**© MAY 2026 | IRE Journals | Volume 9 Issue 11 | ISSN: 2456-8880**
**DOI: https://doi.org/10.64388/IREV9I11-1718000**

# ML-based Road Accident Hotspot Detection

KAVINESH V S[1], PROF. RAKSHITHA B S[2]
[1]*Research Scholar, Department of M.Sc. Computer Science & IT Jain (Deemed-To-Be) University Bangalore, India*
[2]*Department of Computer Science & IT Jain (Deemed-To-Be) University Bangalore, India*

***Abstract—Road traffic accidents remain one of the leading causes of fatalities and injuries worldwide, imposing significant social, economic, and public health burdens. Identifying accident-prone locations, commonly referred to as hotspots, is a critical step toward reducing road casualties through targeted safety interventions. However, traditional statistical and rule-based methods struggle to capture the complex, non-linear relationships among the numerous factors contributing to road accidents.***

***This paper presents a comprehensive review of machine learning methods applied to road accident hotspot detection. We survey existing approaches spanning statistical models, kernel density estimation, supervised classifiers, ensemble methods, deep learning architectures, and explainable AI techniques. We compare their reported performance, datasets used, feature spaces explored, and key limitations. The review identifies critical research gaps, including the lack of systematic multi-model comparisons on balanced hotspot datasets, insufficient use of explainability techniques such as SHAP, and limited integration of diverse feature groups spanning spatial, temporal, road geometry, and weather attributes. The findings of this review provide a structured foundation for future research toward accurate, efficient, and interpretable road accident hotspot detection systems.***

***Index Terms—Road Accident Hotspot Detection, Machine Learning, Random Forest, XGBoost, SHAP, Traffic Safety, Spatial Analysis, GIS, Predictive Modelling, Explainable AI***

## I. INTRODUCTION

### *A. Background of the Study*

Road traffic accidents represent a global public health crisis. According to the World Health Organization, approximately 1.35 million people die in road crashes each year, with millions more sustaining non-fatal injuries [1]. In developing countries, where rapid urbanisation and growing vehicle populations often outpace road infrastructure improvements, the situation is particularly severe. Beyond the human cost, road accidents impose substantial economic losses through healthcare expenditure, property damage, lost productivity, and emergency response costs.

A road accident hotspot is a location or road segment where crashes occur at a significantly higher rate than expected given traffic exposure. Identifying and mitigating these hotspots is a foundational objective of road safety engineering. Traditionally, hotspot identification has relied on frequency-based rankings, kernel density estimation, and statistical threshold tests applied to historical crash records. While useful, these approaches often fail to account for the multifactorial nature of accident causation, which involves the interaction of human behaviour, vehicle characteristics, road geometry, environmental conditions, and temporal patterns [2].

Machine learning provides a compelling alternative, offering the ability to discover complex, non-linear relationships across high-dimensional feature spaces without requiring explicit model specification. Ensemble methods such as Random Forest and Gradient Boosting have demonstrated strong predictive performance across a wide range of structured data problems, and recent advances in explainable AI have made it possible to interpret their predictions at the feature level [8].

### *B. Problem Statement*

Despite the availability of rich accident datasets maintained by traffic authorities and transportation agencies, effective automated hotspot detection systems remain underdeveloped in many regions. Existing methods either rely on simple frequency counts that ignore spatial and temporal context, or employ computationally expensive deep learning models that sacrifice interpretability. There is a clear need for accurate, efficient, and explainable machine learning models that can identify high-risk road segments and provide actionable insights to road safety planners.

Moreover, class imbalance is a persistent challenge in accident datasets: non-accident or low-risk locations vastly outnumber genuine hotspots, causing many classifiers to un-derperform in the minority (hotspot) class. Addressing this imbalance while maintaining strong overall classification performance is a key technical challenge in this domain.

### *C. Objectives of the Study*

The primary objectives of this review are:

- To survey existing methods for road accident hotspot detection, including statistical, spatial, machine learning, and deep learning approaches.
- To analyse the feature spaces used in existing studies, covering spatial, temporal, road geometry, traffic, and environmental attributes.
- To compare reported performance across different meth-ods and benchmark datasets.
- To examine the role of explainable AI techniques such as SHAP and LIME in road safety prediction models.
- To identify key research gaps and outline directions for future work in interpretable hotspot detection.

