Law Enforcement Resource Allocations: AI solution.

How to provide meaningful crime predictions without bias using macro reporting data and AI

Steven Tran
Nicholas Van Bergen
Anand Ramakrishnan

The problem

Legal system

Person-based predictive models perpetuate systemic racism and are inherently biased.

- State and City level law enforcement have already experimented with presumed ‘objective’ AI generated predictions.

“Racism has always been about predicting, about making certain racial groups seem as if they are predisposed to do bad things and therefore justify controlling them”

-Dorothy E. Roberts. Penn Law

Problem statement

Can we use macro-level predictors that presume to alleviate biased foundations?

We think so.

Challenges deep-dive

Challenge 1

‘Broken Windows’

Attempts to heavily regulate small crimes to prevent larger crimes from happening.

Stop or avert small crimes from happening we get less big crimes.

Challenge 2

Predicting is already racist.

Models trained on demographics of the arrest record

Racism by proxy.

Challenge 3

Data availability

Need to be extremely cautious about data and the potential to be racist by proxy.

Moment of Silence

Thank You.

Inform

Columbia Stanford MIT Penn

Amplify

D4BL A.I. Now AJL AFP Chupadatos

Solution

Predict at a state level with ambiguous/unrelated data is a start

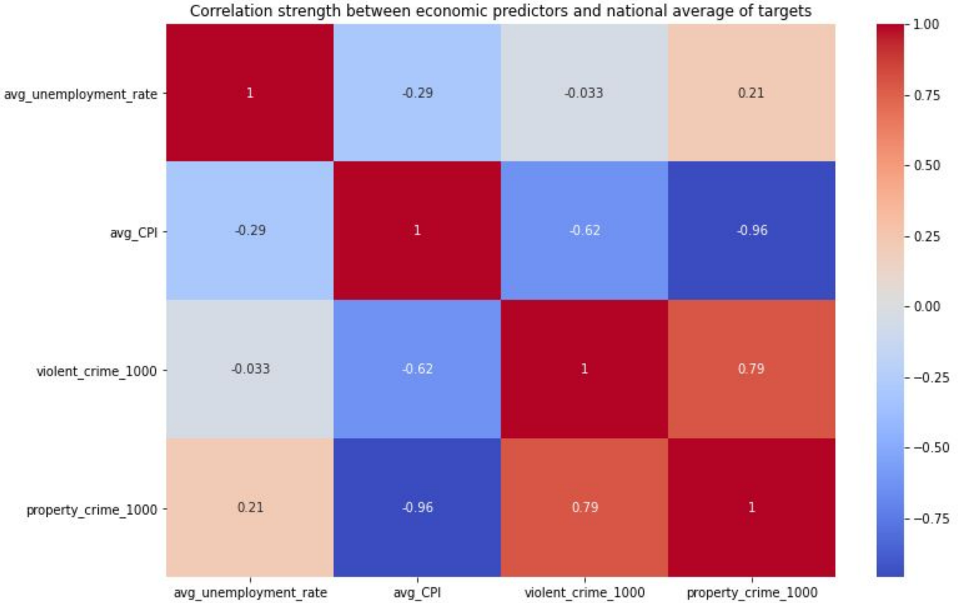

- Predictors:
  - Bureau of Labor Statistics
    - Annualized CPI
    - State Unemployment Rate
  - State Attorney General Political Affiliation
- Targets:
  - FBI crime database:
    - State Per Capita Violent Crime
    - State Per Capita Property Crime

