Snowed In: Measuring Information Exposure for Informing School Closure Decisions

Jeffrey Chen

Roadmap

1. **Motivation**
2. Considerations
3. Data + Methods
4. Results
5. Implications

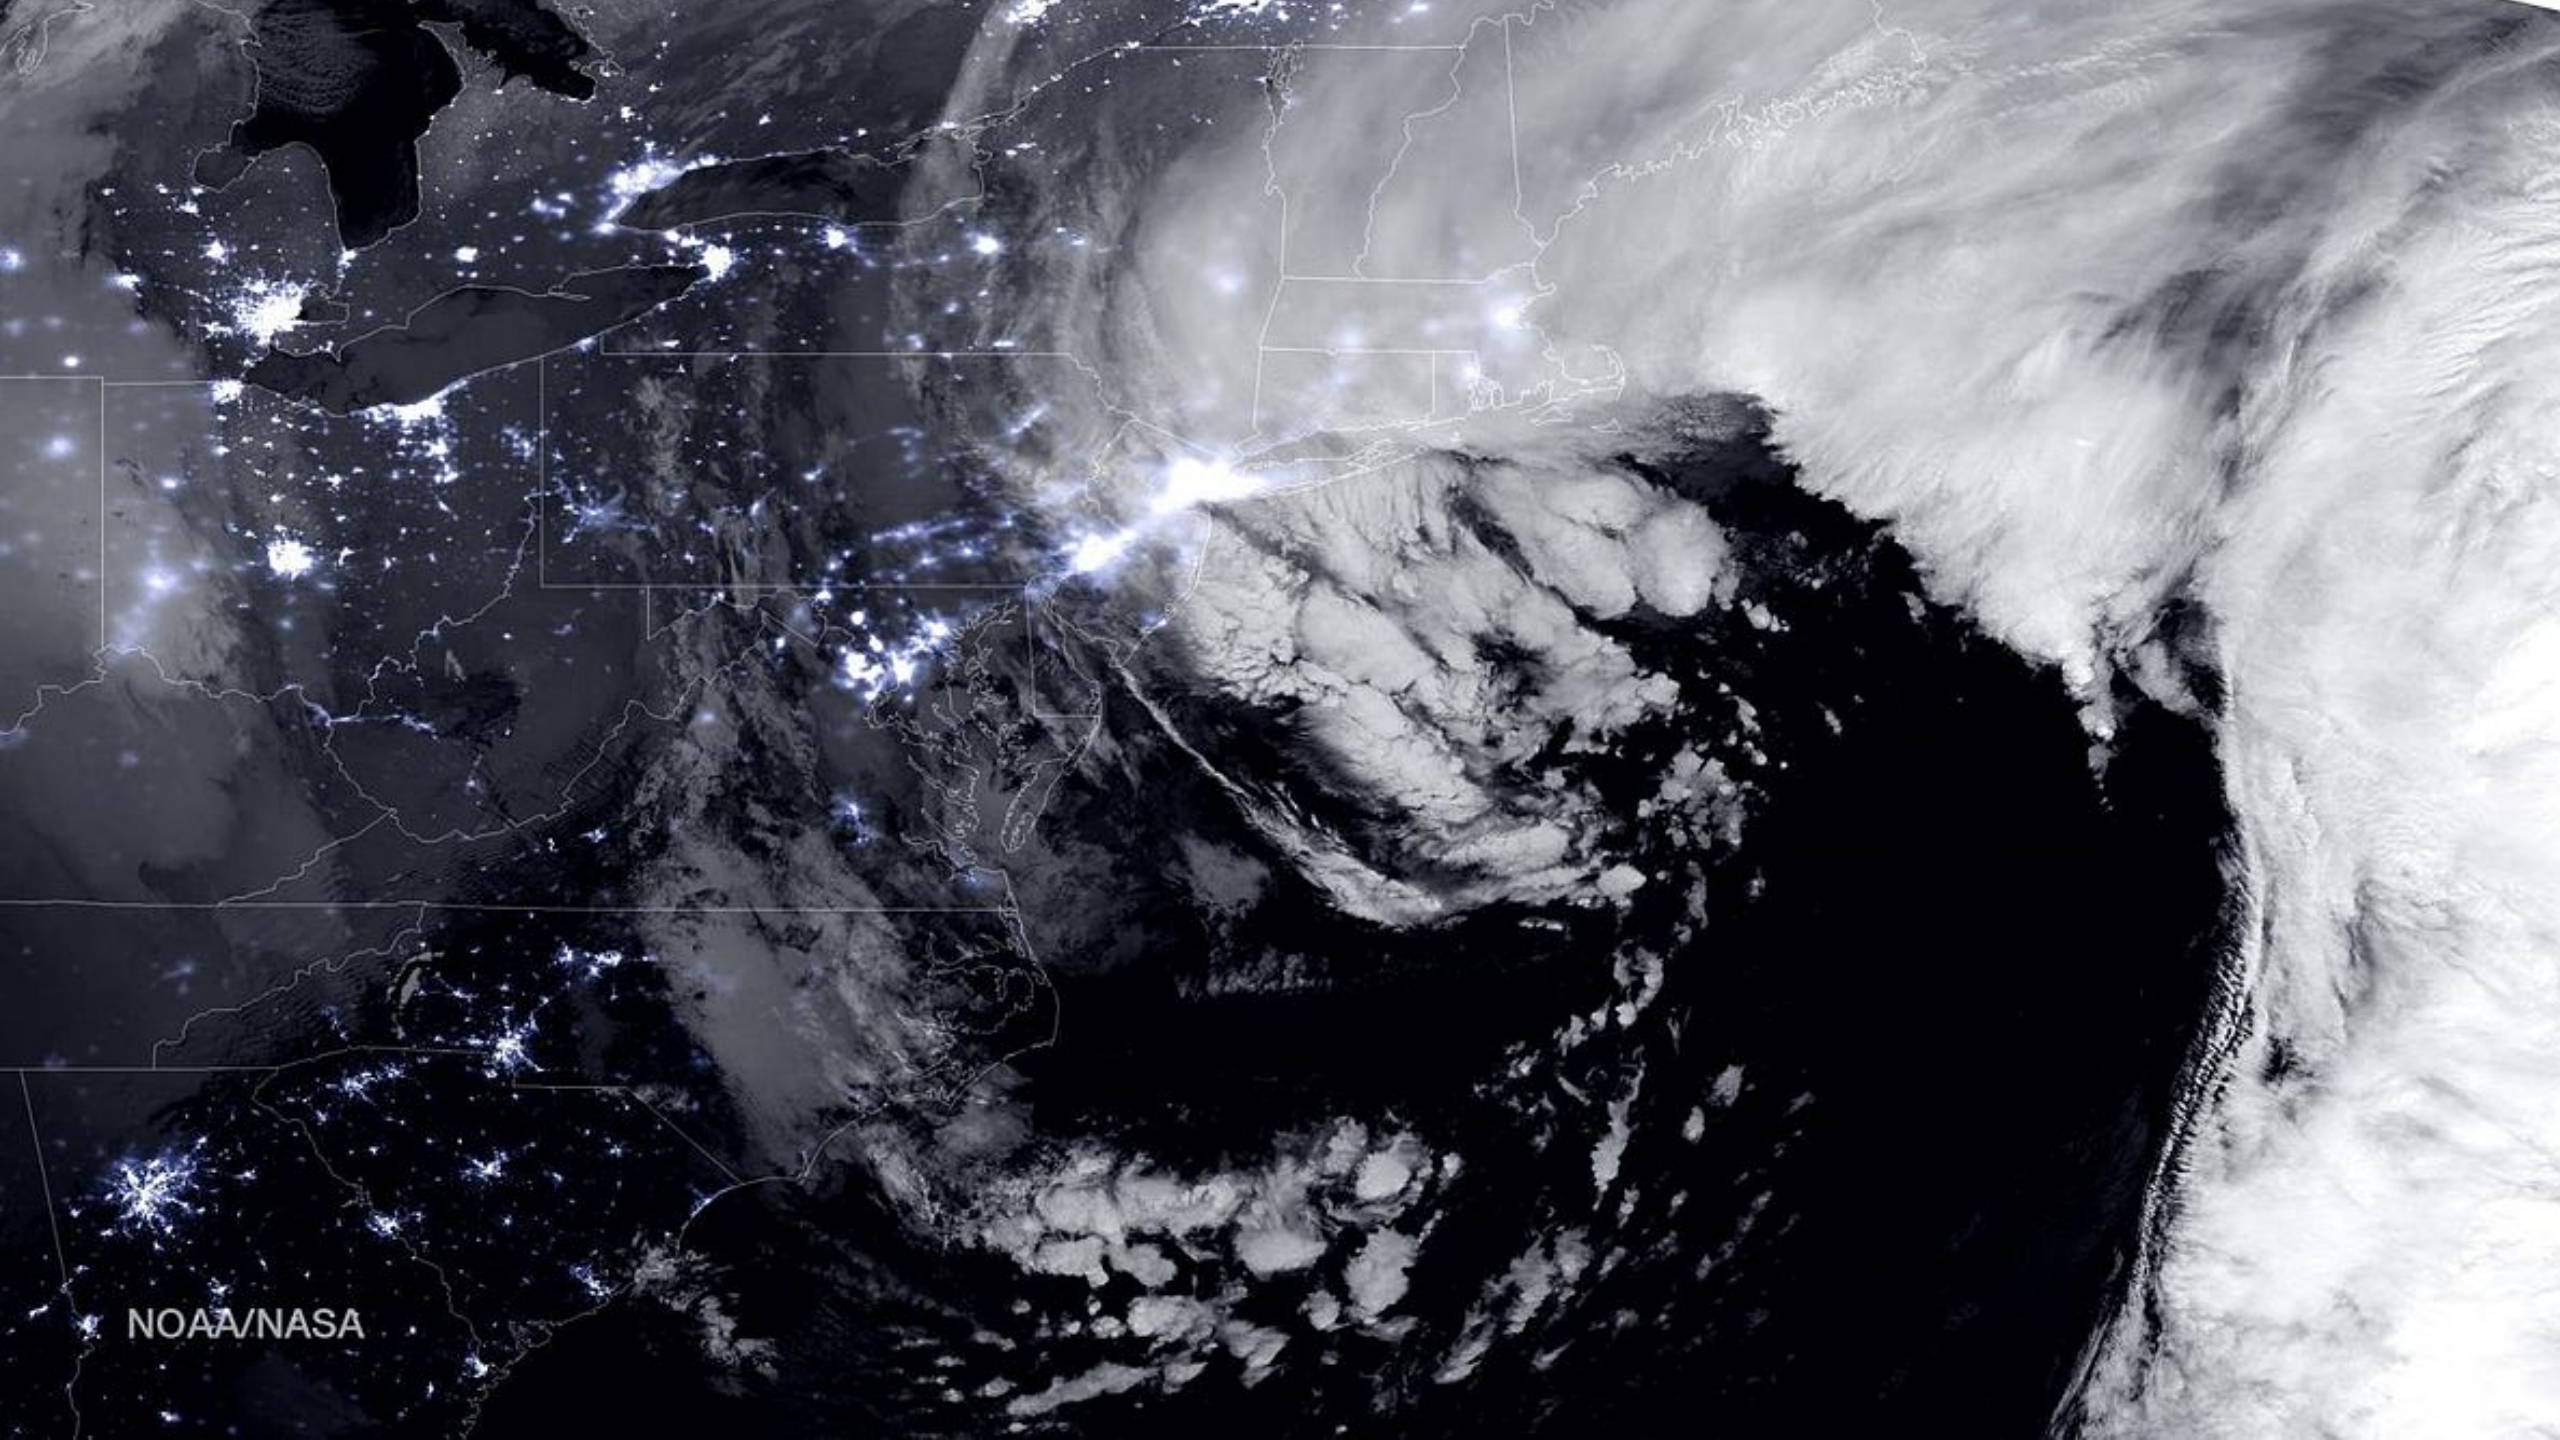
NOAA/NASA

The New York Times

New York Braces for Blizzard Amid Warnings of Closings and Hazards

By Colin Moynihan

Jan. 25, 2015

6

Mayor Bill de Blasio said on Sunday that the storm approaching on Monday was likely to be one of the biggest to ever strike New York City, and he urged people to stay indoors to avoid powerful winds, low visibility and “treacherous” road conditions.

The National Weather Service, which issued a blizzard warning for the greater New York City area, forecast gusts of wind up to 50 miles per hour and snow accumulation of “at least one to two feet.”

