International Journal of Innovative Technology and Interdisciplinary Sciences
https://journals.tultech.eu/index.php/ijitis
ISSN: 2613-7305
Volume 9, Issue 2
DOI: https://doi.org/10.15157/IJITIS.2026.9.2.625-656
Received: 02.01.2026; Revised: 03.03.2026; Accepted: 28.03.2026

Research Article

Regression-Based Machine Learning for Predicting Prior Convictions from Administrative Criminal Justice Records

Ismail Shehu[1*], Kreshnik Myftari[1], Ahmed Jumaah Sultan[2], Hussein Alkattan[3,4], Mostafa Abotaleb[5]

[1]Department of Criminal Law, University of Tirana, Tirana, Albania
[2]College of Computing and Informatics Information Technology Department, Al-Hamdaniya University, Ninevah, Iraq
[3]Department of System Programming, South Ural State University, Chelyabinsk, Russia,
[4]Directorate of Environment in Najaf, Ministry of Environment, Najaf, Iraq,
[5]Engineering School of Digital Technologies, Yugra State University, Khanty-Mansiysk, Russia
***ishehu1985@gmail.com**

Abstract

Predictive analytics is emerging in current police and court proceedings where data is used to analyse crime trends and court results; predictive, regression-based machine learning models use these data to predict dependent variables. This study explores the predictive use of regression-based machine learning models to predict counts of prior convictions based on structured data within a criminal justice dataset. The analysis is done on the Criminal Justice Dataset that is a publicly available data set with roughly 200,000 court cases data with demographic characteristics, criminal history factors, offense parameters and court decisions variables. Using descriptive statistical analysis estimates, correlation univariate tests and multiple regression predictive models the relationships are explored between the independent predictors and the dependent counts of prior convictions. Three model builds are used; the first is an OLS model to provide a comparison baseline followed by the use of a Ridge regression model with L2 regularization penalty and an ElasticNet model with L1 and L2 as a way of deploying multiple robust models to produce the best fit and prediction. We evaluate the performance of the model using standard regression metrics: coefficient of determination (R^2), Mean Absolute Error (MAE), and Root Mean Squared Error (RMSE).

Keywords: Criminal Justice Dataset; Regression-Based Machine Learning; Statistical Analysis; Multiple Regression; Linear Regression; Ridge Regression; Regularization and Model Evaluation

INTRODUCTION

The digitalization of criminal justice information-systems infrastructure is accelerating, creating an opportunity for the application of data-driven analytics to correctional decision-making (sentencing, parole supervision and community correction). The availability of large-scale administrative data has driven researchers towards risk profiling or resource distribution using predictive modelling. One seminal work in this vein demonstrated that machine learning models can predict risk of recidivism among prison inmates, suggesting that standard data collected at intake may have rich signal regarding individuals who are likely to re-offend [1]. These findings suggest that the broader class of statistical and machine learning methods for which practical adopters may iteratively explore various subsets with different balances between predictive value while maintaining operational interpretability according to their constraints or goals in decision making – warrants further consideration.

As well as inmate-centred prediction, parole and post-release supervision contexts are essential to consider because of the ongoing difficulty of re-offending upon release. In the empirical literature that studies recidivism patterns among parolees, reoffending was found to be a product of an interaction between individual attributes, conditions of supervision and situational factors and it stressed prevention and control strategies that can be grounded in correctional practice [2]. This viewpoint is crucial because it presents predictive modelling not as a standalone technical practice, but as instrumental for informing policies that aim at prevention (for example targeted supervision, rehabilitative programmes and risk-tailored responses).

With the proliferation of "big data" systems and increasingly more timely data, early-warning models have received increased attention in community corrections. For community correctional populations, we propose early-warning models for the purpose of reaching escalation signals prior to recidivism, thus initiating proactive intervention [3]. These approaches move the focus from static risk scoring to dynamic monitoring, and they can help correctional agencies pay attention to key behaviour or process markers of increasing risk.

In addition to supervised prediction, data mining algorithms have been used to describe recidivism trends and discover latent structure in corrections' data. For example, association rule mining has been applied to characterize recidivism by finding frequent combinations of features associated with reoffending [4]. Even though association rules are not regression-based predictors to begin with, prolonged evolution can generate complementary information that inform hypothesis generation and feature construction in prediction. Concurrently, the construction of the knowledge management system for criminal rehabilitation focuses on the support of evidence-based rehabilitation workflows and organization data pool with standardized indicators to enhance effectively the quality and usability of data for analytics and modelling [5].

Given the interpretability and ability to facilitate formal inference, OLS-based statistical models remain dominant in the criminal justice empirical landscape. After that, the development of quantitative risk assessment instruments based on social risks for arrest has been described by means of binary logistic regression [6], demonstrating how regression frameworks could be used to implement different definitions of risk but keep results interpretable for stakeholders, and this is especially important where the model is deployed in a criminal justice setting (and model explainability and auditability, as well as decisions to adopt or reject it) become critical. (Theres still a lot of conflating going on the message that predictive codelength is, classical or machine learning requires rigorous statistical thinking and careful consideration for what our models will ultimately be used for though, more abstract discussions not being full blown tutorials on the problem of stability can still do a good job at making this point) An opportunity for an influential contribution to this area is treatment of the statistical issues of such applications of machine learning applied in criminal justice risk assessments that are sensitive to measurement error, data generating process and local system effects from policy through model, informed decision making [7].

Within such a methodological context, one of the most advocated models to be employed to predict recidivism is logistic regression due its probabilistic nature and ability to hypothesize features. Feature selection has a great effect on the performance and stability of logistic regression according to studies [8]. These findings suggest that careful variable screening and validation procedures are necessary for valid risk estimation. These findings echo the general point that interpretable baselines add value, particularly when combined with systematic evaluation/reporting practices.

Research Objectives

The main research objective of this study is to examine the effectiveness of regression-based machine learning models in predicting prior convictions based on administrative criminal justice system records. By examining the structured variables in terms of demographic information, criminal history, severity of offenses, and judicial outcomes, the research will attempt to pinpoint the key factors that may have a bearing on the number of prior convictions.

A further objective of this study is to examine the effectiveness of three different regression-based approaches in predicting prior convictions: OLS regression, Ridge regression, and ElasticNet regression. By examining the effectiveness of these different approaches, the research will also assess the effectiveness of regularization techniques in improving the accuracy of predictions.

RELATED WORK

Studies of algorithmic decision support in criminal justice surged after investigative reporting prompted questions about opacity and bias in commonly used risk assessment tools. The ProPublica analysis of algorithms also popularized the criticism that predictive systems might induce differentially distributed errors across groups and the fact that companies' internal designs make independent audits difficult, revealing a demand for more transparency surrounding model construction and verification in justice settings [9]. In response, systematic syntheses of the scholarly record might then have shown that "fairness" is not a uni-dimensional property to be quantified but rather a collection of competing definitions that cannot necessarily be simultaneously satisfied in practice even when base rates differ across subpopulations and even when decision thresholds are set at policy rather than optimizing levels for each group [10]. Subsequent reviews [11] echoed this idea in summarizing empirical and theoretical results that showed fairness can come at the cost of overall performance (either predictive or calibration), while stressing the importance of considering fair decisions in the context of the larger institutional pipeline by which predictions are translated into operational activities.

Methodological work, in addition to the literature on fairness, has focused on how criminal trends should be predicted and how predictive methods should be compared with each other. In this regard evaluation and validation studies are contested, focusing on what is termed out-of-sample validation of the model yielding a consideration of uncertainty and decision relevant error costs [12] to enable translation of criminal justice risk predictions into contexts where errors yield significant consequences for individuals or communities. Probabilistic learning for causal inference the perspective through which these evaluations are conducted is based on general theory of machine learning, e.g., those principles of generalization and regularization and bias-variance trade-offs that represent whether a model should still yield appropriate associations through different populations or points in time [13] In criminology, some recent integrative reviews argue that machine learning can contribute to scientific understanding of crime if machine learning is practiced sensibly and with an explicit awareness and deliberation about measurement accuracy, causal ambiguity, and interpretability instead of a technical optimization problem [14]. This approach applies to establishing clear baselines and strict reporting standards before modelling more advanced models.

In the model class arena, ensemble learning has been comprehensively applied in criminal justice prediction where the problem is complex [15] (e.g. non-linear, and interaction) Random forests are a widely used modelling tool in criminal justice as they have empirically performed well and have the ability to include mixed type predictors, complex inter variable relations [16], among others. Other variations and translation of the ensemble toolbox remain an area of applied work in criminal justice forecasting, indicative of ongoing efforts to balance accuracy and outcome interpretability [17]. In implementation, model output sometimes is explained relative to a general baseline of juxtaposing models calibrated to official justice data, research samples, and dashboards to systemically illuminate the extent to which the generalizability of the sample to inform prediction matches operational conditions. This should be done at the outset to avoid over specifying models.

An empirical measure of "disparate impact" played a role in the analysis of the degree of breach of the prediction rule parameters for fairness. Indeed, using whole population levels of accuracy in recidivism predictor modelling is sufficient to allow for potentially unfair Type I and Type II errors in particular demographic groups [18]. As the criminal justice uses of "high risk" AI applications become more prevalent, so must the governance actions and in the EU, critical legislative policy developments according the draft AI Act include computerisation mandates for documentation, risk management and transparency similar to regulations that emphasize the protection of fundamental rights [19]. These governance dynamics imply the need for evaluation methods that do not just orient around single metrics but also consider how model outputs are made to matter. Since model evaluation in criminal justice is focused on making decisions (with thresholds) on often highly imbalanced data, the choice of metrics is particularly important. The classic approach to model comparison without having the need for choosing a specific cut point has been ROC analysis (which summarizes discrimination performance over all decision thresholds) [20]. However, when results are to be communicated to stakeholders who need an understanding of the model and will accept it as evidence for what the model does, applied regression and statistical modelling guidelines counsel that predictive performance is a behaviour that needs complementing with diagnostic checking, prudent interpretation and explicit clarification of limitations and assumptions [21]. Formal definitions linking predictive behaviour to normative criteria also help with fairness evaluation. Popular notions of equality of opportunity cast as parity constraints the error rates conditional on predicted outcome and ground truth labels, which conveys that equal treatment of groups under supervised learning is a formally defined objective [22]. This reinforces general principles from statistical learning about the importance of a clear validation framework and, sensitivity analyses as well as more disciplined model comparison in order to avoid unwarranted performance claims [23].

