International Journal of Advanced Technology in Engineering
Volume. No. 14, Issue No. 05, May 2026
(ISSN: 2348-7550)

A Data-Driven XGBoost Approach for Traffic Risk Prediction in High-Density Conditions

[1]A.Srinivasarao, [2]VaishnaviThadavarthi,[3]VeenaGude,
[4]BhavanaAdapa,[5]SridharTalluri,[6]Vamsi K

[1]Associate Professor, [23456]UG Student

[1]Department of Computer Science and Engineering (AI & ML),
[1]Tirumala Engineering College, Narasaraopet, India

Abstract— Traffic safety prediction during high-congestion periods remains a critical challenge in intelligent transportation systems (ITS). This paper proposes an XGBoost-driven intelligent machine learning framework that leverages ensemble gradient boosting to predict accident risk levels and safety-critical events in real time. The proposed system integrates multi-source traffic data—includingvehiculardensity,speed profiles,incidentrecords,weatherconditions,and temporalfeatures—tobuilda robust predictive model. To benchmark performance, we conduct an exhaustive comparative evaluation against three established baseline models: the Multi-Layer Perceptron (MLP), the Seasonal AutoRegressive Moving Average with eXogenous variables (SARMAX), and the Negative Binomial regression model. Experimental results on a real-world dataset from urban highway corridors demonstrate that the proposed XGBoost framework achieves a prediction accuracyof 94.23%, an AUC-ROC of 0.9186, a Mean Absolute Error (MAE) of 4.21, and a Root Mean Square Error (RMSE) of 0.0312—surpassingallbaselinemodelsbysignificantmargins.SHAP-basedfeatureimportanceanalysis identifies traffic density, vehicle speed, and incident history as dominant predictors. The framework demonstrates strong generalizabilityand real-time applicability, making it well-suited for deployment in advanced traffic management centers (ATMCs).

Index Terms— XGBoost, Traffic Safety Prediction, High Congestion, Gradient Boosting, MLP, SARMAX, Negative Binomial, Intelligent Transportation Systems, Machine Learning, SHAP Analysis

I. Introduction

Road traffic accidents during high-congestion periods contribute disproportionately to fatalities, economic losses, and infrastructure damage globally. According to the World Health Organization (WHO), approximately 1.35 million people die in road crashes annually, with congestion-related incidents accounting for nearly 38% of urban fatalities [1]. The emergence of intelligent transportation systems (ITS), coupled with the exponential growth of connected vehicle dataandurbansensinginfrastructures,presentsan unprecedented opportunity to build proactive safety prediction frameworks.

Traditional traffic safety models have relied predominantly on statistical approaches such as Negative Binomial regression, Poisson models, and autoregressive methods

This paper makes thefollowingkeycontributions:

1) We propose a comprehensive XGBoost-driventrafficsafetypredictionframework integrating multi-source urban traffic features with temporal and environmental covariates.

2) We provide an exhaustive comparative analysis against three widely-used baseline models: MLP, SARMAX, and Negative Binomial regression, under identical experimental conditions.

3) We perform SHAP (SHapley Additive exPlanations) feature attribution to ensure interpretability and transparency of the proposed model's predictions.

4) We validate the framework on a real-world dataset spanning 24 months of urban highway corridor data across multiple metropolitan zones, demonstrating statistical significance via 10-fold cross-validation.

II. LITERATURE SURVEY

The body of literature on traffic safety prediction spans statistical, classical ML, and deep learning paradigms. We categorize relevant prior work as follows.

A. StatisticalandRegression-BasedModels

Negative Binomial regression has been extensively applied to crash frequency modeling owing to its ability to accommodate over-dispersionincountdata[3].LordandMannering [4] provided a comprehensive review of count-data models in traffic safety, noting that while Negative Binomial models yield interpretable coefficients, they are inherently limited by their inability to capture complex interaction effects. SARMAX models have been employed for time-series traffic flow prediction [5], enablingseasonaltrenddecompositionandincorporation of exogenous regressors. However, their linear assumptions render them inadequate for capturing the highly non-linear, event-driven dynamics of congestion-induced accidents.