A Big Snowfall in the Forecast

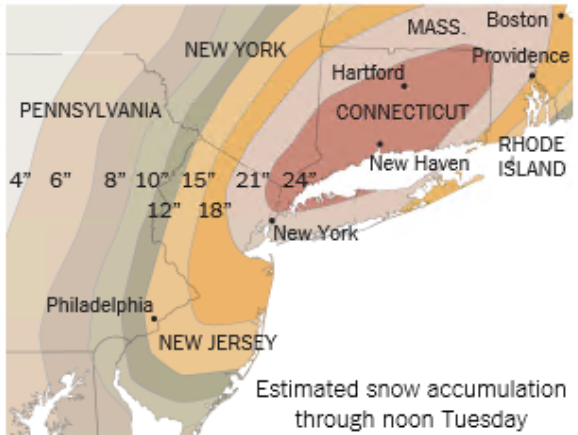

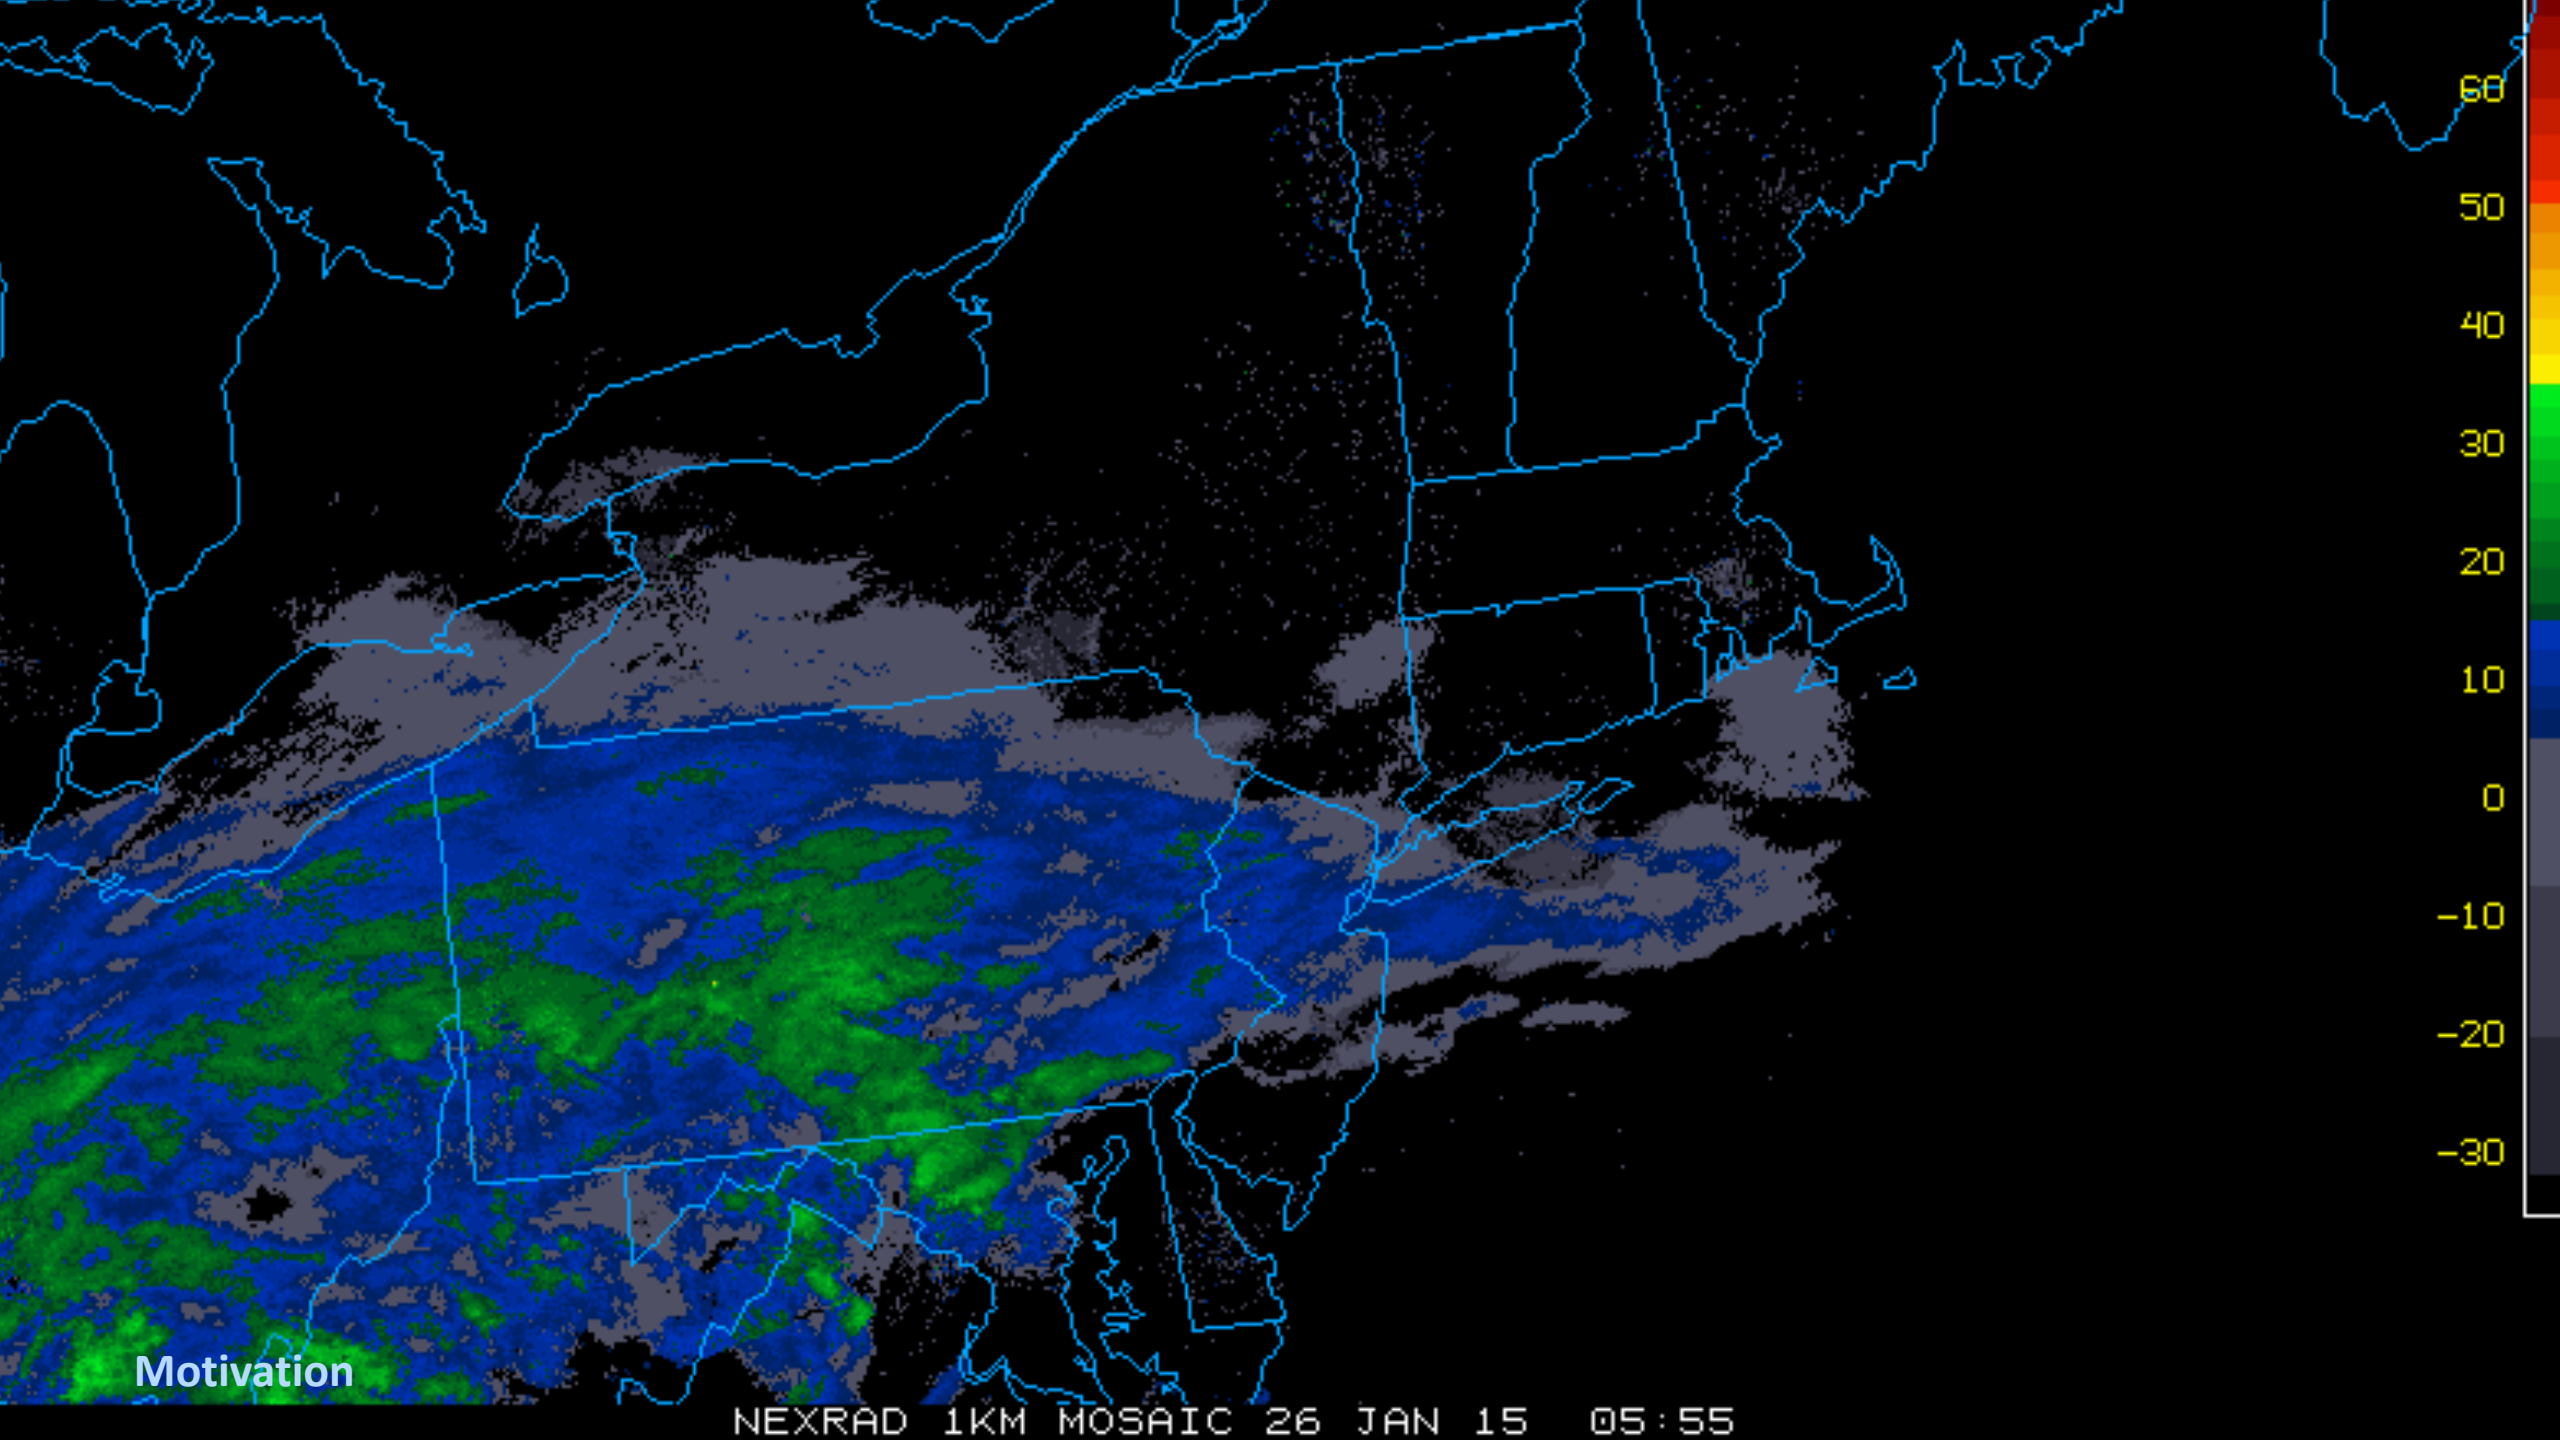
60
50
40
30
20
10
0
-10
-20
-30
Motivation
NEXRAD 1KM MOSAIC 26 JAN 15 05:55

The New York Times

New York City Is Spared the Worst Effects of Snowstorm

Snow plows worked through the night to clear New York City's roads.

3 / 11

By Marc Santora and Emma G. Fitzsimmons

Jan. 26, 2015

850

For the latest on the winter storm affecting the New York region, click here

The New York Times

Leaders in New York and New Jersey Defend Shutdown for a Blizzard That Wasn't

Mayor Bill de Blasio of New York City, with Sanitation Department workers in Manhattan on Monday, when he issued dire warnings about the storm. Yana Paskova for The New York Times

By Matt Flegenheimer

Jan. 27, 2015

1350

Challenges with Anticipating Weather Impacts

- **False positives** and **false negatives** as possible social costs: possible productivity losses and expenses due to alternative child care arrangements

- **Imperfect information**, especially for families, are a barrier to making sound decisions. In other words, making decisions under uncertainty is just plain hard.

Research Goals

- Can school closures be anticipated using winter weather forecasts? (12 hours in advance)
- Relative to NWS warnings, how would a school closure forecast improve the public's information exposure?
- How can information exposure be associated with economic impacts?

