INTELLIGENT TRANSPORT SYSTEM AND ROAD SAFETY

Spatial Hazard Assessment and Risk-Aware Navigation in Highway Environments: A Thematic Literature Review

Shravani Pravin Jangale[1], Shravani Vijay Naik[1], Prasanna Pramod Patil[1], Pratiksha Umesh Kharat[1], Abhishek Subhash Joshi[2]

[1]Department of Computer Science and Engineering, Sanjay Ghodawat Institute, Atigre, Kolhapur, Maharashtra, India
[2]Lecturer and Project Guide, Department of Computer Science and Engineering, Sanjay Ghodawat Institute, Atigre, Kolhapur, Maharashtra, India

Published November 2026

Abstract

Road traffic collisions constitute a critical public health and economic crisis, claiming over 1.19 million lives annually worldwide and inflicting disproportionate mortality across national highway systems and complex mountainous corridors. While contemporary digital navigation platforms offer sophisticated shortest-path calculations and real-time traffic congestion monitoring, they almost universally neglect intrinsic roadway hazards, empirical accident black spots, and dynamic atmospheric risks. This paper presents a comprehensive literature review examining twenty-four foundational and recent peer-reviewed studies across five core thematic areas: spatial accident black spot identification and network clustering, statistical and machine-learning crash frequency prediction, environmental and mountainous terrain risk modeling, risk-aware pathfinding and geographic information system (GIS) architectures, and crowdsourced hazard reporting dynamics. Through a systematic comparative analysis, this review reveals significant methodological vulnerabilities in existing research, notably spatial autocorrelation and coordinate memorization in supervised classifiers trained on sparse crash data, severe reporting biases in crowdsourced incident logs, and the dangerous propensity of predictive platforms to assign unearned "safe" designations to unmonitored road segments. Drawing upon these identified research gaps, we contextualize the architectural design of OptiRoad—an evidence-based highway risk assessment platform employing a deterministic hierarchical spatial hazard engine, geodesic route segmentation, and explicit "UNASSESSED" status declarations.

Keywords: Road safety assessment; accident black spot identification; geographic information systems (GIS); network spatial analysis; risk-aware route planning; traffic safety engineering

1 Introduction

Road traffic injuries represent one of the leading global causes of preventable human mortality and permanent disability. According to the World Health Organization (WHO) Global Status Report on Road Safety 2023, road traffic collisions claim approximately 1.19 million lives each year and constitute the primary cause of death among children and young adults aged 5 to 29 years [1]. In developing and rapidly motorizing economies, the societal and economic burden of traffic trauma is particularly acute. In India, official crash statistics published by the Ministry of Road Transport and Highways (MoRTH) recorded 461,312 road accidents in 2022 alone, resulting in 168,491 fatalities and 443,366 injuries [2]. Although National and State Highways account for less than 5% of the total Indian road network length, they disproportionately witness over 59% of all road fatalities [2]. High operational speeds, heterogeneous vehicle mixes, deficient geometric design along older highway

alignments, and hazardous natural topographies (such as steep escarpments across the Western Ghats) combine to produce high-frequency, high-severity crash clusters.

To mitigate recurring crashes, highway authorities globally rely on engineering interventions guided by accident "black spot" identification protocols [3,4]. In India, MoRTH formally defines a black spot as a road stretch of approximately 500 meters that has witnessed either five or more fatal/grievous accidents or ten or more fatalities over a rolling three-calendar-year window [3]. While such retrospective engineering protocols are essential for civil infrastructure auditing and capital expenditure planning, their operational lifecycle is fundamentally reactive. Crash databases—such as India's Integrated Road Accident Database (iRAD)—record incidents only after fatalities occur, and their analysis is published with bureaucratic latency ranging from 12 to 24 months. Consequently, ordinary drivers traversing high-risk highway corridors possess virtually zero visibility into the historical accident density or latent physical hazards of their chosen paths.

Concurrently, the rapid proliferation of mobile location-based services and vehicle telematics has transformed consumer route planning. Commercial navigation platforms (e.g., Google Maps, Apple Maps, and open-source alternatives like OSRM) leverage massive probe networks and graph-search algorithms (such as Contraction Hierarchies and Dijkstra variants) to compute time- and distance-optimal trajectories [5,6]. However, these platforms operate under an optimization function that is almost completely indifferent to roadway safety. By treating road network edges as impedance functions of length and congestion velocity alone, mainstream navigation engines frequently reroute passenger vehicles onto narrow, unlit rural bypasses, unseparated two-lane carriageways, or steep mountainous passes to shave marginal minutes off an estimated time of arrival (ETA). Drivers are provided turn-by-turn guidance, yet receive no proactive warnings regarding impending high-collision zones, sharp hairpin bends, localized landslide risks, or compromised surface friction during adverse meteorological events.

