

Risk-Assessment: An Art or Science?

Predicting Recidivism at the Time of Sentencing



Jeremy Kim and Miguel Camacho-Horvitz, CS229 2015 Stanford University

Motivation

As a society, we seek to reduce crime. One important facet of this complex goal is the minimization of prisoner recidivism, or a relapse in criminal behavior following release.

One particularly controversial strategy – indeed one that courts in Pennsylvania may soon utilize – would incorporate risk of recidivism as a factor involved determining **“appropriate” incarceration sentences**. This raises both moral as well as statistical questions. For the former, we might ask whether criminals can be sentenced based on what some call “future crimes” as their risk is determined by the behavior of people in some senses “similar” to them in the past?

As far as the latter, **can we even accurately predict which convicted criminals will commit another crime?** For many, the answer to the former will depend on the answer to the latter - the degree of certainty with which we can make predictions on the likelihood of future criminal activity is critical. Thus, in this project, we hope to address the **accuracy of recidivism prediction**.

Objectives

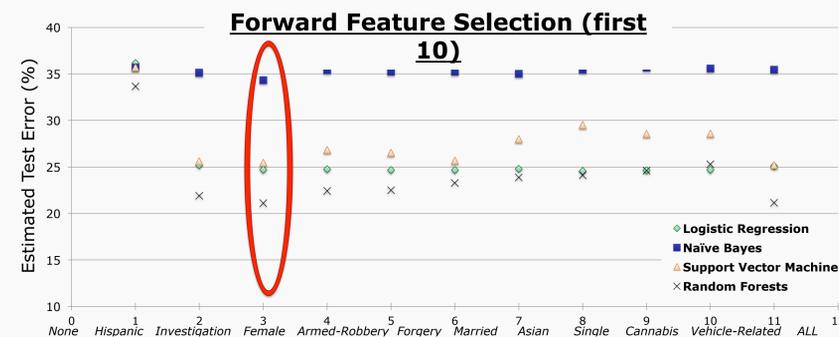
- Develop predictive models that will learn from past criminals and classify (using supervised learning) new criminals into two groups
 - “Positive” Subset – Predict they will reoffend
 - “Negative” Subset – Predict they will not reoffend
- Two main goals:
 - Infer which features are most predictive of recidivism.
 - Optimize for as accurate a predictor as possible, understanding the gravity of false positives.

Data

- Data received from the Bureau of Justice Statistics.
 - Tracked Convicted Felons across the United States from 1986-1989.
 - The dataset included demographic information (race, age, gender) as well as personal history (employment, housing, past crimes).
- For non-ordered, categorical data, we split the feature into many binary features (e.g. type of crime became “was murder”, “was theft”, etc.).
- Finally, the dataset had a binary value corresponding to whether that person recommitted a crime or not.

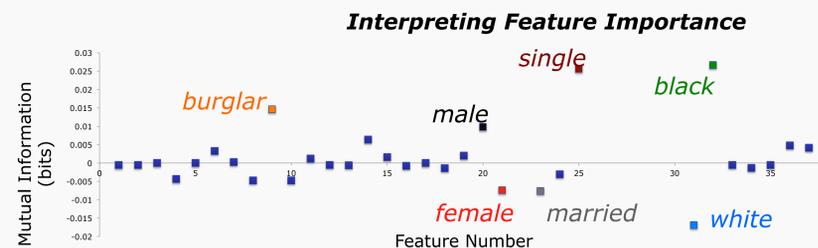
Model Selection

Using a forward selection algorithm on our initial 38-feature data set we found that after adding the top three features, the drop in expected test error rate was minimal. Moreover, the first three features included were not the features we expected.



Inference

- Analyzing the **Mutual Information** statistic for each feature.
- Many interesting dichotomies exist.
 - Black, single, and male were all held significant, “positive” information on recidivism while there “counterparts” of white, married, and female held significant, “negative information.”



- One non-intuitive, certainly not hypothesized important feature was the whether the conviction was due to burglary.
 - Burglary held more information on future criminal activity than did murder, rape, and other (typically considered) more heinous crimes.

Prediction

Baselines: We established some baseline predictions to compare our models against

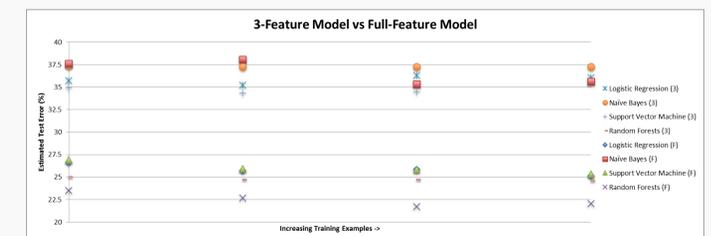
Error if you assume everyone recommit	Assume no-one recommit	Assume everyone under 30 recommit	Assume all violent criminals recommit	Assume all with previous records recommit
87.69 %	12.31%	56.49%	30.10%	24.67%

*The fact that the data was so skewed towards not recommitting proved to be a challenge

Models Used: Multinomial Naïve Bayes, Logistic Regression, Support Vector Machine, Random Forest Classifier

Results: With cross-validation, the best predictors we could get were barely better than the null-hypothesis (no-one recommit)

MNB	LR	SVM	RFC
61.59%	12.28%	12.27%	12.28%



Conclusions

- With the limited data that we have in our set, we could not predict recidivism with much accuracy beyond the null hypothesis.
- Some of the features identified as most predictive, race and gender, are fairly problematic to base sentencing off of and probably tied to other circumstances/features blind to our dataset.
- Without more data and a better model, we don't see the merit in basing sentencing off of recidivism prediction in this manner.

Acknowledgements

- CS229 Staff
- National Archive of Criminal Justice Data
- Bureau of Labor Statistics