

Algorithms as Prosecutors: Lowering Rearrest Rates Without Disparate Impacts and Identifying Defendant Characteristics ‘Noisy’ to Human Decision-Makers

Daniel Amaranto, Elliott Ash, Daniel Chen, Lisa Ren, Caroline Roper

June 25, 2017

Abstract

We investigate how machine learning might bring clarity to a human decisions made during the criminal justice process. We created a model that predicts a defendant’s risk of being rearrested after their charges are dropped. We used a database from the office of the Orleans Parish District Attorney that covers cases from 1988-1999. Applying strategies identified by past research that compared prediction models to judicial decision makers, our model selected higher risk individuals to prosecute than its human counterparts did. For a set charge rate, our model would reduce the rearrest rates between 5% and 9%. Developed further, such a model could have several important policy implications: it might identify defendant characteristics that are particularly ‘noisy’ to prosecutors; it could suggest ways of alleviating criminal caseloads without increasing crime rates; and it might provide important insights into how a prosecutor’s background relates to the quality and nature of their charging decisions.

Introduction

It has been claimed that the US criminal justice system incarcerates too many and generally does so unfairly. Experts have identified certain factors that have particularly strong effects on the system’s purported inefficiencies and inequities, including the aggressiveness of prosecutors and the practice of plea bargaining. [1, 3]. It has been suggested that a reduced emphasis on plea bargaining and a heavier emphasis on the “screening” process — the stage when prosecutors decide whether or not to charge a defendant — could improve the fairness and efficiency of the system.[3] In our investigation we are not concerned with weighing in on the validity of such claims. We aim instead to augment the understanding of screening decisions made by prosecutors.

After a person is arrested and before a trial begins, prosecutors (screeners) can decide to either accept those charges and proceed to a trial or to drop them. In order to assess whether or not the decision to drop charges was made correctly, we use rearrest as our target; that is, if an individual who had charges dropped enters the arrest registry again within a certain time frame, we consider the screen decision to have been wrong. To optimize this prediction problem we use gradient boosted trees, a forward stagewise additive modeling algorithm that averages decision trees that are sequentially improved. After optimizing the model, we employed techniques described by Kleinberg et al [2] to assess its performance

compared to screeners. A reduction in rearrest rate model by the model would allow us to critique the way that screeners select defendants to charge.

Data

Our data describe over a decade of arrests in a federal prosecutor’s office from The Orleans Parish District Attorney’s office. The current data set is from 1988 to 1999 and provides detailed information on approximately 430,000 charges and 280,000 cases (involving 145,000 defendants) filed or adjudicated during this timeframe. The data collected also contains detailed information regarding each individual offender, such as information about the prosecutor, arrest register, and defendant.

Data Cleaning

When possible, we imputed missing values for continuous-valued features based on the values of other non-missing features (i.e., imputing the value of the screening date from date of arrest and number of days between arrest and when the arrestee was seen by a screener). For continuous-valued features where this was not possible, missing values were imputed as the mode. Using the codebook and lookup tables associated with the dataset, we flagged invalid values of binary and categorical features. Binary and categorical features were then transformed using one-hot encoding.

Our model only considers cases where the arrestee was not charged. We found that there were sometimes multiple arrests for a defendant on the same day. Some of these entries were duplicates while others took on different values for the screening outcome code. We flagged the multiple arrests so that there would only be one arrest per defendant per day, and if charges for that defendant were accepted for any of the arrests on that day, we set the screening outcome code to indicate that charges were accepted. Those arrestees were then removed from the dataset fed into the model.

We split the data into training, validation, and test sets using a split of 64/16/20. The split was stratified along the year of arrest so that the distribution of arrests over time was consistent among the training, validation, and test sets.

Target Variable

To construct the target variable, we had to identify arrests where the arrestee was rearrested within a certain number of years in the future. In order to evaluate whether an arrestee was rearrested within a certain number of years, we truncated the data by that number of years from 1999, the last year for which we have arrest data. For example, if we wanted to create a target variable that indicates rearrest within 2 years, we would create that variable for defendants arrested between 1988 and 1997 so that we could conclusively determine whether an arrestee in 1997 was rearrested within two years. We created target variables for rearrests within one to five years, removing 36,246 records (45%) from our training data.