### *D. Contributions of the Paper*

This review paper contributes to the literature by providing a structured and critical survey of machine learning methods applied to road accident hotspot detection. It categorises existing approaches by methodology and identifies consistent patterns in feature importance, model performance, and dataset characteristics. It highlights the persistent gap in explainability research within this domain and proposes a set of criteria for evaluating future hotspot detection systems. By synthesising findings across diverse studies, this paper serves as a useful reference for researchers, engineers, and policymakers working on data-driven road safety improvement.

### *E. Organisation of the Paper*

This paper is organised as follows. Section I presents the introduction and background. Section II reviews related work on accident prediction and hotspot detection, covering statisti-cal, spatial, machine learning, deep learning, and explainable AI approaches. Section III discusses the key concepts and techniques commonly used in this domain. Section IV presents a comparative analysis of existing methods and identifies research gaps. Section V discusses applications and use cases. Section VI concludes the paper and outlines future research directions.

## II. LITERATURE REVIEW

### *A. Statistical Approaches to Hotspot Identification*

Objective: Early works on hotspot identification applied negative binomial regression and empirical Bayes methods to model accident frequency at intersections and road segments, accounting for traffic exposure through vehicle kilometres travelled [13].

Methods: Negative binomial models were fitted to historical crash counts, with road geometry, speed limits, and traffic volume as covariates. Empirical Bayes estimates corrected for regression-to-the-mean bias by combining site-specific observed counts with predicted values from the fitted model [2].

Results: These methods produced reliable baseline rank-ings of high-risk locations but required careful distributional assumptions and performed poorly when accident rates were low or datasets were small.

Conclusions: Statistical models provide interpretable risk estimates but are limited by their parametric assumptions and inability to capture non-linear feature interactions. They remain useful as interpretable baselines against which more complex methods can be benchmarked.

### *B. Kernel Density Estimation for Spatial Hotspot Mapping*

Objective: Kernel Density Estimation (KDE) has been widely used in Geographic Information Systems (GIS) to produce continuous spatial risk surfaces from discrete accident point data [3].

Methods: Gaussian or quartic kernel functions were applied to accident coordinates, with bandwidth selection influencing the spatial resolution of the resulting density map. KDE surfaces were overlaid with road network data to assign risk scores to road segments.

Results: KDE-based maps provided intuitive visual rep-resentations of accident clustering and performed well in identifying broad risk zones in urban areas. Studies such as [3] demonstrated that network-constrained KDE, which restricts density estimation to the road network rather than two-dimensional space, more accurately reflects the linear structure of traffic flow.

Conclusions: KDE is effective for exploratory spatial anal-ysis and communicating risk patterns to non-technical audi-ences, but it does not incorporate road attribute or environ-mental covariates, limiting its predictive power for proactive hotspot identification.

### *C. Machine Learning for Accident Severity and Hotspot Pre-diction*

Objective: A substantial body of research has applied super-vised machine learning classifiers to predict accident severity (fatal, serious, slight) and to identify high-risk locations from structured tabular features [4].

Methods: Decision Trees, Random Forests, Support Vector Machines, and XGBoost were trained on labelled accident records. Feature selection and cross-validation were used to prevent overfitting. Studies varied in the feature sets used, ranging from simple road type and speed limit variables to comprehensive sets including junction type, lighting condition, weather, and vehicle characteristics [6].

Results: Random Forest and XGBoost consistently out-performed simpler models across multiple studies, achieving AUC-ROC scores above 0.85 for severity classification tasks [4]. Class imbalance between fatal and non-fatal or hotspot and non-hotspot locations was identified as a major challenge in nearly all reviewed studies.

Conclusions: Ensemble tree models are well-suited to struc-tured accident data. Their ability to handle mixed feature types, capture non-linear relationships, and rank feature importance makes them a natural fit for road safety applications. However, most studies in this category focus on accident severity rather than spatial hotspot classification specifically.

### *D. Deep Learning for Traffic Accident Prediction*

Objective: Recent works have explored deep learning archi-tectures, including Convolutional Neural Networks (CNN) and Long Short-Term Memory (LSTM) networks, for spatiotem-poral accident prediction [5].