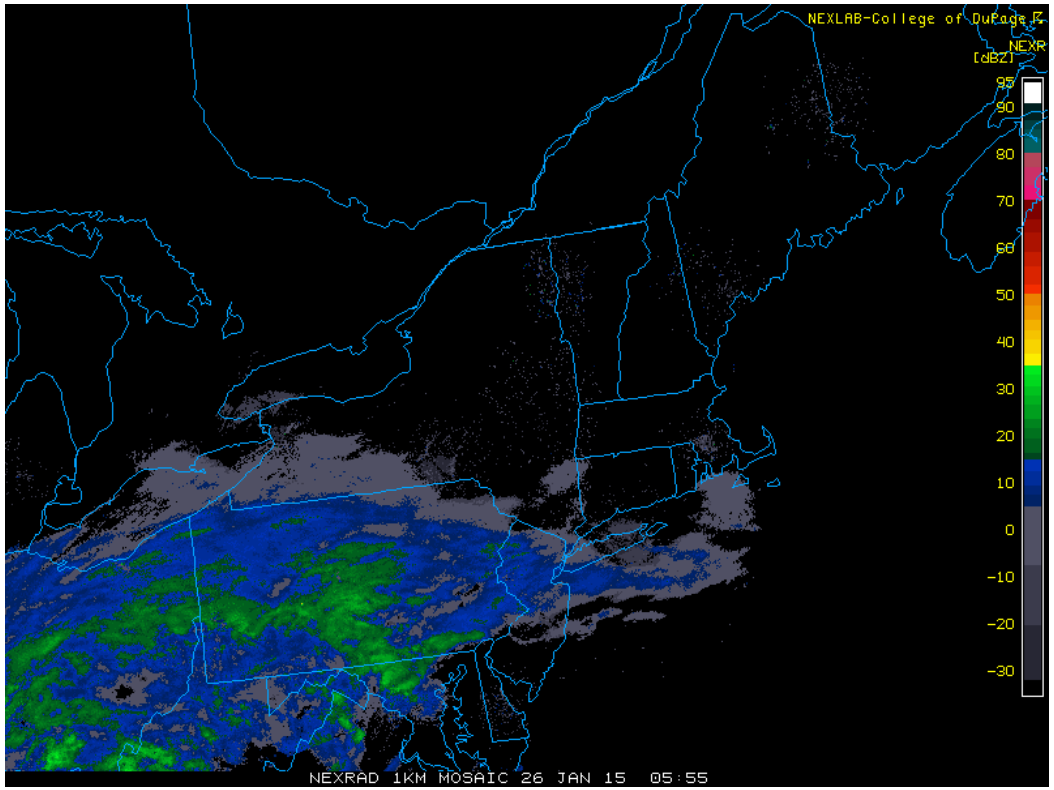

Roadmap

1. Motivation
2. Considerations
3. Data + Methods
4. Results
5. Implications

Considerations: Local Exposure to Weather and Attitudes

Days with Forecasted Snow

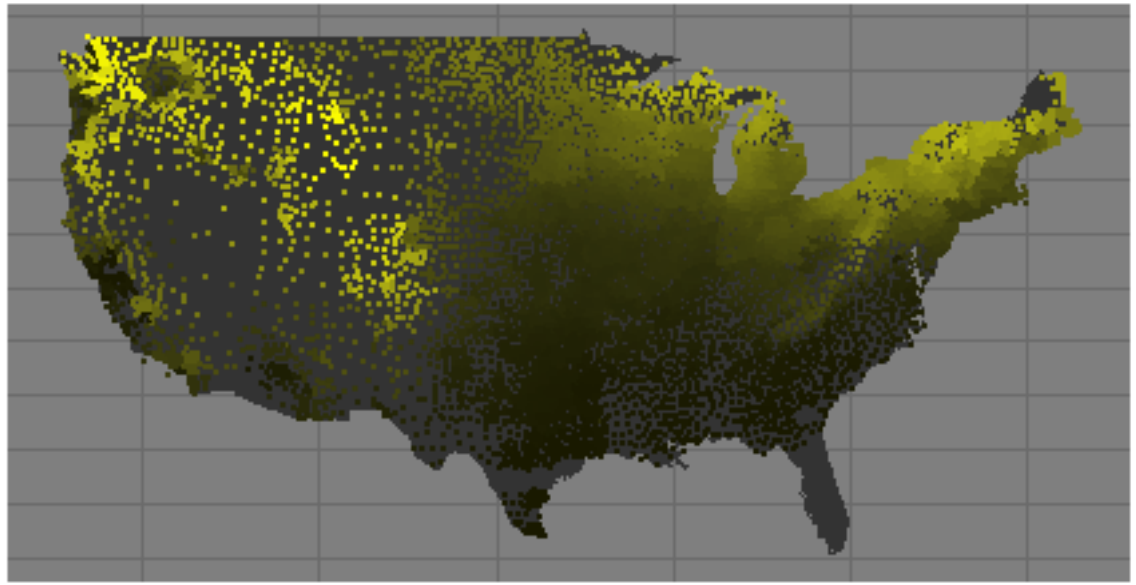

Closure Days

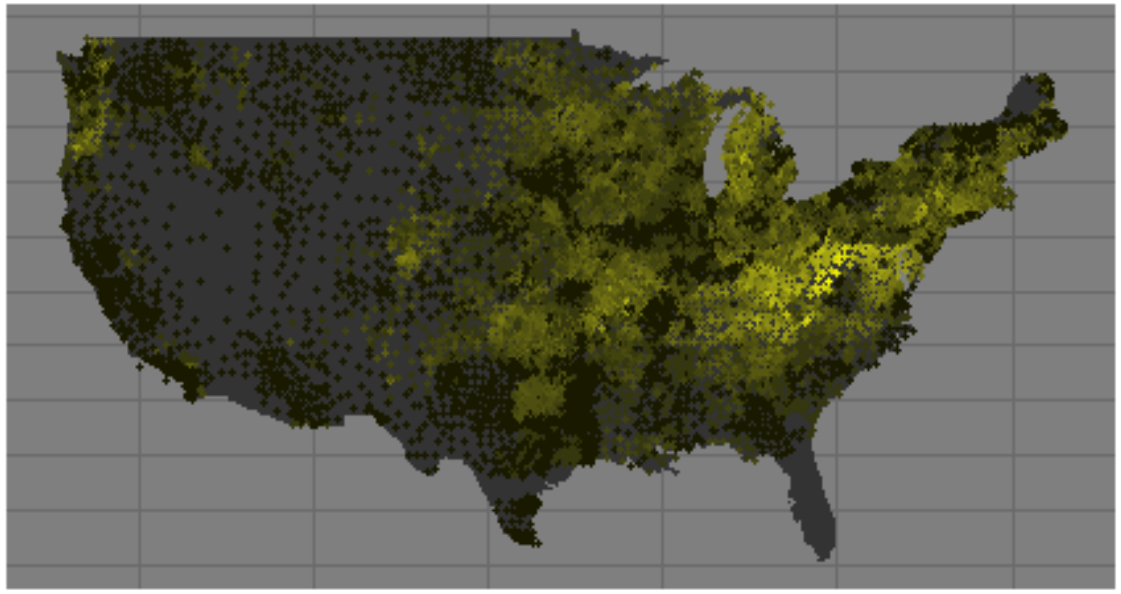

Each region of the country has different exposure to weather.

Thus, how common closure days are will vary from area to area.