Modeling

Our approach starts by creating a model that takes as input instances of arrests in an arrest registry. Each arrest includes details related to both the charge and the arrested individual.

Additionally, each arrest in the data has a field with the number of days, if it exists, until the individual has a subsequent arrest. Using data from which actual outcomes are known and a loss function that penalizes incorrect predictions, the model learns how to optimally predict the outcome if it receives a new input. Thus our model makes a prediction about whether each entry in a registry will or will not be rearrested within a set amount of time. The window of time during which we consider rearrest is variable. The model also generates a probability of rearrest for an instance; this probability, or risk score, is essential to formulating a comparison between the model and human decision makers.

This prediction model was trained on a dataset of 80,217 arrestees. A validation set with 20,055 arrestees was used to evaluate the performance of different permutations of the model. A test set of 25,068 arrestees has still not been touched and awaits our completed model.

Baseline Model - Decision Tree

One of the simplest ways to predict whether someone would be rearrested is to use a decision tree based on age and the severity of the charge. Therefore we constructed a baseline model of maximum depth 4 using these two features. Given a parent node, we fit the model by minimizing the entropy of child nodes. Entropy is defined as

$$-(p_0 \log_2 p_0 + p_1 \log_2 p_1)$$

where p_0 is the proportion of the node comprised of arrestees who were not rearrested and p_1 is the proportion of the node comprised of arrestees were rearrested. The baseline model achieved 60% accuracy on the validation set and an F-score of 65%.

Ensemble method: Gradient Boosted Trees

We refine our baseline model with Gradient Boosted Trees, a sequential ensemble model comprised of decision stumps. We used binomial deviance as our loss function:

$$\log(1 + e^{-2y\hat{y}})$$

where y is the true value of the target variable and \hat{y} is the value predicted by the algorithm. Using this loss function allows us to interpret the score function as a probability.

Model Optimization

Feature Selection

We performed feature selection using the following steps:

1. Create 22 training datasets, one using all features and 21 created by removing one feature at a time.
2. Train a gradient boosted trees model on each of the modified training data sets.
3. Repeat this process for 5 different target variables from rearrest within 1 year to rearrest within 5 years.

4. Evaluate each of the 220 models to identify which feature, target, and model combination results in the highest validation f-score.

We found that the highest F-score resulted from removing Arrest Credit Code, keeping the rest of the features, and predicting rearrest within 5 years.

Parameter Optimization

After selecting rearrest within 5 years as the target variable and choosing a final set of features, we performed parameter optimization. All the models we trained use 50% as the probability threshold for each class, since the classes are balanced in our data (52% of arrestees were rearrested).

For the gradient boosted trees, we trained the model with the following values of hyperparameters for a total of 81 models:

1. number of estimators (100, 300, or 500)
2. learning rate (.05, .1, .5)
3. max depth (3, 5, 10)
4. minimum samples split (2,4,8)

All the gradient-boosted tree models use deviance as the loss function to be optimized and the Friedman mean squared error as the measure of the quality of a split.

Our best-performing final gradient boosted trees model used 500 tree estimators, a learning rate of 0.05, a maximum depth of 5, and a minimum samples split threshold of 4. The associated validation F-score is 0.7694, which is .12 higher than the validation F-score of the baseline model. A comparison of the final model to the baseline shows the substantial improvement achieved.

The final confusion matrix for the optimized model is below.

	Prediction: Not Rearrested	Prediction: Rearrested
Not Rearrested	3905	1391
Rearrested	1265	4433

Note that all results are from the validation set. Because we plan to continue our research, we have not evaluated our model on the test set.

Results and Evaluation

We assessed the actual risk of arrestees compared to the predicted risk returned by our model. Using the validation set, we grouped the arrestees into quintiles by their estimated risk and found that the predicted riskiest arrestees have higher rearrest rates. This shows that the arrestees that were released by a screener and predicted by our model to be risky were in fact risky.