Methods: CNN-based models treated accident risk predic-tion as an image classification problem over spatial grid cells, while LSTM models captured temporal dependencies in accident frequency time series. Hybrid CNN-LSTM architectures were proposed to jointly model spatial and temporal patterns [7].

Results: Deep learning models achieved high accuracy on large datasets with rich spatiotemporal structure, with studies reporting AUC-ROC values above 0.90 in well-resourced settings. However, performance was sensitive to dataset size, and models trained on data from one city generalised poorly to other urban environments.

Conclusions: While deep learning offers strong predictive performance when large labelled datasets are available, the interpretability gap and high computational cost limit its practical adoption in road safety policy contexts, particularly in regions with limited data infrastructure.

### *E. Explainable AI in Road Safety Research*

Objective: Growing recognition of the need for interpretable predictions in safety-critical applications has prompted the application of SHAP (SHapley Additive ex-Planations) and LIME (Local Interpretable Model-agnostic Explanations) to accident prediction models [8].

Methods: SHAP values were computed for Random Forest and XGBoost models trained on accident datasets, producing both global feature importance rankings and local sample-level explanations. LIME was applied to black-box classifiers including neural networks to generate locally faithful linear approximations of model behaviour in the neighbourhood of individual predictions [12].

Results: Road curvature, speed limit, lighting condition, and junction type consistently ranked among the highest-importance features across studies using SHAP. Dependence plots revealed non-linear and interaction effects that would be invisible to traditional feature importance measures such as Gini impurity. Studies employing LIME reported improved stakeholder acceptance of model predictions compared to unexplained black-box outputs.

Conclusions: Explainability analysis using SHAP and LIME enhances stakeholder trust and supports evidence-based safety interventions by making model reasoning transparent. These methods are increasingly recommended as a standard component of any machine learning pipeline deployed in road safety or other safety-critical domains.

### *F. Handling Class Imbalance in Accident Datasets*

Objective: Class imbalance is a pervasive challenge in accident hotspot datasets, where hotspot locations represent a small minority of all road segments [6].

Methods: Reviewed studies addressed this issue through several strategies: oversampling of the minority class using SMOTE (Synthetic Minority Over-sampling Technique) [11], undersampling of the majority class, cost-sensitive learning that assigns higher misclassification penalties to the minority class, and threshold adjustment on classifier output probabili-ties.

Results: SMOTE-based oversampling consistently im-proved recall for the hotspot class without substantially de-grading precision, making it the most widely adopted tech-nique. Studies that ignored class imbalance often reported in-flated overall accuracy metrics driven by correct classification of the majority non-hotspot class, masking poor performance on the target hotspot class.

Conclusions: Rigorous handling of class imbalance is es-sential for meaningful evaluation of hotspot detection sys-tems. Both the resampling strategy and the evaluation metric selection (favouring F1-Score and AUC-ROC over accuracy) must be reported transparently to allow fair comparison across studies.

### G. *Comparative Analysis of Existing Methods*

Existing methods for road accident hotspot detection dif-fer substantially in their input data requirements, modelling assumptions, and interpretability. Statistical models provide principled uncertainty estimates but impose restrictive distri-butional assumptions. KDE-based GIS approaches excel at spatial visualisation but ignore covariate information. Machine learning classifiers, particularly ensemble methods, offer the best trade-off between predictive accuracy and practical ap-plicability for structured tabular accident data. Deep learning models achieve the highest performance on large spatiotem-poral datasets but require substantial computational resources and lack interpretability.

TABLE I
COMPARATIVE ANALYSIS OF EXISTING METHODS

| Study | Method | Dataset | AUC | Limitation |
|---|---|---|---|---|
| Aguero-Valverde (2013) [13] | Neg. Binomial | US Highways | 0.79 | Parametric assumptions |
| Xie & Yan (2015) [3] | KDE + GIS | Urban Roads | – | No covariate modelling |
| Chen et al. (2019) [4] | Random Forest | UK Accidents | 0.87 | No explainability |
| Yuan et al. (2021) [5] | LSTM | China Roads | 0.90 | Black-box; high compute |
| Parsa et al. (2020) [6] | XGBoost + SHAP | Tehran | 0.88 | Single city; limited features |
| Jiang et al. (2023) [7] | CNN-Grid | NYC Open Data | 0.91 | No road geometry features |

### H. *Critical Review*

The reviewed studies demonstrate clear progress in road accident hotspot detection over the past decade. Traditional statistical and spatial methods established the conceptual and methodological foundations of the field, while machine learning approaches have substantially improved predictive accuracy by leveraging richer feature sets and more flexible model architectures. Ensemble tree methods such as Random Forest and XGBoost have emerged as the most reliable choice for structured tabular accident data, combining strong per-formance with manageable computational cost and built-in feature importance measures.

