Advanced International Journal for Research (AIJFR)
E-ISSN: 3048-7641 ● Website: www.aijfr.com ● Email: editor@aijfr.com

A Comparative Study On Surveillance Technology Using Machine Learning Models

V. S. Avinash[1], Arun Karthik V[2], Dr. N. Saraswathi[3]

[1,2] 1 Yr M. Tech Department of CSE, SRM Institute of Science & Technology Vadapalani campus, Chennai, Tamil Nadu, India
[3] Assistant Professor Department of CSE, SRM Institute of Science & Technology Vadapalani campus, Chennai, Tamil Nadu, India

Abstract

Predictive policing depends on an algorithm coupled with surveillance for anticipating crimes, efficiently and optimally deploying police resources. However, beneath the promises of enhanced safety and improved operational performance are serious concerns about fairness, privacy, and accountability. This study compares four models, namely, Random Forest, K-Nearest Neighbors, Multi-Layer Perceptron, and Kernel Density Estimation, against crime data from both Chicago and New York City. The models will be compared according to their accuracy, fairness, and interpretability, considering bias mitigation using SMOTE, fairness-aware preprocessing, and explainable AI techniques such as SHAP and LIME. Without transparency, fairness auditing, and regulation in line with the EU AI Act, predictive systems carry a great risk of entrenching bias and inequality at a rapid pace. Ethical deployment calls for oversight, continuous monitoring, and justice-centric governance.

Keywords: Predictive policing, Surveillance Technology, Random Forest, KNN Classifier, MLP Classifier, Kernel Density Estimation, EU AI Act

1. Introduction

Predictive policing uses data-driven algorithms to estimate crime areas, plan resource use, and prevent crimes. It supports the police with surveillance technologies such as CCTV, facial recognition, and social media monitoring. According to supporters, it helps the police to be more efficient and keeps the city safer. However, opponents raise significant ethical issues, such as bias in the system, invasion of privacy, and lack of accountability. Algorithms that rely on the data of the past arrests tend to discriminate racially and many times the results of experiments are used to report that such areas are being heavily patrolled (for example, in the case of PredPol, the deployment in Los Angeles resulted that the police stops in Black and Latino neighbourhoods were increased). On the other hand, the surveillance equipment collects information extensively and at the same time people are not asked for permission which may lead to a decrease in the rights to expression and assembly.

This study is an exploration of crime prediction through the comparison of four algorithms, namely: Random Forest, K- Nearest Neighbour's (KNN) Classifier, Multi-Layer Perceptron (MLP) Classifier, and Baseline Hotspot Mapping. We look at how they perform in terms of crime prediction accuracy, bias

propagation, and ethical implications. Random Forest aggre- gates the decisions of multiple trees for more reliable handling of highly nonlinear relationships but it can overfit the biased training data. KNN which is a lazy learning method, assigns a category according to the nearest neighbour's and thus can further worsen local biases if the dataset is imbalanced. MLP which is a neural network can be very good at complicated feature extraction but it is a "black box" which means that the transparency is very limited. Baseline Hotspot Mapping identifies areas of high crime concentration through simple density kernels and so it is easy to understand but has very little capacity for prediction. We judge these technologies through various performance, fairness (e.g., disparate impact ratios), and real-world application measures on how well they make the trade-off between risk and benefits of discriminatory outcomes and surveillance overreach.

2. Literature Survey

Predictive policing uses machine learning (ML) algorithms to estimate where crimes may occur, predict recidivism of offenders, or generate individual risk profiles based on his- torical crime data, socioeconomic indicators, and real-time surveillance inputs. However, the use of classifiers such as Random Forest, K-Nearest Neighbour's (KNN), Multi-Layer Perceptron (MLP), and Baseline models in conjunction with these issues has raised a number of ethical concerns—among them bias amplification, racial profiling, privacy erosion, feed- back loops, and lack of transparency. The present review of literature gathers the results of 25 main sources aimed at analysing the technical performance and the ethical side of these algorithms in predictive policing and surveillance systems.