Considerations: Local Exposure to Weather and Attitudes

Percent of district-days with a weather notification

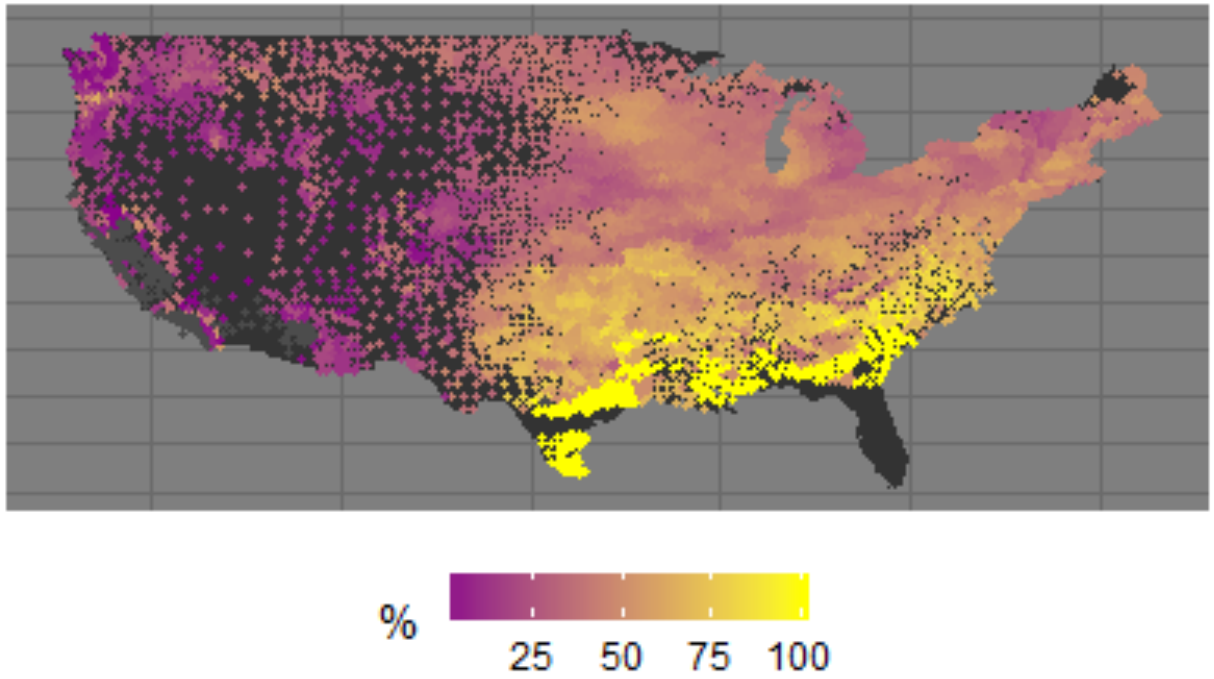

Number of district-days with a weather notification

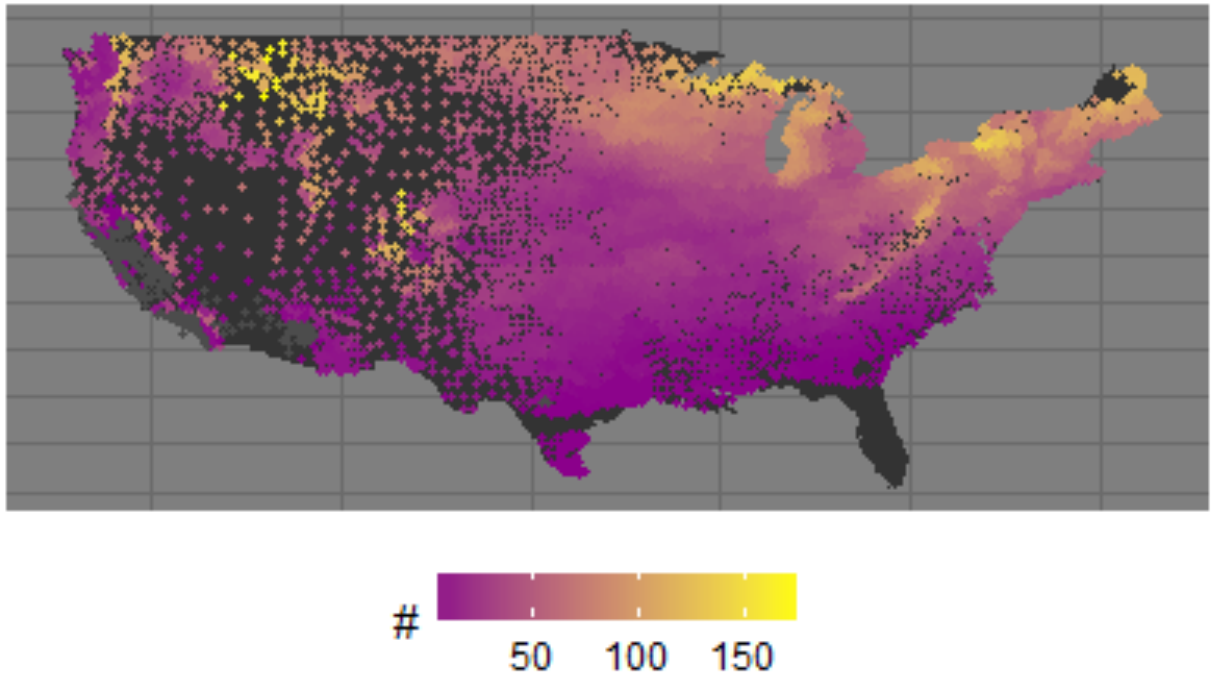

Considerations: Local Exposure to Weather and Attitudes

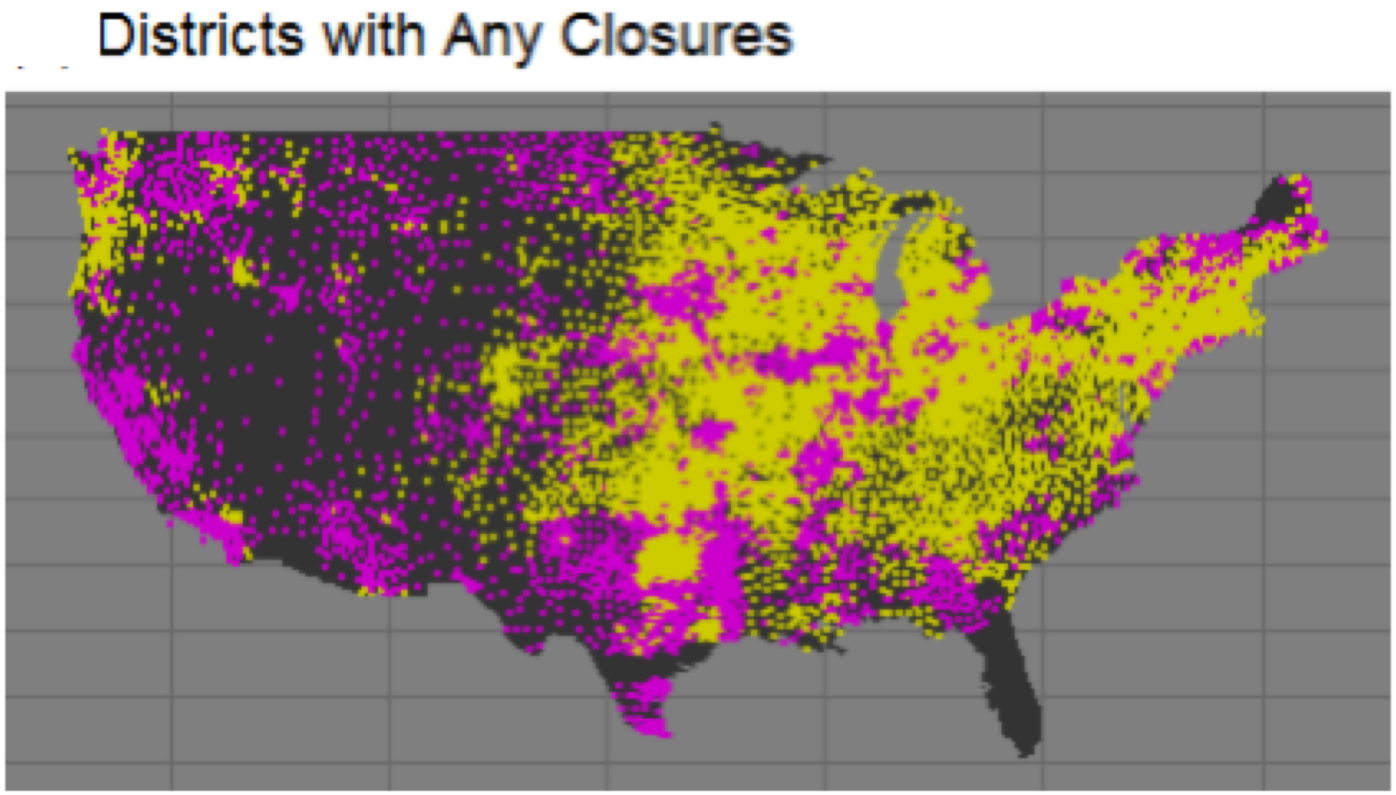

This also meant that some parts of the country have very infrequent potential snow-related weather closures.

Considerations: Benchmark to Beat

Percent of district-days with a weather notification

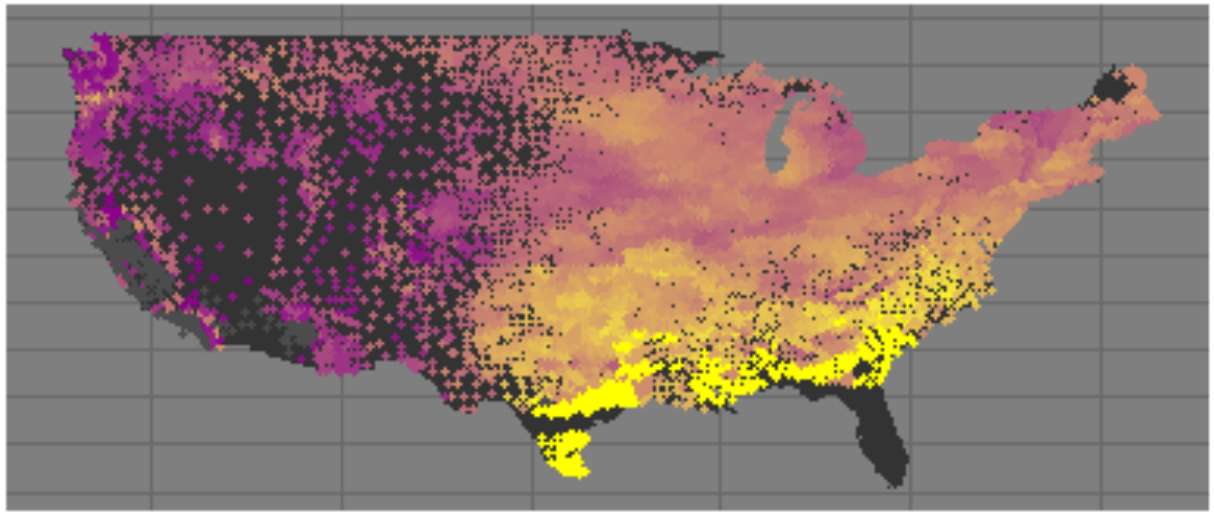

School Closures Per Forecasted Snow Days

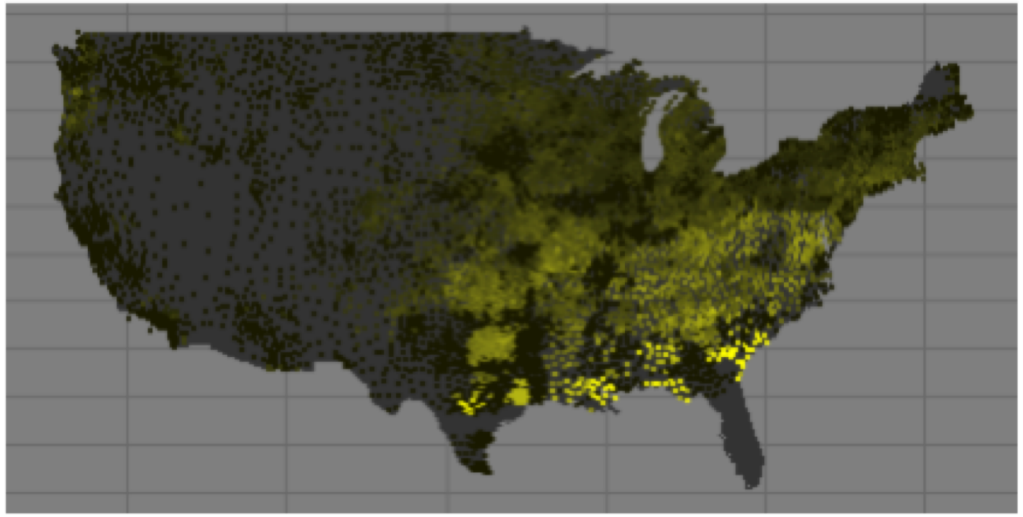

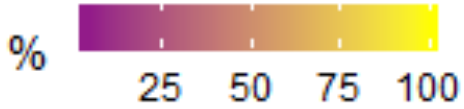

We should ask then if a warning is sufficient to inform a naïve observer about whether a school closure will happen?

Roadmap

1. Motivation
2. Considerations
3. **Data + Methods**
4. Results
5. Implications

Plan of Attack

1. **Construct District-Day Data set**
2. **Train Classifier using 10-folds cross validation**
3. **Use Cross-validated predictions to set notification thresholds**

Generic Modeling Approach

$$Y_{t,i} = f(W_{t-p,i}, Q_{t,i}, R_{t,i}, D_{t,})$$

- $Y_{t,i}$: Binary Target (School Closures)
- $W_{t-p,i}$: Forecasted Weather
- $Q_{t,i}$: Geography
- $R_{t,i}$: Social
- $D_{t,}$: Time

Data Collection

Gridded Weather Forecasts (12h horizon: Rain, Temp Min, Snow)

Local population demographics

School district closures (Wong et al. 2016) (Target)

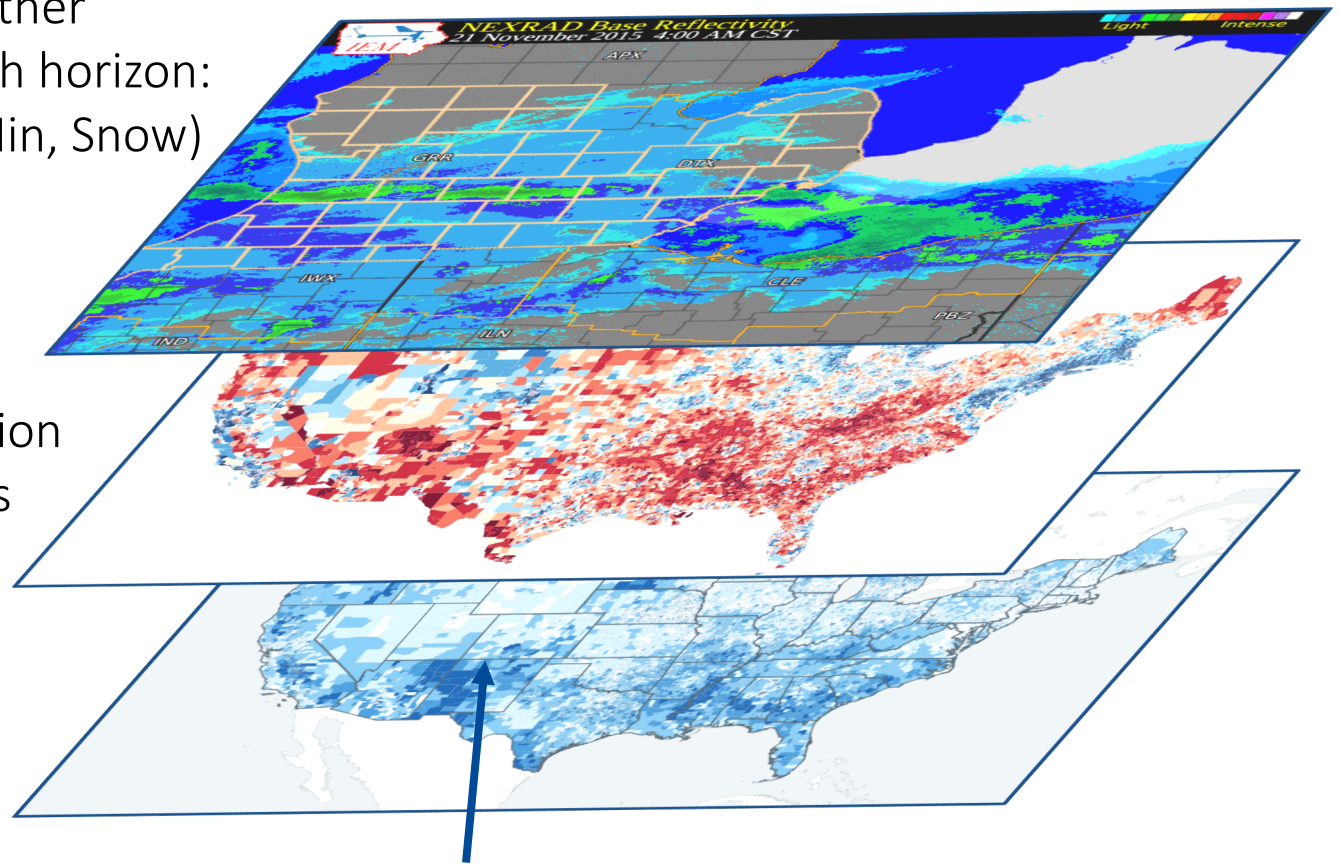

Unit of Analysis: "District-Day with Forecasted Snow"

Plan of Attack

1. Construct District-Day Data set
2. Train Classifier using 10-folds cross validation
3. Use Cross-validated predictions to set notification thresholds

