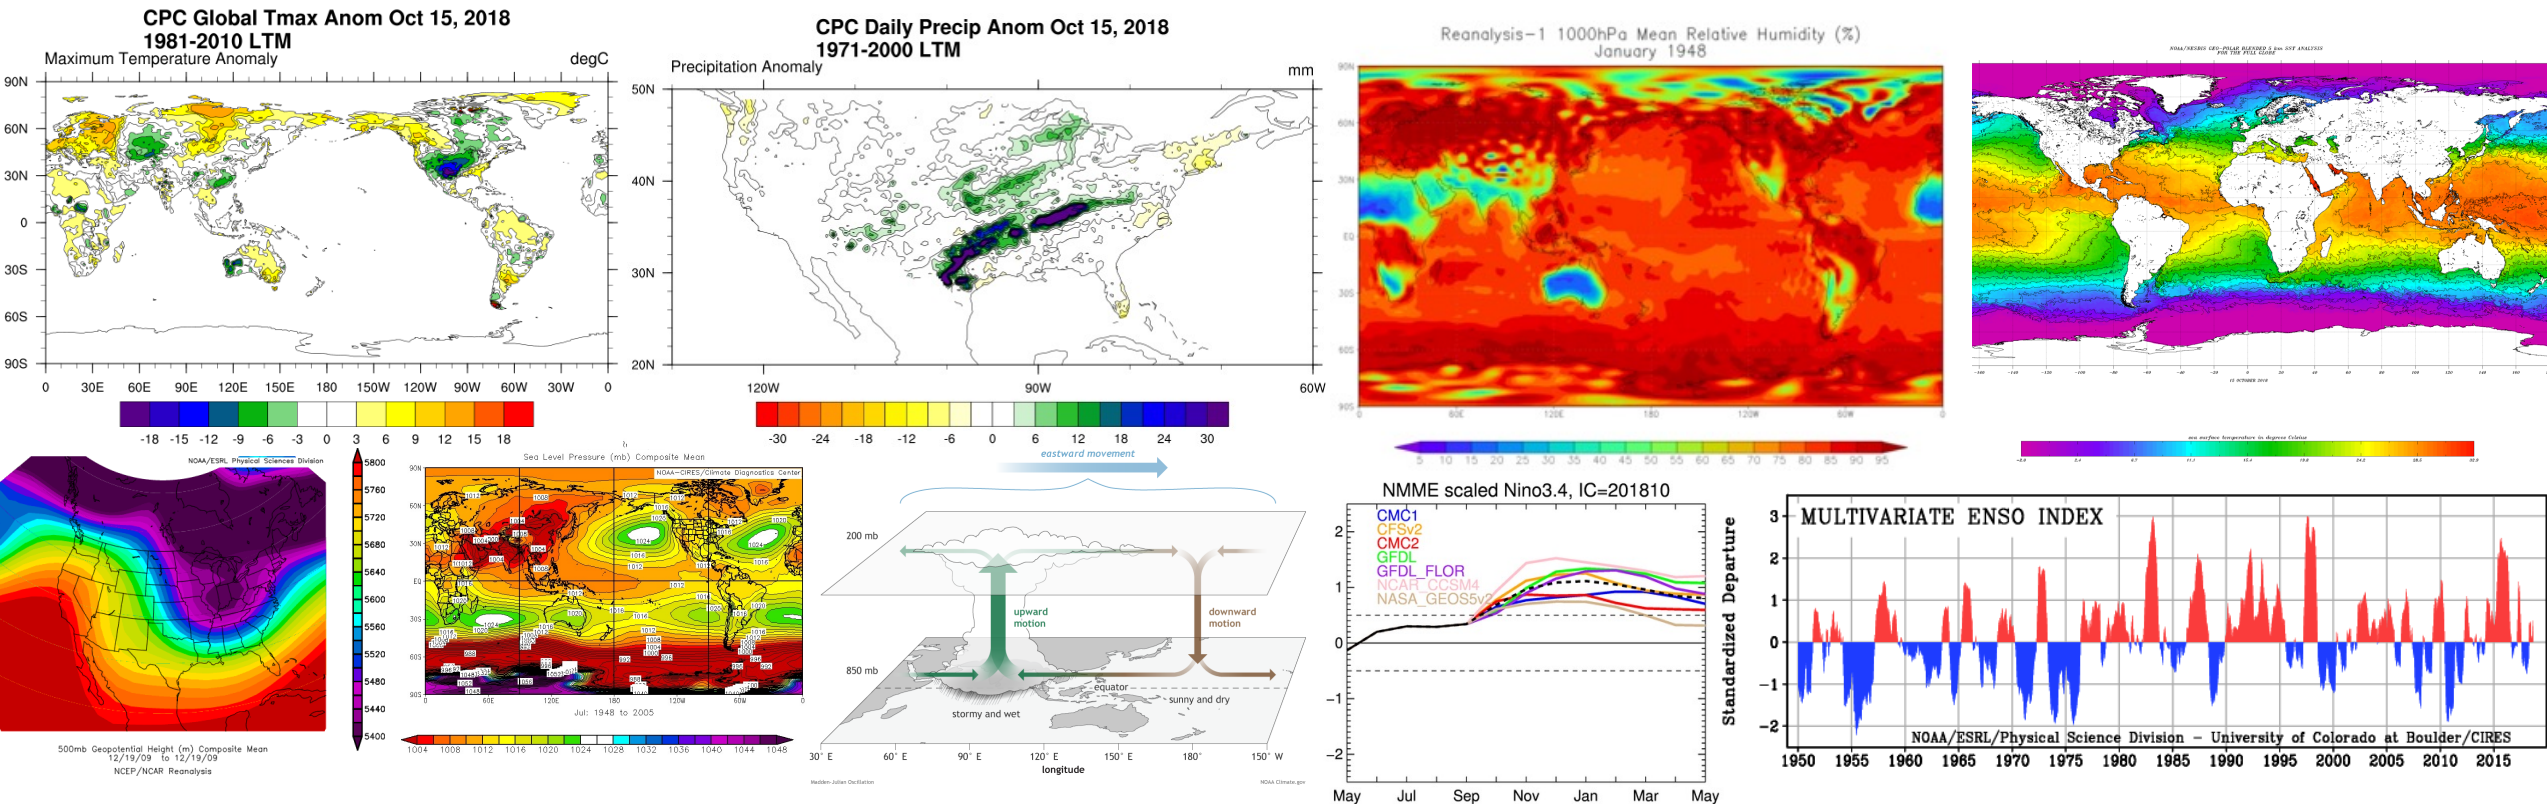


Saddle up for the Sub-Seasonal Climate Forecast Rodeo!

usbr.gov/research/challenges


Our SubseasonalClimateUSA Dataset

- To train and evaluate our models, we constructed a **SubseasonalClimateUSA dataset** from diverse data sources
- Updated daily + accessed via [subseasonal_data](#) Python package




Week 3-4 Forecast submitted 20180109, verifying 20180205

Our skill: 0.8383




Observed




Week 3-4 Forecast submitted 20170905, verifying 20171002


Our skill: -0.0077




Observed




CFSv2 skill: -0.3961





Damped Persistence




CFSv2 skill: -0.1971



Damped Persistence



Skill =
Geographic cosine similarity



Our Forecasting Models

- **Skill** = Geographic cosine similarity between observations and forecasts
- **MultiLLR: Multitask Locally Linear Regression**
 - Incorporates lagged measurements from all data sources
 - Prunes irrelevant regressors using skill-based multitask feature selection
- **AutoKNN: Multitask Nearest Neighbor Autoregression**
 - Identifies dates most similar to target using skill-based similarity measure
 - Regresses onto observed temperature or precipitation of similar dates and fixed lags
- **Ensemble**: Averages the normalized predicted anomalies

$$\hat{\mathbf{a}}_{\text{ensemble}} \triangleq \frac{1}{2} \frac{\hat{\mathbf{a}}_{\text{multillr}}}{\|\hat{\mathbf{a}}_{\text{multillr}}\|_2} + \frac{1}{2} \frac{\hat{\mathbf{a}}_{\text{autoknn}}}{\|\hat{\mathbf{a}}_{\text{autoknn}}\|_2}.$$

- **Proposition** *If the average of the individual model skills is positive, then the ensemble skill is strictly greater than the average of the individual skills.*