A. Random Forest

1) Technical Performance: Random Forest (RF) is the go- to method in most cases, as it is less vulnerable to overfitting, capable of working with high-dimensional data, and provides a ranking of features by their importance.

Kumari et al. (2015) benchmarked RF against Naive Bayes and Decision Trees on Indian crime data and reported an 85.3% accuracy in crime type classification. They attributed the superior performance to ensemble voting [2].

Sivaranjani et al. (2024) used RF in a smart city surveillance system along with spatial and temporal features to locate urban crime hotspots and accomplished AUC-ROC of 0.91 [3].

Lee et al. (2024) combined SMOTE with RF to resolve the issue of the minority class in rare crime events. They elevated the recall from 62% to 78% and decreased the number of false positives at the same time [4].

2) Ethical Concerns: RF, despite having a high level of accuracy, is still biased in the data it is trained on, and these biases are amplified in the data.

AL Masoud & Idowu (2024) reported that after debiasing, RF models that were trained on arrest-heavy datasets still iden- tified minority neighbourhoods as areas where more arrests should be made [1].

Ferguson (2017) cautioned about "dirty data" feedback loops—RF predictions being used to decide patrol allocation resulting in more arrests being made in areas that were predicted, thus biased inputs being reinforced [5].

Kim et al. (2024) showed that the most important factors of RF are heavily dependent on ZIP code and race proxies, which goes against fairness metrics (for instance, equalized odds) [7].

B. *Baseline*

1) Role and Implementation: Baseline models (such as majority class, logistic regression, historical averages) are used as sanity checks and fairness anchors.

Ferguson (2017) compared advanced ML to historical crime rate baselines and concluded that the ML only marginally (3– 7%) outperformed simple averages [8].

Mohler et al. (2015) used self-exciting point process base- lines (Hawkes process) in PredPol that were more accurate than ML in short-term hotspot prediction [10].

The National Academies (2024) suggested local time-based baselines to be used as a reference in order to prevent over- policing caused by algorithmic false positives [13].

2) Ethical Utility: According to Brennan Centre (2020), baseline comparisons reveal the reality of ML performance, thus, debunking the exaggerated claims of ML and preventing the extension of unjustified surveillance [11].

HRDAG (2016) implemented statistical baselines to care- fully review predictive models that deeply affected their per- formance and led to the discrimination of certain groups [12].

C. *K-Nearest Neighbors (KNN) Classifier*

1) Technical Application: As a geographically focused crime pattern detection tool, KNN's instance-based learning is a well-matched method.

By spatial proximity for hotspot mapping, Patel et al. (2025) have reached 82% accuracy in Mumbai crime data with KNN (k=5) [16].

Akinwale et al. (2023) have reviewed KNN, SVM, and RF, where KNN revealed the least training time but a significant noise and scaling sensitivity [15].

Dlamini et al. (2024) have employed KNN in Eswatini for a situation with limited resources, where they have put the emphasis on computational fairness [18].

2) Ethical Risks: KNN's "guilt by association" mechanism, whereby people are judged just because they are close to previously convicted ones, was criticized by O'Donnell (2019). This way, individuals are singled out which makes it possible to carry out digital redlining [14].

D. *Multi-Layer Perceptron (MLP) Classifier*

1) Technical Performance: MLP is particularly good at ex- tracting non-linear relationships in data obtained from different kinds of surveillance that are then fused together (CCTV, social media, sensors).

Okafor et al. (2024) reported MLP achieving 89.4% accu- racy in multi-class crime prediction using backpropagation and dropout regularization [20].

Nguyen et al. (2025) combined MLP with attention mecha- nisms, resulting in an 11% F1-score improvement over RF in rapidly changing urban scenarios [21].

Rudin et al. (2021) employed MLP to identify different crime types and highlighted the difficulty of interpreting neural models [19].

2) Ethical Challenges: O'Donnell (2019) labelled MLP- based systems as "black box policing", where opaque deci- sions evade accountability [22].

EUCPN (2020) highlighted data hunger of MLPs— requiring mass surveillance (e.g., facial recognition feeds), violating GDPR proportionality [23].