Data Collection & Exploratory Data Analysis

Data Sources

Bureau of Labor Statistics Public Data API

Wikipedia & NAAG Web Scrape

F.B.I. Crime Data API

https://www.bls.gov/developers/

https://www.crummy.com/software/BeautifulSoup/

https://github.com/goldsmith/Wikipedia

https://crime-data-explorer.fr.cloud.gov/pages/docApi

Box Plot - Violent Crimes

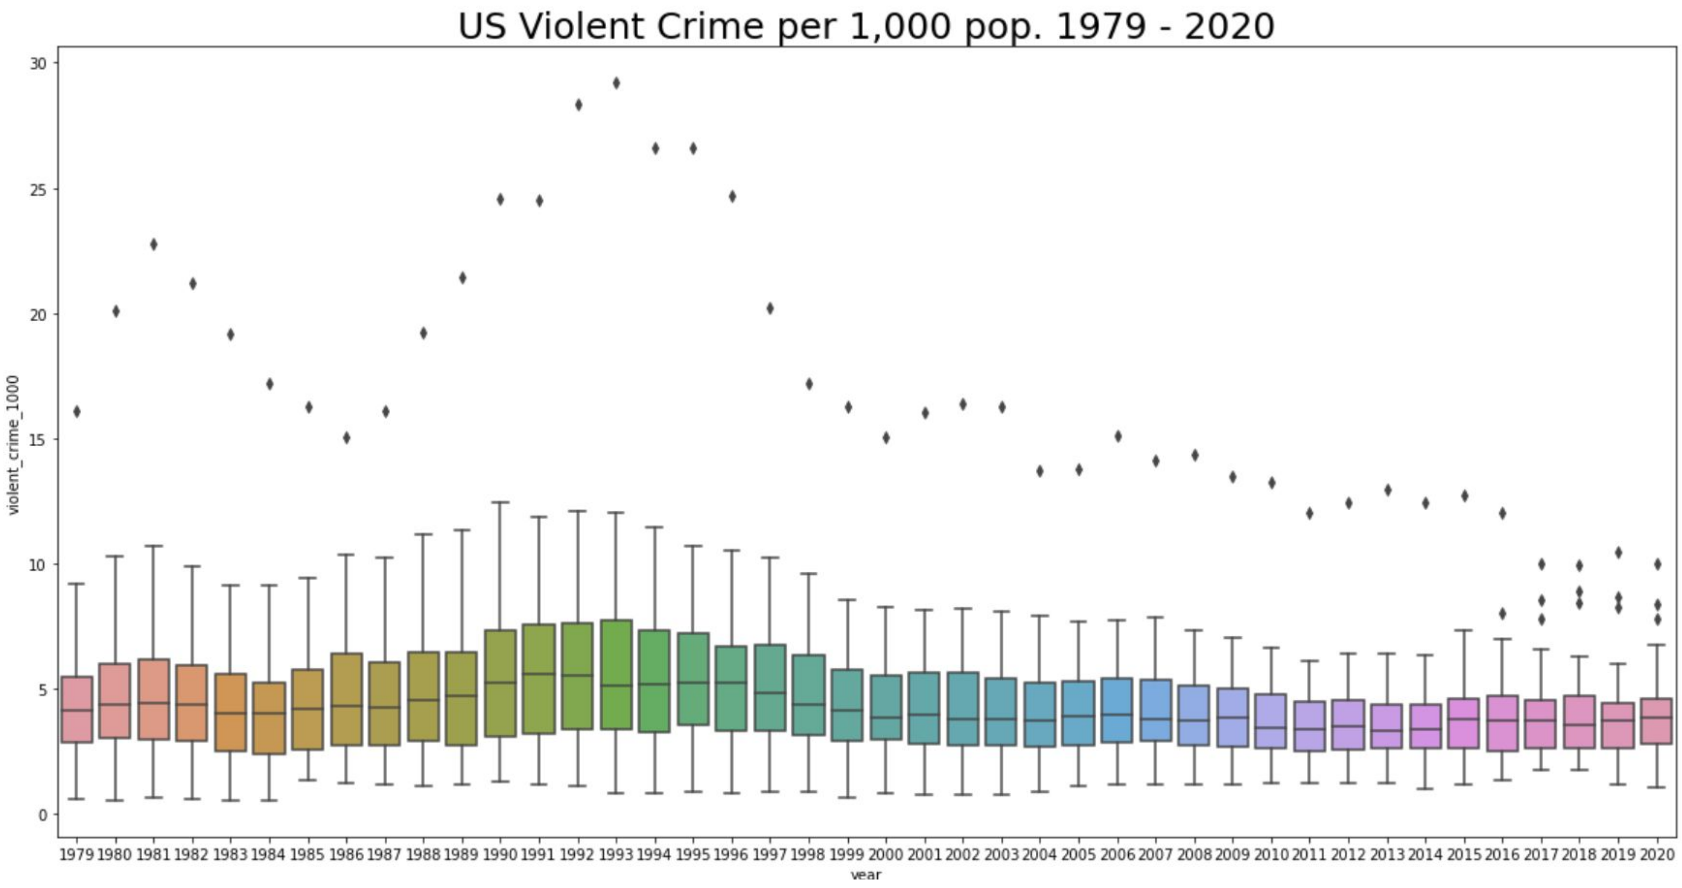

Box Plot - Property Crimes

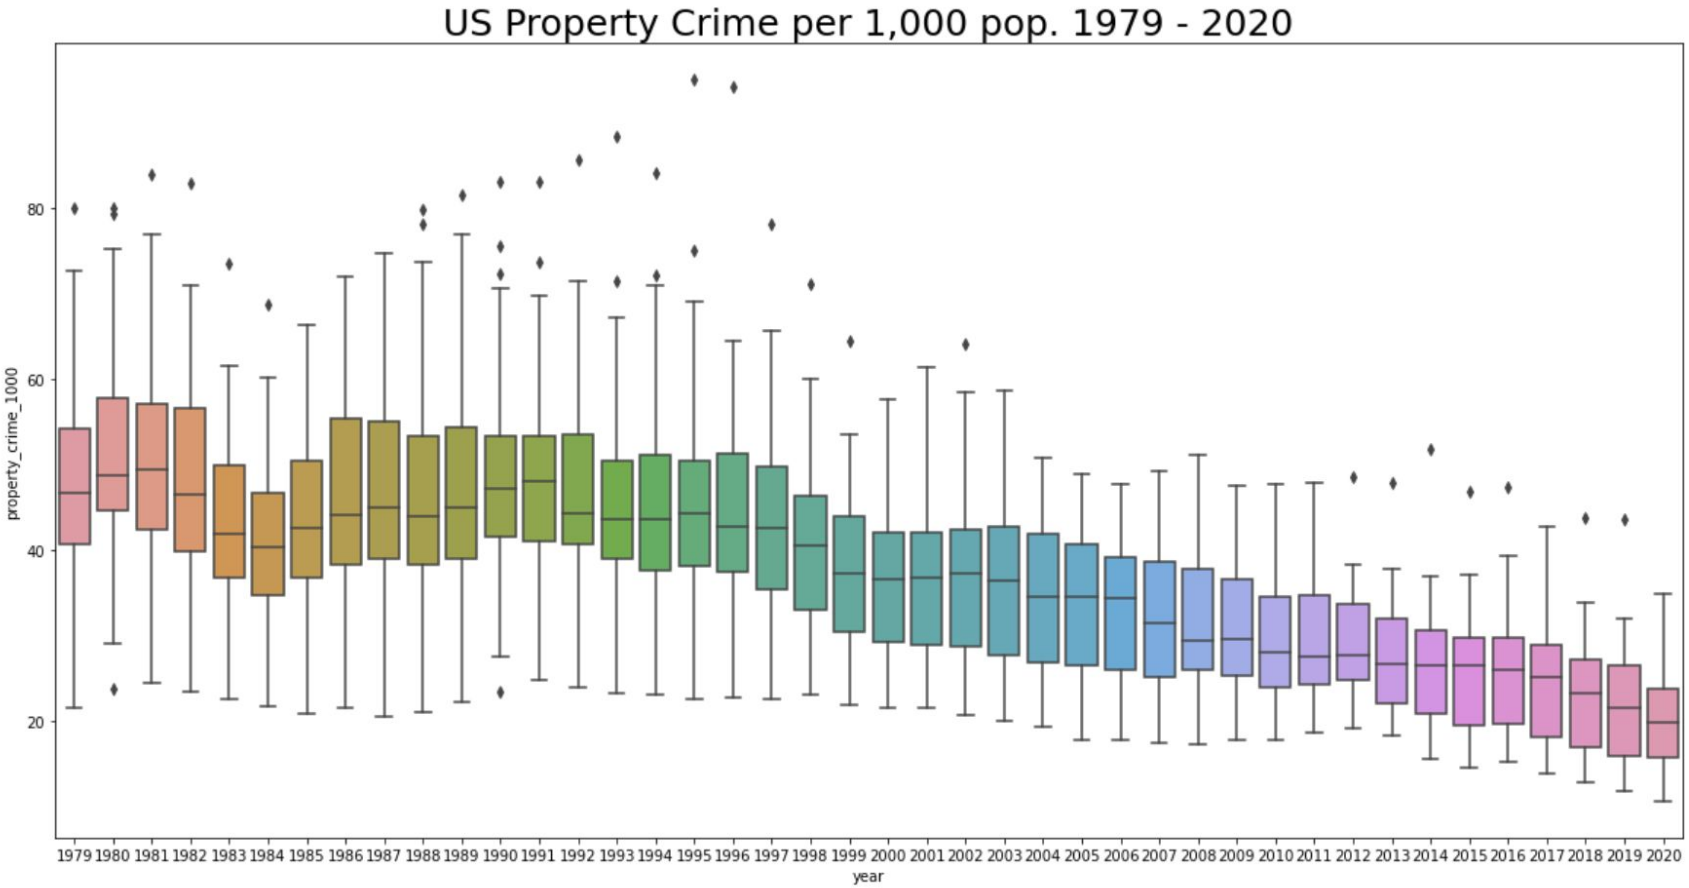

Seasonality?

The aggregated *'National'* data does not appear to be seasonal.

We took a closer look at State level data.

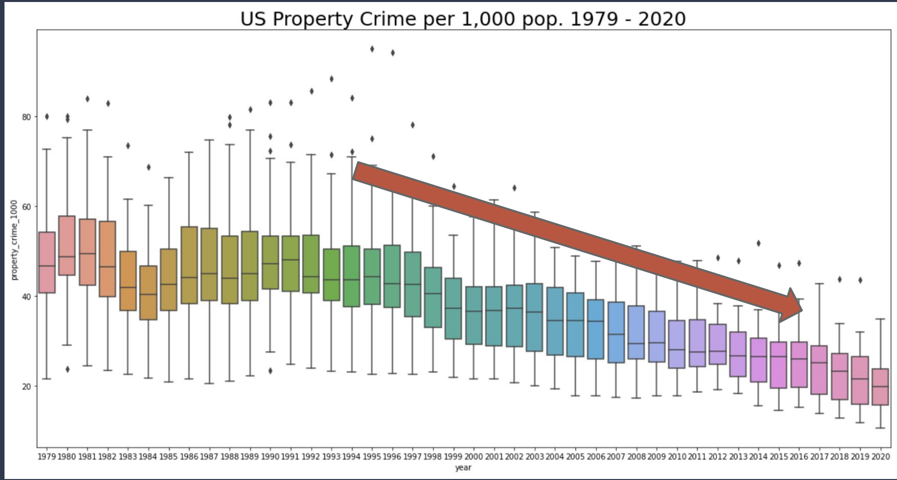

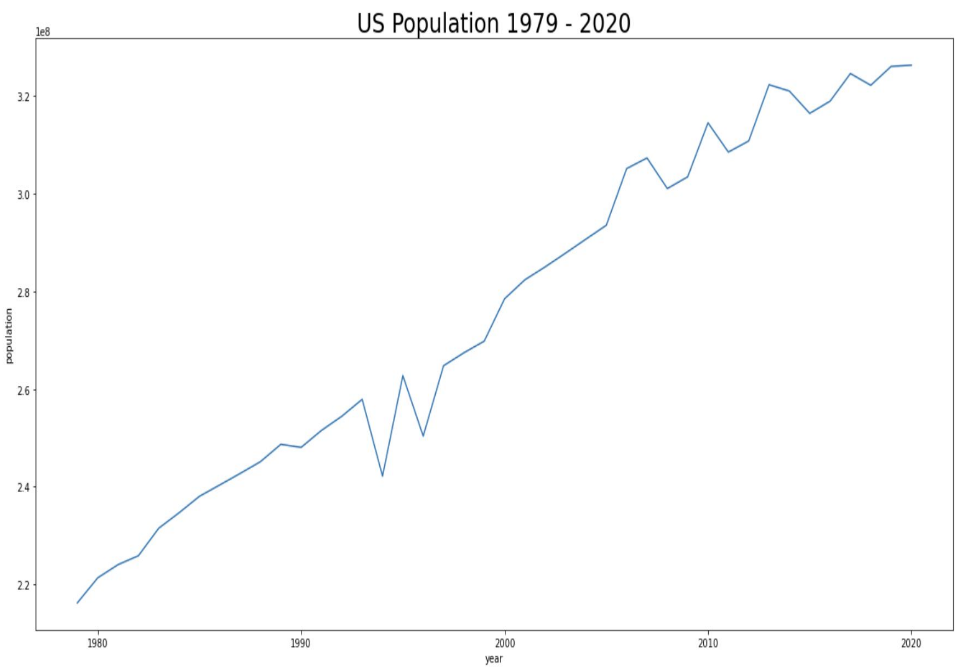

Consistent rise in population growth between 1970 and 2020. May contribute to the decline in Property Crime.