However, several persistent weaknesses are evident across the reviewed literature. The majority of machine learning studies focus on accident severity prediction at the individual crash level rather than spatial hotspot classification at the road segment level, limiting the direct applicability of their findings to infrastructure safety planning. Explainability re-mains underutilised: very few studies apply SHAP or LIME to their models, and fewer still report class-wise or segment-level explanations. Class imbalance is frequently acknowledged but inconsistently addressed, leading to inflated accuracy figures that obscure poor hotspot recall. Finally, most studies use a narrow feature set, typically limited to road type and basic weather variables, without systematically integrating the full range of spatial, temporal, geometric, traffic exposure, and environmental factors known to influence accident risk.

## III. RELATED CONCEPTS

### A. *Road Accident Hotspot*

A road accident hotspot is defined as a location, road segment, or zone where the observed number of accidents is statistically significantly higher than expected given the traffic volume and road characteristics. Various operational definitions exist in the literature, ranging from absolute fre-quency thresholds (e.g., five or more accidents within three years) to rate-based measures that normalise accident counts by exposure (vehicle kilometres travelled) [2]. The choice of definition directly affects which locations are labelled as hotspots and therefore the training data available to machine learning classifiers.

### *B. Machine Learning for Classification*

Supervised machine learning involves training a model on labelled examples to learn a mapping from input features to output class labels. For hotspot detection, the input features de-scribe each road segment (geometry, traffic, environment) and the output label indicates whether the segment is a hotspot or not. Commonly used classifiers in this domain include Logistic Regression, Decision Trees, Random Forests, Support Vector Machines, Gradient Boosting methods (XGBoost, LightGBM), and Multi-Layer Perceptron neural networks [9], [10].

### *C. Ensemble Methods*

Ensemble methods combine the predictions of multiple base learners to produce a more accurate and stable final prediction. Random Forest builds an ensemble of decision trees on bootstrap samples of the training data, using random feature subsets at each split to decorrelate individual trees [9]. Gradient Boosting methods such as XGBoost construct trees sequentially, with each tree correcting the residual errors of the previous ensemble [10]. Both methods are robust to overfitting, handle mixed feature types naturally, and provide built-in feature importance measures.

### *D. Explainable AI (XAI)*

Explainable AI refers to methods and techniques that make the predictions of machine learning models interpretable to human users. In the road safety context, explainability is essential for translating model outputs into engineering rec-ommendations that practitioners can act upon with confidence. SHAP (SHapley Additive exPlanations) provides theoretically grounded, consistent feature attributions based on cooperative game theory, assigning each feature a signed contribution value for each individual prediction [8]. LIME generates locally faithful linear approximations of black-box model behaviour for individual instances [12].

### *E. Spatial Analysis and GIS*

Geographic Information Systems (GIS) are used to manage, visualise, and analyse geospatial data. In road safety research, GIS tools support the spatial joining of accident records to road network segments, the computation of spatial features such as cluster membership and proximity to junctions, and the visualisation of predicted risk surfaces [3]. Kernel Density Estimation is a widely used GIS technique for producing continuous accident density maps from discrete point records.

### *F. Class Imbalance*

In binary hotspot classification tasks, the number of non-hotspot road segments typically far exceeds the number of hotspot segments, creating a class imbalance that biases stan-dard classifiers toward the majority class. SMOTE addresses this by generating synthetic minority-class examples in the feature space rather than simply duplicating existing samples, improving classifier sensitivity to the hotspot class [11].