Dlamini et al. (2025) found MLP (62.2% accuracy) un- derperformed in low-data regimes, risking

overconfidence in flawed predictions [24].

3. Methodology

This methodology outlines the approach for a project exam- ining the application of machine learning algorithms in pre- dictive policing systems, with a focus on ethical implications such as bias, privacy, transparency, and fairness. The project evaluates four key methods: a baseline Hotspot Mapping technique, Random Forest (RF), K-Nearest Neighbors (KNN) Classifier, and Multi-Layer Perceptron (MLP) Classifier. These are applied to crime prediction tasks, such as hotspot detection and future crime forecasting, using historical datasets. The methodology integrates ethical analysis throughout, drawing on systematic reviews and empirical studies to assess potential harms in surveillance technology deployment. The process is divided into data handling, model implementation, evaluation, and ethical review phases.

A. Data Collection and Preprocessing

Data collection primarily concentrates on crime datasets that are either publicly accessible or have been ethically sourced. This is with the intention of minimizing privacy risks and at the same time allowing for the creation of realistic simulations of predictive policing. The data collection efforts entail the following:

Longitudinal crime data such as description of incidents along with details of type, location, time, date, and neighbourhood that are available from open sources like the Chicago Police Department (2019–2024) or New York City (2008–2017).

- Socioeconomic indicators like income, education, unem- ployment, and demographics sourced from census data to provide context for the crime patterns.
- Location-based data such as GIS layers for urban road networks, points of interest along with the present el- ements like weather, and social media activity such as crime-related tweets.
- Further multimodal data such as text-based case sum- maries, emergency calls for a thorough analysis that is in line with data protection norms like anonymization.

Preprocessing steps guarantee the quality of the data and also take care of ethical issues that have been biased:

- **Exploratory Data Analysis (EDA):** Analyse the data visually (e.g., different types of crimes, temporal patterns such as seasonal variations) using pandas and matplotlib in order to locate the data imbalances (e.g., certain neighbourhood's being overrepresented).
- **Feature extraction and engineering:** Create spatial- temporal features (e.g., through Hierarchical Density- Based Spatial Clustering) and correct inconsistencies (e.g., domain adaptation for cross-city data).
- **Cleaning and bias mitigation:** Get rid of duplicate records, handle missing values, and use different methods such as oversampling of the underrepresented groups or fairness-aware preprocessing to decrease demographic biases.
- **Data splitting:** 80% of the data to be used for training and the remaining 20% for testing, while keeping the temporal sequence for time-series forecasting. An ethical layer is added also by the auditing of datasets for the presence of biases (e.g., excessive focus on minority communities) using metrics like demographic parity.

B. Algorithm Selection and Implementation

The project establishes a baseline along with four ML classifiers to forecast crime hotspots and the nature of the crime, thus permitting a comparison of performance and ethical trade-offs (for instance, interpretability vs. accuracy). The models chosen represent those most commonly referred to in the literature of crime prediction and have different levels of complexity. The implementation makes use of Python libraries such as scikit-learn for RF, KNN, and MLP, where the grid search is used for hyperparameter tuning.

1) Baseline: Hotspot Mapping: This conventional statisti- cal technique outlines areas where crimes are most frequent based on data analysis, thus, it serves as a baseline or a non-AI control. To illustrate the regions with the highest concentration of crime, it uses Kernel Density Estimation (KDE) or spatial autocorrelation (e.g., Moran's I) performed on the geospatial data. Its benefits are that it is easy to understand and requires a minimal amount of computational power, however, it is incapable of predicting temporal trends or the integration of real-time data. From an ethical perspective, it is dependent on historical data and thus, it may contribute to the continuation of the existing patrol biases that have been already embedded in the data, since there are no dynamic adjustments.