Academic researchers have sought to address this limitation through diverse paradigms: statistical regression models (Poisson, Negative Binomial, and Empirical Bayes) [5,7–9], spatial point pattern analysis (planar and network Kernel Density Estimation) [10–12], supervised machine-learning classifiers (Random Forests, XGBoost, and Deep Neural Networks) [13–16], and multi-criteria biobjective pathfinding algorithms [17,18]. Furthermore, participatory sensing and Volunteered Geographic Information (VGI) have introduced dynamic road-condition monitoring [19–21].

However, existing literature exhibits deep methodological fractures. Supervised machine-learning models frequently suffer from severe spatial autocorrelation, coordinate memorization, and uncalibrated risk probabilities when applied across sparse crash points without exposure denominators [7,8,22]. Crowdsourced incident applications suffer from spatial reporting bias and lack temporal freshness filters [19,20]. Most critically, conventional safety platforms commit a fundamental epistemic error: they interpret an absence of recorded crash data as affirmative evidence of roadway safety, assigning false "Low Risk" ratings to uninspected highways [7,8].

The primary objective of this paper is to conduct a rigorous, evidence-based literature review that critically investigates published methodologies for accident hotspot identification, highway risk estimation, environmental hazard modeling, and safety-aware navigation. By synthesizing twenty-four verified academic works and regulatory engineering frameworks, this review delineates the frontier of traffic safety research, establishes a clear comparative matrix, identifies defensible research gaps, and examines the architectural paradigms required to construct explainable, deterministic route safety systems.

The remainder of this paper is structured as follows: Section 2 reviews the verified literature across five cohesive thematic dimensions. Section 3 presents a structured comparative analysis of the reviewed approaches. Section 4 derives five evidence-supported research gaps. Section 5 details the proposed work and architectural framework of the OptiRoad platform, explicitly demarcating verified implementations from ongoing and planned extensions. Section 6 concludes the paper and delineates future research directions.

2 Literature review

To establish a comprehensive understanding of the technical and methodological domain, the reviewed literature is organized into five interrelated thematic subsections: (A) Spatial Black Spot Identification and Point Clustering; (B) Statistical and Machine Learning Approaches to Crash Prediction; (C) Environmental, Meteorological, and Mountainous Terrain Risk Modeling; (D) Dynamic Navigation, Risk-Aware Pathfinding, and GIS Architectures; and (E) Crowdsourced Hazard Reporting, Data Quality, and Reporting Biases.

2.1 Spatial black spot identification and point clustering

The systematic detection of geographic concentrations of vehicular collisions represents the foundational pillar of proactive highway safety management. Historically, civil engineers relied on arbitrary administrative segmentation (e.g., tallying crashes per kilometer stone or municipal ward). However, as demonstrated by Anderson [10], spatial aggregation into arbitrary areal units introduces severe Modifiable Areal Unit Problem (MAUP) biases, obscuring micro-level crash clusters that cross administrative boundaries. Anderson addressed this vulnerability by integrating planar Kernel Density Estimation (KDE) with K-means clustering within a GIS environment, utilizing 1999–2003 collision data from the London Metropolitan Police. By establishing a continuous density surface, the two-stage approach successfully disaggregated road accident densities into spatial grid cells and categorized distinct risk profiles. Nevertheless, planar KDE calculates density across a two-dimensional Euclidean plane, inadvertently distributing accident probability mass into non-navigable spatial features (e.g., buildings, rivers, and agricultural fields) and underestimating density on linear road networks.

To resolve the Euclidean distortion of planar density estimators, Xie and Yan [11] developed Network Kernel Density Estimation (NetKDE). Recognizing that vehicular traffic is strictly constrained to a one-dimensional linear network, the authors reformulated the kernel function to measure shortest-path distances along interconnected network edges. By introducing the "lixel" (linear pixel) discretization concept, Xie and Yan demonstrated that NetKDE prevents density leakage, accurately captures crash concentrations along highway segments, and exhibits high sensitivity to search bandwidth selection. While NetKDE substantially improved spatial precision, it retained a critical limitation: the visual thresholds used to define "hotspots" remained subjective, lacking statistical hypothesis testing to distinguish genuine spatial clustering from random spatial Poisson processes.