Historical Forecast Skill (2011-2018)


task	multillr	autoknn	ensemble	cfsv2	ens-cfsv2
temperature, weeks 3-4	0.2230	0.3111	0.3073	0.2557	0.3508
temperature, weeks 5-6	0.2204	0.2810	0.2962	0.2142	0.3279
precipitation, weeks 3-4	0.1573	0.1513	0.1893	0.0860	0.1964
precipitation, weeks 5-6	0.1312	0.1403	0.1703	0.0691	0.1755

Temperature: 37-53% improvement over US operational forecasting system (CFSv2)

Precipitation: 128-154% improvement



Jessica Hwang



Paulo Orenstein



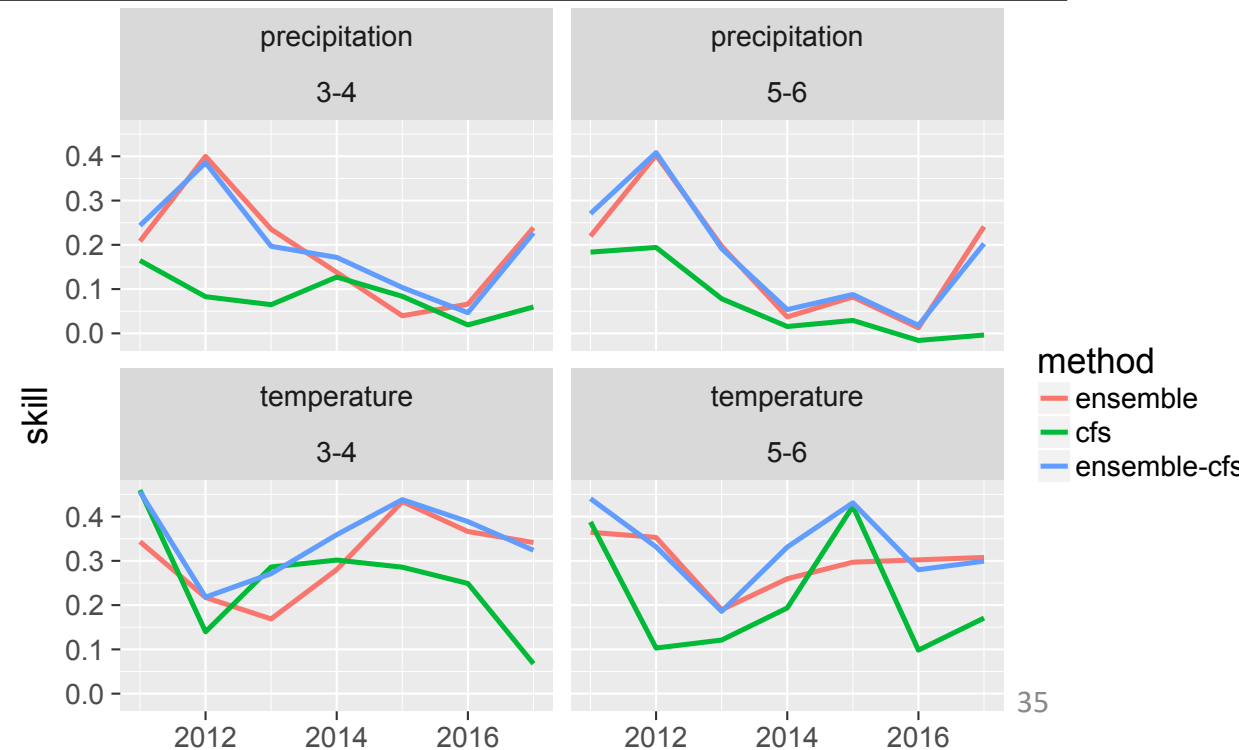
Judah Cohen



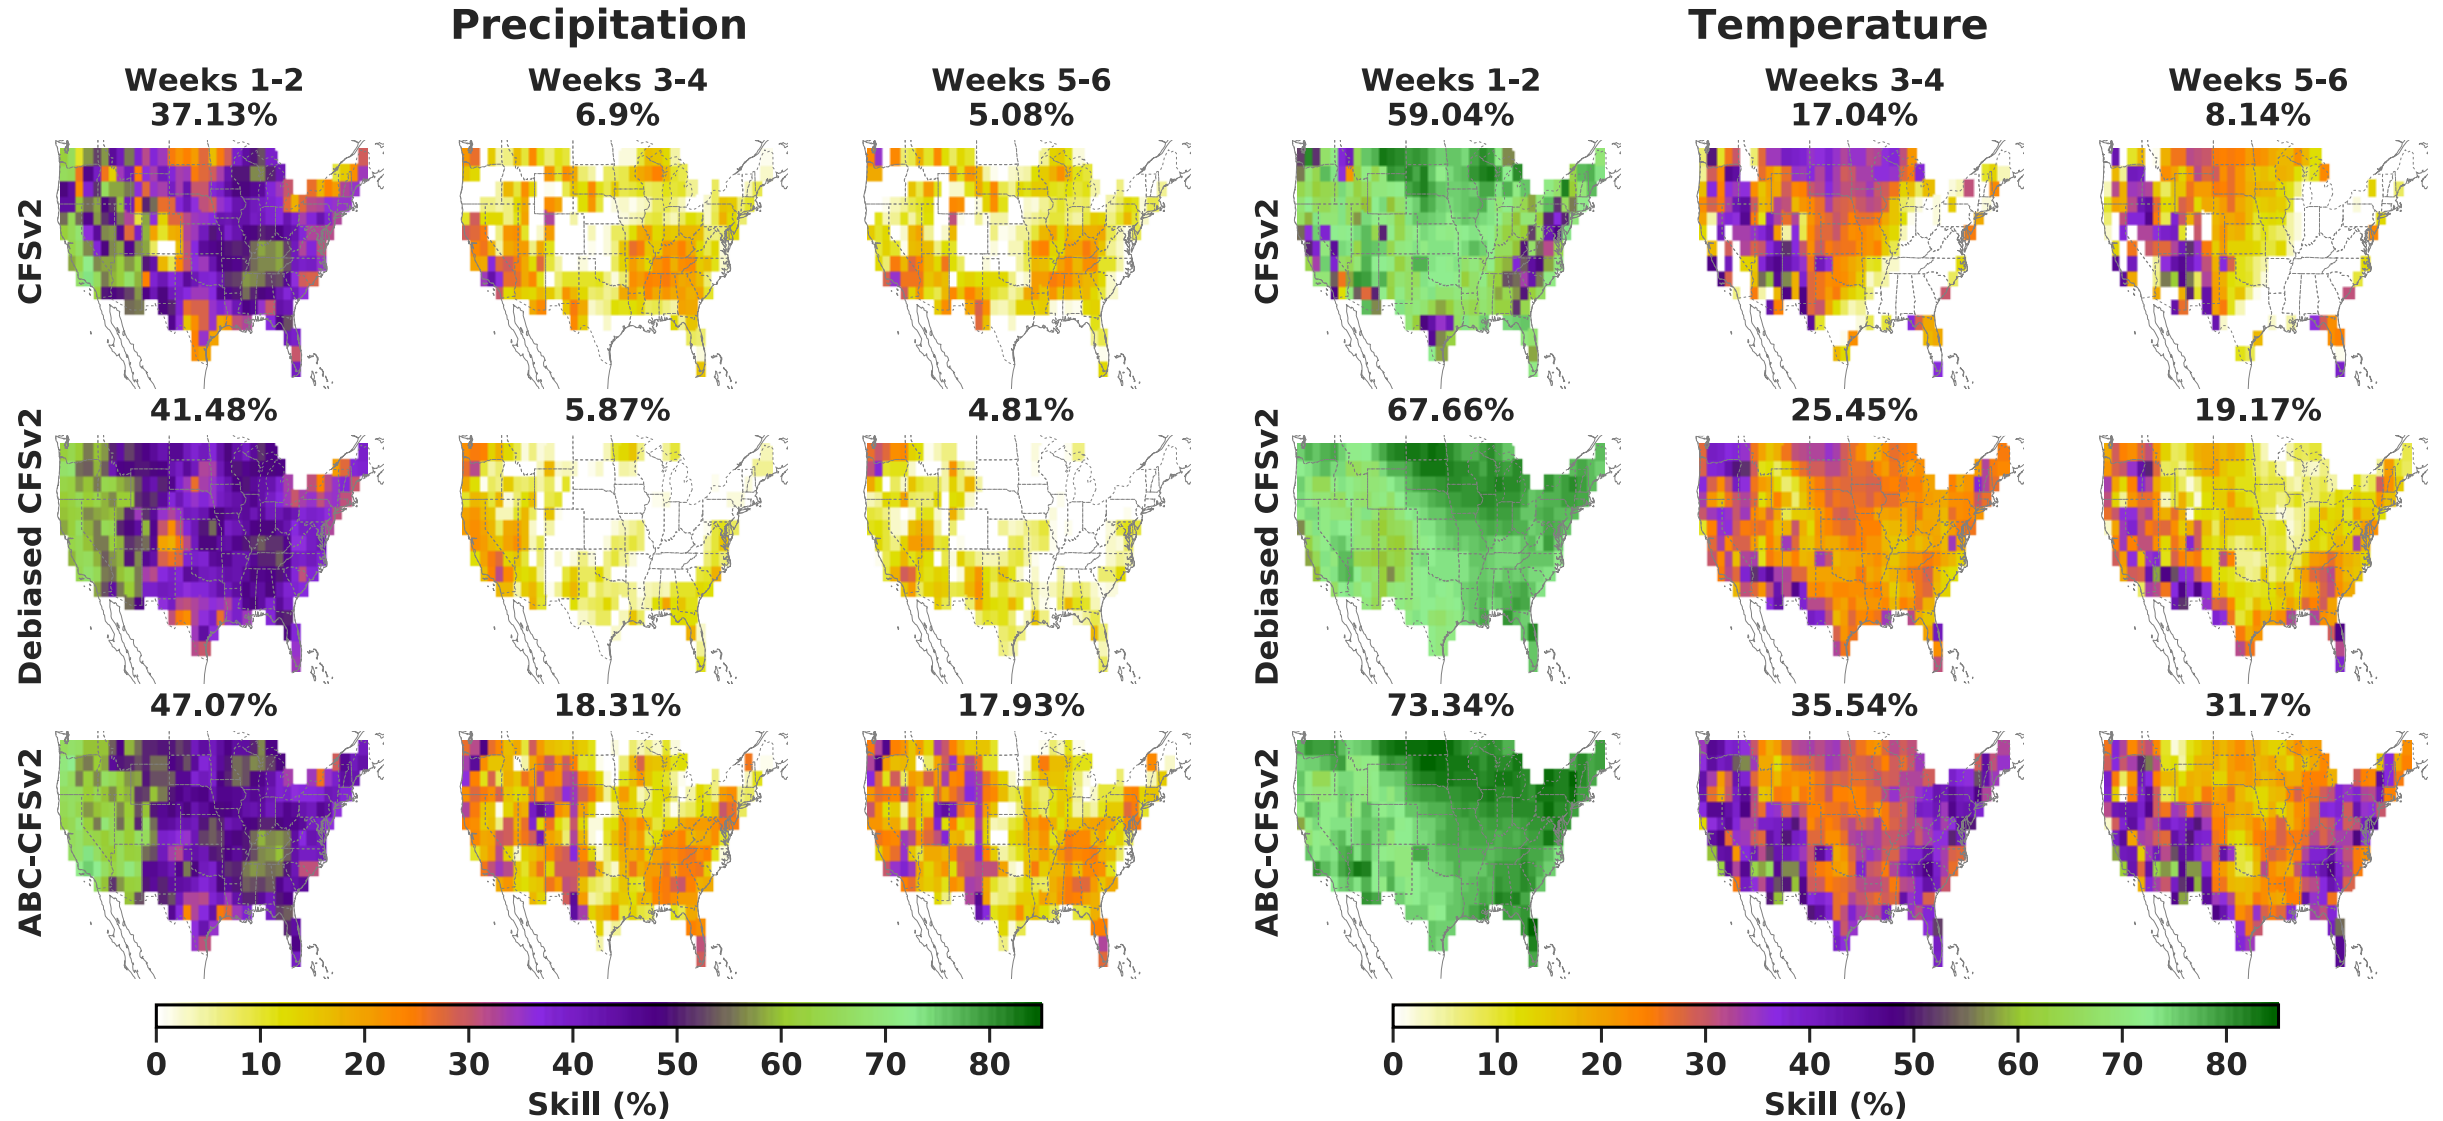
Karl Pfeiffer



Ernest Fraenkel



Adaptive Bias Correction (ABC): Hybrid Physics + Learning Model



- Doubles or triples the forecasting skill of US operational model (CFSv2) [[Nature Communications, 2023](#)]
- Outperforms state-of-the-art machine learning and deep learning methods [[NeurIPS, 2023](#)]