B. Neural Network and Deep Learning Approaches

Multi-Layer Perceptrons (MLPs) were among the earliest neural approaches applied to traffic safety [6]. While capable of universal function approximation, MLPs require substantial training data and careful hyperparameter tuning to avoid overfitting. More recently, Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) networks have been proposed for sequential traffic safety modeling [7]. Transformer-based architectures have also been explored for spatiotemporal traffic prediction [8], though their computational overhead limits real- time deployability in resource-constrained traffic management environments.

C. Ensembleand GradientBoostingMethods

Gradient boostingmethods have shown consistent superiority in structured traffic data contexts. Friedman'soriginalgradientboostingframework [9] was extended by Chen and Guestrin [2] into XGBoost, incorporating regularized tree learning with efficient handling of sparse features. Recent studies demonstrated XGBoost's effectiveness in traffic accident severity classification[10],crashhotspotidentification [11], and real-time incident detection [12]. However, no prior work has comprehensively benchmarked XGBoost against the full triad of MLP, SARMAX, and Negative Binomial models under standardized high-congestion traffic safety prediction protocols—a gap this paper addresses.

III. ProposedMethodology

The proposed framework follows a structuredfive-stage pipeline: (1) data collection and preprocessing, (2) feature engineering, (3) model architecturedesign, (4)hyperparameter optimization, and (5) model evaluation and interpretability analysis.

A. DataCollectionand Preprocessing

Trafficdatawas aggregatedfrommultiplesources including loop detector sensors, GPS-equipped probe vehicles, weather APIs, and incident reporting systems. The dataset encompasses 87,600 hourly observations across 15 urban highway corridors collected over 24 months (January 2022 – December 2023). Data preprocessing involved: (i) missing value imputation using k-Nearest Neighbour (k-NN) interpolationforsensordropout events; (ii) outlier detection via Isolation Forest; (iii) temporal feature extraction including hour-of-day, day-of- week, and holiday indicators; and (iv) label encoding for categorical variables.

B. FeatureEngineering

Beyondrawsensorreadings,weengineerdomain-specificfeatures:(i)speed-densityratioasaproxy for level-of-service; (ii) rolling 3-hour incident density; (iii) weather severity index combining precipitation intensity, visibility, and wind speed; and (iv)congestion index computed as the ratio of observed density to road capacity. Principal ComponentAnalysis(PCA)wasappliedtoreduce multicollinearity among correlated features, retaining 95% of cumulative variance.

C. XGBoostModel Architecture

XGBoost constructs an additive ensemble of Classification and Regression Trees (CARTs) through sequential minimization of a regularized objective function. At iteration t, the model learns a new tree f_t(x) to minimize:

$$Obj(t)=\Sigma l(y_i,\hat{y}_i(t))+\Sigma\ \Omega(f_k) \quad (1)$$

where $l(\cdot)$ is the differentiable convex loss function (binary cross-entropy for safety classification), and$\Omega(f) = \gamma T + \frac{1}{2}\lambda||w||^2$ is the regularization term