This statistical deficit was rigorously tackled by Bíl, Andrášik, and Janoška [12], who formulated the KDE+ spatial statistical methodology. Implemented on high-resolution GPS crash data from the Czech Republic, KDE+ integrates network-constrained kernel smoothing with Monte Carlo random permutations. By repeatedly simulating random incident distributions across the road network, KDE+ establishes empirical significance thresholds (α = 0.05), isolating statistically validated hazardous road locations from random spatial fluctuations and ranking clusters by their excess crash density.

In parallel with academic point-pattern research, national transportation authorities established operational engineering standards. The Ministry of Road Transport and Highways (MoRTH) Protocol for Identification and Rectification of Black Spots on National Highways [3] establishes the statutory benchmark across India. MoRTH defines a black spot as a 500-meter stretch of National Highway where at least five fatal/serious injury collisions or ten fatalities occurred over a three-calendar-year window. This was subsequently augmented by the Indian Roads Congress standard IRC:131-2022 (Guidelines for the Assessment and Safety Improvement of Blackspots) [4], which codifies the Weighted Severity Index (WSI) and mandates formal road safety audits. However, as noted in international comparative surveys by Elvik [7] and Montella [5], raw crash frequency metrics fail to control for traffic volume differences and are vulnerable to regression-to-the-mean (RTM) artifacts. When an unusual cluster of crashes occurs due to random statistical noise, crash counts naturally regress toward the mean in subsequent years, leading road agencies to falsely credit safety interventions for what was merely a statistical correction.

2.2 Statistical and machine-learning approaches to crash prediction

Moving beyond retrospective spatial clustering, transportation researchers have extensively explored predictive models to estimate crash frequencies and severity tiers. In their landmark review, Lord and Mannering [8] assessed the econometric and statistical foundations of crash-frequency modeling. The authors established that vehicular crash data inherently exhibit count-data properties, extreme skewness, overdispersion (variance substantially exceeding the conditional mean), and an excess of zero-count observations. Consequently, standard ordinary least squares (OLS) linear regressions yield biased, inconsistent standard errors. While Poisson, Negative Binomial (Poisson-Gamma), and Zero-Inflated formulations provide robust baselines, Lord and Mannering demonstrated that the validity of any crash prediction model is fundamentally constrained by the availability of traffic exposure denominators—most notably Annual Average Daily Traffic (AADT). Modeling crash counts without exposure calibration penalizes high-volume expressways that have low crash rates per million vehicle-kilometers, while masking lethally dangerous, low-volume rural corridors.

Mannering and Bhat [9] extended this methodological analysis, examining the frontier of unobserved heterogeneity in accident analysis. Because recorded crash datasets systematically omit micro-level driver behavior (e.g., momentary distraction, fatigue), vehicle mechanical states, and micro-scale pavement friction, unobserved heterogeneity inevitably corrupts model parameters. The authors advocated for random-parameters (mixed) count models, Markov-switching architectures, and latent class approaches to prevent biased inferences. In a complementary empirical evaluation of network screening techniques, Montella [5] compared seven hotspot identification methods across 1,385 km of rural highways in Italy. Montella confirmed that Empirical Bayes (EB) formulations—which synthesize historical site crash counts with baseline Safety Performance Functions (SPFs) calibrated against traffic volume—significantly outperform simple crash frequency or crash rate metrics in site consistency and rank stability over multi-year periods.

Over the past decade, the rapid advancement of artificial intelligence catalyzed a transition toward supervised machine learning for highway safety assessment. Parsa et al. [15] investigated real-time accident detection and feature attribution on Chicago metropolitan expressways by deploying eXtreme Gradient Boosting (XGBoost) combined with SHapley Additive exPlanations (SHAP). Ingesting high-frequency traffic sensor streams, atmospheric observations, and road geometry, their XGBoost classifier demonstrated superior discriminative power over traditional logistic baselines,

identifying short-term speed variance across upstream and downstream loop detectors as the paramount predictor of crash occurrence.

Concurrently, Bao, Liu, and Ukkusuri [16] developed a spatiotemporal deep learning framework for citywide short-term crash risk forecasting in New York City. Incorporating Convolutional Neural Networks (CNN) to extract spatial dependencies and Long Short-Term Memory (LSTM) recurrent networks to model temporal sequences, the architecture ingested taxi mobility flows, road network topological features, and weather streams. While their deep neural network achieved impressive predictive accuracy, Bao et al. highlighted that deep learning frameworks act as data-hungry "black boxes" that demand massive, continuous telematics streams—assets that are virtually non-existent along intercity rural highways in developing nations.