We also assessed the performance of our model against human screeners (as in [2]) using the concept of implicit risk ranking. We do not directly observe how the screeners rank

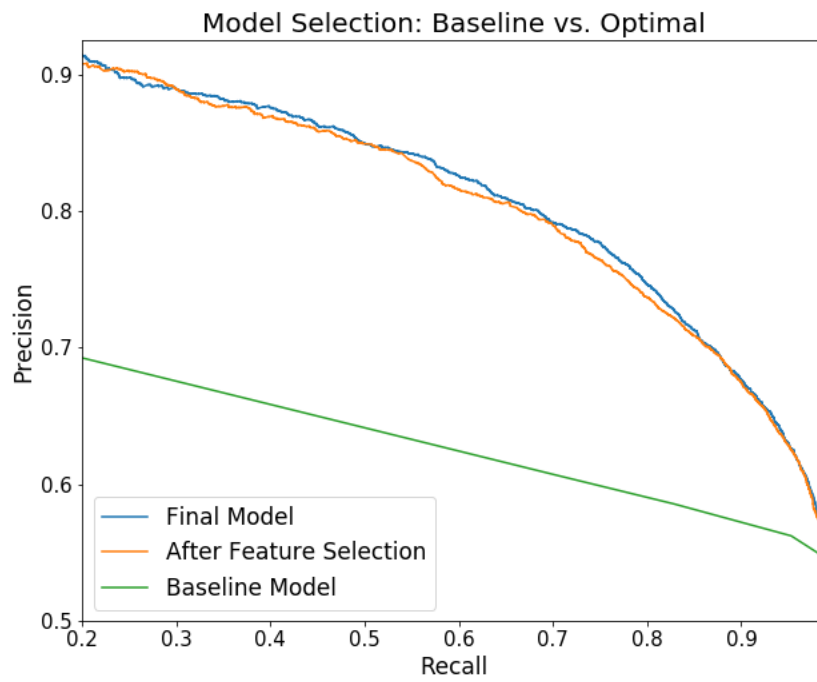


Figure 1: The figure above shows the precision-recall curves of three models: (1) the decision tree baseline model, (2) the gradient boosted trees model with optimal features and default hyperparameters, and (3) the final model with optimized hyperparameters.

the risk of the arrestees that they see. However, we can assess their implicit risk ranking from the variation in charge rates between “strict” and “lenient” screeners by comparing the distribution of predicted risk of the arrestees charged by the “strict” and the “lenient” screeners. The NODA dataset provided sufficient variation in screener charge rates for us to assess the implicit risk ranking (see appendix). We find that the actual risk distribution of risk amongst strict and lenient screeners to differ from what we would expect to see if the screeners were releasing defendants based on their predicted risk.

Furthermore, we assess the potential improvement in outcome from using our model by analyzing the “marginal” defendant. Given a screener (or a group of screeners), we define the marginal defendant as the defendant with the highest predicted risk that was seen and released by that screener (or group of screeners).

We group screeners into two bins based on the percent of arrestees that they charge. Then we calculate the additional number of arrestees that would need to be charged for the “lenient” group of screeners to reach the same charge rate as the “strict” group of screeners. Using the fitted gradient boosted tree model, we choose these “marginal” defendants based on estimated risk. We assess the outcome, in terms of rearrest rates, of charging these additional “marginal” defendants. If we arrive at a lower rearrest rate than the strict human screeners, then our model results in improvements in rearrest rates. The potential improvements in rearrest rates are summarized in the chart below.

To test the robustness of our results, we compared our results with the rearrest rates if we were to choose which the additional “marginal” defendants to charge at random. We conducted 1,000 iterations of this test and were not able to reduce the rearrest rates to the same level as when using the predicted risk ranking.

Discussion

Error Analysis and Time to Rearrest

We conducted error analysis to identify unusual features. We divided the observations according to whether they were correctly predicted and those that were incorrectly predicted and compared the distributions of every feature using box plots for continuous variables and bar plots for categorical variables. The result of the full comparison is in the appendix. We discuss one substantive finding here.

We observed that the incorrectly predicted observations had a similar distribution to correctly predicted observations for almost all features, with the exception of time to rearrest. The box plot below shows the distribution of time to rearrest for correctly predicted and incorrectly predicted observations in the positive class.

