University of Colorado Denver

College of Liberal Arts and Sciences

Vulnerable Populations and Prejudice Propagation:
A Reinforcement Learning Model

Mathematical and Statistical Sciences, University of Colorado Denver

Introduction

- A core tenant of the United States criminal justice system is fair & impartial treatment of all persons
- Deep learning, algorithms based on neural networks, have revolutionized prediction & classification of a variety of data sets, including image & auditory recognition
- Predictive policing, deep learning applied to law enforcement, can determine patterns in crime & weight the probability of future offenses
- While deep learning algorithms are designed to be free of social bias, they can still pick up on its presence in the underlying data & produce skewed results

Motivation

- There are allegations of social bias in predictive policing including concerns about algorithms used to identify potential reoffenders
- These algorithms rank people on sets of questions designed to identify recidivism risk
- None of the questions involve race or ethnicity & the developers have taken steps to account for hidden correlations that might produce skewed results
- Yet there is evidence that these algorithms unfairly rank black & Latino offenders more harshly than white offenders

Reinforcement Learning

- This project uses a form of deep learning called reinforcement learning (RL) to explore how such algorithms deal with underlying social bias
- The RL algorithm learns by making an action & analyzing the result, much like a child learns by exploring the world
- RL is used to train AIs in video games but its adaptability has generated great interest in applying it to complex social problems
- An RL algorithm interacts with its environment & uses the data it collects to continuously update its strategies & policies
- The RL algorithm can be tuned for learning speed, advanced thinking, and exploration (called "epsilon greedy", this parameter is the ratio of random action to ideal action taken)

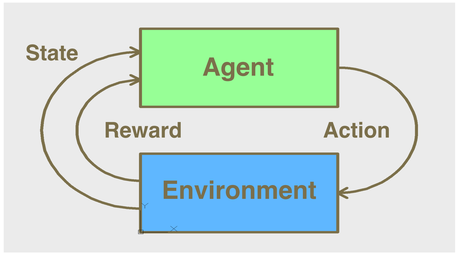

The Population

- A data set containing 1000 synthetic persons was generated
- The people were randomly assigned to two groups, each group was equal in criminality (the person's value as a target) & vulnerability (the person's ability to recover from negative encounters)
- The decision to detain was determined only by a person's suspiciousness level, which was loosely correlated to criminality

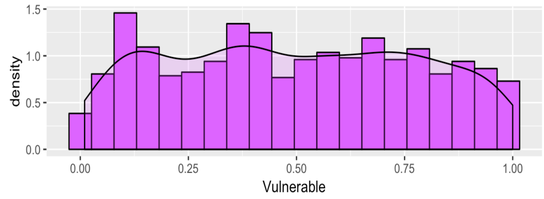

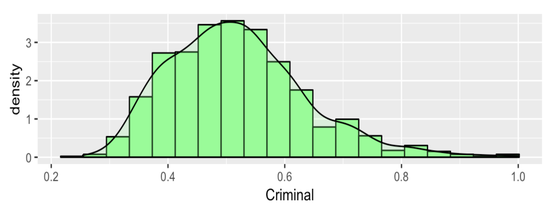

- To model the effect of systematic bias on a group, one set of data added a uniform weight to the target group's suspiciousness level

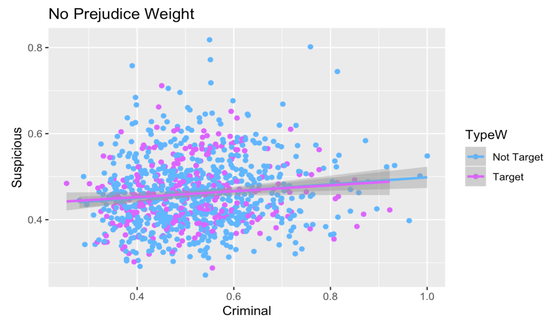

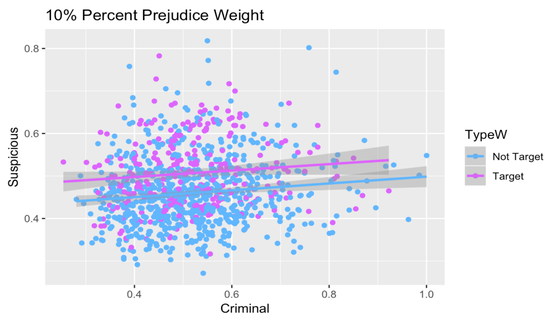

The Environment & Actions

- The environment is filled with a random sample from the population each model day
- At each time step, within each model day, the RL algorithm decides whether to move to a new square or, if a suspect is present, detain them
- A detainment results in a reward equal to the suspect's criminality but at a cost of several time steps so the RL algorithm must carefully weigh tradeoffs

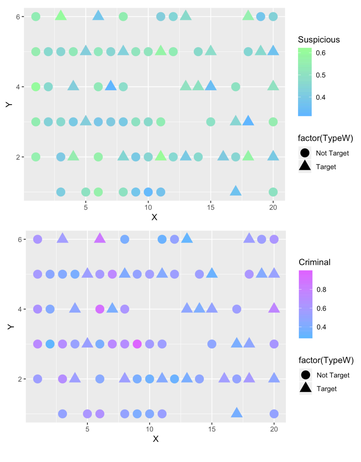

Preliminary Results

- RL applied to the biased and unbiased data sets produced different learning curves
- Both RL algorithms showed sustained improvement over their initial model day results, which were low
- Mean, median, & standard deviation were similar for tests run on both data sets lasting 50+ model days

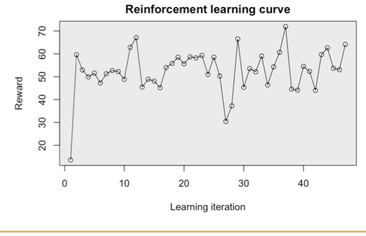

Conclusions & Future Work

- This model is the first stage in ongoing research into social bias propagation in deep learning algorithms
- The model shows promise, but numerous changes will need to be made before it is useful for determining social bias
- The model, as yet, does not return data on an ideal model day nor a list of targets by type, making comparison of the two RL algorithms' choices difficult
- Q learning, RL that does not need a state based environment, will be explored as an alternative
- The effect of feedback loops, where current actions influence future states will be modeled & analyzed
- Strategies will be developed to identifying & minimizing social bias in deep learning
- In addition to criminal justice, applications include improved education & health care deep learning tools

Primary References

- Sutton, R.S. & Barto, A.G. 1998. *Reinforcement Learning*. MIT Press.
- Aggarwal, C.C. *Neural Networks and Deep Learning*. Springer.
- Angwin,J., Larson, J., Mattu, S. & Kirchner, L. (2016, May 23). *Machine Bias*. Retrieved from https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing

For Further Information