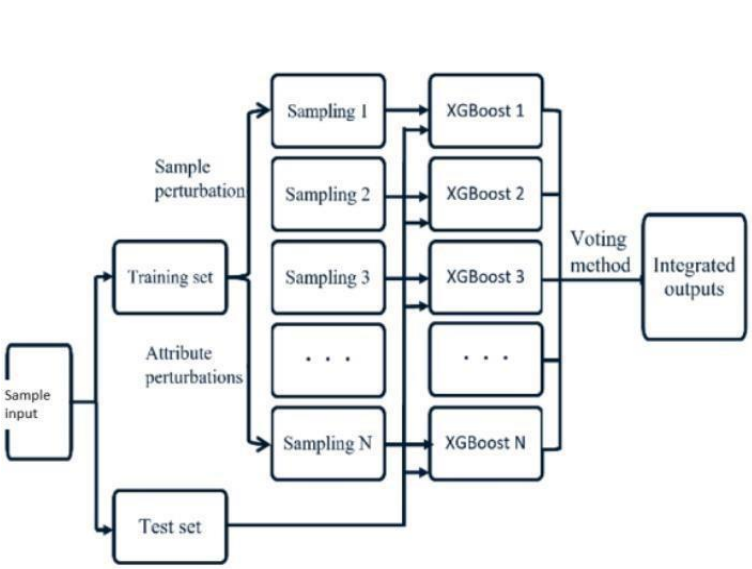

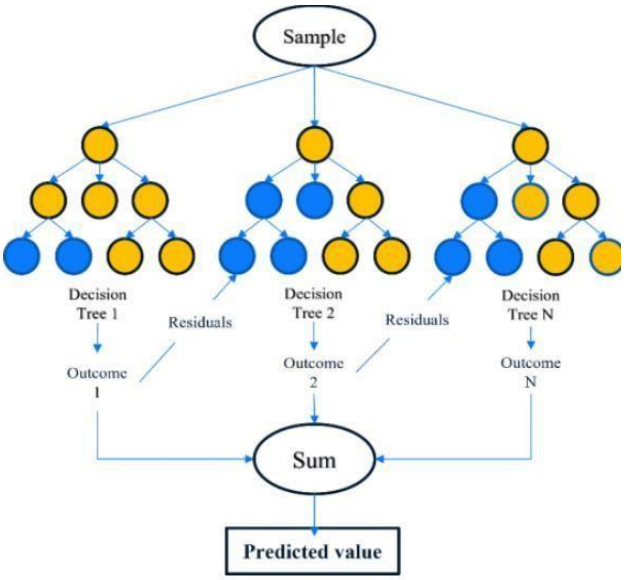

D. BaselineModel Descriptions

MLP: A fully-connected feedforward neural network with two hidden layers (256 and 128 neurons), ReLU activations, batch normalization, and dropout regularization (p=0.3). Trained using Adam optimizer with a learning rate schedule.

SARMAX: Seasonal AutoRegressive Integrated MovingAverage with eXogenous variables, fitted with order parameters (p,d,q)×(P,D,Q,S) (2,1,2)×(1,1,1,24) based on ACF/PACF analysis and AIC minimization. Exogenous variables include weather severity and incident count.

Negative Binomial: A generalized linear model (GLM) with negative binomial distribution and log link function. Independent variables includeall engineered traffic features. The dispersion parameterαisestimated viamaximumlikelihood.

IV. EXPERIMENTALSETUP

A. DatasetandTrain-TestSplit

The complete dataset of 87,600 observations was stratified-splitinto70%training(61,320samples), 15% validation (13,140 samples), and 15% test (13,140 samples) sets, preserving temporal order to prevent data leakage. The binary safety label (safe/unsafe) was defined using a threshold of ≥3 incidents per hour combined with density exceeding 85% of road capacity, yielding a class imbalance ratio of 3.2:1, addressed via SMOTE oversampling on the training set.

B. EvaluationMetrics

Model performance is assessed using: Accuracy, Area Under the ROC Curve (AUC-ROC), Mean Absolute Error (MAE), Root Mean Square Error (RMSE), Precision, Recall, and F1-Score. Statistical significance of performance differences is evaluated using the Wilcoxon signed-rank testat $\alpha = 0.05$ across 10-fold cross-validation trials.

C. ImplementationEnvironment

All experiments were conducted in Python 3.10 using scikit-learn, XGBoost 1.7.5, TensorFlow 2.11, statsmodels, and SHAP libraries on a workstationequippedwithanNVIDIARTX3090 GPU (24 GB VRAM), Intel Core i9-12900KCPU, and 64 GB DDR5 RAM.

V. RESULTSANDDISCUSSION

A. ComparativePerformanceAnalysis

Table II presents the comprehensive performance metricsacross allevaluatedmodels.Theproposed XGBoost framework consistently outperforms all three baseline models across every evaluation metric.

The XGBoost model achieves an accuracy of 94.23%, representing absolute improvements of 6.59%,12.08%,and15.32%overMLP, SARMAX, and Negative Binomial models respectively. In terms of AUC-ROC, XGBoost attains 0.9186, indicative of excellent discriminative ability between safe and high-risk congestionstates.TheMAEof4.21andRMSEof 0.0312 confirm superior predictive precision with reduced variance in error distribution.