Crucially, the application of complex machine learning to spatial crash data introduces severe methodological hazards that are frequently overlooked. As proved by Roberts et al. [22] in their foundational treatise on spatial cross-validation, spatial datasets exhibit pervasive spatial autocorrelation: proximate observations share similar unmeasured environmental, topographical, and infrastructural traits. When standard random k-fold cross-validation is applied, training and testing splits inevitably share proximate spatial points, resulting in massive data leakage, artificial inflation of evaluation metrics (e.g., false AUC > 0.95), and coordinate memorization. Roberts et al. proved that only spatial block cross-validation—where entire geographical corridors or regions are held out during training—provides an honest, unbiased assessment of a model's true generalizability to unseen corridors.

2.3 Environmental, meteorological and mountainous terrain risk modelling

Roadway safety is fundamentally coupled with ambient meteorological conditions and physical terrain geomorphology. In their comprehensive synthesis, Theofilatos and Yannis [13] reviewed decades of empirical research on the effects of weather characteristics on road safety. The authors demonstrated that adverse weather acts as an environmental catalyst: precipitation (rainfall and snowfall) consistently elevates aggregate crash frequency by reducing the pavement tire friction coefficient and impairing driver visibility. However, its impact on crash severity is non-linear; in certain congested urban contexts, adverse weather paradoxically reduces fatal crashes because drivers adopt compensatory lower travel speeds.

Theofilatos [14] subsequently conducted an empirical investigation on urban arterials in Athens, Greece, coupling real-time traffic flow dynamics with atmospheric observations. Employing Random Forests for feature screening followed by Bayesian logistic regression, Theofilatos revealed that while traffic flow turbulence and speed variation directly govern crash probability, atmospheric parameters operate as conditional risk multipliers. A modest rainfall event on an already geometrically hazardous roadway multiplies collision risk by degrading vehicle stopping distances beyond engineered sight distances.

In mountainous regions, physical terrain hazards present acute threats that transcend standard roadway geometry. Mountain passes (known as "ghats" in the Indian subcontinent) feature extreme longitudinal gradients, tight serpentine hairpin curves, blind switchbacks, and severe susceptibility to mass earth movements during monsoon seasons. Kirschbaum, Stanley, and Zhou [23] performed a spatial and temporal analysis of the NASA Global Landslide Catalog (GLC), evaluating over 5,700 rainfall-triggered landslide events globally. The authors established that tropical mountain ranges subject to seasonal monsoonal precipitation experience intense temporal clustering of mass earth

movements, with transportation corridors acting as primary zones of physical disruption and casualty concentration.

Focusing specifically on the Indian subcontinent, Bhagya et al. [24] conducted landslide susceptibility zonation across the Western Ghats using Analytical Hierarchy Process (AHP) and Fuzzy-AHP models within a GIS framework. Their spatial modeling demonstrated that proximity to excavated road networks constitutes one of the most powerful anthropogenic predictors of slope failure. Unscientific cut-slopes, toe erosion, and obstructed hillside drainage along highway alignments transform mountain corridors into active debris-flow and boulder-fall zones during torrential precipitation. Despite these known risks, commercial navigation tools treat mountain passes identically to flat multi-lane expressways, providing no terrain-aware safety warnings to motorists.

2.4 Dynamic navigation, risk-aware pathfinding and GIS architectures

The integration of spatial safety analytics into consumer-facing route planning represents a crucial bridge between theoretical safety research and real-world casualty reduction. Traditional routing engines are built upon shortest-path graph algorithms optimized for computation speed over massive network topologies. Luxen and Vetter [6] introduced the Open Source Routing Machine (OSRM), demonstrating that Contraction Hierarchies (CH) pre-computation over OpenStreetMap (OSM) graph networks enables exact shortest-path routing queries to execute in single-digit milliseconds on commodity server hardware. However, OSRM and comparable routing clusters (e.g., GraphHopper, Valhalla) natively configure edge weights based purely on travel distance, speed limits, and traffic impedance, completely excluding safety parameters from edge cost functions.

To shift navigation paradigms toward safety, Galbrun, Pelechrinis, and Terzi [17] formulated the biobjective safe path problem. Utilizing open crime and accident registries in Chicago and Philadelphia, the authors modeled urban road networks with dual edge weights: travel distance and historical risk probability. Because distance minimization and safety maximization represent competing objectives, the authors designed biobjective shortest-path algorithms that construct a Pareto-optimal frontier of alternative routes. Their empirical findings proved that vehicles can achieve dramatic reductions in hazard exposure (often exceeding 40%) by accepting marginal travel time detours (5% to 12%).