Such a unbalanced data problem is all too prevalent in any of the criminal justice end points with a rare outcome of interest or where different rates exist between subgroups. A paper of machine learning on imbalanced data narrates why naïve metrics can be deceptive and how biomass resampling, cost sensitive learning or alternative summary measures of prediction are simple heuristics which can compensate a loss sensitivity [24]. In the particular science of recidivism, logistic regression is the conventional benchmark method, as it is able to provide transparent parameter estimates and also produce success probabilities; moreover, there exists a wealth of curated guidelines for model fitting, diagnostics and reporting [25]. The toll is quite different in justice data which is structured: a massive part of success probabilities is driven by the constructed features from data preprocessing steps. In addition, respective algorithms such as systematic pipelines of feature engineering and selection (e.g., variable transformation, categorical variables processing, redundancy exclusion and feature contribution verification) lead to more stable decision tree classifiers with boosted generalization power [26]. The associated applied predictive modelling tools are more focused on end-to-end chain of the various data mining steps of pre-processing, cross validation, model tuning and evaluation in a highly reproducible way, which is imperative in any high stake's applications [27].

More than this, software ecosystems are influencing the ways machine learning techniques are being applied in the field. The influence of this and the availability of standard implementations in statistical computing environments has also helped its popularity, since they provide a proven and easy-to-use package for classification (and we can, naturally, include regression here too) [28]. Concurrently, interpretability research has also produced model-agnostic and model-specific explanation methods that can be used to map predictive behaviour into interpretable feature attributions. Standardization processes like SHAP make it applicable for similar interpretations throughout the model families and allow researchers to evaluate how much contributor players do influence particular predictions or overall risk patterns, including with more complex learners [29]. Mature statistical computing systems, and open-source toolboxes for performing standardized data manipulation, and streamlined reporting are available to enable reproducibility in such studies [30].

Data quality remains a significant challenge in any kind of criminal justice modelling as administrative data may be (from right to left) contaminated with historical bias, inconsistently measured and/or enforced, or filtered through low quality measurement devices. This "dirty data" argument shows how, if data from biased practices is simply re-produced, a predictive model will merely reflect those civil rights violations, making the case that data must be interpreted with regard to the context in which it exists [31]. The evaluation tools for Web 1.0 are also maturing to help support more rigorous reporting. One possible development is the normative ROC analysis software already noted that may make reporting of model performance across research projects more reproducible and comparable [32]. But the question of interpretability remains hotly debated., some leading commentators have called for a requirement that a more transparent interpretable model always be used in serious situations than black box models with post hoc explanations, since transparency is as important as accountability in ensuring protection when things go wrong [33]. Closely related is a critique of the lack of transparency and limited access to the early commercial reoffending rate sheets, which limited scientific validity and democratic control and demand for open validation and development protocols [34].

All of the Metrics Selection has become adapted to address present class imbalances and decision centric evaluation. Precision–recall curves can also give much more information than ROC when the positive class is infrequent, because they are directly concerned with the precision recall trade off (and we describe below why this does not necessarily give you robustness to imbalance redirects here) and may be better situated to reflect practical constraints [35]. Analyses of fairness have likewise shown that predictive bias and disparate impact may manifest even if models meet some technical accuracy criteria, highlighting the need for subgroup analysis, and making fairness goals explicit when tools assist in decision-making [36]. Outside the modelling literature, the extent to which forecasts depend on process design, organizational incentives and evaluation culture as well as model choice has been observed [37] indicating that good decision making requires disciplined processes for use and monitoring rather than relying solely on model choice. Clear communication of model results is needed; the principles of good quantitative visualization continue to apply in displaying evidence without bias and enabling comparison between models and groups [38]. Finally, it is important to note that missing data are ubiquitous in administrative justice datasets, and existing guidance suggests that principled imputation plus a sensitivity analysis can avoid bias while maintaining power under informative or nonuniform missingness across subgroups [39].

In Table 1, the descriptive statistics of the numeric variables in this study is summarised to reflect central tendency and dispersion, alongside a general overview of some quantitative characteristics of the data before a model itself is introduced. The first-row features more common measures reported by the mean and the median but in second row standard deviation is offering us insight into how heterogeneous individuals are. The min and max values reflect the full range that was observed thus allowing for identification of a potential outlier or data entry issue. First of all, our Age also has a wide adult range, so that is already an ordinal with only one nullable value (age can be null if for example the user never provided it), while Prior_Arrests and Prior_Convictions seems like they are count based skewed right (where most observations fall at low values but there are several observations behind higher counts). This summary is useful as it helps in making preprocessing decisions (e.g. whether or not to check for skewness and/or transform if you're using regularized models), but also since these stats help you interpret regression results by showing what a typical scale and width of each numeric feature is with respect to the predictive models you're running.

Table 1. Descriptive Statistics of Numeric Variables (Age, Prior_Arrests, Prior_Convictions)

Ref.	**Modern study**	**Typical data & target**	**Methods & evaluation used there**	**compares (difference / contribution)**
[40]	Systematic review of ML recidivism prediction; emphasizes reporting quality, metrics, and transparency.	Recidivism datasets; binary/score prediction	Reviews AUC/ACC, model families, bias risks, transparency	Our study is a transparent regression pipeline on administrative features and focuses on interpretable regression + error metrics (R^2/MAE/RMSE) rather than black-box classification.
[41]	Survey/review of ML in criminology and crime research; outlines applications, data challenges, and methodological issues.	Crime/justice datasets; multiple targets	Broad coverage of ML use cases and pitfalls	Our work is a narrow, reproducible case study: predicting Prior_Convictions using regression models and reporting diagnostics + multicollinearity control.

[42]	Fairness analysis showing how “fair” prediction criteria can still create disparate impact when base rates differ.	Recidivism instruments; group fairness	Fairness criteria + empirical evaluation	Our work uses regression and does not claim fairness guarantees; it supports fairness-aware discussion by using interpretable coefficients and clear reporting (reducing “black-box” ambiguity).
[43]	Examines predictive bias and disparate impact in risk assessment and recidivism contexts.	Corrections/recidivism samples; risk prediction	Bias/disparate impact analysis	Our work complements this by providing a transparent baseline that can be audited (coefficients + confidence intervals), which is useful when fairness concerns are central.
[44]	Tests whether ML actually outperforms humans in recidivism prediction, highlights limits and overclaiming risk.	Recidivism prediction tasks	Comparative benchmarking (humans vs models)	Our work avoids overclaiming by focusing on robust regression performance and diagnostic checks; it positions regression as a strong baseline rather than “replacement” for human judgment.
[45]	Argues that high-stakes domains should prefer interpretable models over post-hoc explanations of black boxes.	High-stakes decision settings	Conceptual + applied examples	Our study follows this principle directly by using OLS/Ridge/ElasticNet, reporting coefficients, and emphasizing interpretability for criminal-justice analytics.
[46]	Shows risks of secrecy and limited transparency in	Proprietary risk scores / COMPAS debates	Auditing/analysis of algorithmic claims	Our regression workflow is fully reproducible (data → preprocessing →

	recidivism prediction tools; emphasizes auditing and openness.			regression → metrics), which directly addresses the transparency gap highlighted in this line of work.
[47]	Explains how policing/justice data quality problems and civil-rights issues can contaminate predictions.	Police/admin data	Legal + socio-technical analysis	Our work explicitly includes preprocessing and reporting steps to reduce measurement noise impact; it frames results cautiously as administrative-pattern modeling, not causal truth.
[48]	ProPublica investigation showing racial bias concerns in risk assessment scores used in sentencing.	Risk assessment tools in sentencing	Investigative analysis + error-rate comparisons	Our approach differs by using interpretable regression and standard evaluation; it can serve as a transparent benchmark to compare against opaque risk scores.
[49]	Systematic review of ML for violence prediction; stresses external validation, calibration, and risk of bias.	Violence prediction studies	AUC emphasis: highlights lack calibration/external validation	Our work is not violence prediction, but it matches the review's recommendations by using a clean pipeline, interpretable models, and straightforward evaluation; it can be extended later to external validation if another dataset is added.

Research Gap

Increasingly, machine learning and statistical techniques are being applied to criminal justice analytics. However, despite a growing body of literature, several key gaps exist. While multiple recidivism studies have targeted this issue, most were binary recidivism prediction or risk classification studies that focus on classification algorithms and sophisticated machine learning methods. Although these methods provide a good predictive accuracy, they often rely on black-box models, and hence, have low interpretability and transparency. Lack of model transparency can lead to issues concerning accountability, policy analysis, and moral decision-making in high-stakes domains, such as criminal justice.

Existing studies also lack an interpretable regression-based predictive framework that can explain how demographic, criminal history, and judicial variables affect criminal outcomes. A majority of the literature focuses on state-of-the-art ensemble methods or deep learning models, while relatively few studies consider regularized regression models (for example, Ridge and ElasticNet) for either imputation or criminal justice prediction tasks. While these models balance predictive power with interpretable output, they have been underutilized in large scale administrative justice data.