Anatomy of Classification Model Predictions

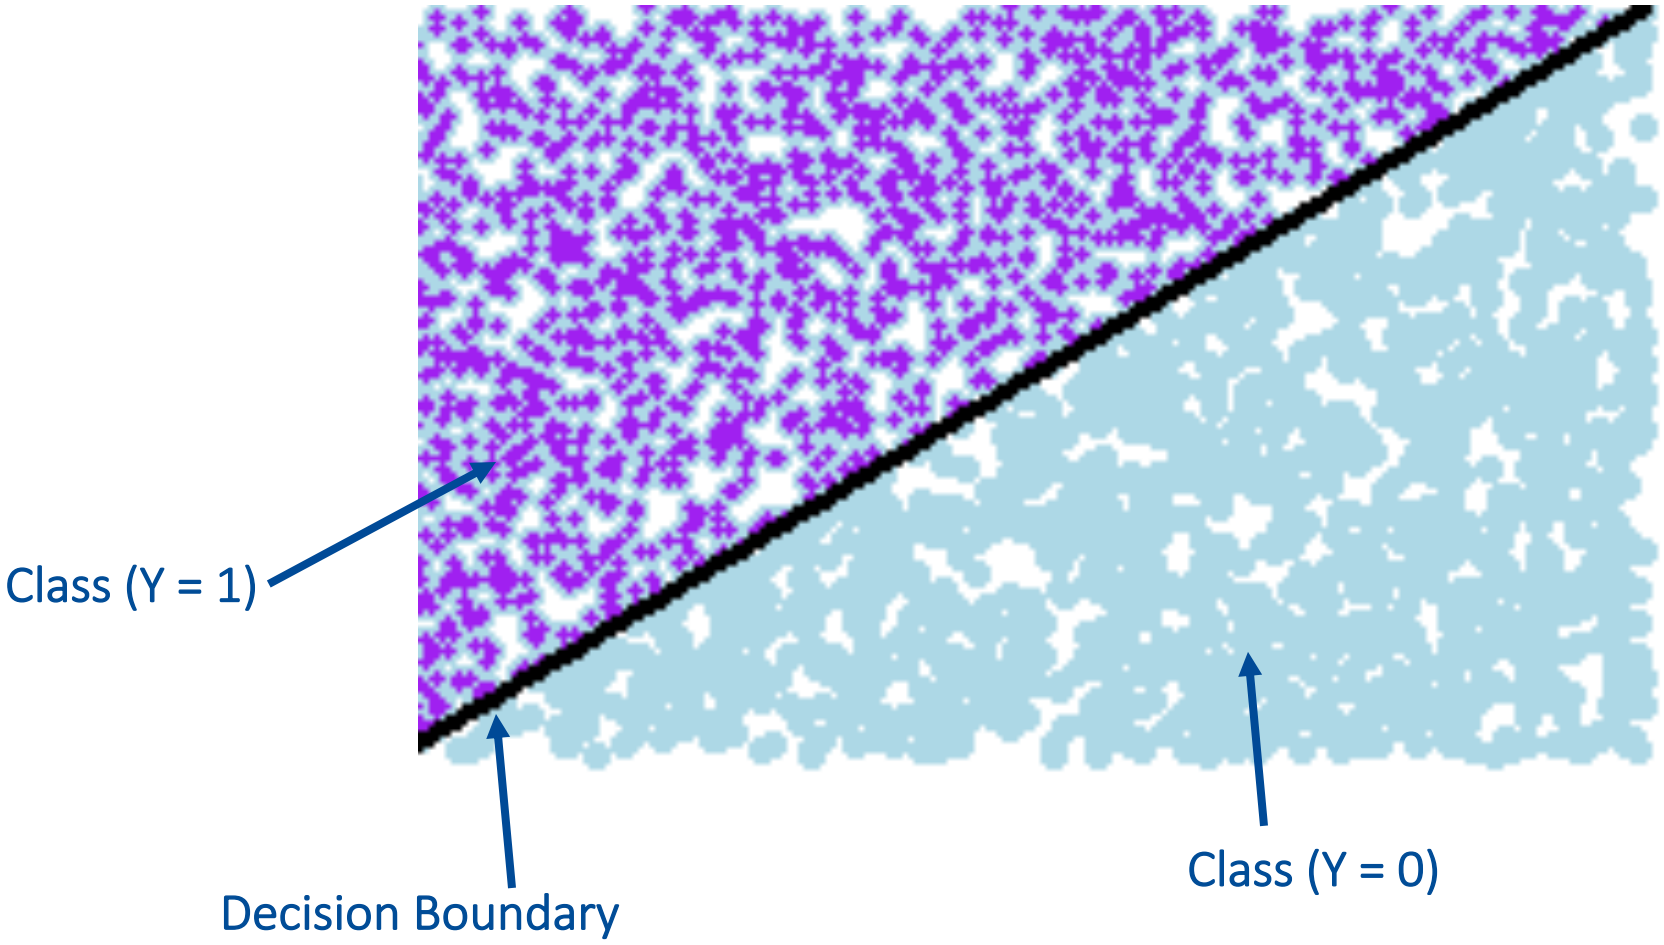

Choice of algorithm is important

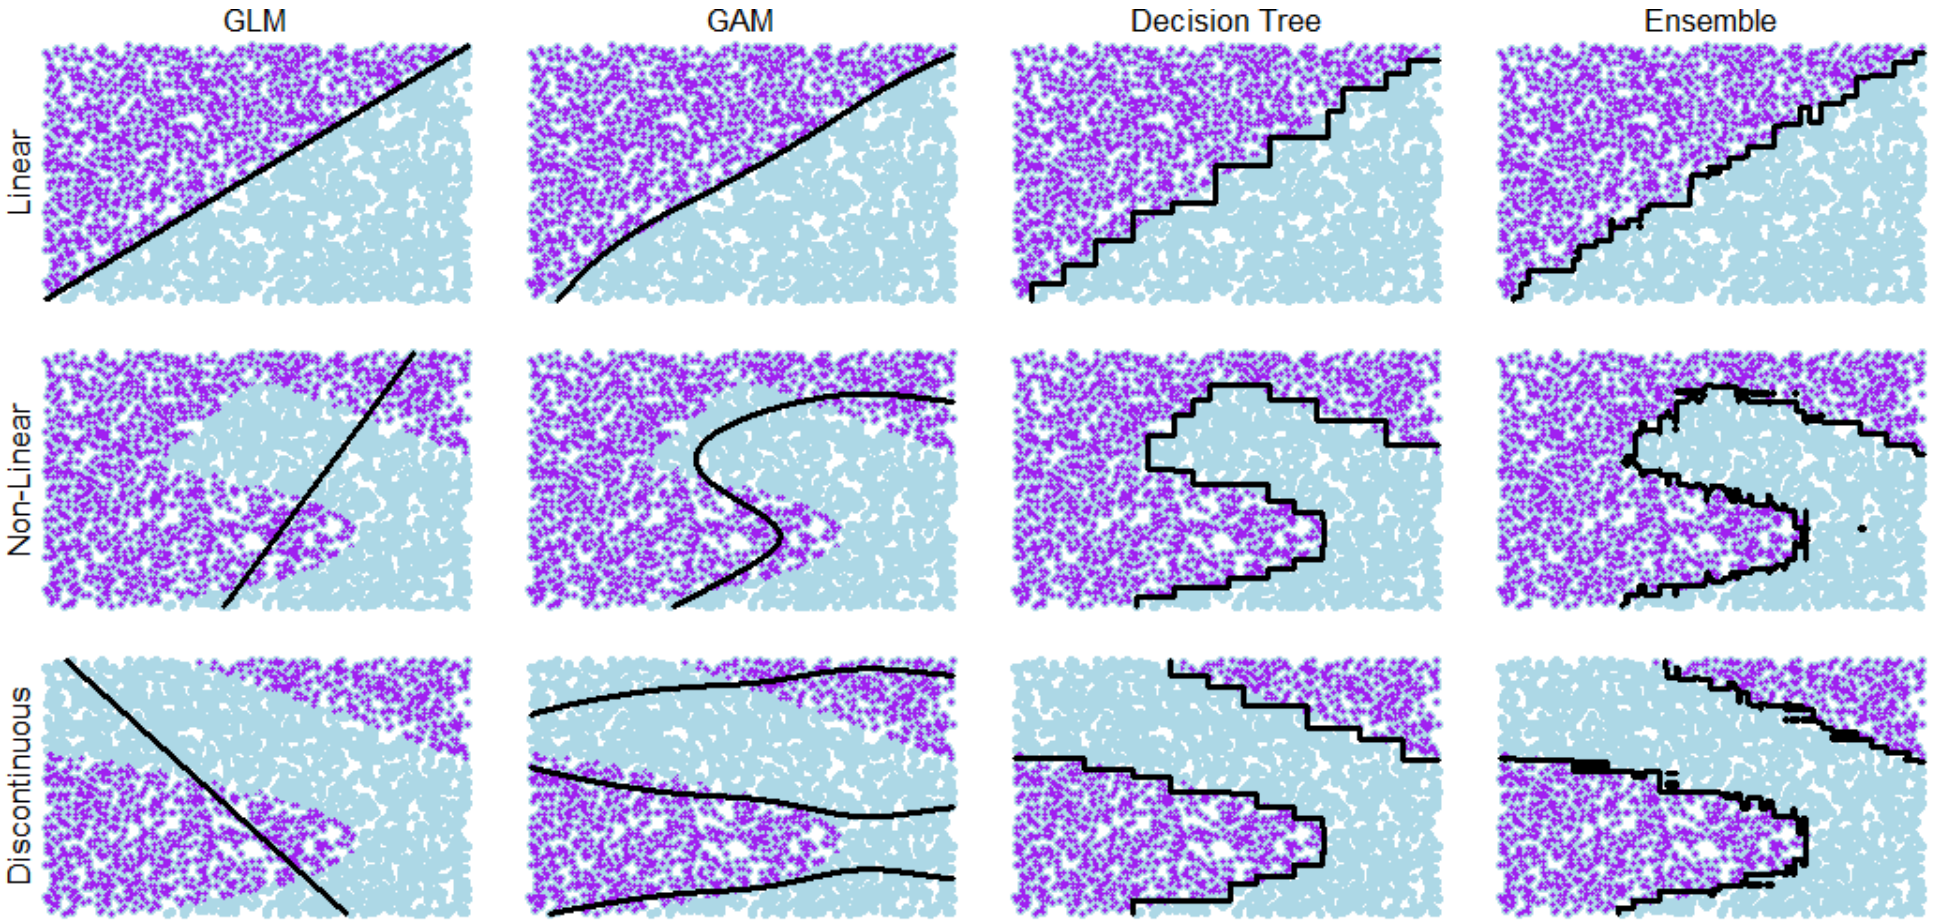

Beede and Chen (2018)

Decision trees as the Foundation of Tree Learning

Decision Boundary

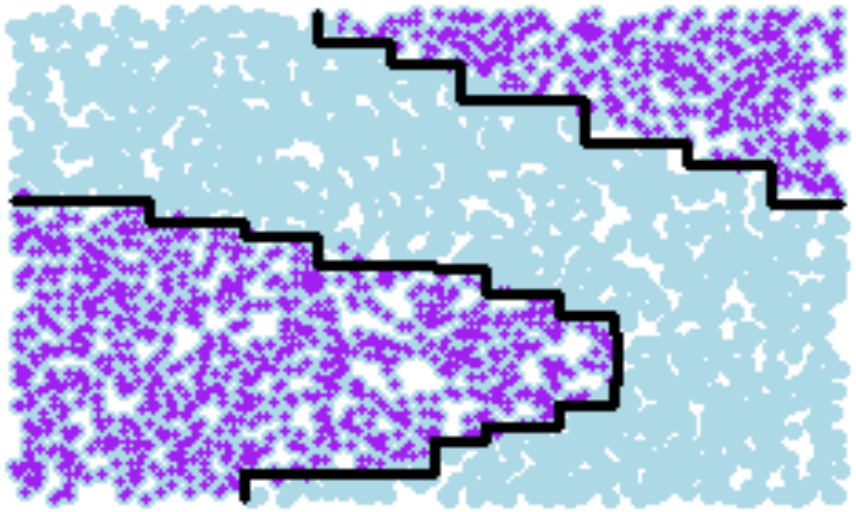

=

Tree Structure

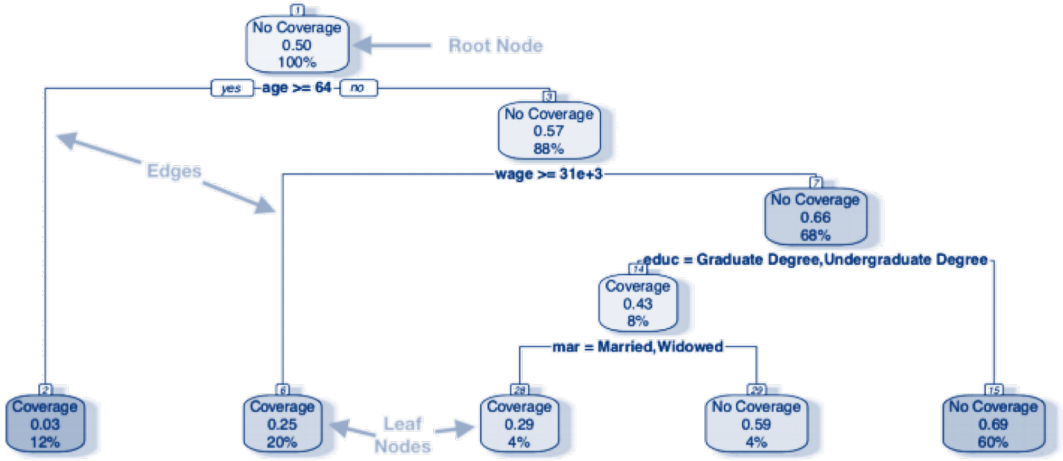

- Jagged edges due to recursive partitioning
- More flexible than regression, but tend to overfit

Decision trees as the Foundation of Tree Learning

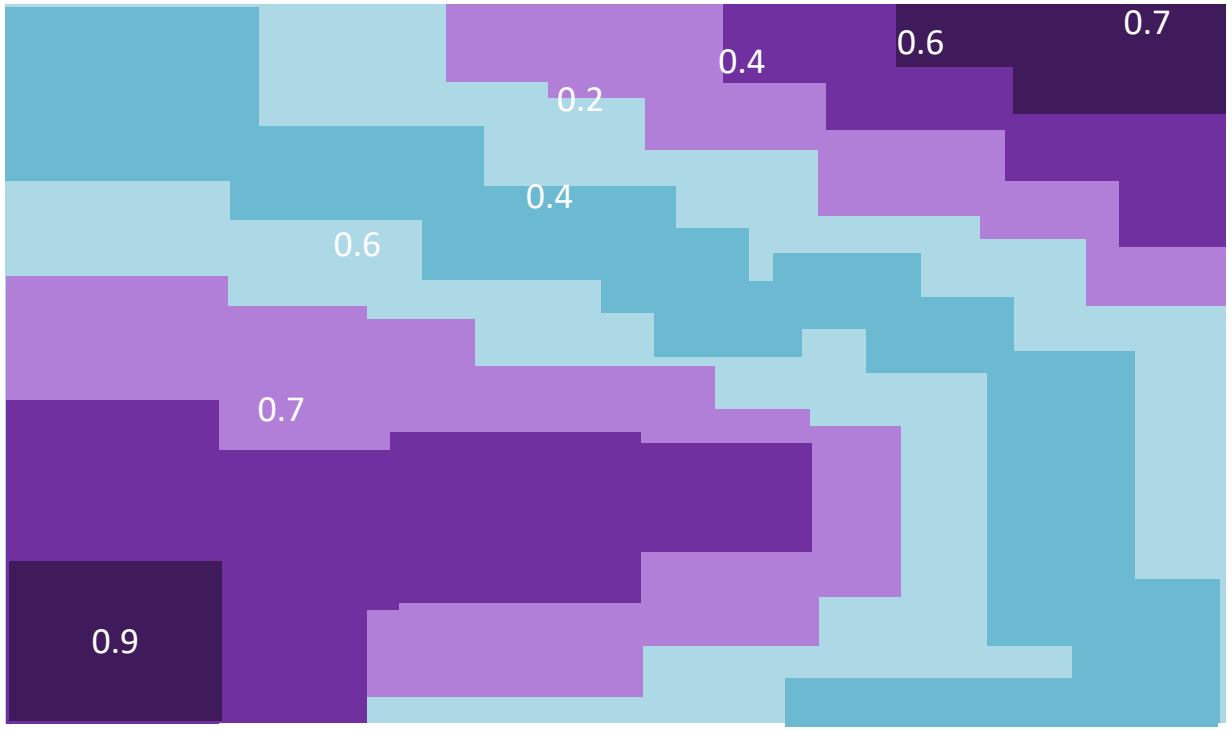

$$T(x;\Theta) = \sum_{j=1}^{J} \gamma_j I(x \in R_j)$$

$$\Theta = \{R_j, \gamma_j\}_1^J$$

Where:

- R_j is a region in the probability space
- γ_j is the mean probability of the outcome in region R_j

Decision trees as the Foundation of Tree Learning

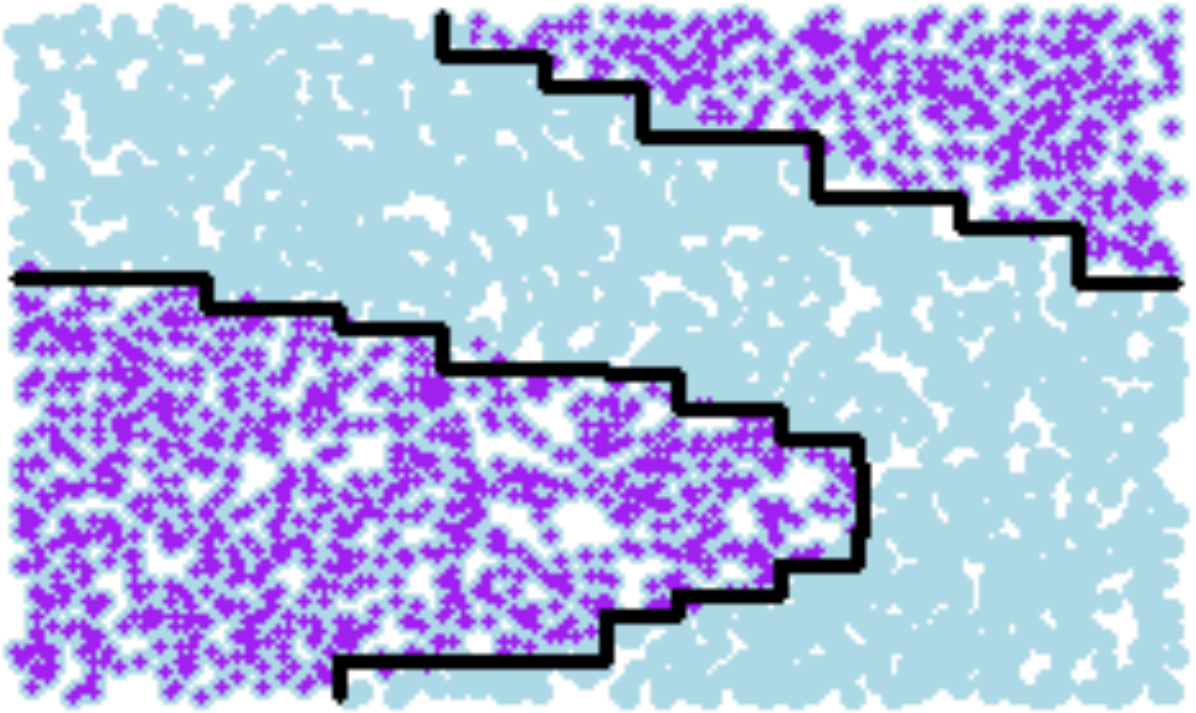

$$T(x;\Theta) = \sum_{j=1}^{J} \gamma_j I(x \in R_j)$$

$$\Theta = \{R_j, \gamma_j\}_1^J$$

Where:

- R_j is a region in the probability space
- γ_j is the mean probability of the outcome in region R_j

Gradient Boosting

A major improvement over decision trees.

$$f_M(x) = \sum_{m=1}^{M} T(x; \Theta_m)$$

Sum of M-number of trees

Base learner (decision stump)

Each iteration is a refinement on the last iteration

$$f_m(x) = f_{m-1}(x) + \eta \sum_{j=1}^{J} \gamma_{jm} I(x \in R_{jm})$$

- $f_m(x)$: Model in m-iteration
- $f_{m-1}(x)$: Model in previous iteration
- η: Shrinkage parameter
- γ_{jm}: Probability for the Region j in iteration m
- $I(x \in R_{jm})$: Identify matrix defined by input features

Gradient boosting loss relies on deviance

$$\Theta_m = argmin \sum_{i=1}^{N} (-g_{im} - T(x_i; \Theta))^2$$

Where:

- g_{im} is the deviance $I(y_i = G_k) - p_k(x_i)$

Train-Test Sample Partition

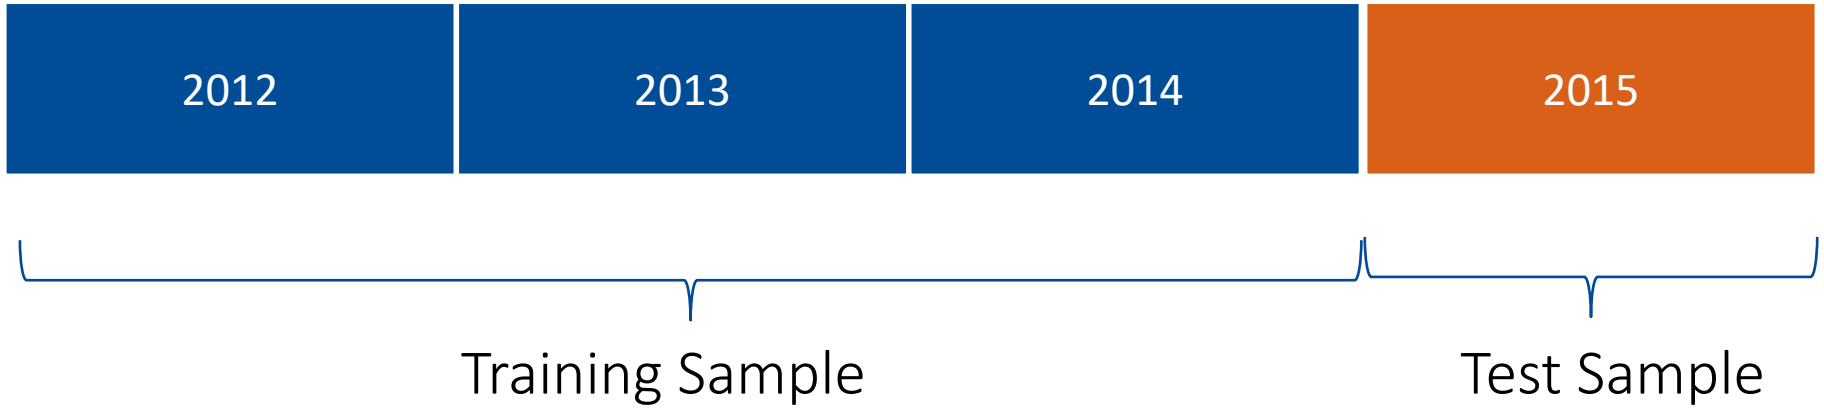

Training: 10-Folds Cross Validation

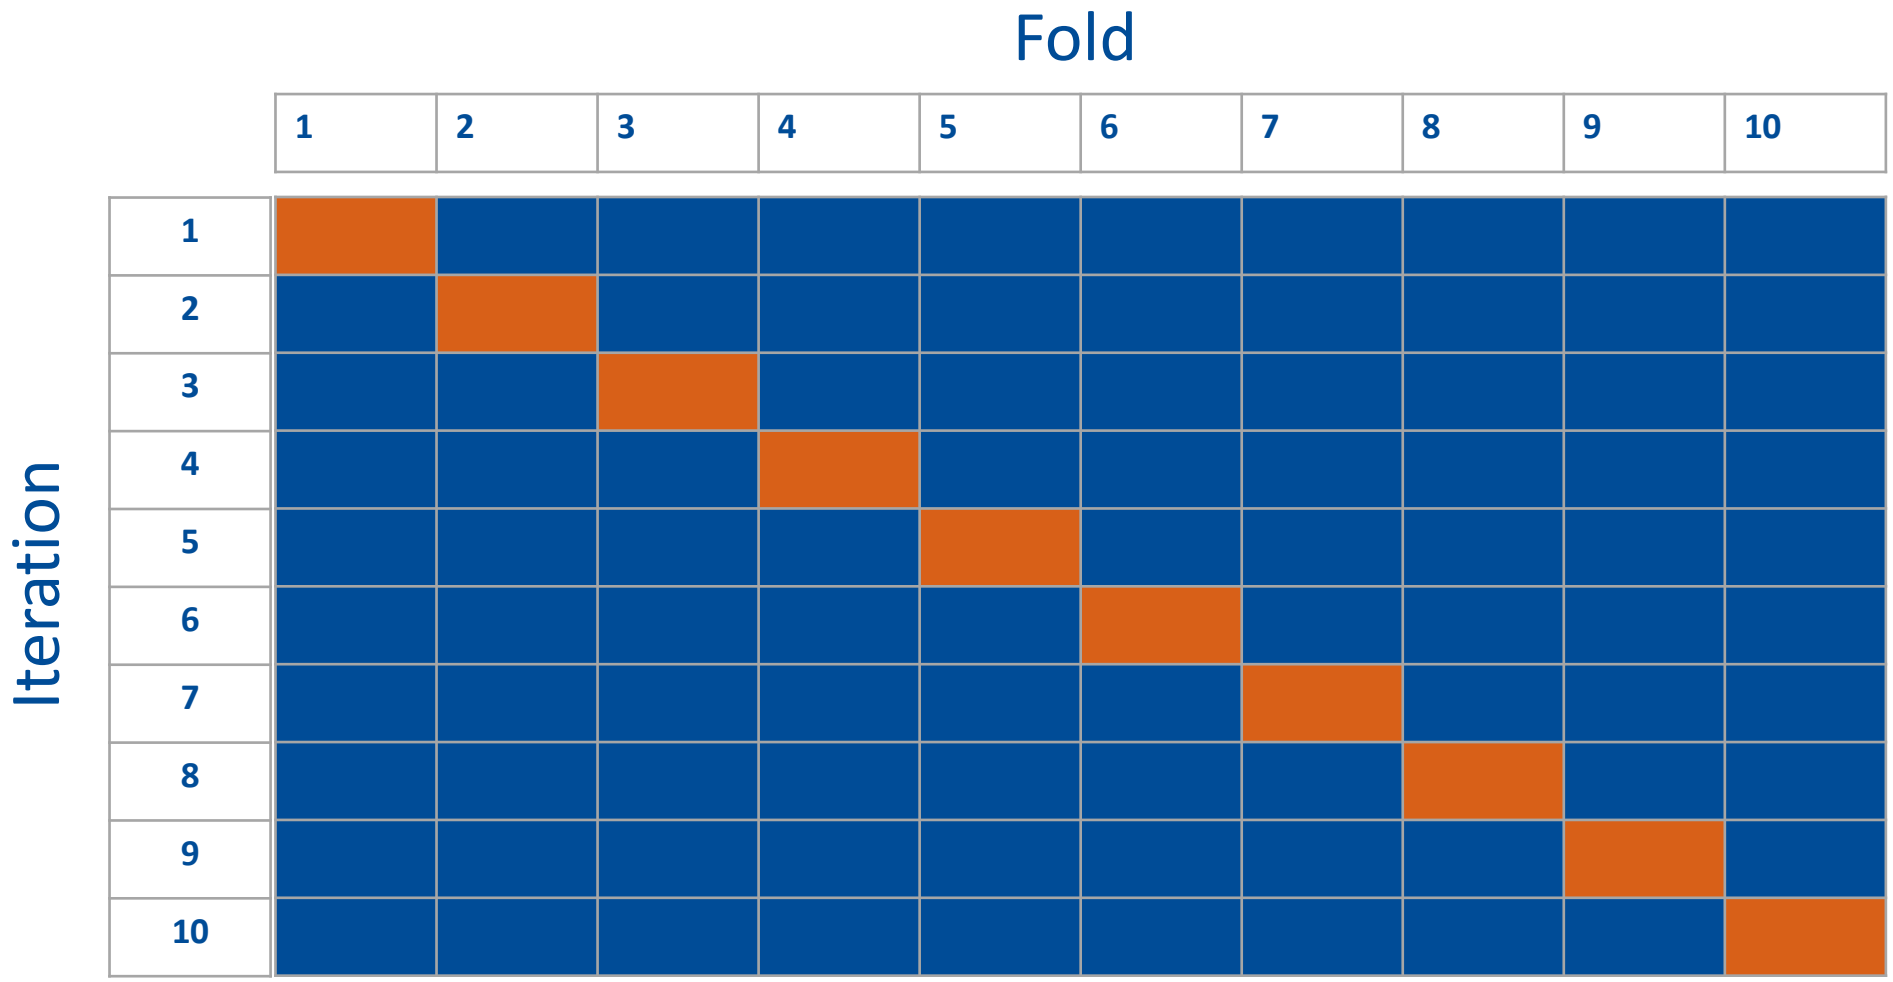

Plan of Attack

1. Construct District-Day Data set
2. Train Classifier using 10-folds cross validation
3. Use Cross-validated predictions to set notification thresholds

Selecting classification threshold

Predicted probabilities from 10-folds CV

Classification Thresholds: Modified F1-Statistic

$$F1 = \frac{2}{\frac{1}{TPR} + \frac{1}{TNR}}$$

s.t. $TPR - TNR \geq d$

Population-Weighted True Positive Rate

Population-Weighted True Negative Rate

Harmonic Mean

TPR – TNR is a positive value with a margin ***d***.