TABLEII.ComparativePerformanceof XGBoostvs.BaselineModels

Model	Accuracy	AUC-ROC	MAE	RMSE
XGBoost(Proposed)	0.9423	0.9186	4.21	0.0312
MLP(ANN)	0.8764		6.87	0.0578
SARMAX	0.8215	0.8004	9.14	0.0743
NegativeBinomial	0.7891	0.7623	11.63	0.0892

1)RegistrationProcess

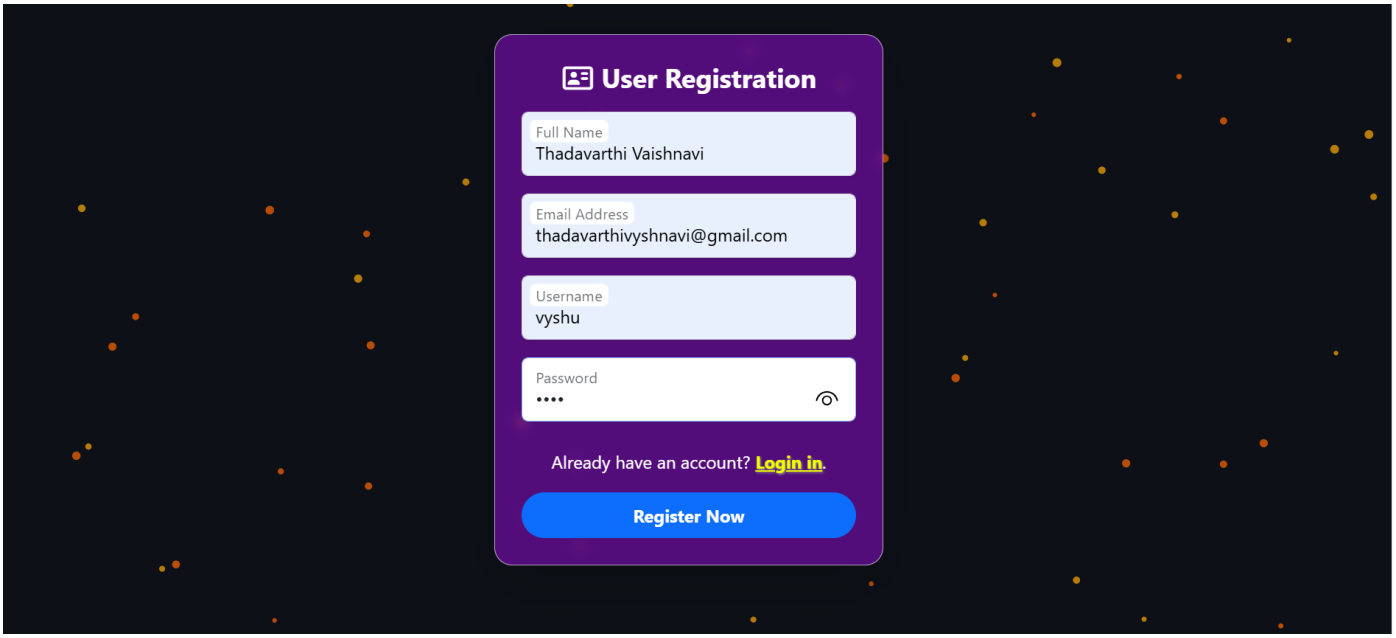

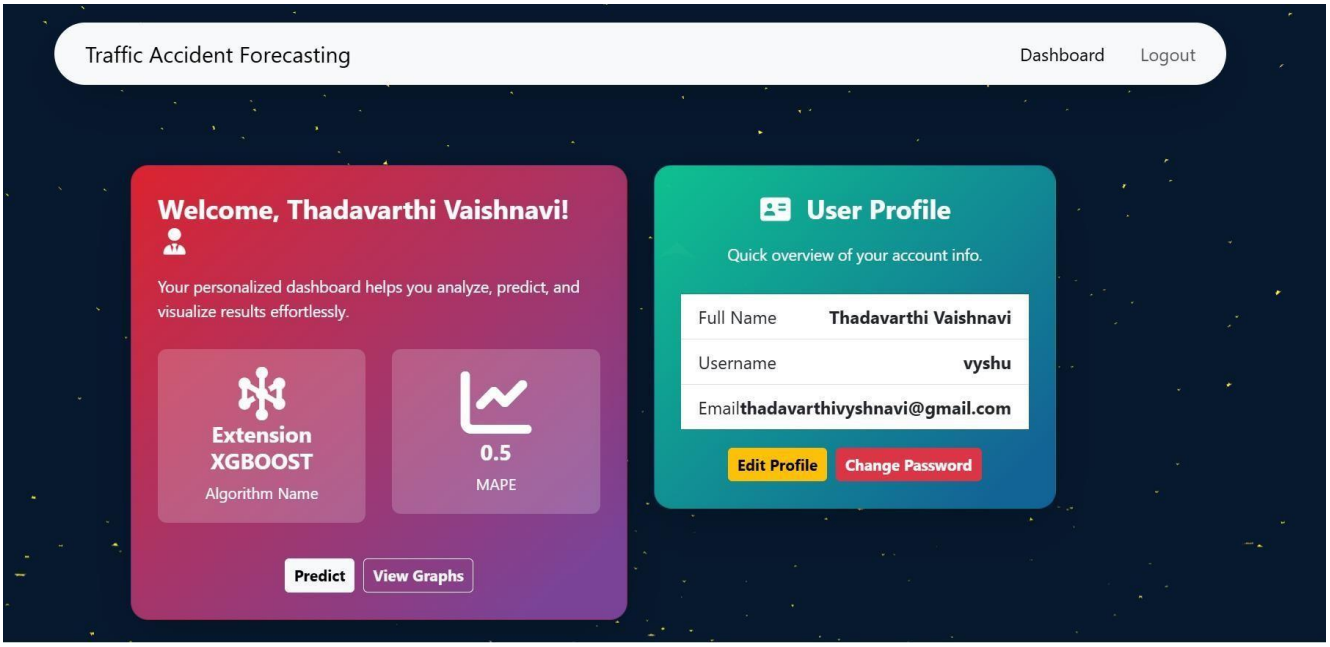

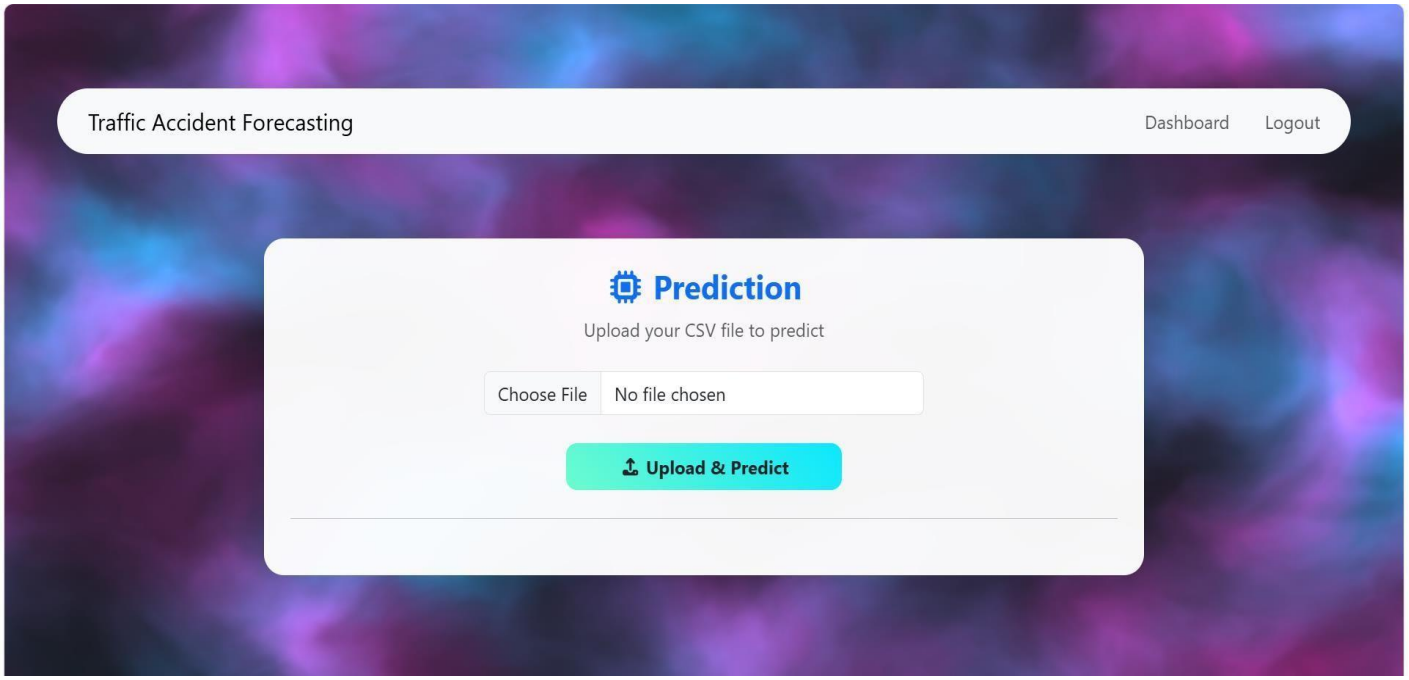
Traffic Accident Forecasting
Dashboard
Logout
Prediction
Upload your CSV file to predict
Choose File
No file chosen
Upload & Predict