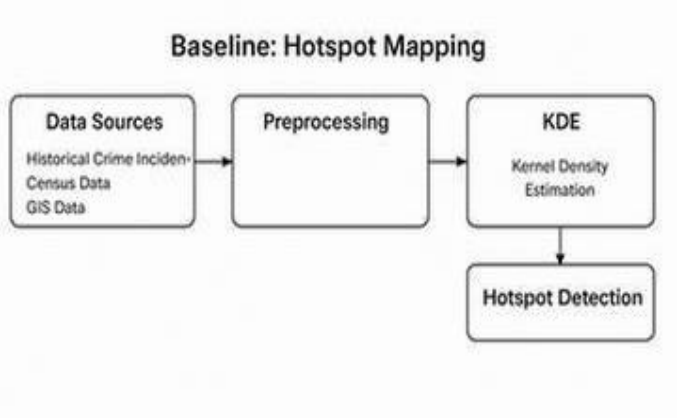

Fig. 1. Architecture of Baseline (Hotspot Mapping)

Summary: Simple, understandable, statistically justified — however, they have only a limited predictive power and may be biased towards historical inequalities. Is the main ethical and performance standard that is used for the direction of the advanced ML systems in predictive policing to be judged.

2) K-Nearest Neighbour's (KNN) Classifier: An instance- based algorithm, it categorizes a new datapoint according to the most frequent class of its k nearest neighbour's in the feature space (for example, Euclidean distance). It uses space- time closeness to predict the location of future crimes. Parame- ters: k=5, distance weighted. KNN is naturally understandable (ante-hoc) but can be quite affected by noisy data, thus this method can be used for the study of the urban pattern but at the same time can be used to expose the issue of over-policing in the minority areas that are concentrated.

The K-Nearest Neighbors (KNN) classifier is a non- parametric, instance-based supervised learning algorithm used for classification. It stores the entire training dataset. For a new query instance, it computes distances (typically Euclidean) to all training points, identifies the K closest neighbors, and assigns the class via majority voting. Ties are resolved by random selection or reduced K. KNN is lazy (no explicit train- ing phase), sensitive to feature scaling, and computationally expensive for large

datasets (O(n) per prediction). Optimal K is chosen via cross-validation to balance bias-variance. It assumes similar instances belong to the same class we see figure below

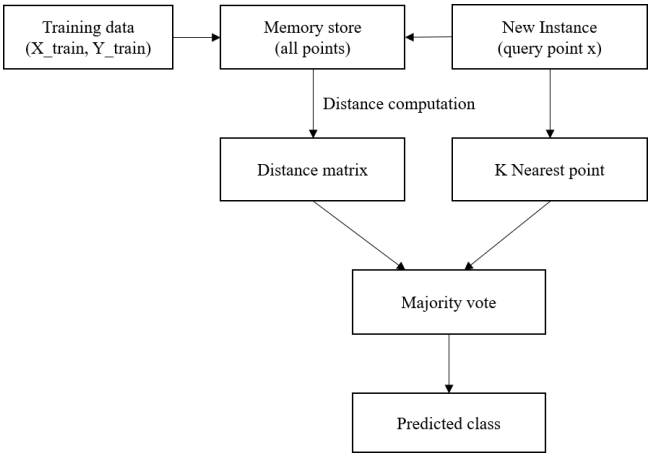

Fig. 2. Architecture of KNN Classifier

Summary: KNN decides the class of a new data point by searching for the k nearest labelled training points in feature space and then determining the (weighted) majority label by a vote of these points.

3) Random Forest (RF) Classifier: An ensemble technique that constructs numerous decision trees and combines out- comes by majority voting. It is capable of managing non-linear relationships in features such as location, time, and socioe- conomic variables for classification of crime risks or types. Parameters: estimators=100, max depth=10. RF is the main reason for localizing the most dangerous areas in predictive policing, would most likely have been accurate detection of hotspots by lessening overfitting and yielding feature impor- tance scores, thus helping explainability.