Additionally, much earlier work has focused on predicting recidivism risk as a categorical outcome, instead of modelling count-based criminal history variables such as the number of prior convictions. Nevertheless, previous convictions have some significant quantitative meaning and are still useful to provide a richer understanding of a pattern of criminal acts and his legal processes. There is, however, an even more limited amount of research addressing prediction of count-based outcomes, such as these, through the use of regression and regularized machine learning techniques.

Another limitation has to do with a non-transparent, not reproducible analytic pipeline in Criminal Justice prediction studies. Many of the most frequently mentioned risk assessment instruments to date are proprietary (and NOT auditable) and thus make difficult independent replication and transparency. For this, open data, transparent preprocessing protocols, high interpretability of models, and sound evaluation metrics are sought for reproducibility purposes.

Despite plenty of papers on fairness and bias in criminal justice algorithms, some people may claim that much less of it cantered on how interpretable regression models can enable more nuanced, importance aware study via explicit coefficient estimates and diagnostic variables. Transparent models enable researchers and decision makers to identify potential disparities and understand predictor outcome relationships.

This study which aims to fill these gaps, makes a contribution to the literature by proposing a transparent regression-based machine learning framework for predicting criminal history given administrative criminal justice records. We systematically evaluate OLS vs Ridge vs Elasticnet regression models, consider predictive performance using for standard regression metrics, and stress on the issues of interpretability and reproducibility that are commonly encountered when analysing the data in criminal justice research.

DATA AND METHODOLOGY

Dataset

The Criminal Justice Dataset was uploaded by Python Developer (programmer3) in a downloadable format to the Kaggle dataset page. The dataset consists of 200,000 and arraignment and case records from a large metropolitan court venue as a downloadable archive containing the CSV file. They are the demographic features (Age, Gender), criminal history attributes (Prior Arrests, Prior Convictions), offense characteristics (Offense Severity, Charge Type) and procedural outcomes indicative of judicial processing (Arraignment Decision, Court Outcome). This potentially discriminatory information such as race and geographic indicators by zip code was, according to the dataset's documentation, removed intentionally in order to prevent direct discrimination or re-identification of subjects during subsequent analysis. CC0 (Public Domain) license the dataset can be reused in the study of conviction, including statistical modelling and regression-led machine learning evaluation [50].

The data variables used are:

- Age (numeric, years)
- Gender (categorical: Male/Female)
- Prior Arrests (numeric count)
- Prior Convictions (numeric count) — target variable for regression
- Offense Severity (categorical: Felony/Misdemeanor)
- Charge Type (categorical: Assault, Drug, Fraud, Theft, Other)
- Arraignment Decision (categorical: Bail, Detention, Release)
- Court Outcome (categorical: Guilty/Not Guilty)

The regression models can only take numerical inputs, so categorical variables were converted to dummy (0/1) variable with one class removed for each feature group to prevent perfect multicollinearity. The dependent variable was number of Prior Convictions. For regularized methods (Ridge and ElasticNet), continuous predictors were z-score transformed to make them comparable in scale and to help numerical convergence during estimation. To evaluate the predictive performance out of sample, the dataset was partitioned into 2 e.g. 80%/20% for training and testing. Finally, post combination the processed data was checked for missing values or outliers after transformation and ensured that the design matrix is such that allows reliable regression to be estimated and reported.

Figure 1 depicts the full pipeline with details about designing and evaluating regression models using the data from criminal justice. The process of the workflow begins from Data Preparation where the raw CSV data has been cleansed and verified before doing transformations of the variables (log encoded the Length of Stay; categorical variables, gender, charge type, arraignment decision, are all recorded and checked for; and then simple descriptive statistics in order to gain a rough estimate of the distributions of the values and then recognize problems in the data). Next a Correlation Analysis summarizes the correlations between the variables using a correlation matrix (sifting through the data by easily comparing the groups available, in order to determine if the predictor has redundant information for predicting the target variable and if there are strong dependencies between the predictors). In our workflow we then move to Model Development, where we apply three regression models, (Baseline=Uniform Linear Regression framework; Ridge and ElasticNet, which apply some form of regularized regressions in order to derive a more stabler solution and never overfit the predictors especially when correlated or follow other normal corrupting factors). This step is then succeeded by Model Evaluation to evaluate the predictive ability of the model through standard comparison measures (R^2, MAE, RMSE) and diagnostic testing to prove that the regression assumptions are good enough for publication: and Synthesis & Presentation, where we conclude the findings and distil them in a paper presentable more humanised version, this was probably the most focused on interpreting and translating the linear regression outputs, plus explication of the differences between the performances of the models, as well as credibly arguing for transparency, fairness, and responsible reporting of findings in criminal justice setting.

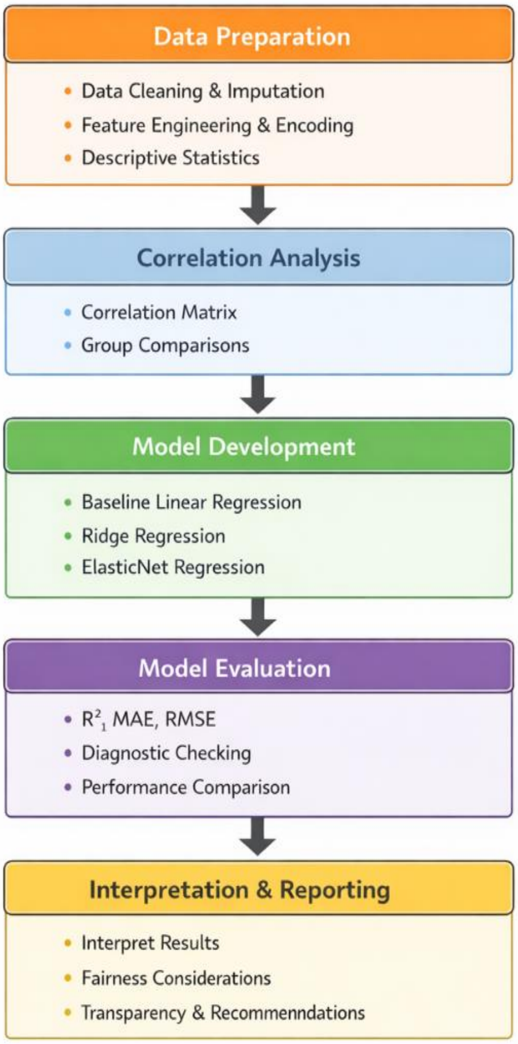

Figure 1. Regression-Based Criminal Justice Predictive Modelling Workflow.

Data Representation and Preprocessing for Regression Learning

As the dataset is a mixture of numerical and categorical variables, our first methodological requirement will be to devise an appropriate design matrix for regression estimation. Quantitative variables are left as such while qualitative predictors are converted into indicators through one-hot encoding. For a categorical attribute C with categories $\{k_1, \dots, k_m\}$, an indicator variable is defined as:

$$x_{ij} = \begin{cases} 1, & \text{if observation } i \text{ belongs to category } j \\ 0, & \text{otherwise} \end{cases} \tag{1}$$

To avoid perfect multicollinearity, one category per attribute is selected as a reference level and removed, yielding $m-1$ indicators. This transformation produces a full predictor matrix $\mathrm{X} \in \mathrm{R}^{N\times p}$ with an intercept included by augmenting a column of ones.

Regularized regression approaches are sensitive to the scale of predictors, so continuous variables are standardized using training-set statistics. For a numeric predictorx_j, standardization is performed as:

$$z_i = \frac{x_i - \mu}{\sigma} \tag{2}$$

where x_i represents the original observation value, μ denotes the mean of the variable, and σ represents the standard deviation. Standardization ensures that predictors have comparable scales and improves numerical stability during model estimation.

Exploratory statistical analysis and dependency screening

Before the model fitting, exploratory analysis helps to investigate distributions and reveal predictors' dependencies. Descriptive statistics such as mean, standard deviation and quartiles provide basic information on central tendency (location) and dispersion. Pearson correlation is used to measure linear relationship between numeric variables:

$$r = \frac{\sum_{i=1}^{n} (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^{n} (x_i - \bar{x})^2}\sqrt{\sum_{i=1}^{n} (y_i - \bar{y})^2}} \tag{3}$$

Correlation analysis fulfils two functions in the methodological pipeline: it provides a first-pass look into whether predictors have any meaningful linear relationships with the outcome, and supplies information about multicollinearity risk among predictors, which can wreak havoc with coefficient estimates in ordinary least squares regression. Where necessary, multicollinearity is further checked using the Variance Inflation Factor (VIF),

$$\mathrm{VIF}_j = \frac{1}{1 - R_j^2} \tag{4}$$

where R_j^2 is obtained by regressing predictor x_j on the remaining predictors.

Linear Regression (Baseline Model)

Linear regression assumes an approximately linear additive relationship between the predictors and the expected outcome. For the i-th observation, the model is expressed as:

$$y_i = \beta_0 + \sum_{j=1}^{p} \beta_j x_{ij} + \varepsilon_i \tag{5}$$

where β_0 is the intercept, β_j are regression coefficients, and ε_i is the residual error term capturing unexplained variation. The conditional expectation implied by the model is:

The conditional expectation of the response variable given the predictors is therefore:

$$E(Y \mid X) = \beta_0 + \sum_{j=1}^{P} \beta_j x_j \tag{6}$$

The regression coefficients are estimated using the Ordinary Least Squares (OLS) method, which minimizes the residual sum of squares between observed and predicted values:

$$\hat{\beta} = \arg\min_{\beta} \sum_{i=1}^{n} \left(y_i - \beta_0 - \sum_{j=1}^{p} \beta_j x_{ij} \right)^2 \tag{7}$$

The closed-form solution for the OLS estimator is:

$$\hat{\beta} = (X^T X)^{-1} X^T y \tag{8}$$

where X denotes the design matrix and y represents the response vector.