## IV. IDENTIFIED RESEARCH GAPS

The literature review reveals several important gaps that limit the current state of machine learning-based road accident hotspot detection:

- Lack of Systematic Multi-Model Comparisons: Most existing studies evaluate one or two machine learning models in isolation. Systematic comparisons of multiple classifiers on the same balanced hotspot classification benchmark, using consistent evaluation metrics, are rare. This makes it difficult to draw general conclusions about which model families are most suitable for this task.
- Insufficient Explainability: The majority of reviewed machine learning studies do not apply any explainability technique to their models. Where feature importance is reported, it is typically limited to global Gini-based rankings that do not convey the direction or magnitude of individual feature effects. Class-wise and segment-level explanations using SHAP or LIME are almost entirely absent from the hotspot detection literature.
- Limited Feature Engineering: Many studies use a nar-row feature set limited to basic road type,

speed limit, and weather condition variables. Comprehensive feature en-gineering pipelines that jointly integrate spatial clustering attributes, temporal accident distribution patterns, detailed road geometry descriptors, traffic exposure metrics, and environmental features are underexplored.

- Inconsistent Handling of Class Imbalance: Class im-balance is a well-known challenge in accident datasets but is inconsistently addressed across studies. Many papers report overall accuracy as the primary metric, which is misleading when the hotspot class is a small minority. Transparent reporting of recall and F1-Score for the hotspot class, alongside robust resampling strategies, is needed.
- Limited Generalisability: Most studies are conducted on data from a single city or country, raising questions about the generalisability of trained models to different road environments, driving cultures, and accident reporting standards. Cross-regional validation studies are largely absent.
- Absence of Real-Time and Dynamic Frameworks: Existing approaches are predominantly static, trained on historical data without mechanisms for continuous model updating as new accident records accumulate. Real-time hotspot detection systems that incorporate live traffic and weather feeds are not yet well-developed in the academic literature.

## V. APPLICATIONS AND USE CASES

### *A. Road Safety Engineering*

Machine learning-based hotspot prediction maps can di-rectly support road safety audit processes. Highway agencies can prioritise geometric improvements, signage upgrades, and road surface treatments at high-risk segments identified by the model, optimising the allocation of limited safety improvement budgets. When combined with cost-benefit analysis, such frameworks enable evidence-based infrastructure investment decisions that go beyond simple frequency-based rankings.

### *B. Traffic Management and Enforcement*

Predicted hotspot locations can inform adaptive traffic man-agement strategies, including the deployment of variable speed limit systems, increased traffic enforcement presence during high-risk time windows, and rerouting recommendations for navigation applications. Integration with real-time traffic man-agement centres would allow dynamic risk scoring as traffic conditions change throughout the day.

### *C. Urban Planning and Land Use*

The spatial patterns revealed through machine learning and explainability analysis can inform urban planning decisions. New residential or commercial developments that generate significant additional traffic should be assessed against the predicted risk profile of adjacent road segments before plan-ning approval is granted. Junction design, pedestrian crossing placement, and road lighting decisions can all be guided by model-identified risk factors.

### *D. Insurance and Telematics*

Insurance providers and fleet operators can use predicted hotspot risk scores as one input into route risk scoring for telematics-based pricing and driver coaching applications. Drivers can be alerted to approaching high-risk segments via in-vehicle systems, supporting defensive driving behaviour.

### *E. Policy and Public Health*

National and regional transport authorities can use aggre-gated hotspot predictions to monitor trends in road network risk over time, evaluate the impact of safety interventions, and report progress against road casualty reduction targets. The transparent, explainable nature of machine learning frame-works ensures that findings can be communicated to non-technical policymakers and public stakeholders with confi-dence.

### *F. Academic Research*

This review contributes to academic knowledge in trans-portation engineering, traffic safety, spatial data science, and explainable AI. The identified research gaps provide a struc-tured agenda for future empirical studies, including the design of standardised benchmark datasets for hotspot detection, the development of novel feature engineering pipelines, and the investigation of causal inference methods that go beyond predictive association to support intervention design.

## VI. CONCLUSION

This paper presented a comprehensive review of machine learning methods for road accident hotspot detection, syn-thesising findings from statistical approaches, kernel density estimation, supervised classifiers, ensemble methods, deep learning architectures, and explainable AI techniques. The review identified ensemble tree methods, particularly Ran-dom Forest and XGBoost, as the most consistently well-performing approaches for structured tabular accident data, offering a strong balance of predictive accuracy, computational efficiency, and feature interpretability.