To explain this pattern, we examined the top features by feature importance. Below, we include the top 5:

Column Name	Importance
AGE	0.281481
BAR_ADMIT_DAYS	0.167391
ARREST_TO_SCREEN	0.115987
SCREENING_DAYS	0.072958
CRIMINAL_FLAG	0.048115




Figure 2: The top graph shows the actual distribution of risk of the defendants released by screeners of increasing strictness (from left to right). The bottom graph shows what we would expect the distribution to be if screeners were releasing defendants based on their predicted risk. If screeners were to release defendants at random, we would expect to see an even distribution of predicted risk for each set of screeners.

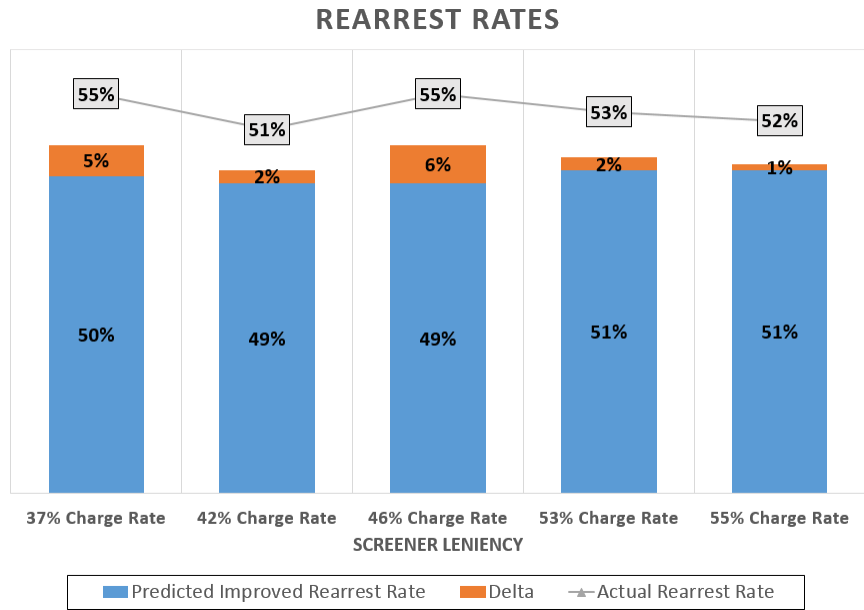


Figure 3: The graph shows the improved outcome, without changing the percent of defendants charged, if defendants are charged based on their predicted risk.

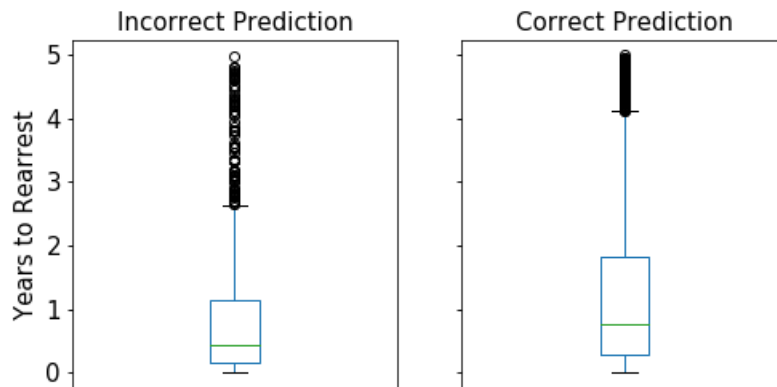


Figure 4: This figure represents the distribution of years to rearrest for arrests followed by a subsequent arrests. The left side of the figure corresponds to instances incorrectly classified and the right side of the figure corresponds to instances correctly classified.