Act V: Pandemic

- **Mar. 2020:** Massachusetts issues **COVID-19** stay-at-home advisory
- **Apr. 2020:** Email from Ameet Talwalkar about Ryan Tibshirani's COVID effort



From Flu to COVID



Ryan Tibshirani

- Co-lead Carnegie Mellon's flu forecasting group, [Delphi](#)
- **2019:** Named Influenza Forecasting Center of Excellence by the U.S. Centers for Disease Control and Prevention (CDC)



Roni Rosenfeld

- **2020:** CDC: What about COVID-19 forecasting?


Delphi COVID-19 Response Team



Initial Mission

- Collect, analyze, and release US-wide indicators of COVID-like symptoms
- Build a website (covidcast.cmu.edu) for public health officials to monitor

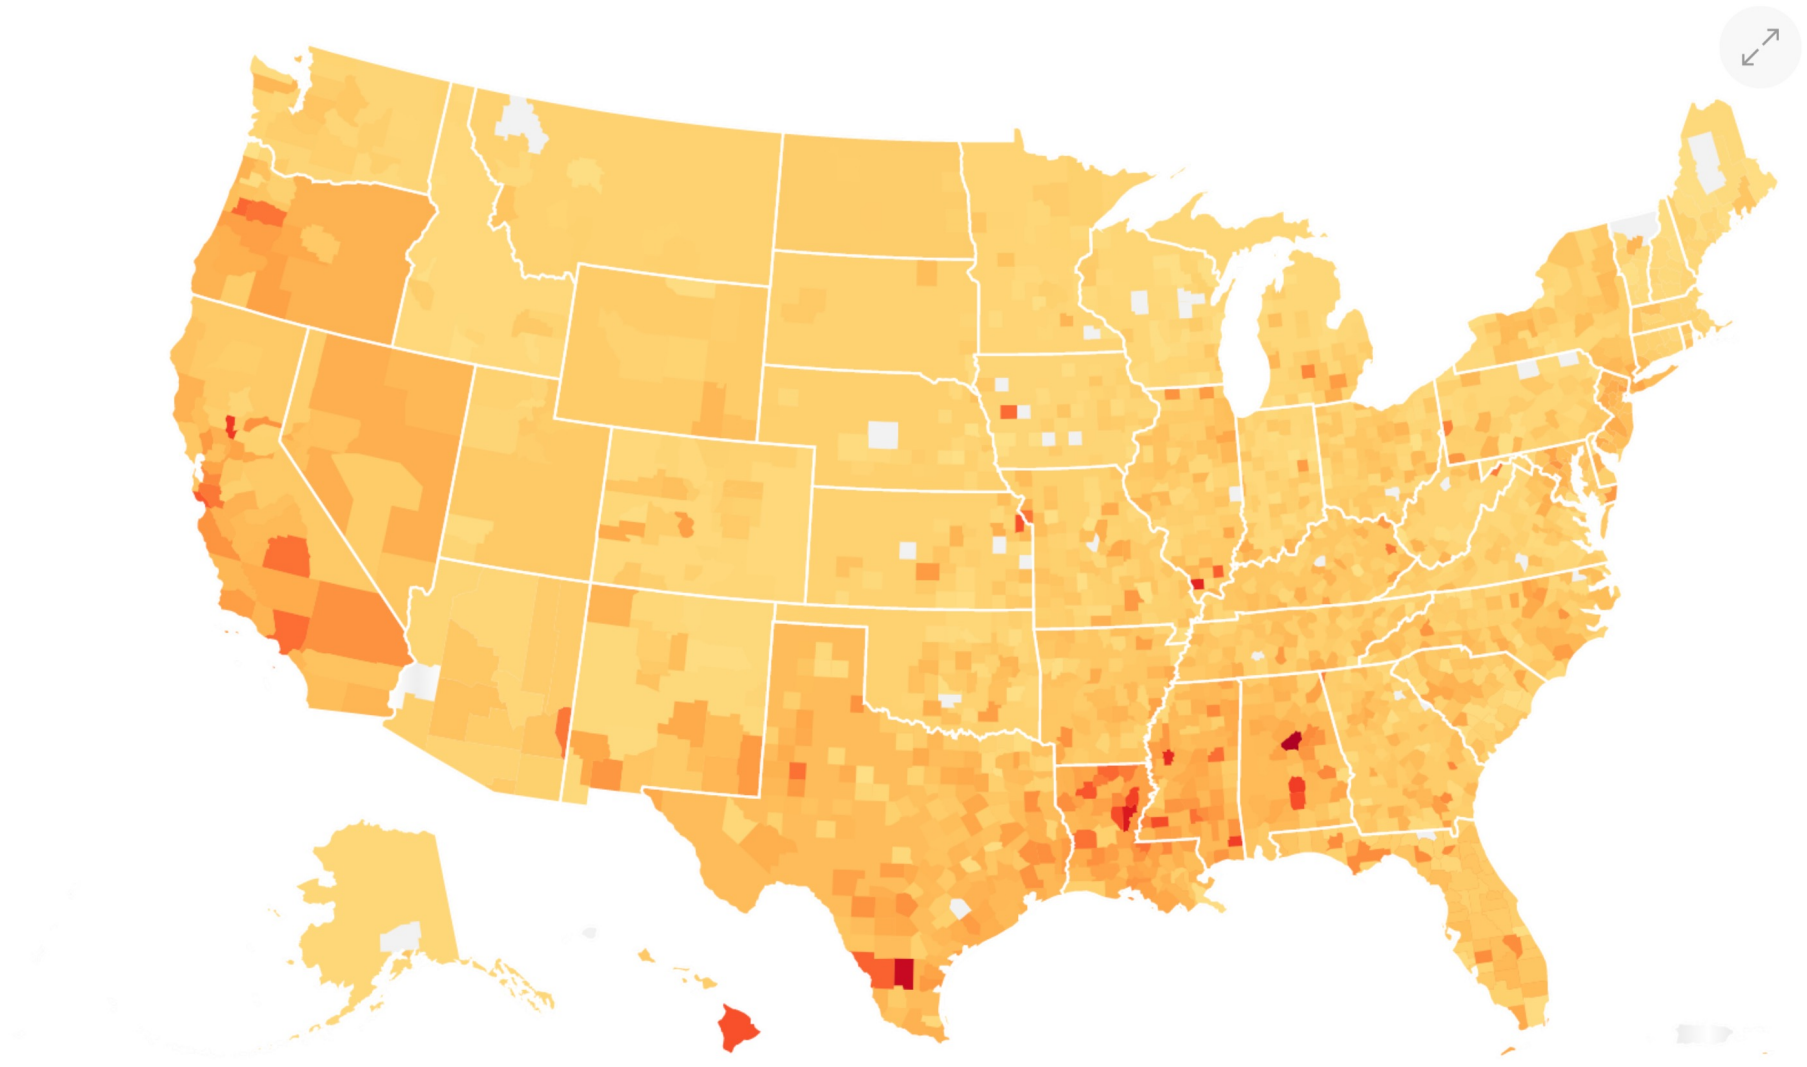
COVID-RELATED DOCTOR VISITS MAP

 Click on a state to explore further

Beehive Grid



Choropleth Map




COVID-LIKE SYMPTOMS IN COMMUNITY MAP

 Click on a state to explore further

Beehive Grid




Choropleth Map



PEOPLE WEARING MASKS MAP


Click on a state to explore further

Beehive Grid Choropleth Map



PEOPLE WEARING MASKS CHART

People Wearing Masks in United States per 100 people



Region
Middlesex County, MA

Wed, Jul 20 2022

Indicator
Symptom Searches (Smell and Taste) on Google


Middlesex County, MA population

1,612,000

people

SYMPTOM SEARCHES (SMELL AND TASTE) ON GOOGLE CHART

Symptom Searches (Smell and Taste) on Google in Middlesex County, MA
scaled search volume