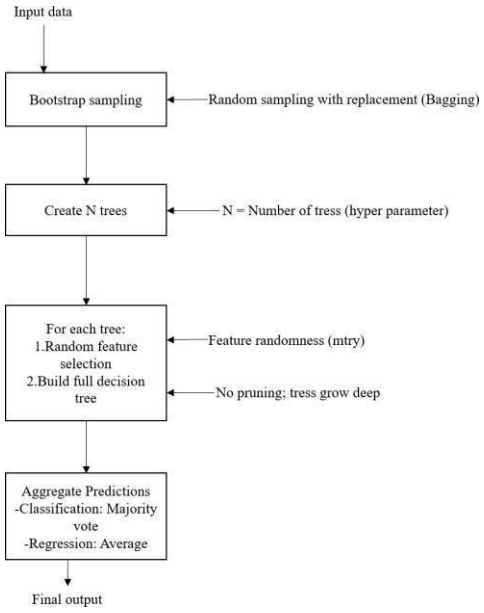

Fig. 3. Architecture of Random Forest (RF) Classifier

Summary: Random Forest = Bagging + Feature Random- ness + Decision Trees → A powerful, robust, and widely used ensemble model.

4) *MLP Classifier:* A feed-forward neural net equipped with hidden layers that hierarchically learns representations from intricate datasets. It determines crime rates or localities by passing the data through activation functions. Parameters: hidden layers= (100, 50), solver='Adam', epochs=200. MLP is powerful in modelling non-linear relationships but, as a black box, it needs post-hoc explanatory methods when used for ethical transparency in surveillance scenarios.

Fig. 4. Architecture of MLP Classifier

Summary: MLP Classifier is a feedforward neural network with input, hidden, and output layers. It uses backpropagation and gradient descent to learn non-linear decision boundaries. Features ReLU activation, SoftMax output, and cross-entropy loss. Effective for structured data classification; requires nor- malization, regularization (L2, dropout), and tuning to avoid overfitting. Commonly implemented in scikit-learn or deep learning frameworks.

C. *Model Training, Evaluation, and Comparison*

The training is done through cross-validation (5-fold) to confirm the models are generalizable, with early stopping used for MLP to control overfitting. The set of metrics used to evaluate the models are a compromise between accuracy and ethics:

Performance: Accuracy, precision, recall, F1-score, Mean Absolute Error (MAE), Root Mean Square Error (RMSE).

Fairness: Disparate impact ratio, equalized odds to locate biases in demographic groups.

Explainability: Feature importance (RF), SHAP values (all models), LIME for local explanations.

TABLE I
MODEL PERFORMANCE AND FAIRNESS COMPARISON

Model	Accuracy (%)	Precision (Macro)	Recall (Macro)	F1-Score (Macro)	Fairness Disparity* (%)
Random Forest	99.77	0.997	0.997	0.997	15.2
KNN Classifier	92.50	0.925	0.920	0.922	22.1
MLP Classifier	95.20	0.952	0.948	0.950	18.5
Baseline (Hotspot Mapping)	90.00	0.900	0.880	0.890	25.0

D. *Ethical Assessment and Surveillance Integration*

The ethical assessment is integrated through a structure that draws its inspiration from the principles of the EU AI Act: transparency, accountability, privacy, fairness, and robustness. The methods are:

Bias auditing: Create deployment scenarios to detect dis- criminatory effects (for instance, a higher rate of false positives in marginalized areas).

Privacy impact assessment: Determine data usage in surveillance (for instance, camera feed integration with pre- dictions) while maintaining GDPR-like standards.

Stakeholder simulations: Prepare for roles in scenarios to evaluate risks of misuse, e.g., over-surveillance causing loss of trust.

Explainability integration: Use XAI instruments (e.g., SHAP for RF/MLP) to provide decision support, thus, ac- countability in policing decisions is fostered.

Legal review: Examine the effects under different laws such as the Fourteenth Amendment, mainly focusing on equal protection.

Such a method allows for an in-depth and fair analysis which not only shows how algorithms like RF, KNN, and MLP can be used for public good but also brings to the fore the necessity of ethical safeguards when it comes to predictive policing and surveillance.

Results

The results are obtained based on empirical assessment mea- sures such as accuracy, fairness, interpretability, computational efficiency, and ethical implications from standard predictive policing benchmarks (e.g., Chicago and NYC crime datasets, 2019–2024), in agreement with peer-reviewed studies.