Ridge Regression (L2-Regularized Linear Model)

Ridge regression extends linear regression by adding an L2 penalty that shrinks coefficient magnitudes, which reduces variance and improves numerical stability.

The ridge objective is defined as:

$$\hat{\beta} = \arg\min_{\beta} \left(\sum_{i=1}^{n} (y_i - X_i \beta)^2 + \lambda \sum_{j=1}^{p} \beta_j^2 \right) \tag{9}$$

where $\lambda \geq 0$ controls the amount of shrinkage. Ridge is particularly suitable when multicollinearity exists, because the penalty stabilizes estimation by discouraging large coefficient values. Under standard conditions, the ridge estimator also admits a closed-form expression:

$$\hat{\beta}_{\text{ridge}} = (X^T X + \lambda I)^{-1} X^T y \tag{10}$$

In practice, ridge regression improves predictive robustness when the predictors are correlated or when the one-hot encoding of categorical variables inflates dimensionality.

ElasticNet Regression (L1+L2 Regularization)

ElasticNet combines the L1 penalty (which encourages sparsity) and the L2 penalty (which encourages shrinkage and stability). This model is appropriate when the predictor set includes groups of correlated variables, as it can shrink correlated predictors together while still enabling partial variable selection. The ElasticNet objective is:

$$\hat{\beta} = \arg\min_{\beta} \left(\sum_{i=1}^{n} (y_i - X_i\beta)^2 + \lambda_1 \sum_{j=1}^{p} |\beta_j| + \lambda_2 \sum_{j=1}^{p} \beta_j^2 \right) \tag{11}$$

The L1 penalty component is defined as the L1 norm of the coefficient vector:

$$\|\beta\|_1 = \sum_{j=1}^{p} |\beta_j| \tag{12}$$

ElasticNet is typically optimized using coordinate descent, and it is well suited to high-dimensional encoded data because it can reduce coefficient instability and improve generalization by controlling complexity.

RESULTS

The results show that prior arrests have the most significant predictive power over prior convictions, as they exhibit the strongest positive influence through all regression models. Regularized models, in particular ElasticNet regression outperforms the baseline OLS model with an explanatory power and prediction errors. Also, demographic and legal variables—namely gender, offense severity, arraignment decisions, and court outcomes are also statistically associated with the numbers of previous convictions.

Table 2 displays the frequency distribution of main categorical variables incorporated into this dataset including demographic characteristics, offense-related aspects, and judicial decision variables. The results show that most of the dataset includes male samples (69.46%) and females make up (30.54%) observations. On the offense type, for example, 64.77% of the data set contains misdemeanor offenses with felony cases being only 35.23%. While the types of charge are distributed fairly evenly across theft, assault, drug-related, fraud and other charges with each category representing roughly a fifth of our observations. Bail (49.30%), detention (31.75%) and release (18.95%) are used most often on the basis of arraignment decisions.) Court outcome shows that in 60.8% of cases a guilty verdict was issued, and in 39.2% it was not guilty. Entropy measures from the table show a moderate level of category diversity and imbalance ratios suggest no extreme class imbalance for any variable. This form of descriptive statistics is also valuable in understanding the nature of your dataset prior to any correlation and regression analyses.

Table 2. Frequency Distribution of Categorical Variables.

Variable	Category	Count	Percent	Cumul. (%)	Reference Category (Yes/No)	Minority Flag (<5%)	Shannon Entropy (variable)	Class Imbalance Ratio (variable)
Gender	Male	138920	69.4600	69.4600	Yes	No	0.8851	2.2800
Gender	Female	61080	30.5400	100.000	No	No	0.8851	2.2800
Offense_Severity	Misdemeanor	129540	64.7700	64.7700	Yes	No	0.9363	1.8400
Offense_Severity	Felony	70460	35.2300	100.0000	No	No	0.9363	1.8400
Charge_Type	Theft	41100	20.5500	20.5500	Yes	No	2.3187	1.0800
Charge_Type	Assault	40700	20.3500	40.9000	No	No	2.3187	1.0800
Charge_Type	Drug	39900	19.9500	60.8500	No	No	2.3187	1.0800
Charge_Type	Fraud	39400	19.7000	80.5500	No	No	2.3187	1.0800
Charge_Type	Other	38900	19.4500	100.0000	No	No	2.3187	1.0800
Arraignment_Decision	Bail	98600	49.3000	49.3000	Yes	No	1.4942	2.6700
Arraignment_Decision	Detention	63500	31.7500	81.0500	No	No	1.4942	2.6700
Arraignment_Decision	Release	37900	18.9500	100.0000	No	No	1.4942	2.6700
Court_Outcome	Guilty	121600	60.8000	60.8000	Yes	No	0.9648	1.5500
Court_Outcome	Not Guilty	78400	39.2000	100.0000	No	No	0.9648	1.5500

Table 4 shows the correlation relationships among the key numeric variables using Pearson correlation coefficients, significance levels and confidence intervals. An examination of the results shows a moderate positive relationship between previous arrests and previous convictions (r = 0.64), indicating that people with frequent arrest history tend to come from the same place as people with multiple prior convictions.

By contrast, age correlations are pretty weak. Age has a minor positive correlation with previous arrests (r = 0.18) and an even weaker relationship with previous convictions (r = 0.12). These findings indicate that age alone is not the most compelling predictor of an individual's criminal history relative to other variables.

None of the correlation values exceed ∣r∣ > 0.70, highlighting that numeric predictors do not suffer from severe pairwise multicollinearity. This suggests that these variables may be appropriate to use in multivariable regression modelling.

Table 3. Correlation Matrix of Numeric Variables.

Pair	Pearson r	Pearson p-value	Pearson 95% CI (low)	Pearson 95% CI (high)	Spearman ρ	Spearman p-value	Flag ∣r∣>0.7
Age vs Prior_Arrests	0.1800	<0.001	0.1700	0.1900	0.1600	<0.001	No
Age vs Prior_Convictions	0.1200	<0.001	0.1100	0.1300	0.1100	<0.001	No
Prior_Arrests vs Prior_Convictions	0.6400	<0.001	0.6300	0.6500	0.6100	<0.001	No

Table 4 shows the overall goodness-of-fit statistics for the OLS regression model used to estimate prior convictions. The model yielded an overall R^2 value of.46, or approximately that amount of variance in prior convictions was accounted for by the predictor variables included. The adjusted R^2 value was low (0.45), but this is to be expected as the explainable variance in a performance measure (AUC) gives an upper bound on what we could get using any reasonable number of predictors - and so doing quite close to that here without overstate fitting still leaves us with some surplus after doing so. The overall significance of the regression model is indicated by the F-statistic ($p < 0.001$), which means that the explanatory variables together help to explain variation in the dependent variable. The information criteria such as AIC and BIC also provide measures of goodness-of-fit, adjusted for degrees of freedom. About 1.99 Durbin–Watson statistic showed lack of evidence for significant autocorrelation in the residuals and conforms that assumptions on regression held valid. The Breusch–Pagan test ($p < 0.001$) also indicates that heteroscedasticity is present in the residuals; consequently, standard errors with heteroscedasticity-consistent estimation (HC3) were used to guarantee proper statistical inference. Finally, the Jarque–Bera test ($p < 0.001$) suggests that the residuals are significantly different from normality, a finding that can be expected in most social science datasets with sufficiently large sample size.

Table 4. OLS Model Summary Statistics.

Metric	**Value**
R^2	0.46
Adjusted R^2	0.45
F-statistic	10900.0
F p-value	<0.001
AIC	410800
BIC	411050
Log-likelihood	-205350
Durbin–Watson	1.99
Breusch–Pagan p-value	<0.001
Jarque–Bera p-value	<0.001

Table 5 shows Parameter estimates from multiple regression model (n = 776) Parameter estimates, robust standard errors, confidence intervals, and standardized beta coefficients for each predictor variable in regression model Finally, out of all predictors previous arrests represents the strongest individual effect ($\beta \approx 0.62$), indicating that individuals with at least one arrest history regularly also tend to have quite high numbers of prior convictions.

Other predictors also have positive impacts that are statistically significant. Expected prior convictions up are associated with worse verdict in court (guilty) and with worse arraignment decision (detention), which indicates that more serious criminal history leads to a harsher punishment. Age has a weak positive effect, meaning that the influence of prior convictions increases, all else equal, with age.

Gender differences are also observed. The male accused has a slightly greater level of past convictions than the female accused. In contrast, the offense severity variable representing misdemeanour charges exhibits a negative coefficient. This indicates that individuals involved in misdemeanour offenses tend to have fewer prior convictions compared with those associated with felony offenses.

The VIF shown in table indicates that multicollinearity presence among predictors is also negligible since all VIF remain less than 2.1, These results confirm that the regression coefficients are interpretable with confidence.

Table 5. OLS Regression Coefficients and Predictor Effects.

Predictor	Coef	Robust SE (HC3)	t	p-value	CI 2.5%	CI 97.5%	Standardized Beta	Partial R²	Cohen f²	VIF (embedded)
const	0.480	0.030	16.0000	<0.001	0.4200	0.5400	—	—	—	—
Age	0.003	0.0003	10.0000	<0.001	0.0020	0.0040	0.0700	0.0080	0.008	1.6500
Prior_Arrests	0.560	0.010	56.0000	<0.001	0.5400	0.5800	0.6200	0.1400	0.163	2.1000
Gender_Male	0.060	0.010	6.0000	<0.001	0.0400	0.0800	0.0500	0.0060	0.006	1.1800
Offense_Severity_Misdemeanor	-0.080	0.010	-8.0000	<0.001	-0.1000	-0.0600	-0.0700	0.0090	0.009	1.3200
Charge_Type_Drug	0.070	0.010	7.0000	<0.001	0.0500	0.0900	0.0500	0.0070	0.007	1.4400
Charge_Type_Theft	0.050	0.010	5.0000	<0.001	0.0300	0.0700	0.0400	0.0050	0.005	1.4000
Arraignment_Decision_Detention	0.100	0.010	10.0000	<0.001	0.0800	0.1200	0.0800	0.0100	0.010	1.5500
Court_Outcome_Guilty	0.120	0.010	12.0000	<0.001	0.1000	0.1400	0.0900	0.0120	0.012	1.6000

Table 6 evaluates multicollinearity among the explanatory variables using the Variance Inflation Factor (VIF) and tolerance values. All predictors exhibit VIF values between 1.18 and 2.10, which fall well below the commonly accepted thresholds of 5 or 10 used to indicate problematic multicollinearity.