Least Violent State - Maine

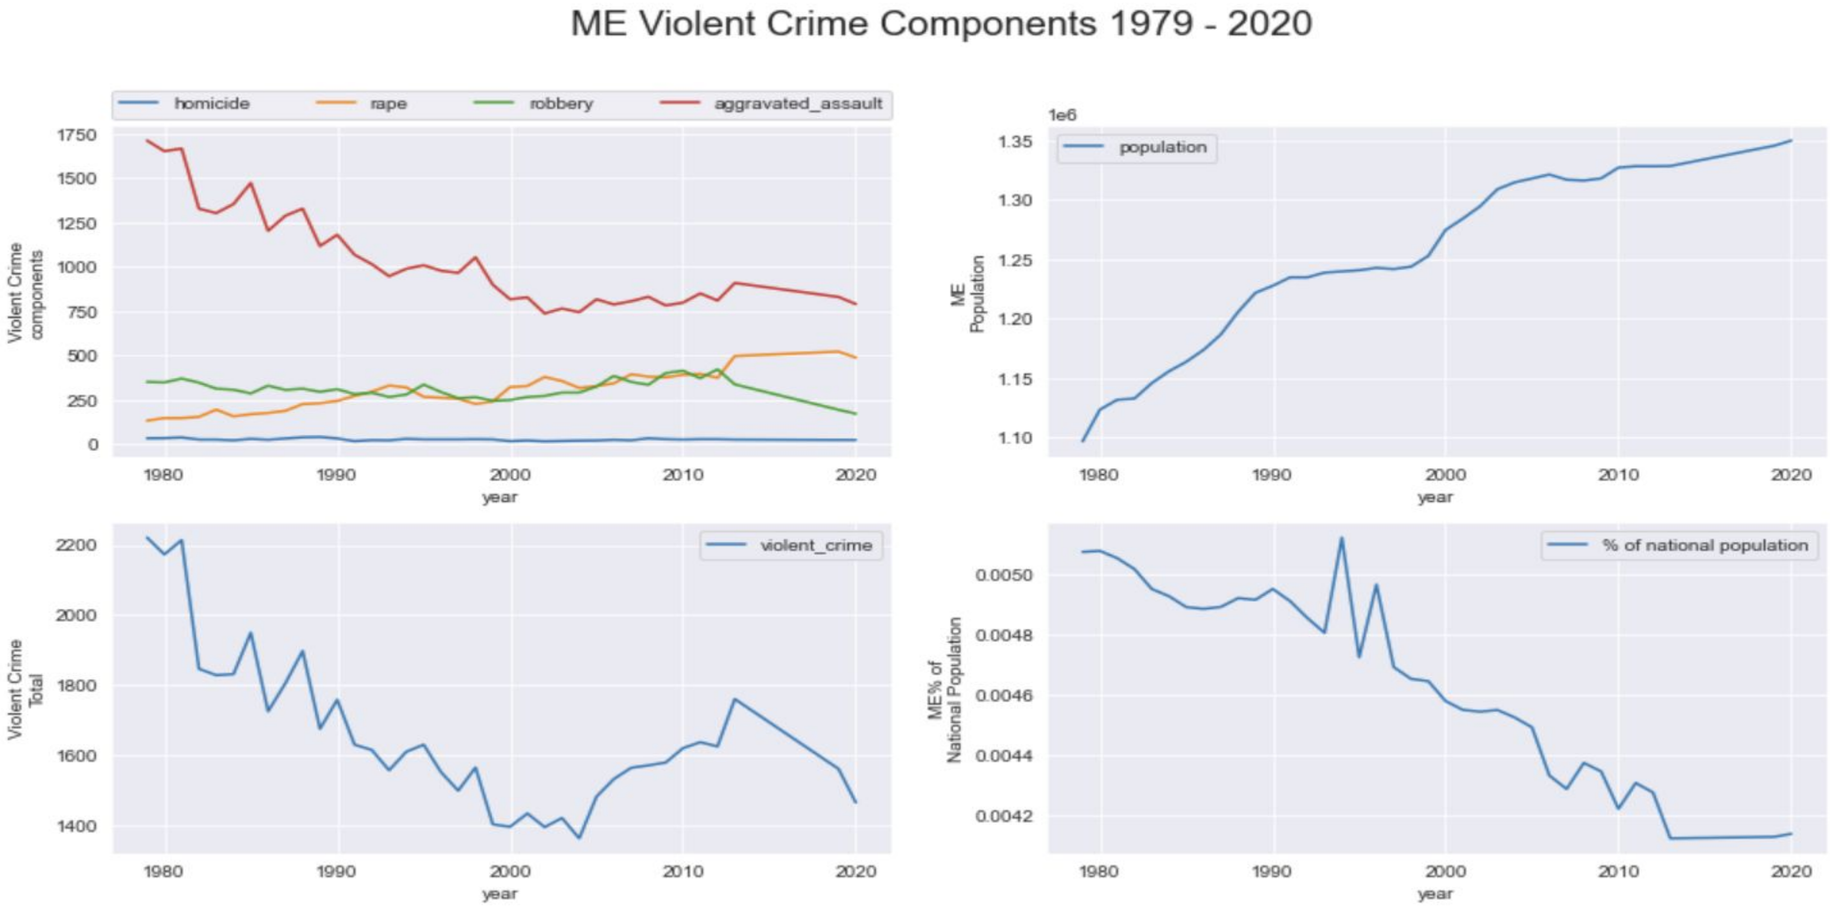

Least Property Crime State - Massachusetts

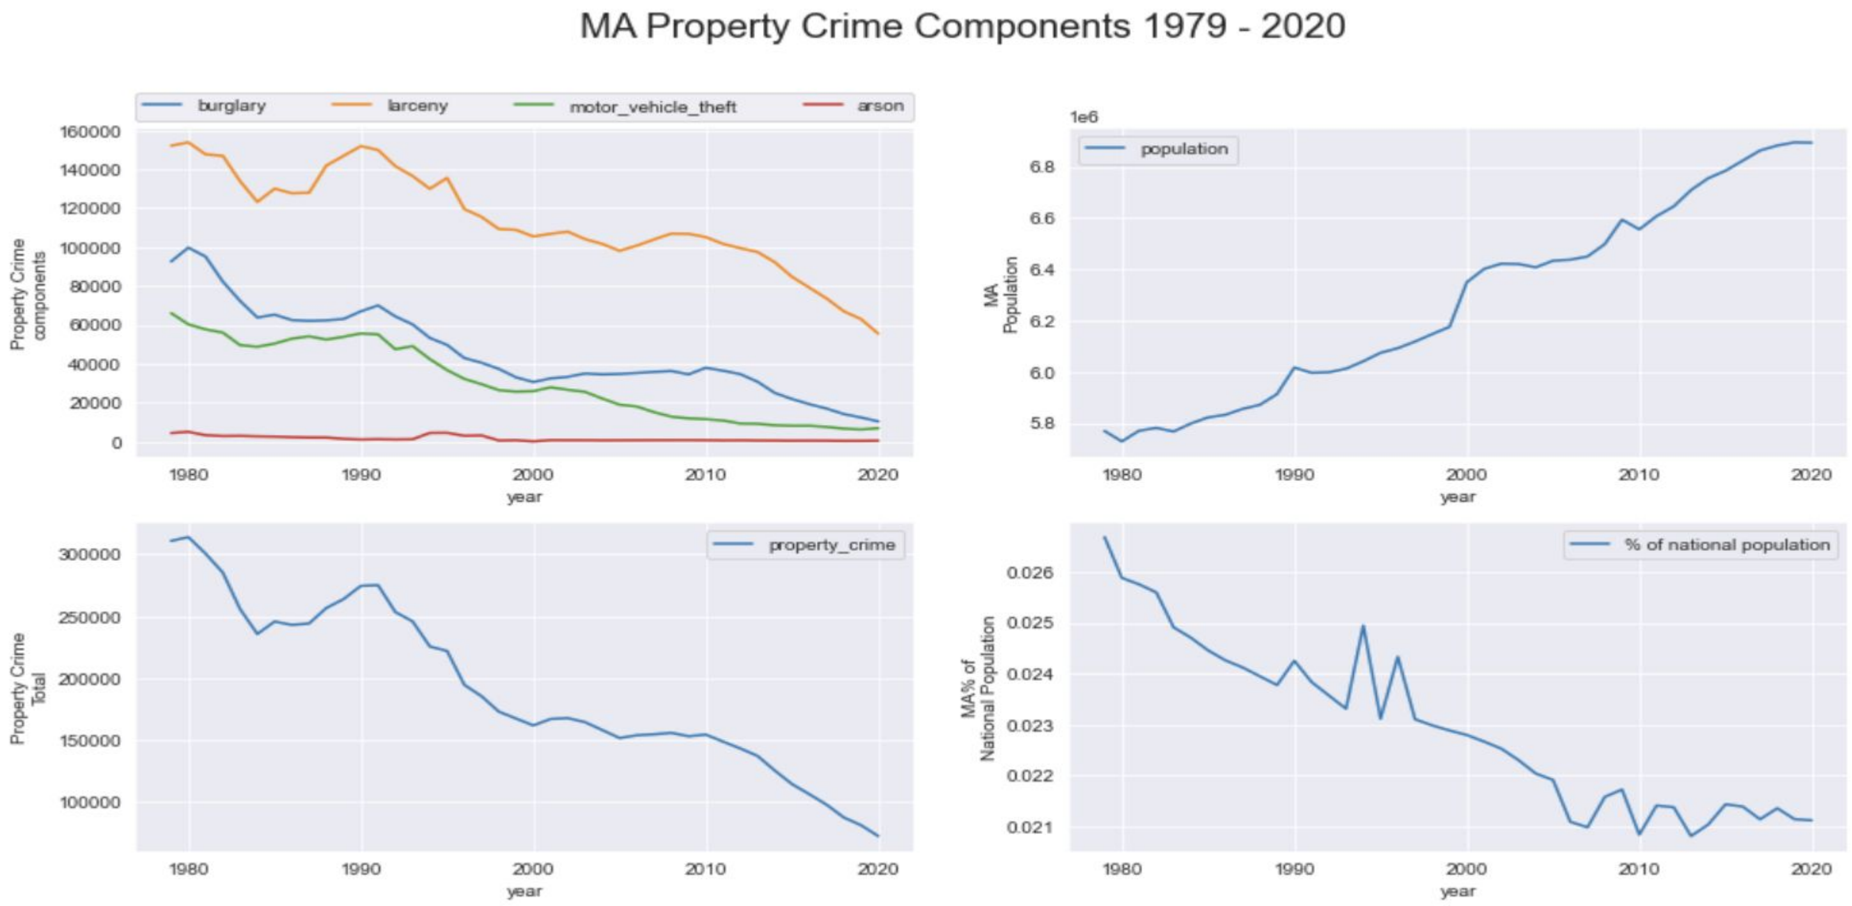

Most Violent State - Washington DC

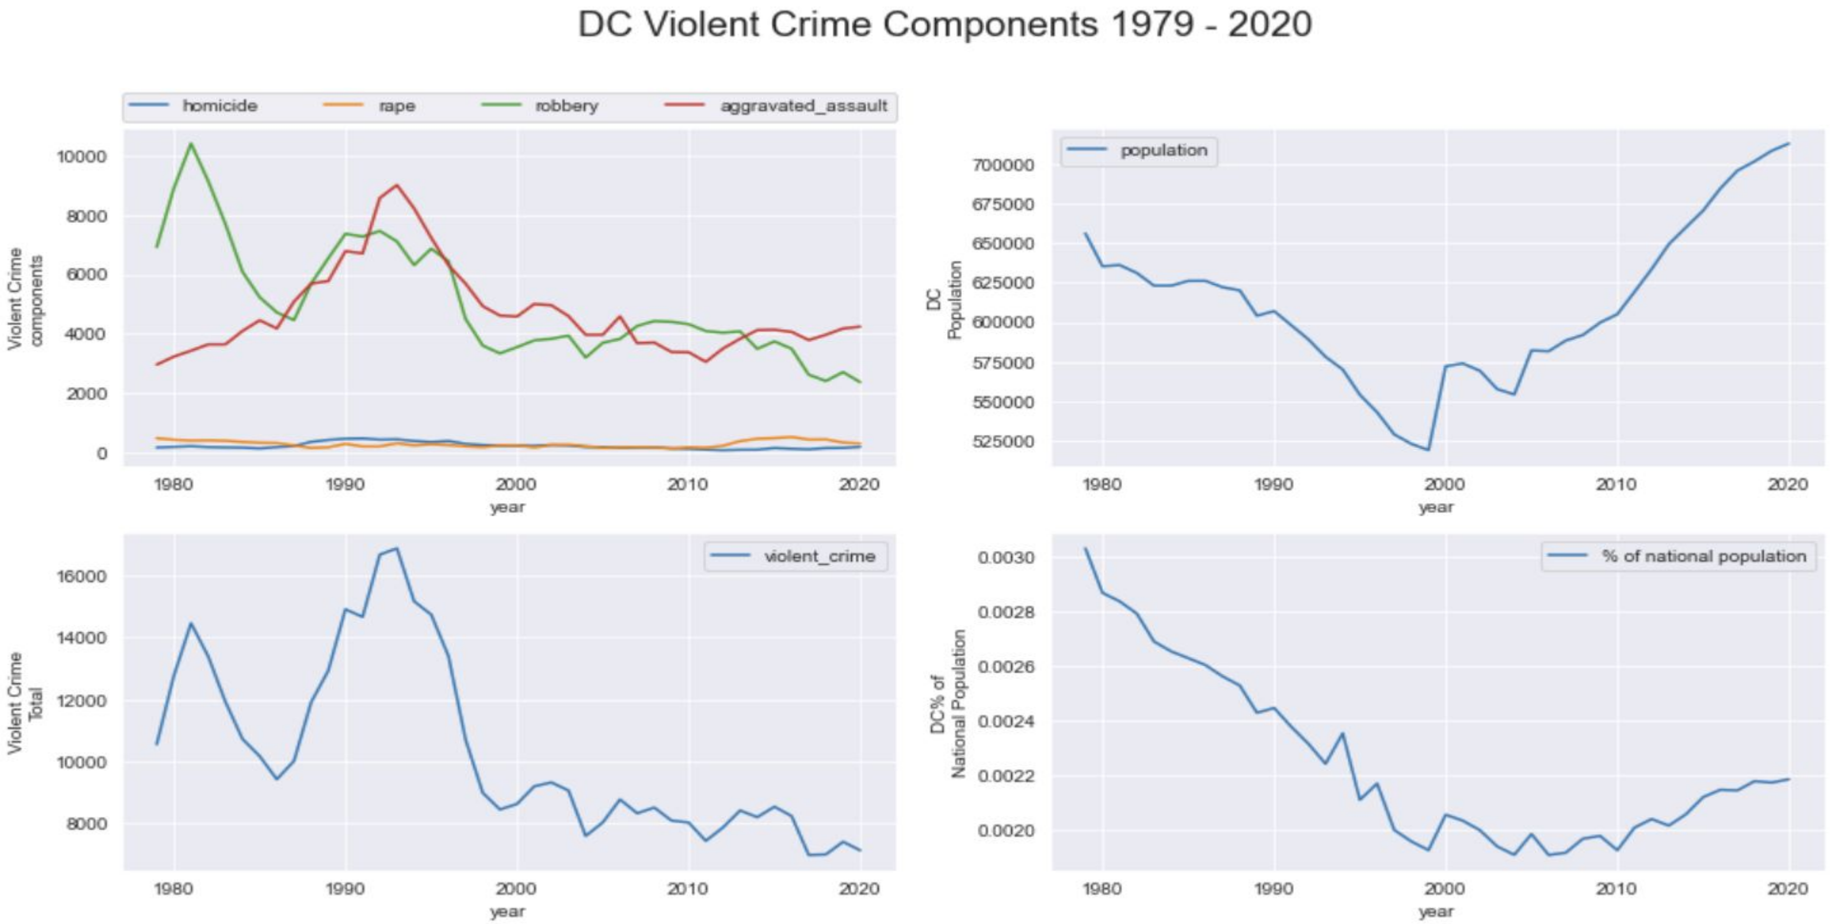

Most Property Crime State - Washington DC

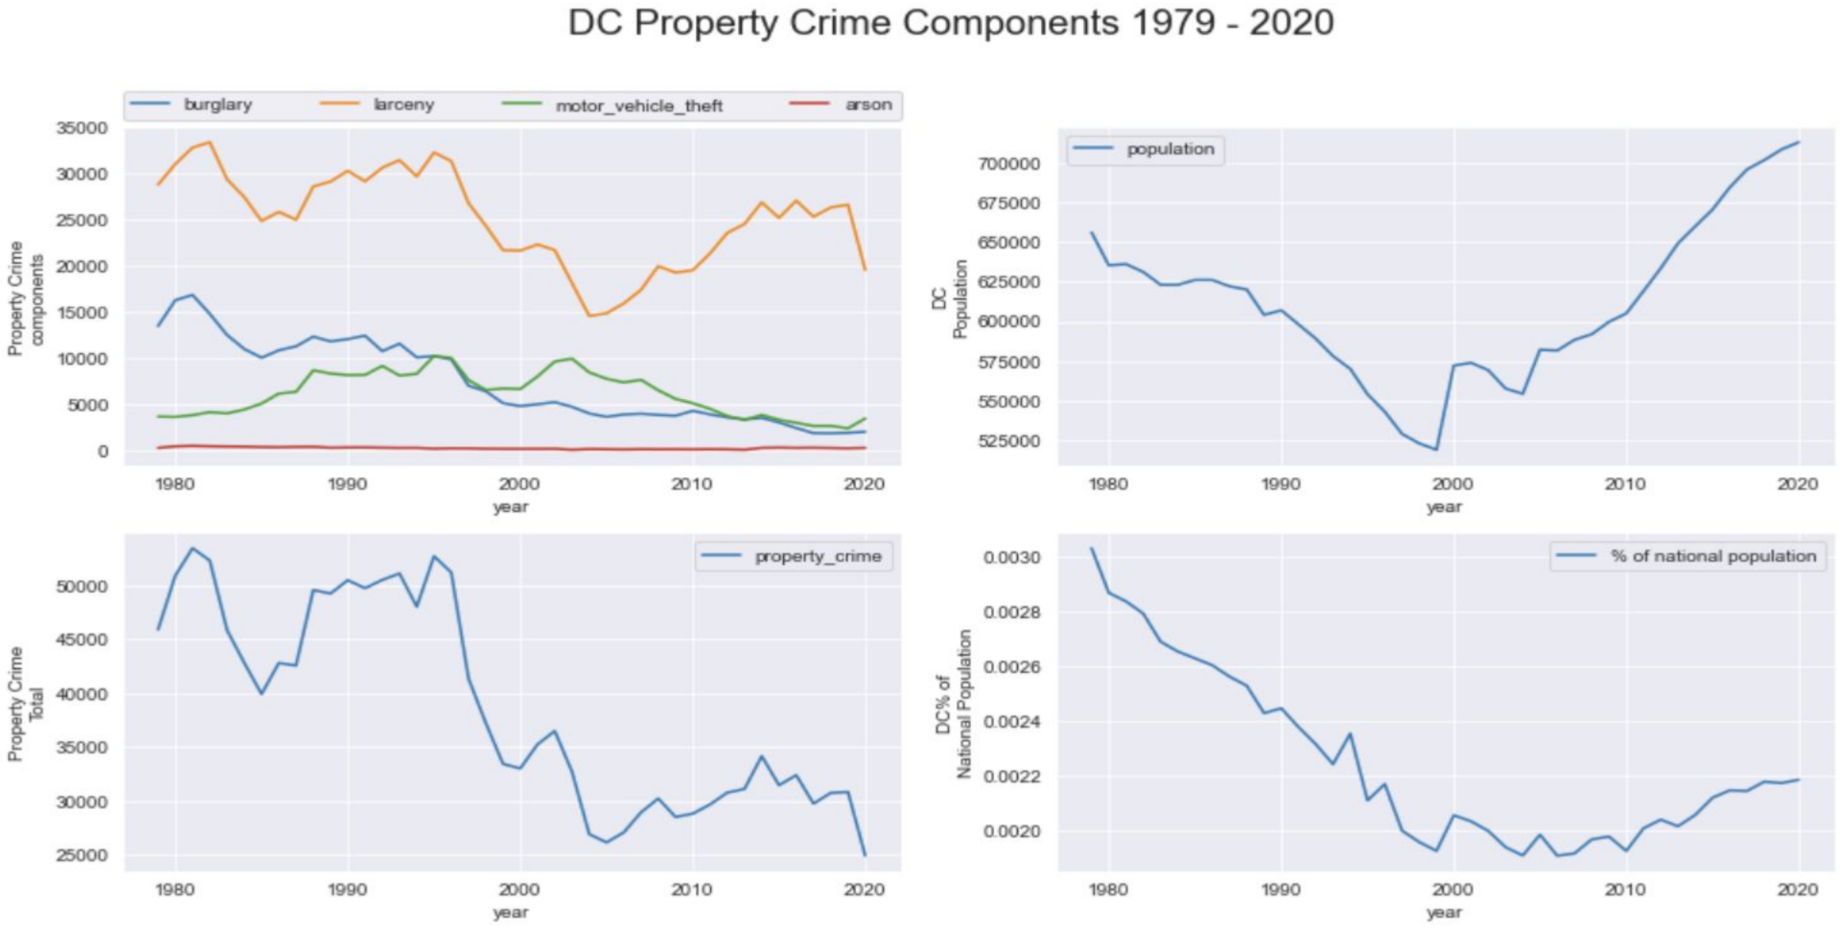

Rankings