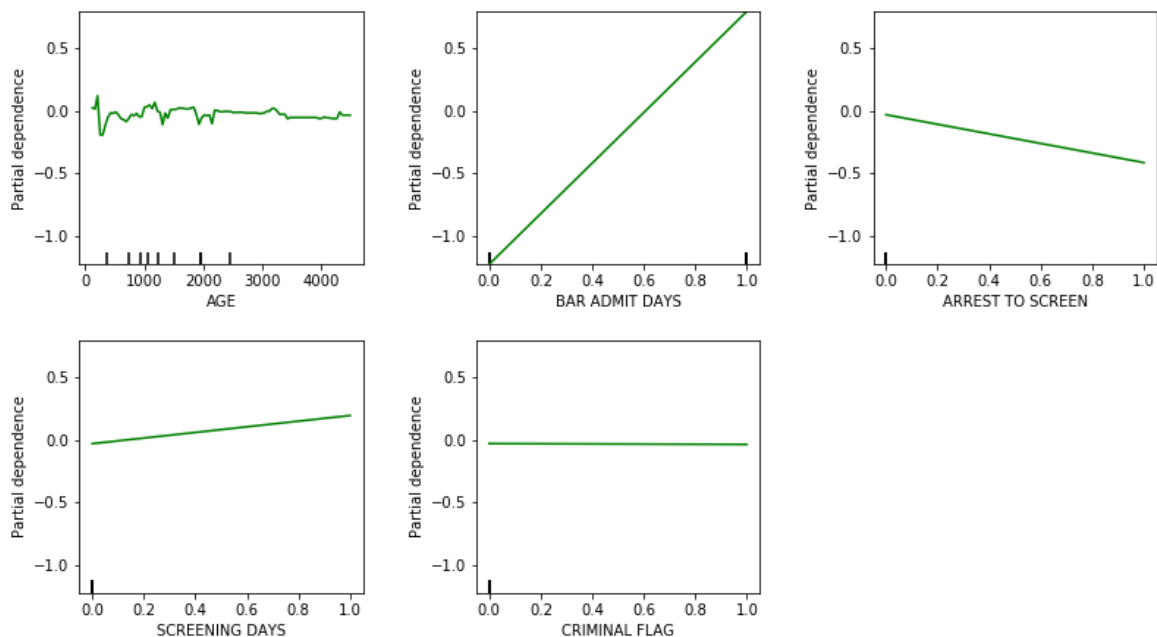


Figure 5: The partial dependence plots of the top 5 features provide a clearer sense of the direction of the relationship.

We also examine the non-categorical features correlated with time to rearrest among the observations that were rearrested within five years. Below are the correlations between time to rearrest and the five non-binary, non-categorical features that have the strongest correlation with time to rearrest:

Column Name	Correlation with Time to Rearrest
CRIMINAL_FLAG	0.229738
JUVENILE_FLAG	-0.184845
AGE	0.135949
INITIAL_DETENTION_FLAG	-0.037101
ARREST_TO_SCREEN	0.034602

Time to rearrest is positively correlated with criminal flag, and criminal flag is associated with a higher predicted probability of rearrest, which may explain in part why recall is higher among observations with a higher time to rearrest. The correlation between time to rearrest and criminal flag is statistically significant according to the Pearson correlation coefficient's p-value. Our results suggest that those arrested for criminal charges who were rearrested were rearrested later than those arrested for non-criminal charges. This finding may seem slightly counterintuitive, but further investigation may explain why the two are correlated.

Demographic Analysis

In order to choose which features would optimize the model (including both categorical and non-categorical variables), we removed individual features one at a time and observed how the F-score changed. Sex, age, and race were consistently among the features whose removal had the highest negative impact on model performance. That criminality is associated with male gender and youth might not be especially surprising or concerning, but it is necessary to assess whether the model is exacerbating racial disparities by, for example, ascribing disproportionately less risk to white defendants.

While we did find that the algorithm tended to predict higher riskiness for black defendants (see below), the false positive rate (i.e., the percent of defendants predicted by the algorithm to be rearrested but were not rearrested) for white defendants was 24% while it was 21% for black defendants. Furthermore, we found evidence that screeners may have some racial bias that may be improved upon by using our model. 75% of the white defendants that were released were rearrested within 5 years while 42% of black defendants released were rearrested within 5 years. Out of these defendants that were released, our model predicted that 88% of the white defendants to be rearrested and 38% of the black defendants to be rearrested within 5 years. This suggests that it is possible that the human screeners may be too lenient on white defendants and too strict on black defendants. We also note that in our training data, only 11% of the defendants are white and 86% of the defendants are black. It is possible that our model would have improved performance if provided more instances of white defendants.

It is all but certain that there are unobserved variables at play which add bias to our model; educational background, socioeconomic status, and a host of other factors may affect the rearrest outcome but are only indirectly observable to the model through the race variable to which such class markers are related. Finding ways to include such unknowns would be an important way to develop this analysis.