Calibrating the School Closure Classification Threshold

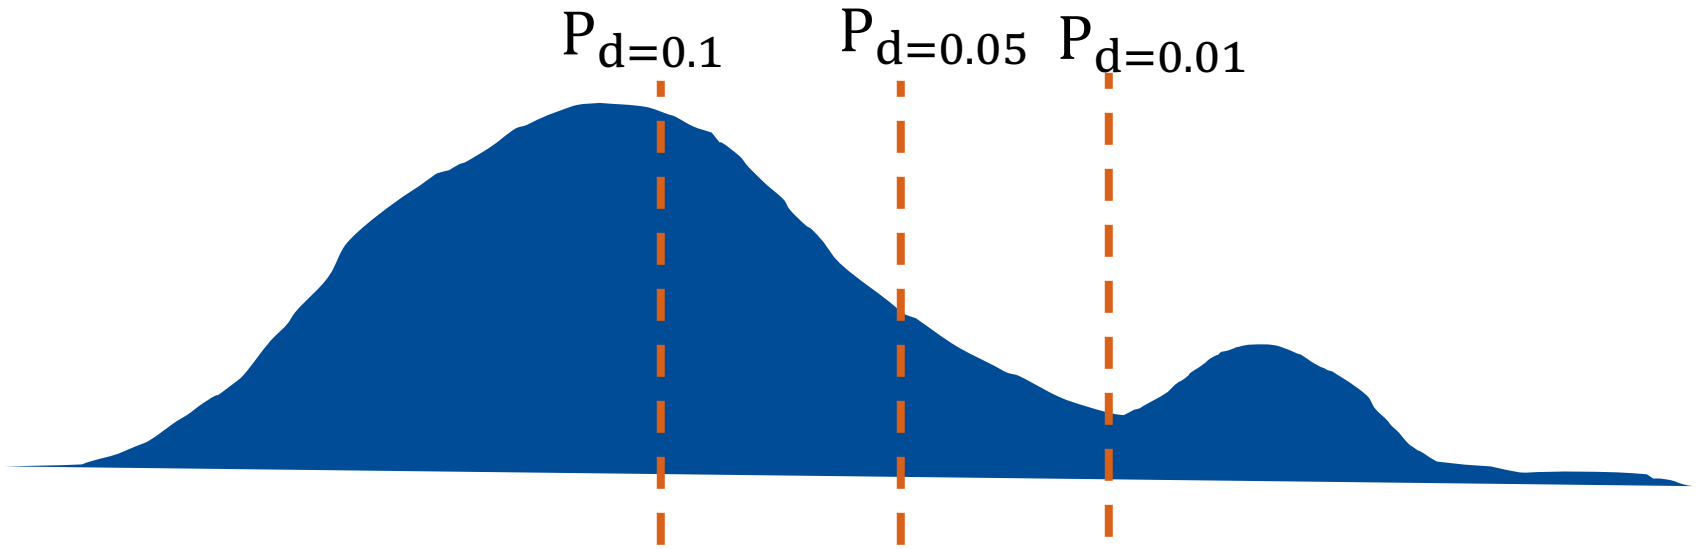

1. For $\hat{p}$, identify a probability threshold P_i that maximizes a population-weighted accuracy measure;

2. Split each fold along P_i; then

3. For the subset where $\hat{p} > P_i$, apply step 1 again until a desired number of tiers have been defined.

Roadmap

1. Motivation
2. Considerations
3. Data + Methods
4. **Results**
5. Implications

Accuracy across geography

TPR by School District

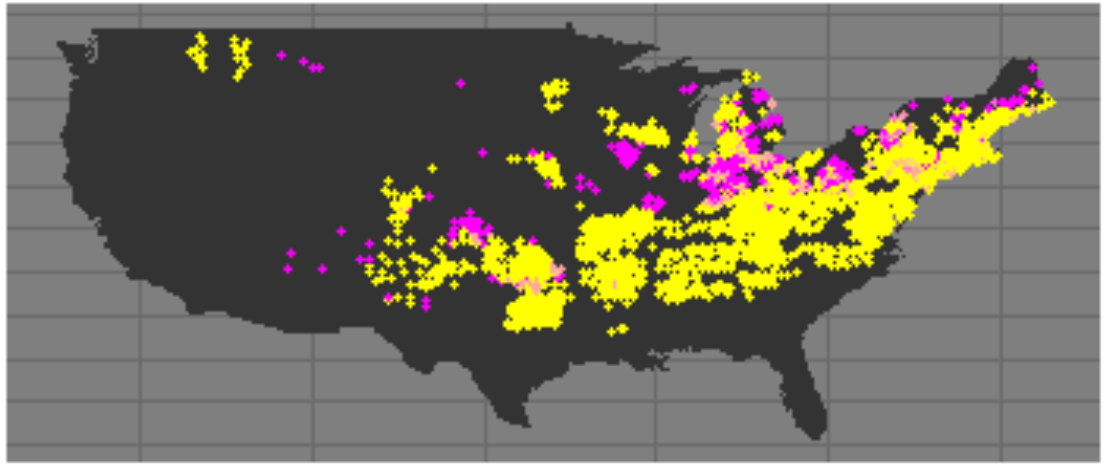

TNR for ML Model by School District

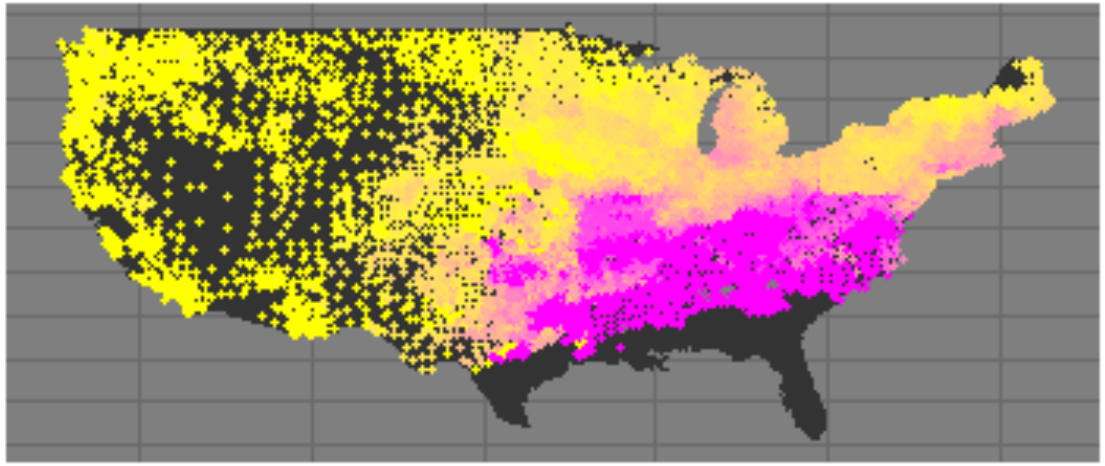

Predictions are best outside of the US Southeast.

Learned School Decisions versus Winter Warnings

TPR by School District

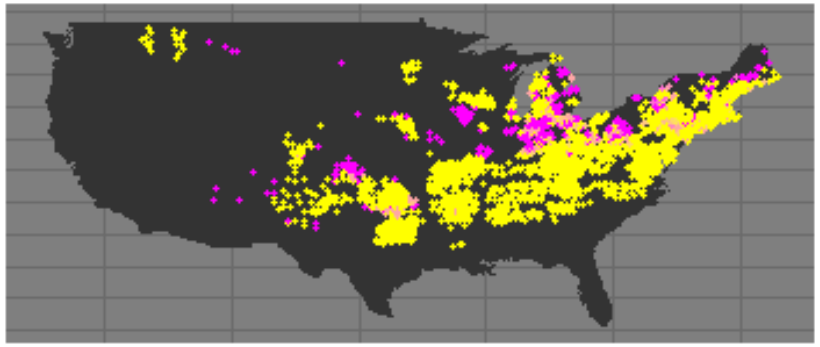

TNR for ML Model by School District

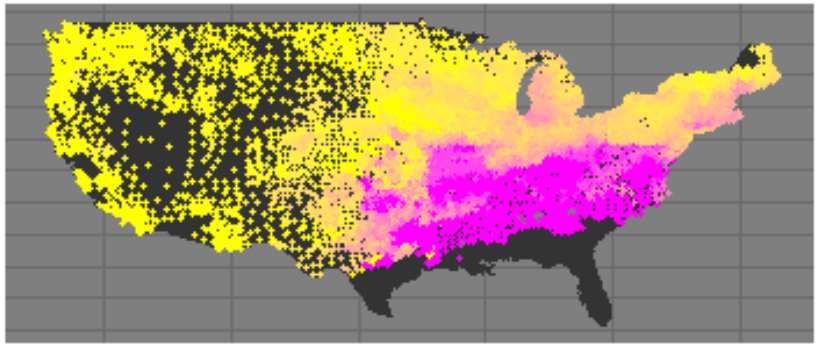

TPR for NWS Notifications

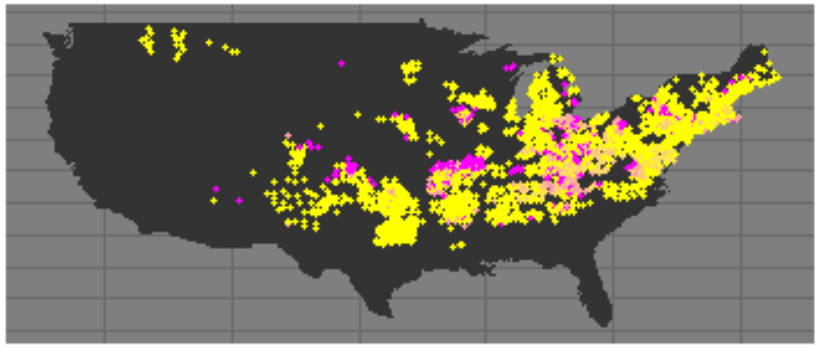

TNR for NWS Notifications

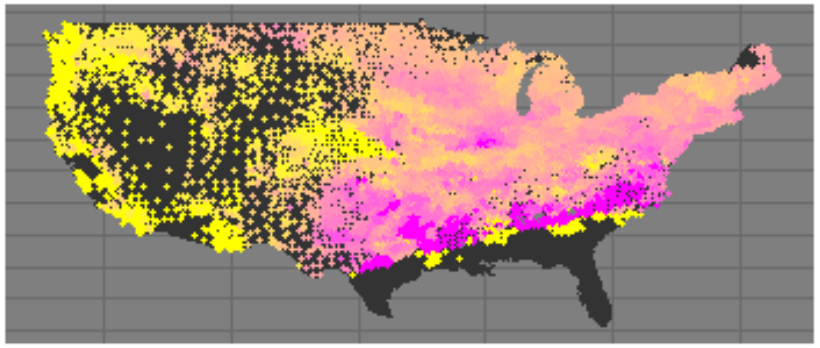

Roadmap

1. Motivation
2. Considerations
3. Data + Methods
4. Results
5. **Implications**

Maximum Possible Reduction of False Alarms and Missed Alarms

Margin	FN	FP
Model: d = 1%	0	-146 M
Model: d = 5%	-0.7 M	-126 M
Model: d = 10%	-1.6 M	-83 M

Day One of Blizzard (January 26, 2015)