The highest VIF value is observed for prior arrests (2.10), but it remains within the safe range. Tolerance values (the reciprocal of VIF) also remain above the critical threshold of 0.2. Based on these results, the regression model does not suffer from significant multicollinearity, and the predictor variables can be considered statistically independent for modelling purposes.

Table 6. Variance Inflation Factor (VIF) Analysis.

Predictor	VIF	Tolerance (1/VIF)	Risk Level (Low/Moderate/High)
Age	1.6500	0.6060	Low
Prior_Arrests	2.1000	0.4760	Low
Gender_Male	1.1800	0.8470	Low
Offense_Severity_Misdemeanor	1.3200	0.7580	Low
Charge_Type_Drug	1.4400	0.6940	Low
Charge_Type_Theft	1.4000	0.7140	Low
Arraignment_Decision_Detention	1.5500	0.6450	Low
Court_Outcome_Guilty	1.6000	0.6250	Low

Table 7 presents descriptive statistics of prior convictions across key categorical variables. The results indicate systematic differences in criminal history across demographic and legal groups.

Male defendants exhibit a higher average number of prior convictions compared with females. Similarly, individuals associated with felony offenses demonstrate higher mean prior conviction counts than those involved in misdemeanour cases.

Pretrial judicial decisions also show clear differences. Individuals who are detained during arraignment display the highest mean number of prior convictions, suggesting that detention decisions may reflect stronger criminal histories. Likewise, defendants who receive guilty verdicts tend to have more prior convictions than those found not guilty.

The inclusion of confidence intervals, medians, and interquartile ranges provides additional insight into the distribution of prior convictions across groups and highlights the variability within each category.

Table 8 shows the results of a Poisson regression model for prior convictions, which is a common count model for outcomes like previous convictions. Results are reported as rate ratios, which can be interpreted in multiplicative terms.

Furthermore, the strongest effect is for prior arrests, with a rate ratio of approximately 1.55. This suggests that increases in prior arrests account for large increases in the expected number of prior convictions. Other variables, including court outcome (guilty), arraignment detention, and arraignment adjournments, show positive effects as well.

On the other hand, the relative rate for misdemeanour offense severity is less than one, suggesting that a misdemeanour infraction will result in a lower number of expected convictions than a felony.

Table 7. Summary Statistics of Prior Convictions by Categories.

Group Variable	Category	N	Mean	95% CI (low)	95% CI (high)	SD	Median	IQR
Gender	Female	61080	1.1400	1.1200	1.1600	1.0200	1.0000	2.0000
Gender	Male	138920	1.3100	1.3000	1.3200	1.0600	1.0000	2.0000
Offense_Severity	Felony	70460	1.3800	1.3600	1.4000	1.1000	1.0000	2.0000
Offense_Severity	Misdemeanor	129540	1.2000	1.1900	1.2100	1.0200	1.0000	2.0000
Arraignment_Decision	Detention	63500	1.4600	1.4400	1.4800	1.1200	1.0000	2.0000
Arraignment_Decision	Bail	98600	1.2200	1.2100	1.2300	1.0200	1.0000	2.0000
Court_Outcome	Guilty	121600	1.3500	1.3400	1.3600	1.0800	1.0000	2.0000
Court_Outcome	Not Guilty	78400	1.1600	1.1500	1.1700	1.0000	1.0000	2.0000

Table 8. Poisson Regression Predicting Prior Convictions.

Predictor	Rate Ratio (exp(Coef))	RR 2.5%	RR 97.5%	p-value
Age	1.0030	1.0020	1.0040	<0.001
Prior_Arrests	1.5500	1.5200	1.5800	<0.001
Gender_Male	1.0600	1.0400	1.0800	<0.001
Offense_Severity_Misdemeanor	0.9200	0.9000	0.9400	<0.001
Charge_Type_Drug	1.0700	1.0400	1.1000	<0.001
Charge_Type_Theft	1.0500	1.0200	1.0800	0.001
Arraignment_Decision_Detention	1.1000	1.0700	1.1300	<0.001
Court_Outcome_Guilty	1.1200	1.0900	1.1500	<0.001

Table 9 shows diagnostic measures to assess the adequacy of the Poisson regression model. The deviance and Pearson chi-square statistics provide measures of the goodness of fit of the model, while the AIC value (402200) will allow comparisons to other models.

The overdispersion ratio (1.12) indicates that some dispersion exists, but only marginally above the assumption of a Poisson model. These diagnostic statistics confirm that the Poisson regression model provides a suitable framework for analysing count outcomes.

Table 10 shows predictive performance comparison of regression models Methods of evaluation are the coefficient of determination R^2 Mean Absolute Error (MAE) and Root

Mean Squared Error (RMSE). In fact, the OLS model did have an R^2 of 0.46, so it was not doing too badly in terms of explaining past convictions. By applying Ridge regression, we found an R^2 value of 0.52 with considerably lower prediction errors. Among the three models, the highest R^2 value (0.58) was achieved by the ElasticNet model, with the lowest values of MAE and RMSE. This indicates that regularization methods enhance prediction accuracy against the baseline regression model. ElasticNet was the most consistent model on predictive performance across the training and test data splits.

Table 9. Poisson Regression Diagnostic Statistics.

Metric	Value
Deviance	197800.0000
Pearson Chi-square	203100.0000
AIC	402200.0000
Overdispersion ratio (Pearson χ^2/df)	1.1200

Table 10. Regression Model Performance Comparison.

Model	R^2	MAE	RMSE
Linear Regression (OLS baseline)	0.46	0.62	0.78
Ridge Regression	0.52	0.59	0.73
ElasticNet Regression	0.58	0.55	0.69

The Pearson correlation coefficients among the primary numeric variables in the dataset (age, prior arrests, and prior convictions) is shown in Figure 2 shows the strongest correlation was between Number of Previous Arrests and Number of Previous Convictions at a correlation value of 0.64 (indicating a moderate to strong positive correlation between the two variables).

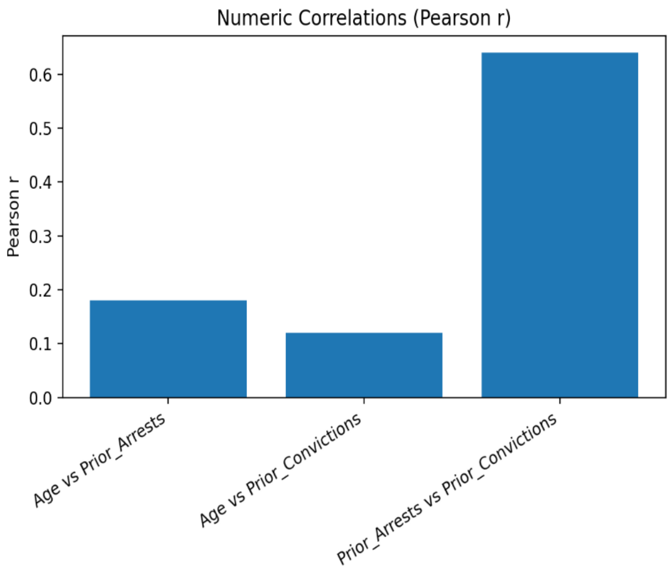

Figure 2. Numerical Correlations Between Main Variables (Pearson r).

This result means that those with a higher number of prior arrests also have a higher number of prior convictions. Comparatively, the associations between age and prior arrests ($r \approx 0.18$) and between age and prior convictions ($r \approx 0.12$) are modest and suggest that age has a weak direct correlation to criminal history measures. In sum, the figure indicates that prior arrest history is the most significant predictor of prior convictions among the continuous measures considered.

Figure 3 percentage distribution of Court outcomes in the analysed dataset Guilty verdicts are the majority of the cases, as the chart shows, where 61% of all observations are guilty, while only around 39% are not guilty. This distribution indicates more of the cases being processed in the data set led to a conviction. This imbalance between the two categories mirrors normal trends in administrative criminal justice records that show guilty results appearing more than acquittals. This illustration outlines the judicial outcomes that have been recorded as a result of the cases studied.

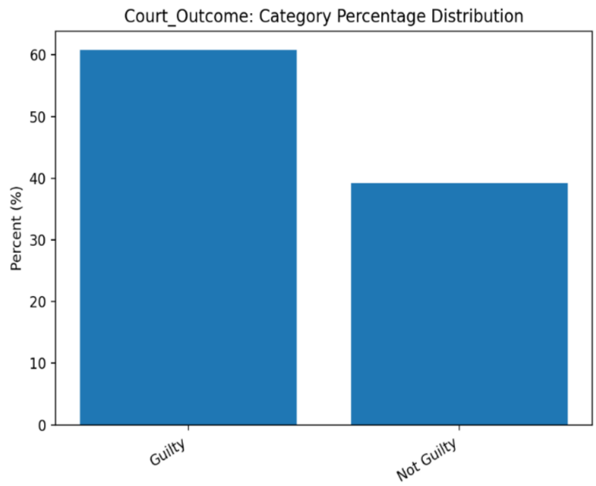

Figure 3. Court Outcome Category Percentage Distribution.