Aligning the screener’s and the algorithm’s prediction

An important distinction between screeners and the algorithm is that while our model is classifying declined cases that result in rearrest to be ‘wrong’ decisions, a screener might not have been making the same prediction. It is possible that different predictions (e.g., whether or not a case can be tried and won) are actually driving the decisions at this node.^[4] This does not, however, invalidate the comparison that we present. If a prosecutor’s office cares about rearrest rates, and one would assume that they all do, then the results of a successful rearrest algorithm should still be relevant to them.

Potential Impact and Additional Steps

Given that the algorithm we developed predicted rearrest rates with relatively more accuracy than screeners, a potential implication would be to make it available to screeners as they consider declinations in real time. As pointed out by Kleinberg et al. in seminal research on the analysis of judicial bail decisions, such a model could provide decision makers with a risk score or flag for individual cases and serve as an aid, though not a replacement, for their judgment. Alternatively an algorithm could be used to rank entire populations of defendants for larger-scale recommendations. ^[2] Large-scale ranking could be important when districts have to make assessments about the feasibility of caseloads and prison populations. While

Histogram of Predicted Risk for Released White Defendants



Histogram of Predicted Risk for Released Black Defendants




Figure 6: These two figures show the distribution of predicted risk by for black defendants and white defendants that were released by the screeners.

our model's success rate is promising to this end, an extensive amount of further research is required before a practical application could materialize.

One way to make the model more versatile would be to make it capable of identifying if lower charge rates can result in the same rearrest rate. This would allow policy makers to lessen the burden of cases with minimal impact on the level of rearrests.

Finally, we were not able to analyze the relationship between screener traits and the characteristics of their declinations. Analyzing variation across screeners could identify whether or not biases exist based on screeners' political affiliation, age, experience, race, etc. One way of addressing this problem is to fit a model to predict screener declinations and to compare the feature most predictive in that model with the predictive features in the rearrest model.

Limitations

Incomplete data was a limitation in our study. We were unable to track the arrest registry beyond 1999 and unfortunately our requests for more recent data were not fruitful. Having access to the arrest records of the Orleans Parish jurisdiction would also have improved the model; it is highly likely that at least some of the defendants in the registry were rearrested in the future in another district but were identified by our model as correct decisions.

Conclusion

The decision to decline or pursue charges against a defendant has been identified as an potentially underemphasized point in the criminal justice process [1, 3]. Using eventual rearrest as an indicator of a successful declination decision, we created a model that outperformed human screeners on data from the NODA database. Applying this model to decisions at the declination node would have achieved lower rearrest rates between 5% and 9%, depending on the strictness level of the screeners we compared. Underlying biases in the model need to be addressed by including more explanatory variables, particularly with respect to demographic data of defendants and screeners. Using machine learning prediction algorithms to assess human decisions is a promising field of research in the court context and beyond, and further research such as this will hopefully have meaningful impacts on policy discussions.

Appendix

Fall in rearrest rate as charge rate increases

To get a sense of the volume of decisions made by screeners of a given strictness level, we determined the cumulative percentage of cases seen as strictness increased. This revealed that the bulk of cases (roughly 60%) are seen by screeners with charge rates between 35% and 60%. Narrowing our focus to those screeners, we looked at the rearrest rates of their declinations across various time frames. As shown in the following plot, the rearrest rate reaches a peak at a charge rate of around 40% across several rearrest time frames. Then, remarkably, the rearrest rate falls before rising again and steadying. A possible explanation for this outcome is that cases are not randomly assigned to screeners, but rather they might be divided based on type of crime or some other factor. Screeners at charge levels from 0.4 to approximately 0.48 might be focusing on types of crime that are not associated with repeated criminal behavior.

Aside from the question of random assignment, this decrease in rearrest rate might be explained by separating these plots according to type of crime.




Figure 7: The first graph shows the cumulative percentage of cases that are seen as more strict screeners are included. The bottom graph has the same x-axis, and shows the rearrest rate vs screener charge rate for 5 different rearrest time frames. Notice that over 65% of all cases are seen by screeners with charge rates between 30% to 60%.