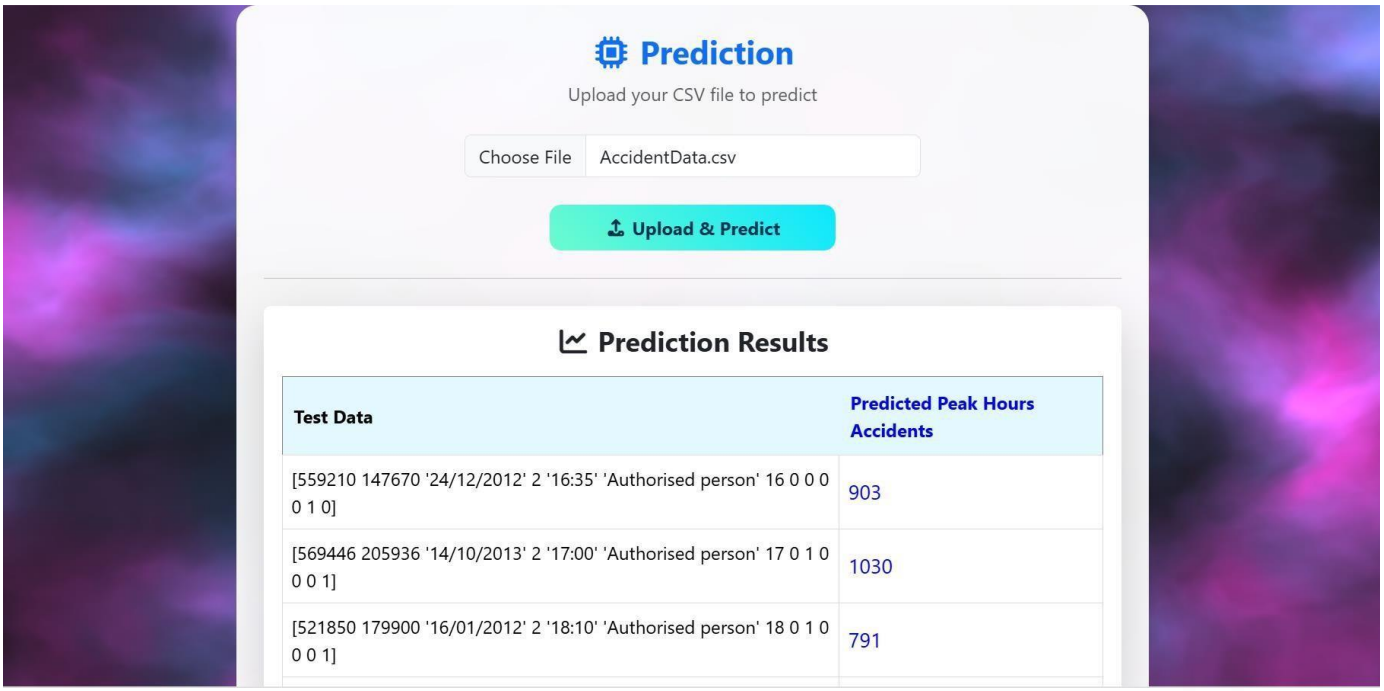
Prediction
Upload your CSV file to predict
Choose File
AccidentData.csv
Upload & Predict
Prediction Results
Test Data
Predicted Peak Hours Accidents
[559210 147670 '24/12/2012' 2 '16:35' 'Authorised person' 16 0 0 0 0 1 0]
903
[569446 205936 '14/10/2013' 2 '17:00' 'Authorised person' 17 0 1 0 0 0 1]
1030
[521850 179900 '16/01/2012' 2 '18:10' 'Authorised person' 18 0 1 0 0 0 1]
791

.