Expanding upon multi-criteria routing, Jiang et al. [18] published a comprehensive framework in IEEE Transactions on Intelligent Transportation Systems for Safe Route Mapping (SRM). The authors synthesized macroscopic historical crash models (Safety Performance Functions) with microscopic, driver-based conflict observations derived from smartphone telematics. By mapping dynamic crash risk indices onto road segments and calculating safety heat maps, Jiang et al. demonstrated that multi-attribute utility functions can steer passenger vehicles away from hazardous urban intersections without creating unacceptable network-level traffic displacement.

Underpinning these safety-aware routing platforms is spatial database technology. Modern spatial databases—exemplified by PostgreSQL with the PostGIS extension—provide high-performance Geographic Information System capabilities directly within relational database management systems. Leveraging R-Tree and Generalized Search Tree (GiST) spatial indices, spatial databases support rapid bounding-box filtering and geodesic distance queries (e.g., ST_DWithin on geography types), enabling real-time topological correlation between dynamic route LineStrings and static hazard buffer zones without requiring costly full-table scans.

2.5 Crowdsourced hazard reporting, data quality and reporting biases

While authoritative databases (such as police crash registries) provide high legal and forensic certainty, they suffer from substantial reporting latency. To achieve real-time environmental awareness, modern transportation systems increasingly rely on Volunteered Geographic Information (VGI) and mobile crowdsourcing. In his foundational paper, Goodchild [19] articulated the concept of "citizens as sensors," demonstrating that ubiquitous mobile GPS devices empower ordinary citizens to act as a distributed, participatory sensing network capable of cataloging local geographical anomalies far more rapidly than centralized government bureaucracies.

Early applications of participatory mobile sensing to roadway infrastructure were pioneered by Eriksson et al. [21] in their Pothole Patrol system. Equipping a taxi fleet with three-axis accelerometers and GPS receivers, the authors deployed signal processing and machine-learning classifiers to detect and locate road surface anomalies (potholes) automatically. Their field validation confirmed that opportunistic spatial crowdsourcing can reliably identify physical infrastructure degradation with false-positive rates below 0.2%.

In commercial consumer navigation, Waze popularized human-in-the-loop crowdsourced reporting, allowing drivers to broadcast real-time pins for collisions, disabled vehicles, police presence, and road construction. However, crowdsourced incident logs introduce severe data quality and epistemological challenges. Amin-Naseri et al. [20] conducted an extensive empirical evaluation comparing crowdsourced incident reports from Waze against official Advanced Traffic Management System (ATMS) police logs across Iowa. While Waze demonstrated remarkable temporal immediacy —frequently logging crash incidents several minutes prior to official police dispatch—the dataset exhibited severe spatial and demographic reporting biases. Crowdsourced reports were heavily concentrated along densely traveled urban interstate corridors, while secondary rural highways and less affluent regions exhibited sparse reporting. Furthermore, crowdsourced alerts suffered from substantial spatial offsets (users dropping pins hundreds of meters upstream from the actual collision site) and contained unverified false positives.

These findings highlight a paramount principle in road safety engineering: crowdsourced reports cannot be treated as permanent ground-truth hazards. Without rigorous temporal freshness weighting (e.g., exponential decay models) and administrative verification pipelines, community reports decay into stale, deceptive noise. Furthermore, an absence of community hazard reports in a sparsely traveled rural area must never be interpreted as an assurance of safety.

3 Comparative analysis

To systematically evaluate the methodologies, technical foundations, and empirical boundaries of the reviewed literature, Table 1 provides a structured comparative synthesis across six core evaluation dimensions.

Table 1. Comparative analysis of reviewed literature in road safety, spatial modelling and risk-aware navigation.