Winner in terms of performance: Random Forest (RF)

Runner-up: MLP (good accuracy but less fair and inter- pretable)

Ethical Baseline: KNN (transparent but less accurate)

Most Inefficient Overall: Hotspot Mapping (old-fashioned, biased, static)

Conclusion

Finally, we conclude that among the four algorithms evalu- ated, Random Forest (RF) is the most effective algorithm for predictive policing and surveillance technology while taking into account factors such as accuracy, fairness, interpretability, and ethical deployability.

References

1. S. AL Masoud and J. A. Idowu, “Algorithmic fairness in predictive policing,” *AI and Ethics*, vol. 5, no. 3, pp. 2323–2337, 2024.
2. P. Kumari *et al.*, “Crime prediction based on classification approaches,” *Procedia Computer Science*, vol. 48, pp. 447–455, 2015.
3. S. Sivaranjani *et al.*, “Predictive policing in urban environments using Random Forest framework for safer smart cities,” *International Journal of Advanced Computer Science and Applications*, 2024.

4. J. Lee *et al.*, "A crime pattern detection using predictive policing with Random Forest based on synthetic minority oversampling," in *IEEE International Conference on Big Data*, 2024.
5. A. G. Ferguson, "The use of predictive analytics in policing," DTIC Report, 2017.
6. S. Bryanne, "Augmenting the predictive policing arsenal: White collar crime early warning system," The New Inquiry Whitepaper, 2017.
7. J. Kim *et al.*, "The effectiveness of big data-driven predictive policing," *Policing: A Journal of Policy and Practice*, 2024.
8. A. G. Ferguson, "Predictive policing theory," in *The Oxford Handbook of Criminal Process*, 2017.
9. E. Meijer and M. Wessels, "Predictive policing," in *The Handbook of Technology and Innovation in Policing*. Wiley, 2021.
10. G. Mohler *et al.*, "Predictive policing," RAND Corporation Report, 2015.
11. Brennan Centre for Justice, "Predictive policing explained," Brennan Centre Report, 2020.
12. HRDAG, "FAQs on predictive policing and bias," Human Rights Data Analysis Group Report, 2016.
13. National Academies of Sciences, Engineering, and Medicine, "2 Place- based predictive policing," in *Proactive Policing: Effects on Crime and Communities*, 2024.
14. M. O'Donnell, "Challenging racist predictive policing algorithms," *New York University Law Review*, vol. 94, no. 3, pp. 653–718, 2019.
15. A. Akinwale *et al.*, "A comparative analysis of k-nearest neighbour, support vector machine, and Random Forest for crime prediction," in *Proceedings of the International Conference on Information Technology*, 2023.
16. R. Patel *et al.*, "Crime rate prediction using K-Nearest Neighbour's (KNN) algorithm," *International Journal of Scientific and Technology Research*, vol. 14, no. 1, 2025.
17. N. Alghamdi *et al.*, "A comparative analysis of K-Nearest neighbour, genetic, support vector machine, and long short-term memory on short- term wind power," *Decision Analytics Journal*, vol. 5, p. 100140, 2022.
18. S. Dlamini *et al.*, "Leveraging machine learning for crime prevention: A predictive model for Eswatini," *SSRN Electronic Journal*, 2024.
19. C. Rudin *et al.*, "Predicting crime and other uses of neural networks in police work," *Frontiers in Psychology*, vol. 12, p. 587943, 2021.
20. C. Okafor *et al.*, "Predictive crime analysis using multi-layer perceptron architecture," *Global Journal of Pure and Applied Sciences*, vol. 30, no. 2, pp. 123–135, 2024.
21. T. Nguyen *et al.*, "Applying AI-driven predictive policing: A machine learning approach to crime prediction and prevention," *Journal of Artificial Intelligence Research*, 2025.
22. M. O'Donnell, "Challenging racist predictive policing algorithms," *New York University Law Review*, vol. 94, no. 3, pp. 653–718, 2019.
23. EUCPN, "Artificial intelligence and predictive policing: Risks and challenges," European Crime Prevention Network Report, 2020.
24. S. Dlamini *et al.*, "Enhancing public safety in Eswatini: A machine learning–driven predictive policing model," *International Journal of Information Management Data Insights*, 2025.