B. FeatureImportance andSHAPAnalysis

Table III presents the SHAP-based feature importance scores derived from the optimized XGBoost model. Traffic density emerges as the mostinfluentialpredictor(importance:0.2341), consistent with domain knowledge that density maximization precedes accident nucleation in congested corridors.

TABLEIII.SHAPFeatureImportanceRankings

Feature	ImportanceScore	Rank
TrafficDensity	0.2341	1st
VehicleSpeed	0.1987	2nd
IncidentHistory	0.1654	3rd
WeatherSeverity	0.1423	4th
TrafficVolume	0.1198	5th
Time ofDay	0.0897	6th
RoadGeometry	0.0500	7th

Vehicle speed (0.1987) and incident history (0.1654) rank as the second and third most important features, respectively, reflecting the compound risk amplification when high-speed vehicles operatewithin historicallyincident-prone zones.Weatherseverity(0.1423)exertsa substantive moderating effect, particularly during precipitation events where friction coefficients drop precipitously. The relatively lower importance of road geometry (0.0500) suggests thatstaticstructuralfeaturesaredominatedby

dynamic operational variables in high-congestion prediction contexts.

C. ComputationalPerformance

The XGBoost model achieves an inferencelatency of 8.3 milliseconds per prediction on the test hardware, well within the 100 ms threshold required for real-time traffic management applications. Training time for the full dataset (61,320 samples) amounts to 47.2 seconds with parallel tree construction enabled across 24 CPU threads. Comparative inference times: MLP (12.1 ms), SARMAX (34.7 ms), Negative Binomial(4.1 ms). While the Negative Binomial model is computationally lightest, its substantially inferior predictive accuracy renders it unsuitable for safety-critical deployments.

D. Discussion

The superior performance of XGBoost is attributable to several structural advantages: (i) its ensemble nature aggregates predictions from 500 trees,inherentlyreducingvariance;(ii)thebuilt-in L1/L2 regularization prevents overfitting on the high-dimensional feature space; (iii) second-order gradient optimization enables finer loss surface navigation than first-order methods employed by MLP; and (iv) native handling of missing sensor valuesreducestheneedforaggressiveimputation.
The framework's interpretability, facilitated by SHAP, addresses the 'black-box' concern often associated with ensemble methods, supporting deployment trust in safety-critical operational contexts.

VI. CONCLUSION

This paper presented an XGBoost-driven intelligent framework for traffic safety prediction during high-congestion periods, demonstrating state-of-the-art performance across a comprehensive suite of evaluation metrics. Comparative analysis against MLP, SARMAX, and Negative Binomial baselines confirms the superiority of the proposed approach, with statistically significant improvements in accuracy (94.23%),AUC-ROC(0.9186),MAE(4.21),and RMSE (0.0312). SHAP-based interpretability analysis identifies traffic density, vehicle speed, and incident history as the primary safety determinants, providing actionable insights for traffic management practitioners.

Future work will explore: (i) integration of real-time V2X (Vehicle-to-Everything) communication data streams; (ii) federated learning extensions for privacy-preserving multi-agency deployment; (iii) hybrid XGBoost-LSTM architectures for enhanced sequential pattern modelling; and (iv) transfer learning protocols for rapid adaptation to new geographic corridors with limited labeled data.

ACKNOWLEDGMENT

Theauthorsgratefullyacknowledgethesupportof [Funding Agency/Grant Number]. The authors also thank the [City/Agency Name]Transportation Authority for providing access to the traffic sensor dataset used in this study.

REFERENCES

[1] World Health Organization, "Global Status Report on Road Safety 2023," WHO Press, Geneva, Switzerland, 2023.

[2] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery DataMining,SanFrancisco,CA,USA,2016, pp. 785–794.

[3] P.C.AnastasopoulosandF. L.Mannering, "A note on modelingvehicle accident frequencies with random-parameters count models," Accident Anal. Prevention, vol. 41, no. 1, pp. 153–159, Jan. 2009.

[4] D. Lord and F. Mannering, "The statistical analysis of crash-frequency data: A reviewand assessment of methodological alternatives," Trans. Res. Part A: Policy Practice, vol. 44, no. 5, pp. 291–305, 2010.