(1) Actual School Closures

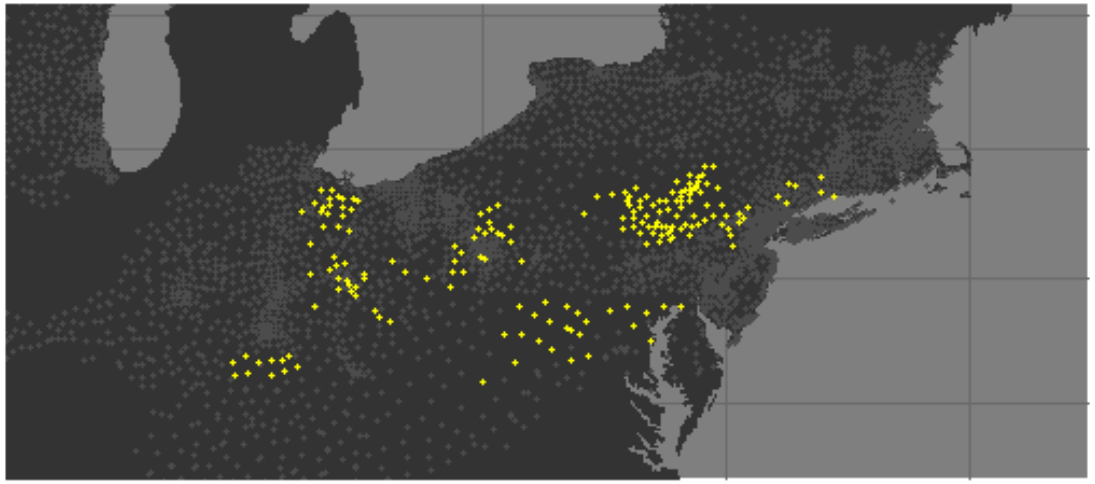

(2) Any Warning Notification

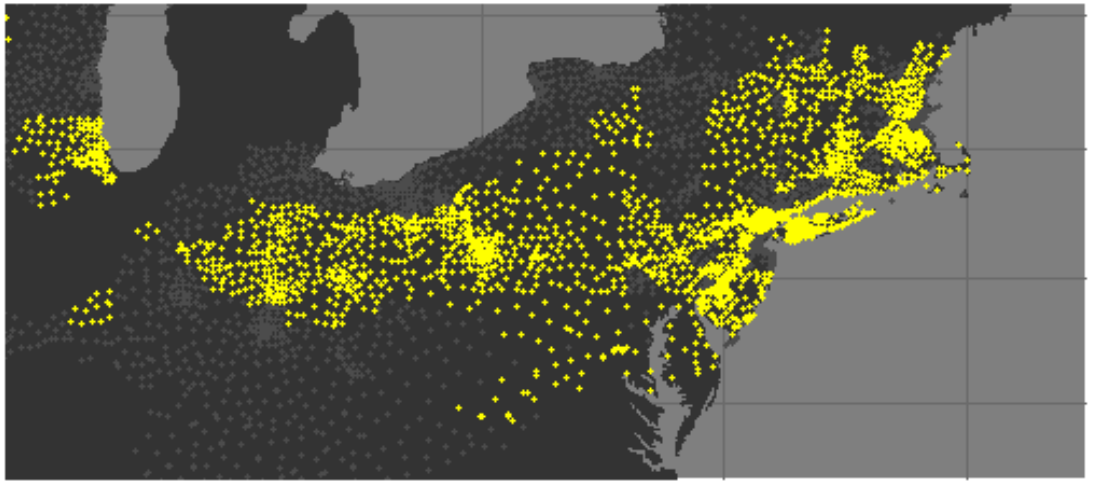

(3) Lagged Notification

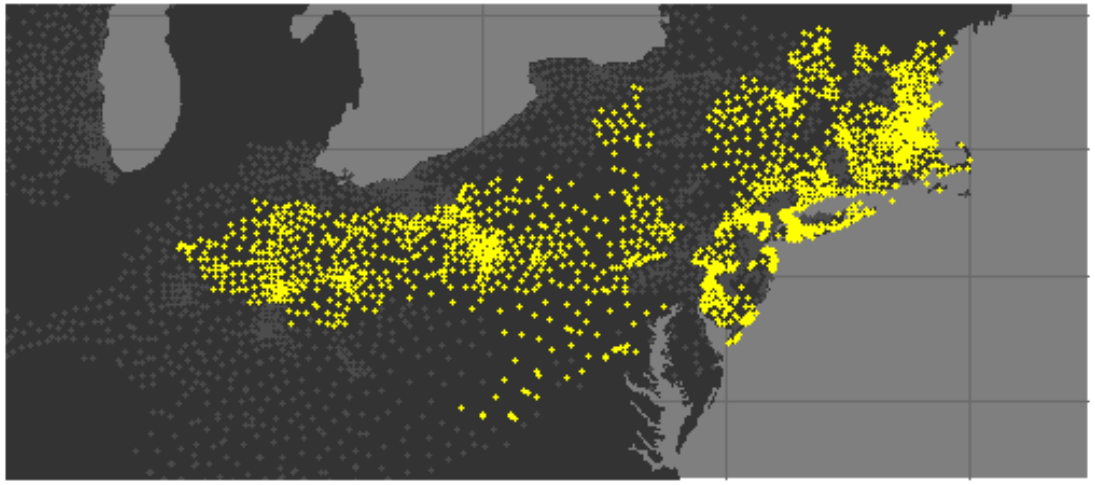

(4) ML Model

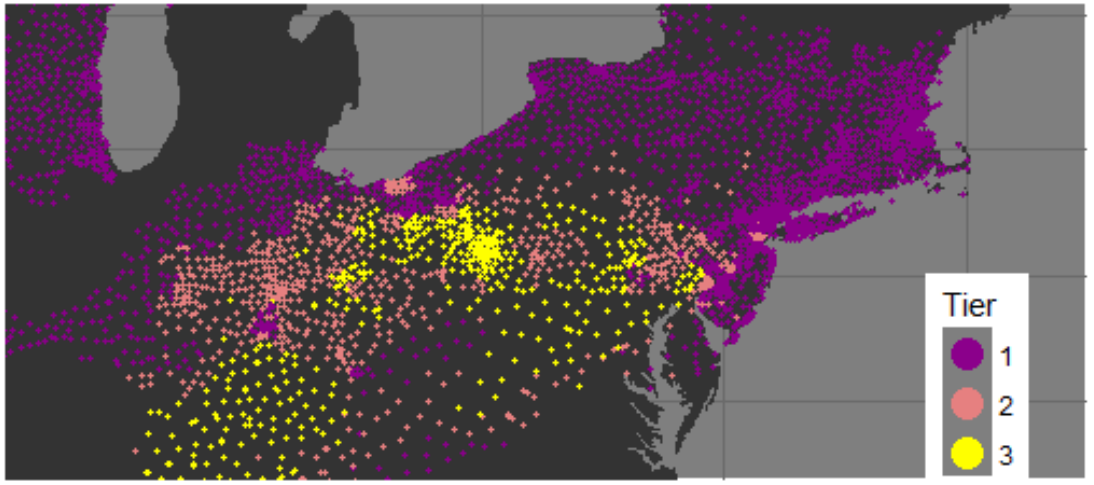

Day Two of Blizzard (January 27, 2015)

(1) Actual School Closures

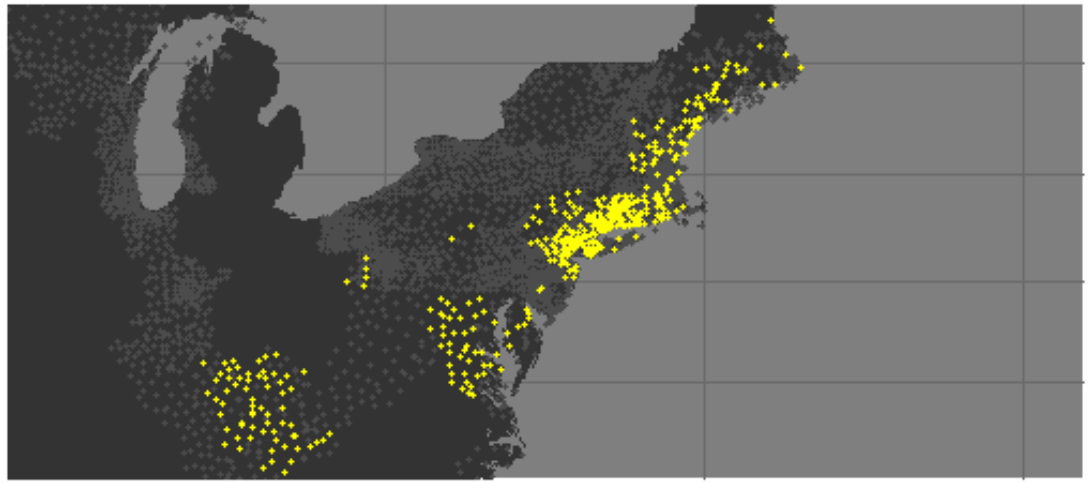

(2) Any Warning Notification

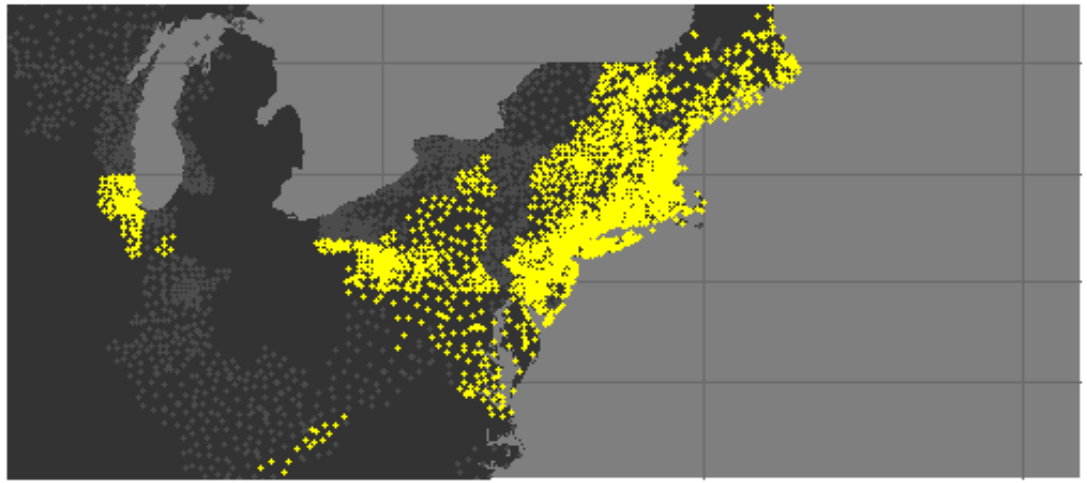

(3) Lagged Notification

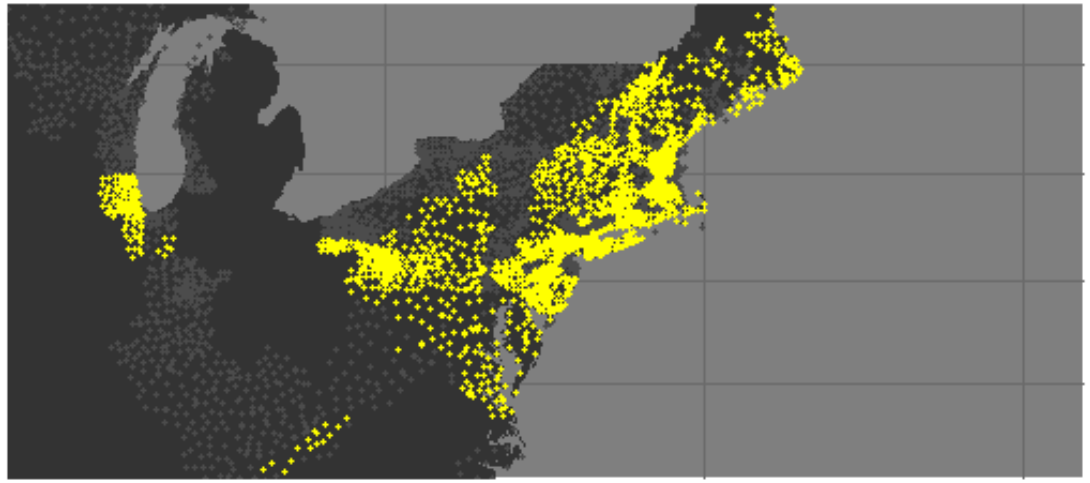

(4) ML Model

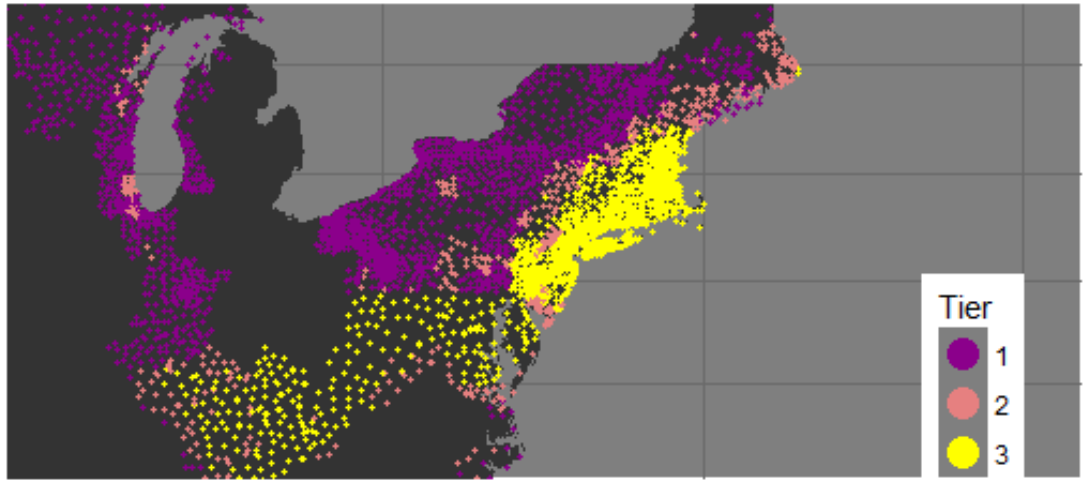

Wrapping up

- Possible to anticipate closures using weather information
- Technique could refine current weather warnings with a specific target audience
- Demonstrates that highly granular weather information can be linked to economic processes – useful for economic measurement.

Jeffrey.Chen@bea.gov