The review identified six key research gaps that limit the current state of the field: the absence of systematic multi-model comparisons on balanced hotspot benchmarks, insuf-ficient application of SHAP and LIME explainability, limited feature engineering pipelines, inconsistent handling of class imbalance, restricted generalisability across geographic re-gions, and the lack of real-time dynamic detection frameworks. Addressing these gaps represents a clear and important agenda for future research.

The practical value of machine learning-based hotspot de-tection is substantial. By translating raw accident records into actionable spatial risk predictions, these systems can support road safety engineers, urban planners, transport authorities, and policymakers in making evidence-based decisions that save lives and reduce the social and economic costs of road traffic accidents. Future work should focus on developing generalised, interpretable, and continuously updating hotspot detection pipelines that can be deployed across diverse road environments and data ecosystems.

## REFERENCES

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

[2] R. Elvik, “The predictive validity of empirical Bayes estimates of road safety,” *Accident Analysis & Prevention*, vol. 40, no. 6, pp. 1964–1969, 2008.

[3] Z. Xie and J. Yan, “Kernel density estimation of traffic accidents in a network space,” *Computers, Environment and Urban Systems*, vol. 32, no. 5, pp. 396–406, 2008.

[4] C. Chen, G. Zhang, J. Huang, J. Liu, and Y. Qiu, “Examining driver injury severity outcomes in rural non-interstate roadway crashes using a hierarchical ordered logit model,” *Accident Analysis & Prevention*, vol. 129, pp. 196–204, 2019.

[5] C. Yuan, Y. Li, M. Huang, and S. Wang, “Using traffic flow character-istics to predict real-time conflict risk: A novel method for trajectory data analysis,” *Analytic Methods in Accident Research*, vol. 35, art. no. 102217, 2022.

[6] A. B. Parsa, A. Taghipour, S. Derrible, and A. K. Mohammadian, “Real-time accident detection: Coping with imbalanced data,” *Accident Analysis & Prevention*, vol. 129, pp. 202–210, 2019.

[7] S. Jiang, W. Chen, and X. Wang, “Deep learning for traffic accident risk prediction: A convolutional approach on urban grid data,” *Transportation Research Part C: Emerging Technologies*, vol. 147, art. no. 104004, 2023.

[8] S. M. Lundberg and S.-I. Lee, “A unified approach to interpreting model predictions,” in *Proc. 31st Conf. Neural Information Processing Systems (NeurIPS)*, Long Beach, CA, USA, 2017, pp. 4765–4774.

[9] L. Breiman, “Random forests,” *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.

[10] T. Chen and C. Guestrin, “XGBoost: A scalable tree boosting system,” in *Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*, San Francisco, CA, USA, 2016, pp. 785–794.

[11] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, “SMOTE: Synthetic minority over-sampling technique,” *Journal of Artificial Intelligence Research*, vol. 16, pp. 321–357, 2002.

[12] M. T. Ribeiro, S. Singh, and C. Guestrin, “Why should I trust you? Ex-plaining the predictions of any classifier,” in *Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*, San Francisco, CA, USA, 2016, pp. 1135–1144.

[13] J. Aguero-Valverde, “Full Bayes Poisson gamma, Poisson lognormal, and zero inflated random effects models: Comparing the precision of crash frequency estimates,” *Accident Analysis & Prevention*, vol. 50, pp. 289–297, 2013.

[14] A. Theofilatos and G. Yannis, “A review of the effect of traffic and weather characteristics on road safety,” *Accident Analysis & Prevention*, vol. 72, pp. 244–256, 2014.

[15] J. Lee, B. Abdel-Aty, and I. Yue,

"Development of zone system for macro-level traffic safety analysis," *Journal of Transport Geography,* vol. 25, pp. 13–21, 2021.

[16] S. Moosavi, M. H. Samavatian, S. Parthasarathy, and R. Ramnath, "A countrywide traffic accident dataset," *arXiv preprint,* arXiv:1906.05409, 2019.

[17] Department for Transport, "Road Safety Data," UK Government Open Data, 2022. [Online]. Available: https://data.gov.uk/dataset/ cb7ae6f0-4be6-4935-9277-47e5ce24a11f/road-safety-data

[18] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[19] F. Pedregosa et al., "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

[20] S. M. Mohammad and P. D. Turney, "Crowdsourcing a word–emotion association lexicon," *Computational Intelligence,* vol. 29, no. 3, pp. 436–465, 2013.