Rankings can be difficult since the scales are not the same from one state to another.

Though, DC did have population growth and a general decline in property crime, they just had the most property crime relative to other states.

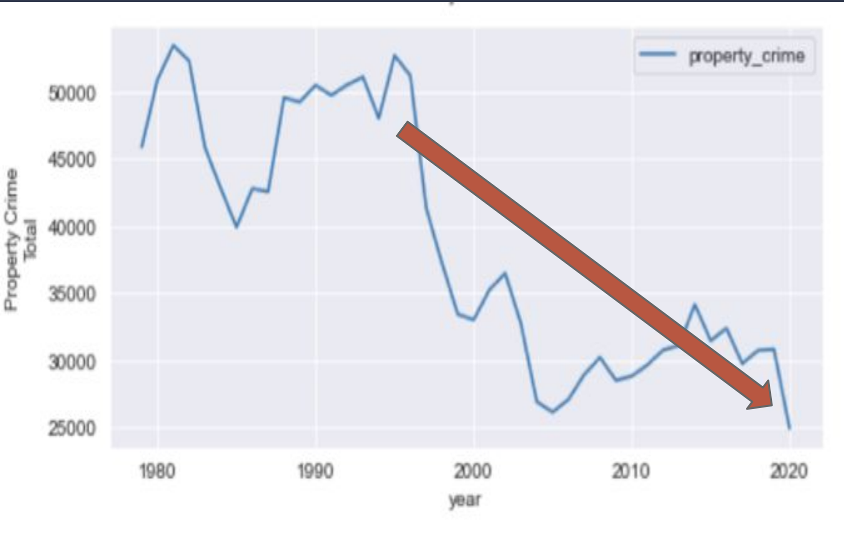

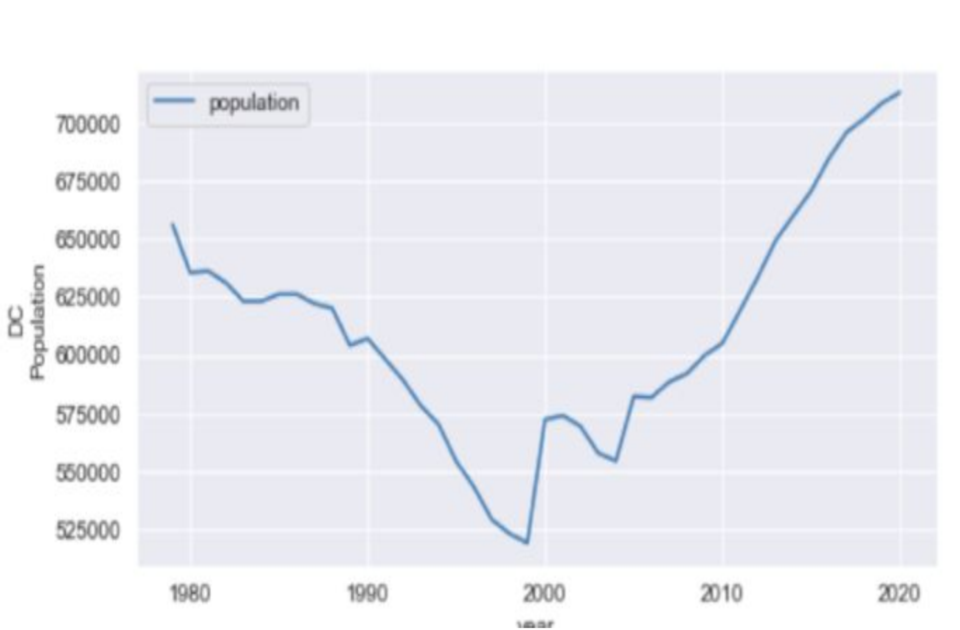

Consistent rise in population growth between 1970 and 2020. May contribute to the decline in Property Crime.

Stationarity

Violent Crime: States that exhibit stationarity

- 7.84%
- 1st difference: 82.35%
- 2nd difference: 94.12%

Property Crime: States that exhibit stationarity

- 0%
- 1st difference: 84.31%
- 2nd difference: 92.16%

Second Degree differencing stationarity states. Where p > than alpha.