Delphi Group, delphi.cmu.edu/covidcast

All Dates Raw Data


● Middlesex County, MA
0.08

● Average of 4 Neighboring Counties
0.06



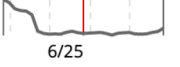






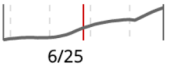

● Massachusetts
0.08

● United States
0.09

COVID-19 INDICATORS TABLE

 Click on an indicator name to explore further



	INDICATOR	VALUE	RELATIVE CHANGE TO PREVIOUS WEEK	HISTORICAL TREND	
				6/11	7/09
Public Behavior	People Wearing Masks (Delphi US COVID-19 Trends and Impact Survey)	29.00 _{/100}	-2.48%		>
	Vaccine Acceptance (Delphi US COVID-19 Trends and Impact Survey)	84.59 _{/100}	+0.01%		>
	Symptom Searches (Smell and Taste) on Google (Google Symptoms Search Trends)	0.09	-0.02%		>
	Symptom Searches (Common Cold) on Google (Google Symptoms Search Trends)	2.15	+0.57%		>
Early Indicators	COVID-Like Symptoms (Delphi US COVID-19 Trends and Impact Survey)	3.26 _{/100}	+7.31%		>
	COVID-Like Symptoms in Community (Delphi US COVID-19 Trends and Impact Survey)	26.81 _{/100}	+2.15%		>
	COVID-Related Doctor Visits (Doctor Visits From Claims)	2.32 _{/100}	+5.36%		>
Cases and Testing	COVID Antigen Test Positivity (Quidel Inc.)	21.85 _{/100}	+6.47%		>
	COVID Cases (Johns Hopkins University)	30.7 _{PER 100K}	-1.97%		>
Late Indicators	COVID Hospital Admissions (U.S. Department of Health & Human Services)	1.5 _{PER 100K}	+10.08%		>
	COVID Deaths (Johns Hopkins University)	0.1 _{PER 100K}	+11.83%		>

COVIDcast

- All data available through a public API <https://cmu-delphi.github.io/delphi-epidata/api/covidcast.html>


delphi-epidata

An open API for Epidemiological Data, from the Delphi research group.

- Used as indicators for forecasting deaths and cases for the CDC

Reinhart et al. An open repository of real-time COVID-19 indicators, Proceedings of the National Academy of Sciences, 2021.

<https://doi.org/10.1073/pnas.2111452118>



The End?

Doing Some Good: Volunteer

DataKind

<https://www.datakind.org/>



STATISTICS
WITHOUT BORDERS

Statisticians in compassionate service.

<https://swb.wildapricot.org/>



What is **Solve for Good**?

Solve for Good is an online platform for social good organization to post projects they need help with, for volunteers to help scope those projects into well-defined problems, and to help solve those problems.

<https://www.solveforgood.org>

Doing Some Good: Data Science Contests



INNOCENTIVE[®]

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

DRIVEN**DATA**

<https://www.drivendata.org/>

DREAM
CHALLENGES

The **DREAM Challenges** are
**crowdsourcing challenges examining
questions in biology and medicine.**

<http://dreamchallenges.org/>

kaggle

<http://kaggle.com/>

Doing Some Good: Summer Programs

Data Science for Social Good Summer Fellowships

- Carnegie Mellon University <https://www.dssgfellowship.org/>
- Stanford University <https://dssg.stanford.edu/>
- University of Washington <https://escience.washington.edu/dssg/>
- Alan Turing Institute <https://www.turing.ac.uk/collaborate-turing/data-science-social-good>

Doing Some Good: 4 Challenges

1. Teach more ML for good

Data Impact Lab

<https://datalab.stanford.edu/>

Where advanced students tackle important, high-impact problems with external partners

Improving healthcare outcomes

Alleviating poverty in the developing world

2. Publish more ML for good

ACM Conference on Computers and Sustainable Societies (COMPASS)

<https://acmcompass.org/>

AI for Social Impact Track: “Submitted papers are expected to present an argument for the (either realized or potential) social impact of the work.”

Doing Some Good: 4 Challenges

3. Incentivize more ML for good

LATHAM & WATKINS LLP

- All pro bono work counts for evaluation and bonuses
- Firm commits to at least 60 pro bono hours per lawyer per year
- Dedicated team for vetting pro bono matters

4. Prioritize more ML for good

The End