Data Lost To Create Target

In order to create the target variable, we removed arrests from the last 5 years of the dataset, because we did not have data on whether the arrestee was rearrested within the next 5 years. This removed 36,246 records (45%) from our training data, which originally consisted of 80,217 records. The table below summarizes the year in which each arrest occurred, and the number of records we'd remove to set the target variable based on a number of years:

ARREST_DATE_y	Records	Cumulative Total
1999.0	7329	7329
1998.0	9432	16761
1997.0	8450	25211
1996.0	5929	31140
1995.0	5106	36246
1994.0	6093	42339
1993.0	4770	47109
1992.0	6598	53707
1991.0	7280	60987
1990.0	9769	70756
1989.0	7783	78539
1988.0	1678	80217

Error Analysis

We presented substantive findings from the error analysis in the Discussion section. These plots demonstrate that time to rearrest was the only variable for which the distribution among correctly predicted observations noticeably differs from the distribution among incorrectly predicted observations.

Figures

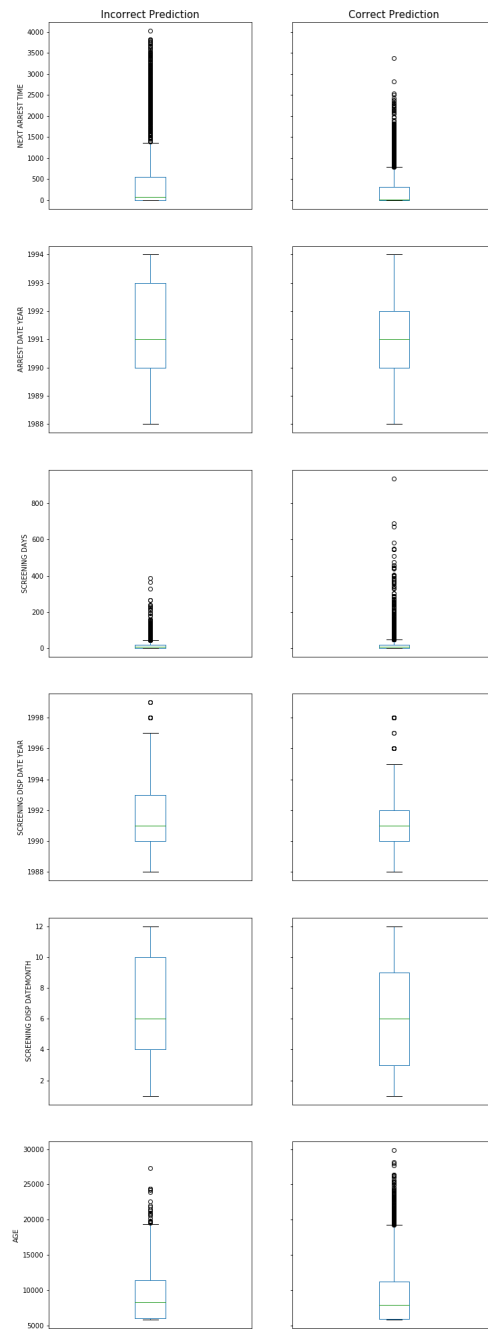


Figure 8: Box plots of the distributions of each continuous variable for correctly and incorrectly classified instances.