Violent Crime States that lack stationarity

- Indiana
- Michigan
- Oregon

Property Crime States that lack stationarity:

- Iowa
- Michigan
- Tennessee
- West Virginia

Modeling

Model Considerations

ARIMA

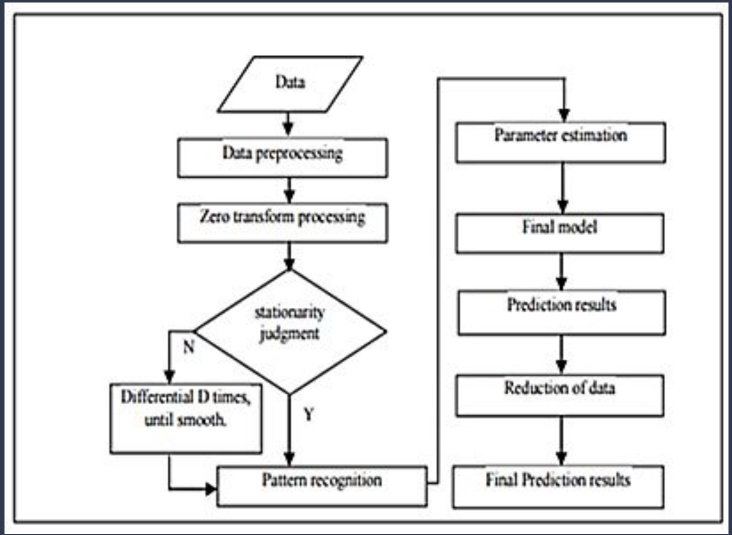

Flow Chart image Credit:
Delima, Allemar Jhone. (2019). Predicting Scholarship Grants Using Data Mining Techniques. International Journal of Machine Learning and Computing. 9. 513-519. 10.18178/ijmlc.2019.9.4.834.

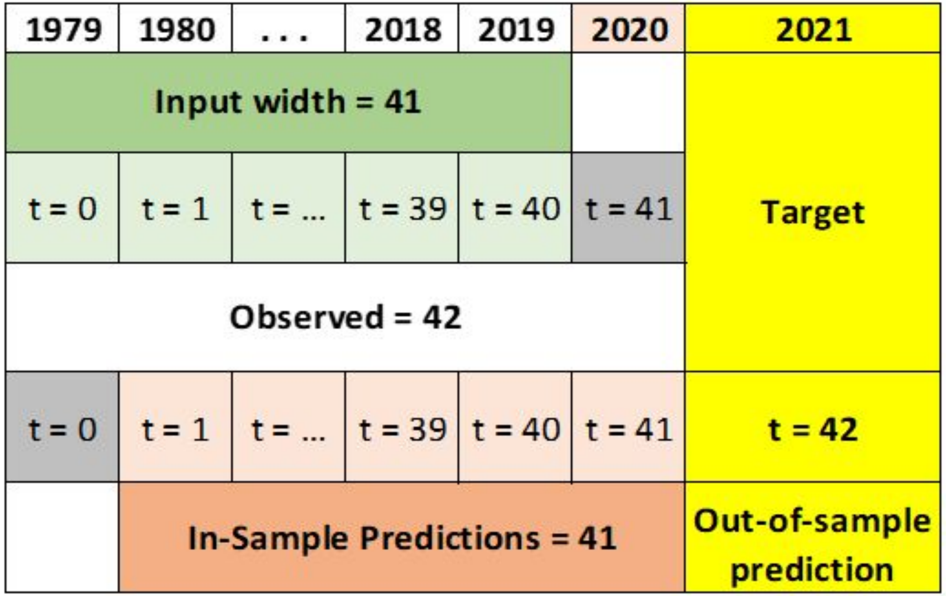

Recurrent Neural Net with LSTM (Long Short-Term Memory)

Each model was run 51 times for each target variable, for each state.

Performance

ARIMA

Parameters:

- Endogenous variable (violent or property crime)
- Exogenous variables (predictors)
- Best order (calculated using auto_arima)

Results:

Violent Crime States

- MAE avg: 0.369
- RMSE avg: 0.795
- R2 avg: -0.531

Property Crime States

- MAE avg: 2.243
- RMSE avg: 4.772
- R2 avg: 0.545

Recurrent Neural Network with Long Short-Term Memory

- Predictors: economic vars and political measure
- Too few observations per sample to do a true train and testing split
- LSTM was trained on predictors for 1979 through 2020.
- Model predicted targets for years 1980 through 2021
- Evaluation was measured by errors compared to observed 1980 - 2020

<table>
<tr><td>1979</td><td>1980</td><td>. . .</td><td>2018</td><td>2019</td><td>2020</td><td>2021</td></tr>
<tr><td colspan="5">Input width = 41</td><td></td><td rowspan="3">Target</td></tr>
<tr><td>t = 0</td><td>t = 1</td><td>t = ...</td><td>t = 39</td><td>t = 40</td><td>t = 41</td></tr>
<tr><td colspan="6">Observed = 42</td></tr>
<tr><td>t = 0</td><td>t = 1</td><td>t = ...</td><td>t = 39</td><td>t = 40</td><td>t = 41</td><td>t = 42</td></tr>
<tr><td></td><td colspan="5">In-Sample Predictions = 41</td><td>Out-of-sample prediction</td></tr>
</table>