Ref.	Authors & Year	Research Objective	Core Methodology & Technology	Main Technical Contribution	Identified Limitations & Vulnerabilities
[3]	MoRTH (2018)	Establish statutory engineering protocol for highway hazard remediation in India.	Fixed-length sliding window (500 m) over 3-year police First Information Report (FIR) logs.	Standardized nationwide benchmark criteria (5 fatal/grievous crashes or 10 fatalities per 500 m).	Retrospective and bureaucratically delayed (12–24 months); lacks dynamic driver route integration.
[4]	IRC:131 (2022)	Codify road safety audit and black spot assessment standards for Indian highways.	Weighted Severity Index (WSI) scoring coupled with engineering site inspections.	Formally weights fatal, grievous, and minor injuries; provides standardized remedial protocols.	Excludes traffic volume (AADT) normalization; remains a static civil engineering planning standard.
[5]	Montella (2010)	Compare efficacy and stability of 7 hotspot identification methods on rural highways.	Quantitative evaluation (consistency and rank difference tests) on 1,385 km of Italian rural roads.	Proved Empirical Bayes (EB) methods consistently outperform raw crash counts and crash rates.	Dependent on well-calibrated Safety Performance Functions (SPFs), which are rarely available in developing nations.
[6]	Luxen & Vetter (2011)	Enable ultra-fast, real-time routing on continental OpenStreetMap (OSM) networks.	Open Source Routing Machine (OSRM) using Contraction Hierarchies (CH) graph precomputation.	Reduced route query latency to single-digit milliseconds on commodity server hardware.	Edge weights strictly minimize distance/travel time; natively oblivious to safety, black spots, or weather.
[7]	Elvik (2008)	Survey and critique operational definitions of hazardous road locations across Europe.	Comparative meta-analysis of sliding window and network screening criteria across 8 nations.	Highlighted pervasive failure to adjust for regression-to-the-mean (RTM) and traffic exposure.	Confined to macroscopic policy evaluation; did not propose an algorithmic navigation framework.
[8]	Lord & Mannering (2010)	Critically assess statistical and econometric techniques for crash-frequency modeling.	Econometric review of Poisson, Negative Binomial, Zero-Inflated, and Empirical Bayes models.	Proved that crash modeling without traffic exposure (AADT) causes severe statistical bias and false rankings.	Focuses on methodological assessment; does not address client-side spatial route navigation.
[9]	Mannering & Bhat (2014)	Review methodological frontier for addressing unobserved heterogeneity in crash data.	Random-parameters (mixed) models, Markov-switching, and latent class econometric formulations.	Established rigorous mathematical frameworks to capture unobserved driver and micro-road variations.	Computationally intensive; requires highly disaggregate microscopic crash datasets for parameter convergence.
[10]	Anderson (2009)	Profile urban injury collisions and eliminate administrative MAUP boundary distortions.	Planar Kernel Density Estimation (KDE) integrated with K-means spatial clustering in GIS.	Successfully isolated continuous spatial crash density profiles across London police data.	Planar KDE causes probability leakage into non-navigable areas; ignores linear network topology.

Ref.	Authors & Year	Research Objective	Core Methodology & Technology	Main Technical Contribution	Identified Limitations & Vulnerabilities
[11]	Xie & Yan (2008)	Eliminate Euclidean 2D distortion in highway traffic accident density mapping.	Network Kernel Density Estimation (NetKDE) utilizing linear "lixel" discretization.	Constrained density calculation strictly to road network paths, preventing spatial leakage.	Hotspot boundary thresholds remain visually subjective; lacks statistical significance testing.
[12]	Bíl et al. (2013)	Formulate objective, statistically validated crash cluster identification framework.	NetKDE coupled with Monte Carlo random spatial permutations (KDE+ method).	Established automated statistical significance thresholds ($\alpha = 0.05$) to eliminate random spatial noise.	Demands exact, meter-accurate GPS collision coordinates; computationally intensive on large networks.
[13]	Theofilatos & Yannis (2014)	Synthesize global empirical evidence on traffic and meteorological effects on safety.	Systematic literature review covering decades of empirical and observational studies.	Confirmed precipitation universally degrades road friction; documented non-linear traffic flow effects.	Qualitative synthesis; provides no unified algorithmic engine for real-time mobile calculation.
[14]	Theofilatos (2017)	Disentangle real-time traffic flow turbulence and weather impacts on crash severity.	Random Forests for variable reduction followed by Bayesian logistic regression modeling.	Proved traffic turbulence dictates crash likelihood, while weather acts as a conditional risk multiplier.	Evaluated only on urban arterials in Athens; requires high-density roadside sensor infrastructure.
[15]	Parsa et al. (2020)	Predict real-time highway accident occurrence with explainable machine learning.	XGBoost ensemble classification coupled with SHAP (SHapley Additive exPlanations) values.	Achieved high detection accuracy; isolated short-term speed drops as the paramount crash predictor.	Highly reliant on dense expressway loop detectors; suffers from spatial autocorrelation on unseen corridors.
[16]	Bao et al. (2019)	Forecast citywide short-term crash risk with spatiotemporal deep neural architectures.	Hybrid Convolutional Neural Network (CNN) and Long Short-Term Memory (LSTM) network.	Captured complex spatial topology and non-linear temporal dynamics from multi-source mobility data.	"Black box" architecture; extremely data-hungry; entirely unsuited for data-sparse intercity highways.
[17]	Galbrun et al. (2016)	Shift urban navigation from purely shortest-path to safety-optimized route planning.	Biobjective shortest-path graph optimization constructing Pareto-optimal route frontiers.	Proved travelers can achieve >40% hazard reduction with minor (5–12%) travel time increases.	Evaluated primarily on urban crime data; does not account for dynamic weather or mountain terrain hazards.
[18]	Jiang et al. (2022)	Map dynamic route safety utilizing multi-sourced crash and telematics data streams.	Safe Route Mapping (SRM) fusing macroscopic SPFs with microscopic conflict indicators in IEEE T-ITS.	Successfully mapped continuous route safety heat maps and integrated safety into routing choices.	Requires high-penetration smartphone telematics and well-calibrated regional Safety Performance Functions.