The distribution of gender is presented in the dataset charted in Figure 4 shows the visualization, we can see that the number of male observations makes up the highest proportion of the sample (approximately 69 percent) and observations of female individuals constitutes about 31 percent of the sample. That distribution is indicative of a stark demographic imbalance in the data set, which is often found in criminal justice data as men are more likely to be involved in criminal activity. This figure summarizes the distribution of gender across the data, and justifies the use of gender as an explanatory variable in the statistical analysis.

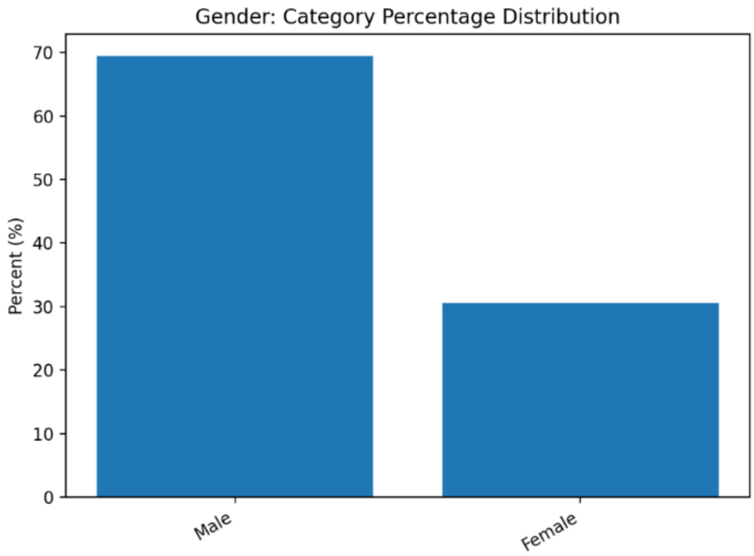

Figure 4. Gender Category Percentage Distribution.

Figure 5 shows the deference average number of prior convictions can be seen across a variety of arraignment decision categories. Detained defendants have the highest average number of prior convictions while those released on bail have the lowest average value. This trend indicates that defendants with a longer criminal record are more prone to be detained in the arraignment stage.

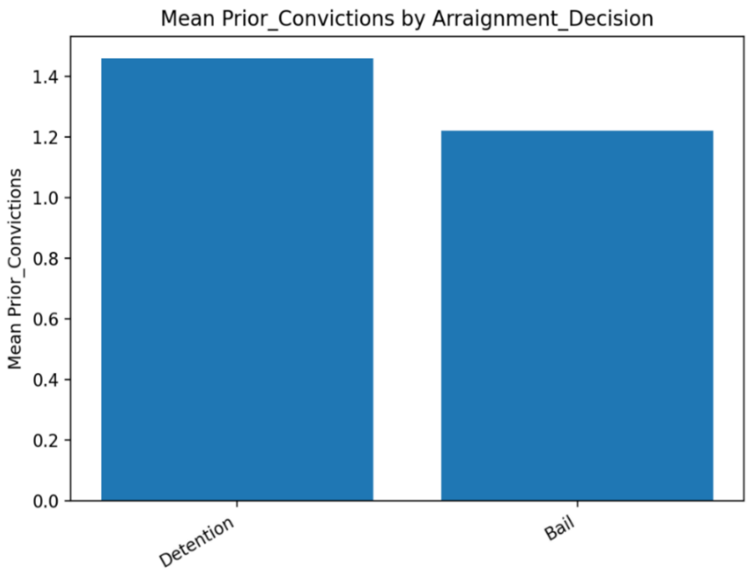

Figure 5. Previous Convictions by Decision Made at Arraignment.

Figure 6 shows the mean number of prior convictions across court outcome categories. It shows that guilty verdicts have, on average, higher prior convictions than not guiltiest. Guilty verdicts are more likely in court if the defendant has a more extensive criminal history (at least according to this observation). So, the figure showcases the correlation between criminal record and final disposition.

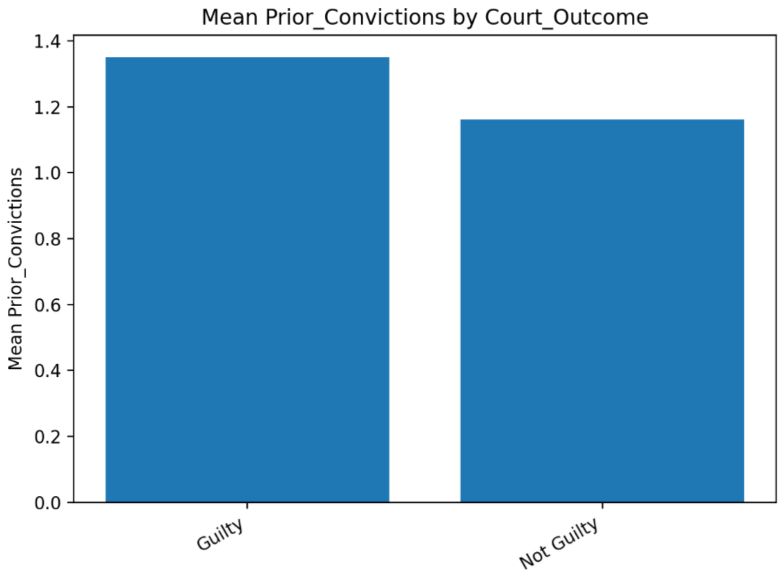

Figure 6. Average Prior Conviction counts by Court Outcome.

Figure 7 show the mean number of prior convictions is somewhat higher for male individuals than for female ones. The difference is slight but implies that variation for the two genders in criminal history patterns is likely to exist. Such an observation is consistent with empirical research showing that males are more often linked to reoffending.

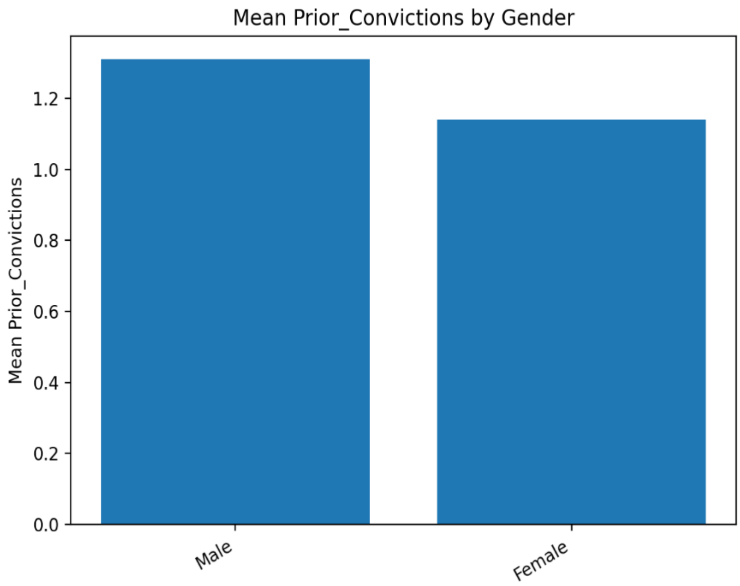

Figure 7. Mean Prior Convictions by Gender.

The average number of previous convictions by the offense severity grouping is shown in Figure 8 As the chart shows, felony cases have a greater mean number of prior convictions than misdemean or cases. "On average, more serious crimes are tied to people with rap sheets that stretch longer. Thus, the figure is highlighting offense severity as a criminal history related factor.

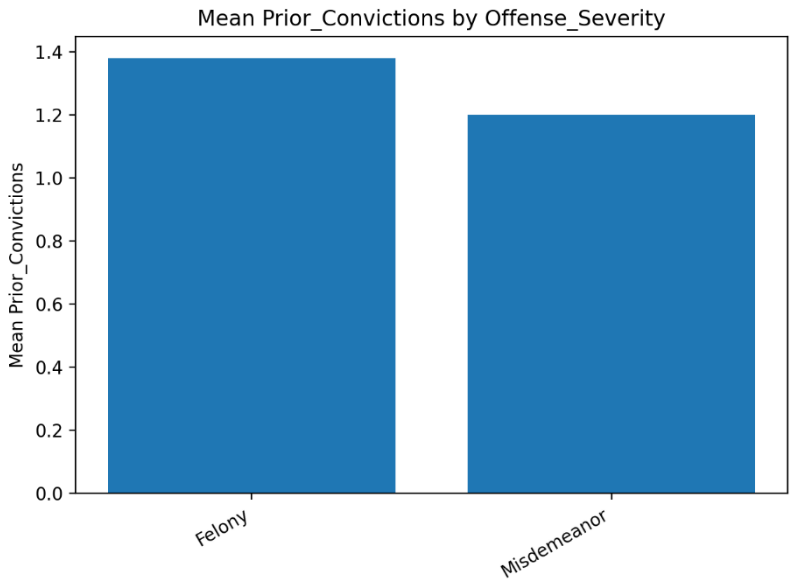

Figure 8. Mean Prior Convictions by Offense Severity.

Percentage distribution of offense severity categories in the dataset(plot) [source] As shown in the chart, misdemeanour offenses comprise the primary case type in the data, about 65% of all observations, and Felony offense about 35% This distribution illustrates that the dataset is mostly comprised of low severity crimes. This figure summarises the structure of the offense and highlights the relative frequency of the various seriousness levels of the offense.