LSTM Performance

LSTM Performance

A case where we outperformed baseline

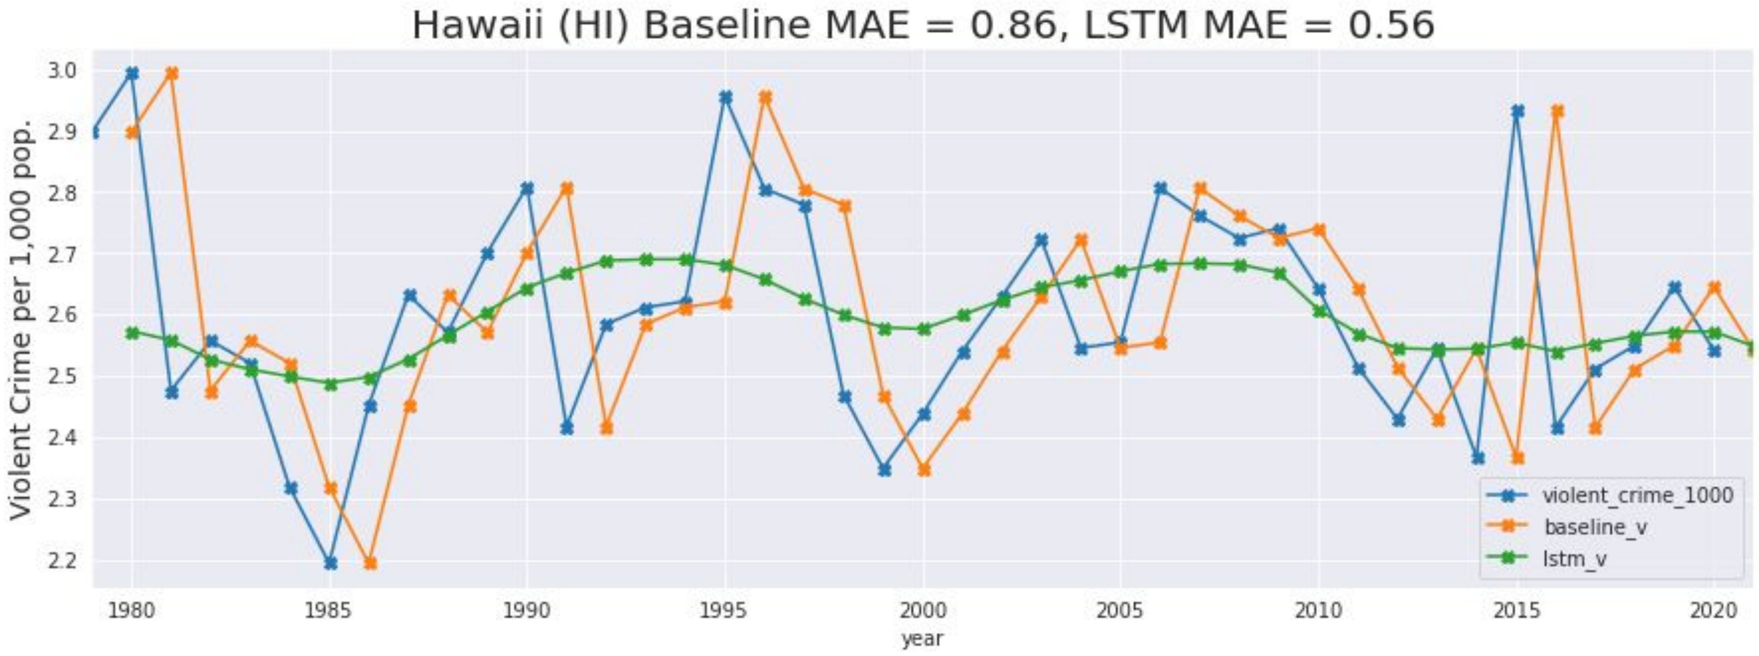

LSTM Performance

. . . and an example of our **worst** performance vs. baseline

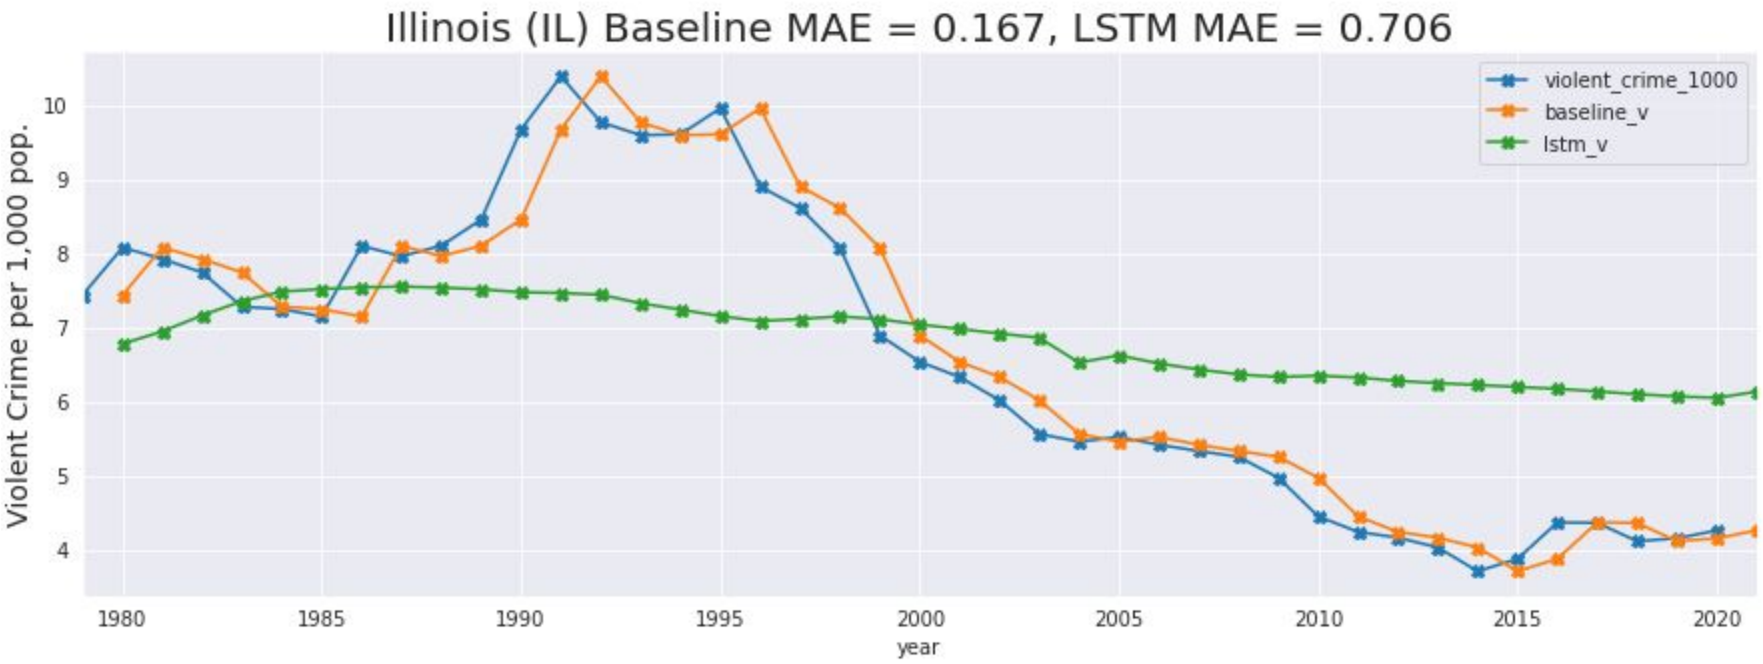

LSTM Performance

Our best in predicting property crime (still underperformed, though):