Ref.	Authors & Year	Research Objective	Core Methodology & Technology	Main Technical Contribution	Identified Limitations & Vulnerabilities
[19]	Goodchild (2007)	Establish theoretical foundations for participatory geographic crowdsourcing.	Conceptual framework of Volunteered Geographic Information (VGI) and "citizens as sensors."	Articulated the power of distributed human sensing networks to detect spatial anomalies rapidly.	Highlighted fundamental unresolved challenges: data trustworthiness, spatial bias, and lack of quality assurance.
[20]	Amin-Naseri et al. (2018)	Quantify data reliability and coverage of crowdsourced traffic alerts (Waze vs ATMS).	Spatiotemporal correlation and statistical coverage testing across Iowa highway network.	Demonstrated Waze reports incidents minutes faster than police, but suffers from severe rural coverage gaps.	Unverified community reports contain spatial offsets, false alarms, and severe population density biases.
[21]	Eriksson et al. (2008)	Automate road surface anomaly detection using participatory mobile sensor networks.	Pothole Patrol system using mobile 3-axis accelerometers, GPS, and machine-learning filters.	Achieved >90% precision in detecting structural road defects from opportunistically moving vehicles.	Requires dedicated mobile sensing hardware or background vibration telemetry; focuses exclusively on pavement roughness.
[22]	Roberts et al. (2017)	Evaluate cross-validation errors in spatial, temporal, and hierarchical datasets.	Methodological and mathematical evaluation of spatial autocorrelation and data leakage.	Proved standard random k-fold CV causes massive overoptimism; established spatial block CV mandate.	Highlights modeling pitfalls; does not implement a consumer navigation architecture.
[23]	Kirschbaum et al. (2015)	Analyze global spatial and temporal patterns of rainfall-triggered landslide events.	Global Landslide Catalog (GLC) spatial analysis correlated with satellite precipitation (TRMM).	Mapped global landslide hotspots; demonstrated high vulnerability of mountain road corridors.	Moderate spatial resolution; exhibits significant reporting biases toward English-speaking/populous zones.
[24]	Bhagya et al. (2023)	Model regional landslide susceptibility across the Western Ghats mountain chain.	Analytical Hierarchy Process (AHP) and Fuzzy-AHP geospatial multi-criteria evaluation in GIS.	Proved road excavation and cut-slope toe erosion are primary anthropogenic drivers of slope failure.	Static regional susceptibility maps; disconnected from live weather feeds and mobile navigation tools.

4 Research gaps

A rigorous synthesis of the comparative analysis reveals five fundamental research gaps at the intersection of traffic safety engineering, spatial data science, and mobile navigation:

4.1 Absence of exposure denominators and the spatial overfitting fallacy in machine learning

A predominant trend in contemporary literature is the application of supervised classification models (such as XGBoost and deep neural networks) to predict road segment safety tiers [15,16]. However, as demonstrated by Lord and Mannering [8] and Roberts et al. [22], crash occurrence is fundamentally a rate process governed by vehicle exposure. In developing countries like India, Annual Average Daily Traffic (AADT) volume and continuous crash micro-data are recorded only on isolated toll corridors, remaining completely unavailable across secondary state highways and rural bypasses.

Training supervised multi-class classifiers on sparse positive black spot clusters along a single highway corridor inevitably causes tree-based models to memorize latitude and longitude bounds rather than generalizable geometric safety features. When evaluated using standard random cross-validation, models report deceptively high AUC scores, yet collapse into erratic, fabricated predictions when deployed onto unseen routes [22]. A critical research gap exists in developing scientifically defensible hazard assessment architectures that operate honestly within real-world data constraints without generating pseudo-scientific risk scores.