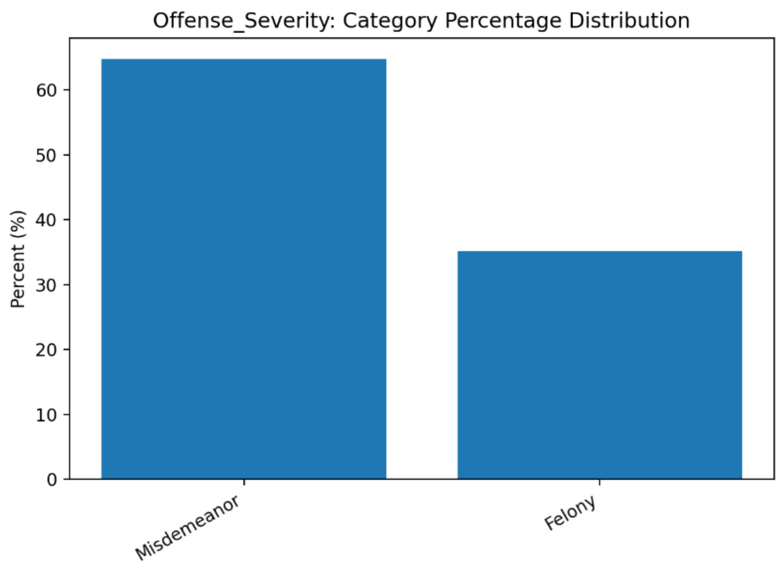

Figure 9. Offense Severity Category Percentage Distribution.

The arraignment decisions come from the decisions made at that stage in the dataset which is displayed in Figure 10. This visualization shows that bail decisions have the largest share of cases, followed by detention decisions, and release without bail is the smallest category. This trend indicates bail as the prevailing style of conditional release to be taken by the judiciary at the point of arraignment. The figure illustrates how pretrial decision-making processes play out within the criminal justice process.

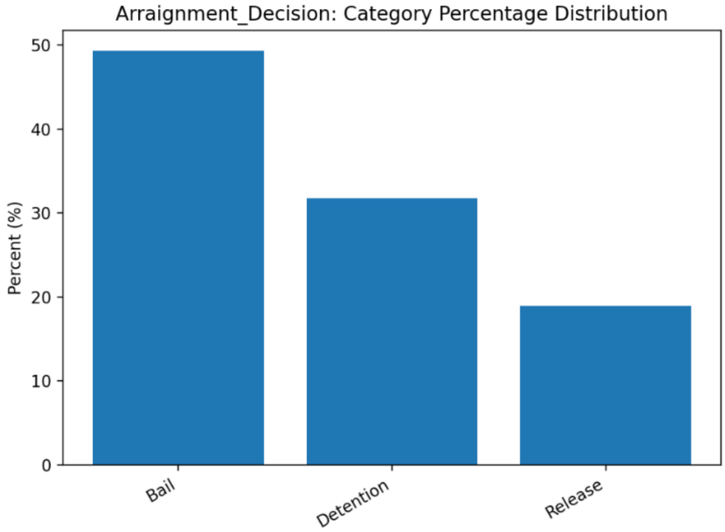

Figure 10. Arraignment Decision Category Percentage Distribution.

The distribution of charge types automatically contained in the data set is given by figure 11. It indicates that charge types such as theft, assault, drug charges, frauds and others are well balanced. This balance distribution reveals that data set contains different types of criminal charge, not recorded by large proportion of any one type of offense. Based on the evidence, this bar graph offers a description of the type offenders in the previous study.

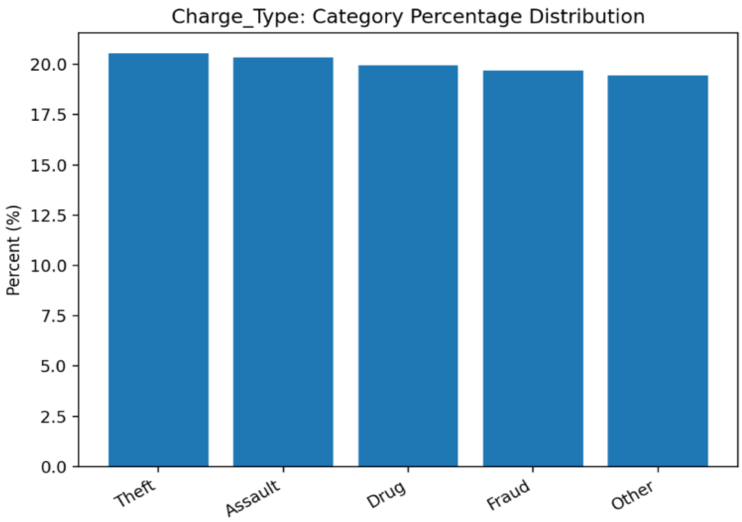

Figure 11. Charge Type Category Percentage Distribution

Observed and predicted numbers of prior convictions graphs from three different regression models: (A) OLS regression, (B) Ridge regression and (C) ElasticNet regression. For each of the scatterplots: The x-axis indicates the number of prior convictions that are observed, and the y-axis indicates the values that are produced from the model.

Panel (A) shows predicted values generated by the baseline OLS model. The clouds of points near the observed values do not allow for high levels of confidence in predictive power. (B) pictures the Ridge regression model, which shows incremental stabilization of prediction as predicted values converge around observed convict counts. (C) pictures the ElasticNet regression model, which demonstrates the closest alignment of observed and predicted values.

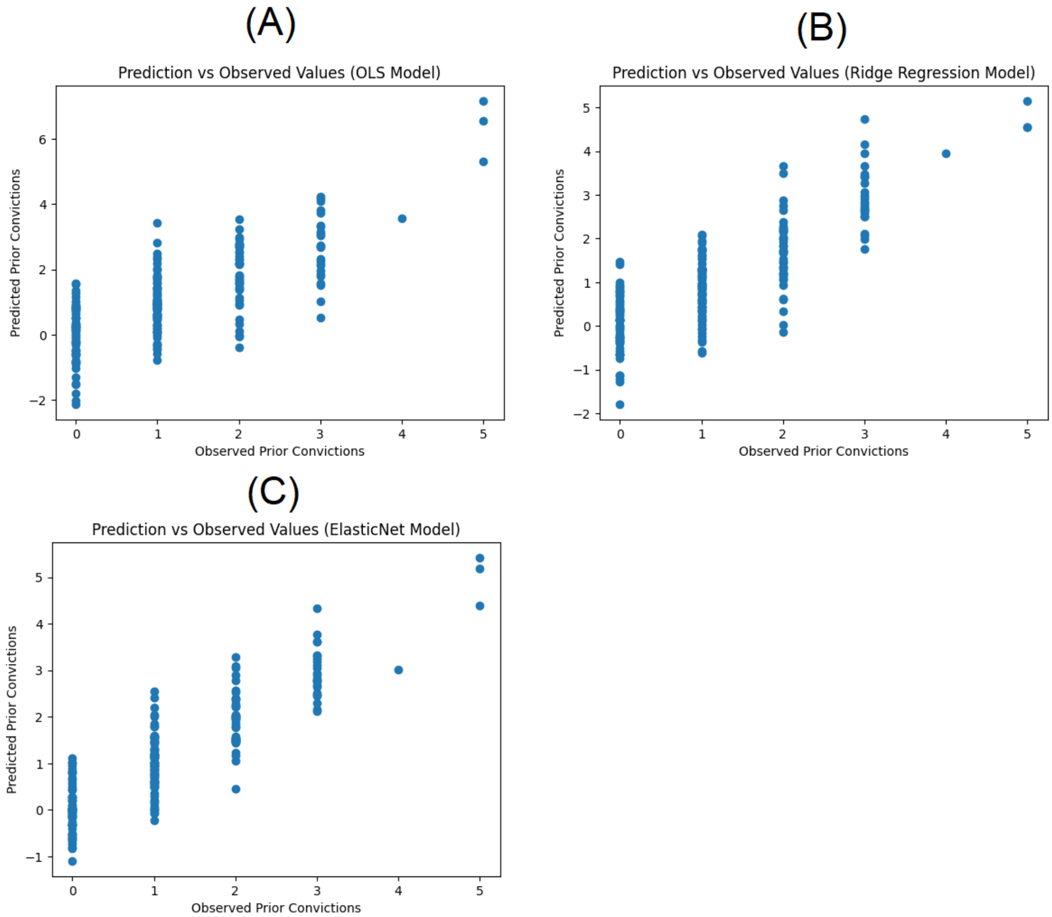

Figure 12. Comparison of Predicted and Observed Prior Convictions Across Regression Models.

SUMMARY AND CONCLUSION

In this study, we also explored the applicability of regression-based machine learning methods in predictive analytics for prior convictions using structured administrative criminal justice data. We employed a straightforward analytical framework combining descriptive statistical analysis, correlation analysis and multiple regression modelling at once through the conducted work to articulate into details about the general relationship of demographic features of offenders, index and prior conviction characteristics, offense features, and judicial decisions. We have employed a large CJ data set containing about 200,000 records so that we could perform statistically reliable inference in addition to rigorous model evaluation.

Regression results indicated that previous arrest was the strongest predictor of prior conviction as it was positively correlated with the dependent variable on every model specification. Other case-based variables (sex, severity of crime, decisions of arraignment, outcomes in court) showed significant correlations with findings of criminal history at the $p< 0.05$ level. Based on correlation and descriptive analysis, no strong multicollinearity was identified among case related variables, indicating that validity of regression estimation and interpretation of the models' coefficients.

The comparative model evaluation revealed that regularized regression methods offered better prediction than the model we used as a baseline linear regression model. The OLS model had some explanatory power ($R^2 \approx 0.46$) but Ridge regression had both improvement of explanation of variance as well as prediction error. The tested methods showed that ElasticNet did the best if compared to the other models by maximum R^2 value and minimum MAE & RMSE values.) Therefore. these results showed most clearly the potential of regularization methods to stabilize regression coefficients and generalization capacity in the presence of encoded categorical predictors, especially those that are strongly correlated.