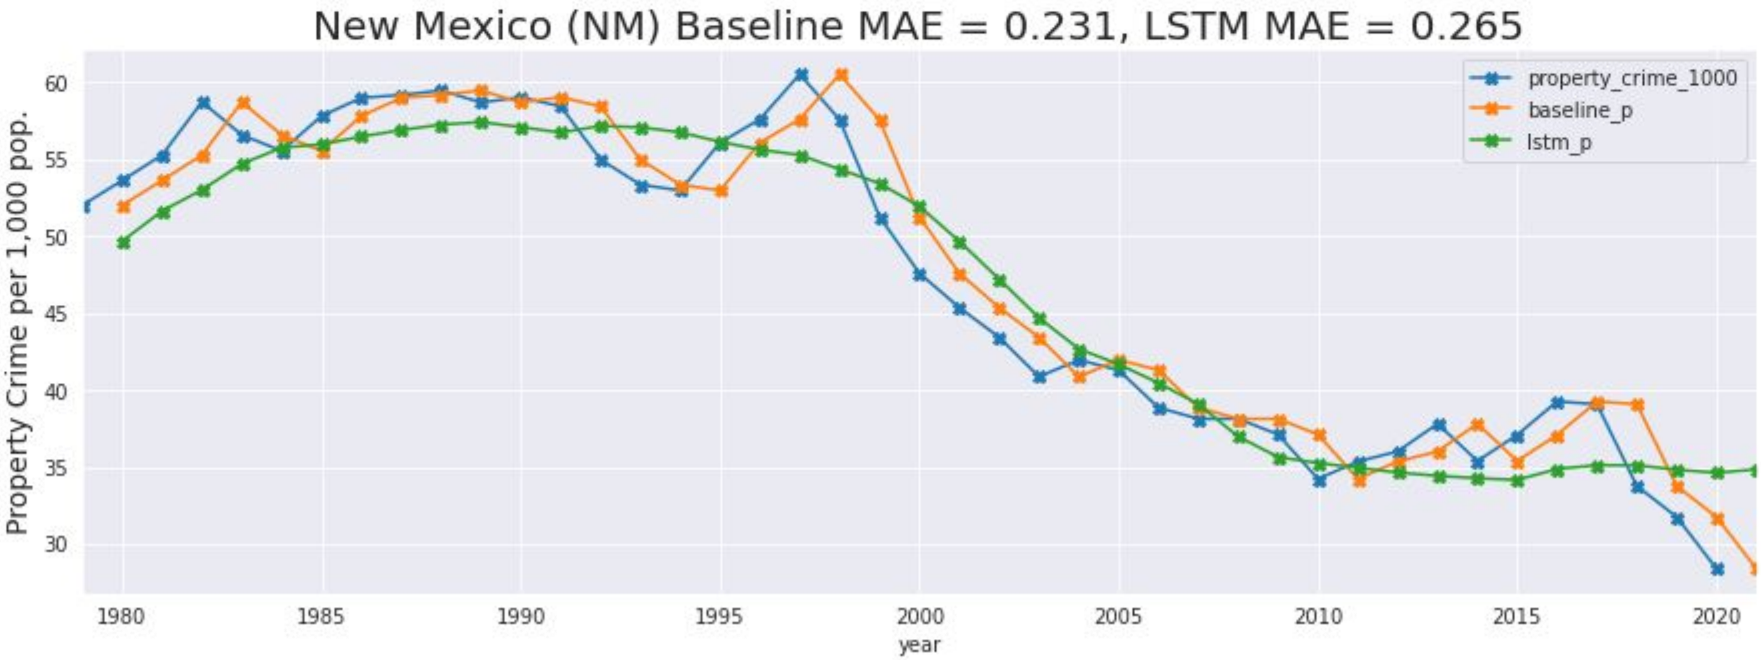

LSTM Performance

Like with violent crime, our worst-performing model for property crime was bad

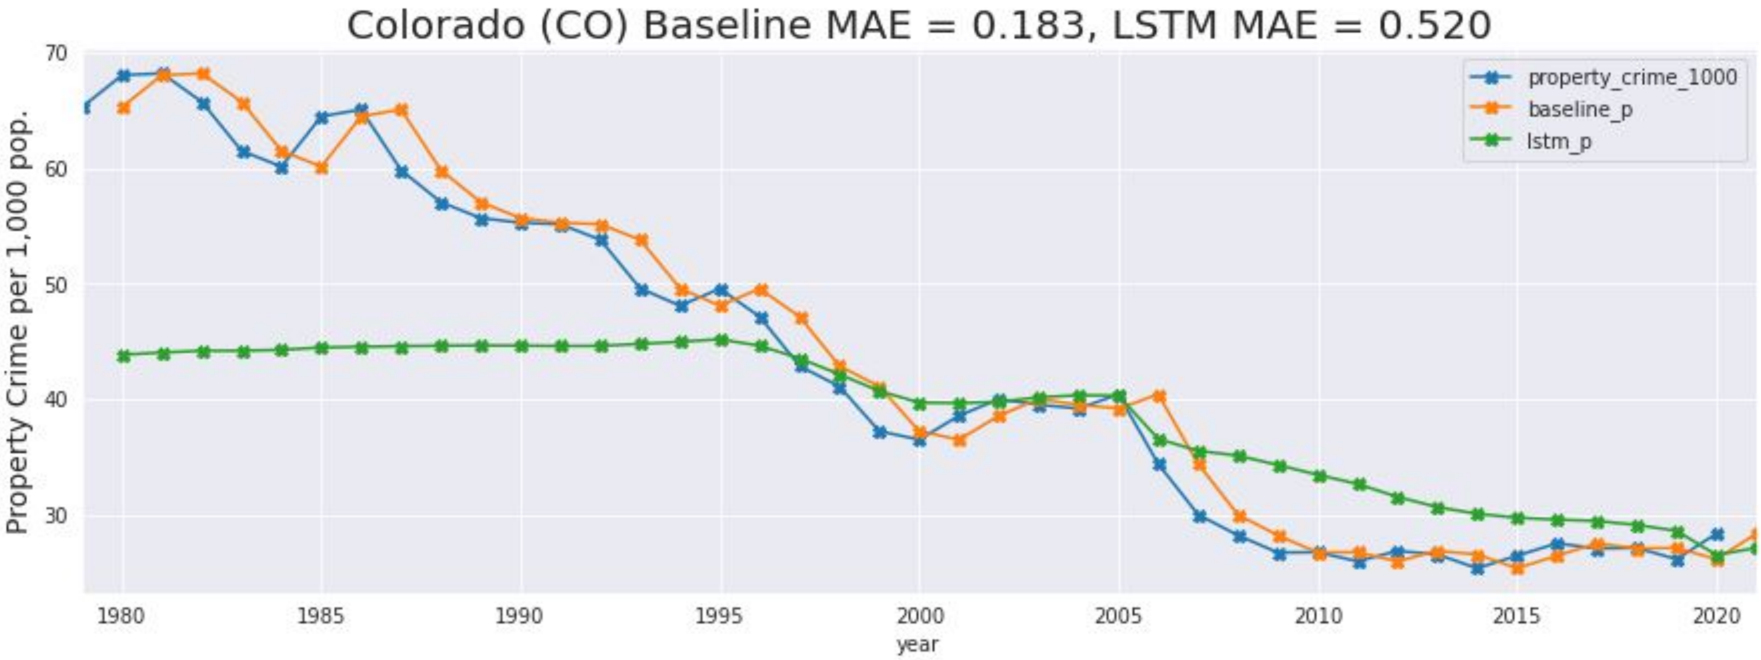

Recurrent Neural Network with Long Short-Term Memory

BRIGHT SPOTS

The RNN method had less stringent requirements for meeting classical modeling assumptions.

The baseline method was a high bar to beat with this kind of data. We think our model could do a lot better with more data to learn from!

LIMITATIONS

- As a black-box deep learning specification, it's difficult to tell which of the features conveyed the most importance to predicting crime rates.
- Limited frequency of data and small number of observations likely does not give the model enough to train on.
- A next attempt should incorporate more data, but perhaps also attempt an advanced technique like autoregressive recurrence.

Recommendations

Conclusions

- This is a very hard problem that has many systemic challenges with it.
- Hawaii might be able to actually use LSTM model.
- Main benefit of our tool presented here today is that we are able to absorb shocks and smooth our predictions as to how much resources are allocated.

We would want to find ‘the right data’ not just more of it.

We want to contribute to the fight against systemic racism in law enforcement and believe that AI could be ethically deployed in the future.

Conclusions

Use Cases:

- State government officials (Treasurer, Governor, Attorney General...) Could benefit to start the resource allocation/policy process.
- Government watch groups can use this public tool to hold elected officials accountable.
- Can expand to include a dashboard

- Could be expanded to include targets or co-targets.
- We could use more training instances to better tune a model.

Future expansion: Dashboard

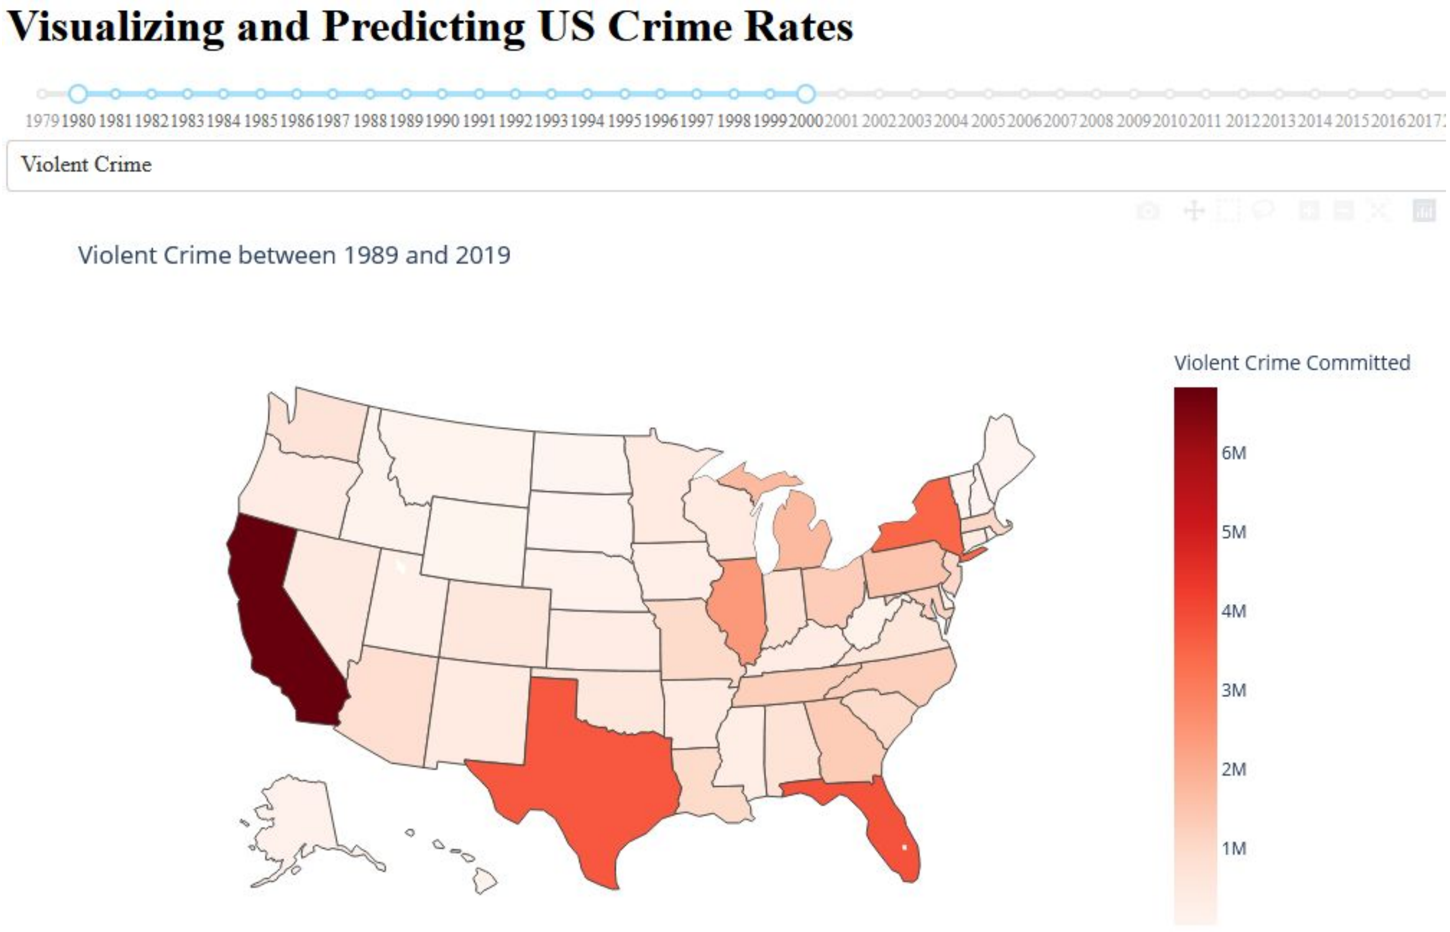

Q/A

Thank You