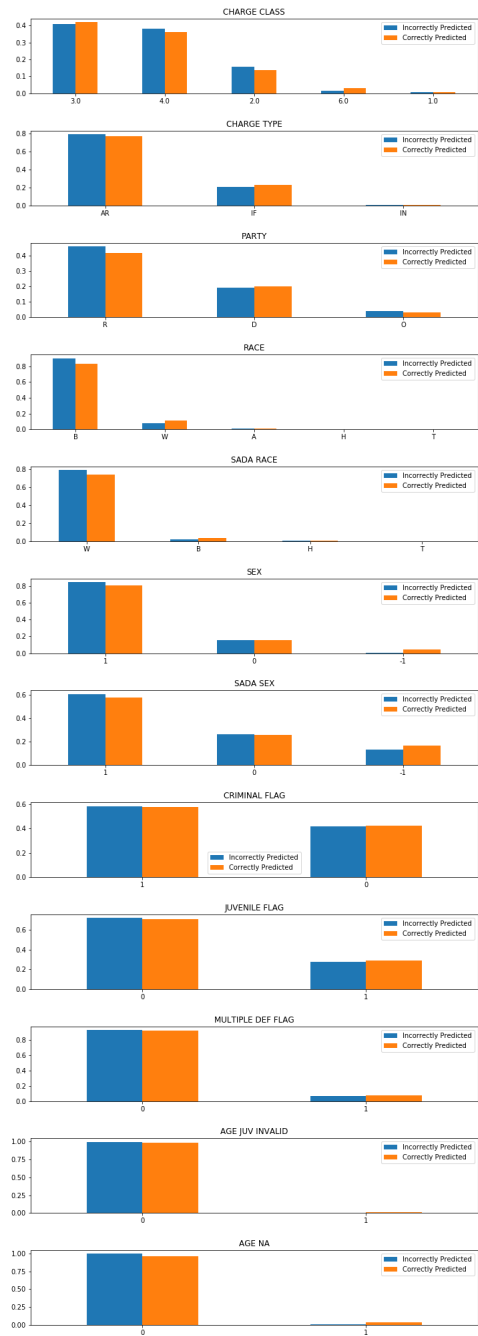


Figure 9: Bar graphs of the distributions of each categorical or binary variable for correctly and incorrectly classified instances, limited to the five most common values of the categorical variable for simplicity.

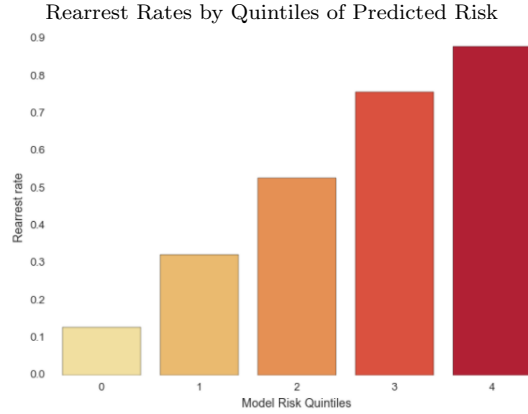


Figure 10: This shows that the arrestees that were released by a screener and predicted by our model to be risky were rearrested at a higher rate

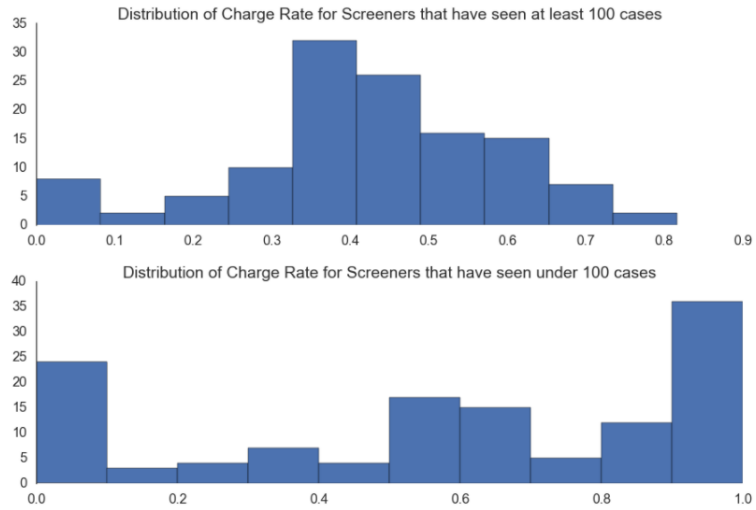


Figure 11: These two charts show the variation in charge rates across different screeners.

References

- [1] Gopnik, A. (2017). How we misunderstand mass incarceration. *The New Yorker*, April 10, 2017.
- [2] Kleinberg, J., Lakkaraju, H., Leskovec, J., Ludwig, J., and Mullainathan, S. (2017). Human decisions and machine predictions. National Bureau of Economic Research Working Paper Series.
- [3] Miller, M. L. and Wright, R. F. (2002). The screening/bargaining tradeoff. *Stanford Law Review*, 55(29):29–118.
- [4] Miller, M. L. and Wright, R. F. (2008). The black box. *Iowa Law Review*, 94(1):125–196.