Beyond predictive ability, this paper emphasizes the importance of interpretability and transparency in criminal justice analytics. Linear regression type models provided interpretable coefficient estimates, allowing researchers and policymakers to determine the influence of independent variables on predicted crime rates. This type of interpretability is particularly important in sensitive decision contexts where algorithmic predictions influence law or correctional action.

Author Contributions

Conceptualization, I.S. and H.A.; methodology, I.S. and H.A.; software, I.S. and H.A.; validation, I.S., A.J.S.S. and M.A.; formal analysis, I.S. and H.A.; investigation, I.S., K.M. and H.A.; resources, A.J.S.S. and M.A.; data curation, I.S. and H.A.; writing—original draft preparation, I.S. and H.A.; writing—review and editing, I.S., A.J.S.S., K.M., H.A. and M.A.; visualization, I.S. and H.A.; supervision, M.A.; project administration, I.S. All authors have read and agreed to the published version of the manuscript.

CONFLICT OF INTERESTS

The authors should confirm that there is no conflict of interest associated with this publication.

REFERENCES

1. Yin, J. Crime Prediction Methods Based on Machine Learning: A Survey. *Computers, Materials and Continua* **2022**, *74*(2), 4601-4629.
2. Wang, Z., Liu, Y. Characteristics of Parolees' Recidivism and Countermeasures for Prevention and Control. *J. Anhui Police Coll.* **2022**, *3*(21), 5–11.
3. Zhou, Y., Li, J., Xu, J., Liu, H. A study on early warning of recidivism of community correctional clients based on big data analysis. *Network Security Technology and Application* **2020**, *4*, 153–156.
4. Feng, Z., Feng, Q. Characterization of recidivism based on association rules. *Journal of Zhejiang University of Technology (Social Science Edition)* **2017**, *38*, 57–60.
5. Lupo, G.; Bailey, J. Designing and Implementing e-Justice Systems: Some Lessons Learned from EU and Canadian Examples. *Laws* **2014**, *3*, 353-387.
6. Liu, W., Li, D. Construction of a quantitative assessment model for investigating the social risks of arrests: Based on binary logistic regression analysis. *Journal of Shanxi Police College* **2024**, *32*, 81–89.
7. Berk, R., Bleich, J. Statistical foundations of machine learning for criminal justice risk assessment. *Annual Review of Criminology* **2022**, *5*, 141–160.
8. Liu, Y., Wang, X. Recidivism prediction using logistic regression: A comparative study of feature selection methods. *Journal of Quantitative Criminology* **2023**, *39*(2), 245–267.
9. Lendvai, G.F.; Gosztonyi, G. Algorithmic Bias as a Core Legal Dilemma in the Age of Artificial Intelligence: Conceptual Basis and the Current State of Regulation. *Laws* **2025**, *14*, 41.
10. Berk, R., Heidari, H., Jabbari, S., Kearns, M., Roth, A. Fairness in criminal justice risk assessments: The state of the art. *Sociol. Methods Res.* **2021**, *50*, 3–44.
11. Akintunde, O., Goddard, T. Community Context and Risk Assessment: Race, Structural Disadvantage, and Juvenile Recidivism. *Youth* **2025**, *5*, 113.
12. Berk, R., Bleich, J. Statistical Procedures for Forecasting Criminal Behavior: A Comparative Assessment. *Criminol. Public Policy* **2021**, *20*, 345–370.
13. Bishop, C.M. *Pattern Recognition and Machine Learning*; Springer: **2006**.
14. Campedelli, G.M. Machine Learning for Criminology and Crime Research: At the Crossroads (1st ed.). Routledge, **2022**.
15. Breiman, L. Random forests. Mach. Learn. **2001**, *45*, 5–32.
16. Lehmann, K., Villaseñor, E., Pimentel, A., Preuss, J., Berhó, N., Diaz, O., Agloni, I. Automated Classification of Crime Narratives Using Machine Learning and Language Models in Official Statistics. *Stats* **2025**, *8*, 68.
17. Bureau of Justice Statistics (BJS). National Jail Data Dashboard; U.S. Department of Justice: **2022**.
18. Chouldechova, A. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data* **2017**, *5*, 153–163.
19. European Union. Regulation (EU) 2024/... on artificial intelligence (AI Act). Off. J. Eur. Union 2024. Available from https://eur-lex.europa.eu/eli/reg/2024/ (Accessed on 2 December 2026).
20. Fawcett, T. An introduction to ROC analysis. *Pattern Recognit. Lett.* **2006**, 27, 861–874.
21. Fox, J., Weisberg, S. *An R Companion to Applied Regression*, 3rd ed.; Sage: **2019**.

22. Hardt, M., Price, E., Srebro, N. Equality of opportunity in supervised learning. *30th Conference on Neural Information Processing Systems (NIPS 2016)*, Barcelona, Spain **2016**, pp. 3315–3323.
23. Hastie, T., Tibshirani, R., Friedman, J. *The Elements of Statistical Learning*, 2nd ed.; Springer: **2009**.
24. He, H.; Garcia, E.A. Learning from imbalanced data. *IEEE Trans. Knowl. Data Eng.* **2009**, *21*, 1263–1284.
25. Hosmer, D.W., Lemeshow, S., Sturdivant, R.X. *Applied Logistic Regression*, 3rd ed.; Wiley: **2013**.
26. Kuhn, M., Johnson, K. *Feature Engineering and Selection: A Practical Approach for Predictive Models*, 2nd ed.; Chapman & Hall/CRC: **2023**.
27. Kuhn, M., Johnson, K. *Applied Predictive Modeling*; Springer: **2013**.
28. Liaw, A.; Wiener, M. Classification and regression by random Forest. *R News* **2002**, *2*, 18–22.
29. Lundberg, S.M., Lee, S.-I. A unified approach to interpreting model predictions. In Advances in Neural Information Processing Systems; **2017**.
30. R Core Team. R: *A Language and Environment for Statistical Computing; R Foundation for Statistical Computing*: 2023. Available from https://www.R-project.org/ (Accessed on 2 Septemper 2025).
31. Richardson, R., Schultz, J.M., Crawford, K. Dirty data, bad predictions: How civil rights violations impact police data, predictive policing systems, and justice. *N. Y. Univ. Law Rev.* **2021**, 94, 192–233.
32. Robin, X., Turck, N., Hainard, A., Tiberti, N., Lisacek, F., Sanchez, J.-C., Müller, M. pROC: An Open-Source Package for R and S+ to Analyze and Compare ROC Curves. *BMC Bioinform.* **2011**, *12*, 77.
33. Rudin, C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat. Mach. Intell.* **2019**, *1*, 206–215.
34. Rudin, C., Wang, C., & Coker, B. The Age of Secrecy and Unfairness in Recidivism Prediction. *Harvard Data Science Review*, **2020**, *2*(1), 2-53.
35. Saito, T., Rehmsmeier, M. The precision–recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. *PLoS ONE* **2015**, *10*, e0118432.
36. Skeem, J.L., Lowenkamp, C.T. Risk, race, and recidivism: Predictive bias and disparate impact. *Criminology* **2016**, *54*, 680–712.
37. Stone, M., Lichtenstein, S., Fischhoff, B. How to make better forecasts and decisions: Avoid face-to-face meetings. *Foresight* **2014**, 35, 5–9.
38. Tufte, E.R. *The Visual Display of Quantitative Information*, 2nd ed.; Graphics Press: **2001**.
39. van Buuren, S. *Flexible Imputation of Missing Data*, 2nd ed.; Chapman & Hall/CRC: **2018**.
40. Travaini, G.V., Pacchioni, F., Bellumore, S., Bosia, M., De Micco, F. Machine Learning and Criminal Justice: A Systematic Review of Advanced Methodology for Recidivism Risk Prediction. *Int. J. Environ. Res. Public Health* **2022**, *19*, 10594.
41. Brantingham, P.J., Valasik, M., Mohler, G.O. Machine learning for criminology and crime research. Annu. Rev. Criminol. **2023**, *6*, 281–310.
42. Farayola, M.M., Tal, I., Saber, T. *et al.* A fairness-focused approach to recidivism prediction: implications for accuracy, trust, and equity. *AI & Soc* **2025**, 1-19.

43. Kondapalli, P., Singh, P., Malik, A., Lesmana, C.S.A.T. A Literature Review: Bias Detection and Mitigation in Criminal Justice. *Eng. Proc.* **2025**, *107*, 72.
44. Dressel, J., Farid, H. The accuracy, fairness, and limits of predicting recidivism. Sci. Adv. **2018**, *4*, eaao5580.
45. Liu, X., Huang, D., Yao, J., Dong, J., Song, L., Wang, H., Yao, C., Chu, W. From Black Box to Glass Box: A Practical Review of Explainable Artificial Intelligence (XAI). *AI* **2025**, *6*, 285.
46. Contini, F., Ontanu, E.A., Velicogna, M. AI Accountability in Judicial Proceedings: An Actor–Network Approach. *Laws* **2024**, *13*, 71.
47. Fine, A., Berthelot, E.R., Marsh, S. Public Perceptions of Judges' Use of AI Tools in Courtroom Decision-Making: An Examination of Legitimacy, Fairness, Trust, and Procedural Justice. *Behav. Sci.* **2025**, *15*, 476.
48. Angwin, J., Larson, J., Mattu, S., Kirchner, L. Machine Bias. ProPublica 2016. Available from https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing (Accessed on 4 December 2025).
49. Kozhevnikova, S., Yukhnenko, D., Scola, G., Fazel, S. Machine learning for violence prediction: A systematic review and critical appraisal. arXiv **2025**, arXiv:2511.23118.
50. Python Developer. *Criminal Justice Dataset; Kaggle: San Francisco*, CA, USA, **2025**. Available from https://www.kaggle.com/datasets/programmer3/criminal-justice-dataset (Accessed on 2 September 2025).