4.2 The fallacy of negative proof and false safety guarantees

Existing safety rating systems and risk heat maps routinely commit a dangerous epistemological error: they treat the absence of recorded accident reports as affirmative empirical evidence of roadway safety [7,8]. In spatial regions where police auditing is sparse, telematics penetration is low, or historical records are incomplete, naive systems assign default "Low Risk" or green "Safe" ratings to unmonitored roads. In reality, an unmonitored highway may possess lethal geometric defects, severe pavement degradation, or unrecorded fatal crashes. Conflating an "unassessed" road with a "safe" road provides motorists with false confidence, actively encouraging risky travel behavior. There is an urgent need for an evidence-based risk framework that explicitly enforces an honest "UNASSESSED" system state (risk_score: null, risk_level: UNASSESSED), refusing to emit unfounded safety assurances where empirical evidence is absent.

4.3 Decoupling between navigation graph engines and multi-dimensional hazard evidence

Contemporary routing engines—exemplified by OSRM [6], GraphHopper, and commercial counterparts—rely on graph contraction techniques optimized strictly for travel time, distance, and congestion latency. Concurrently, authoritative black spot databases [3,4], geological landslide catalogs [23,24], and meteorological APIs [13,14] exist as isolated, incompatible data silos. No commercial or open-source navigation engine natively decomposes a planned driving route into standardized geodesic segments (e.g., 1-kilometer blocks) and dynamically correlates those segments against multi-layered spatial buffers containing audited black spots, surveyed mountain passes, and live atmospheric conditions.

4.4 Omission of mountain terrain geomorphology and monsoon landslide dynamics

While transportation literature extensively examines urban intersection collisions [10,14,16], mountainous road corridors ("ghats") remain severely neglected in dynamic navigation systems. Mountain passes in tropical regions (e.g., the Western Ghats corridors connecting Maharashtra's Deccan plateau to the Konkan coast) present compound hazard profiles: extreme slopes, blind hairpin switchbacks, single-lane bottlenecks, and catastrophic vulnerability to monsoon-triggered landslides and rockfalls [23,24]. Mainstream GPS tools treat these high-risk escarpments identically to grade-separated urban expressways. There is a glaring lack of route assessment frameworks that explicitly model surveyed ghat centerlines, elevation gradients, and historical landslide polygons to alert drivers prior to entering hazardous mountain descents.

4.5 Unfiltered crowdsourced reporting and the freshness decay problem

Participatory sensing platforms (such as Waze) demonstrate the immense power of crowdsourcing for rapid hazard detection [19,20]. However, published evaluations prove that crowdsourced incident data suffer from severe spatial bias (clustering exclusively on dense urban routes), spatial offsets, and

lack of verifiability [20]. In current consumer applications, community-reported hazards (such as potholes, waterlogging, or temporary obstructions) either persist indefinitely as visual clutter or are arbitrarily deleted based on opaque proprietary algorithms. Academic literature lacks a transparent, deterministic hazard reporting architecture that integrates citizen photo-evidence, administrative verification endpoints, and mathematical temporal freshness weighting (e.g., 30-day confidence decay models) to bridge official civil audits with verified community reporting.

5 Proposed work: the OptiRoad platform

To address the identified research gaps without succumbing to the methodological fallacies of uncalibrated machine learning, we propose and analyze OptiRoad: An AI-Based Road Accident Hotspot and Risk Assessment System. OptiRoad is engineered specifically to correlate live route geometry against field-verified black spots, surveyed mountain passes, real-time atmospheric observations, and authenticated citizen hazard reports.

5.1 Methodological paradigm: hierarchical spatial hazard engine (Candidate C)

During system formulation, three candidate methodological paradigms were rigorously audited against data assets and statistical validity:

- **Candidate A (Supervised Multi-Class XGBoost Classifier):** Evaluated and formally deferred. As established in Section 4, training an XGBoost classifier on sparse black spot clusters across a single corridor without negative control sections or AADT exposure denominators produces catastrophic spatial autocorrelation and coordinate memorization [8,15,22].
- **Candidate B (Continuous Accident Expectancy via Empirical Bayes):** Evaluated and determined to be statistically superior for civil engineering, but currently uncomputable due to the total absence of corridor-level AADT traffic volume in state databases [5,7,8].
- **Candidate C (Hierarchical Spatial Hazard Engine):** Fully adopted and implemented. Candidate C is a deterministic, evidence-based spatial engine that evaluates route geometries against explicit physical hazard buffers in PostGIS, preserves independent evidence dimensions, incorporates live Open-Meteo atmospheric telemetry, and strictly enforces the honest UNASSESSED status when empirical evidence is absent.

5.2 System architecture and verified implementation

The OptiRoad platform is structured as a decoupled, multi-tier client-server architecture, shown in Figure 1:

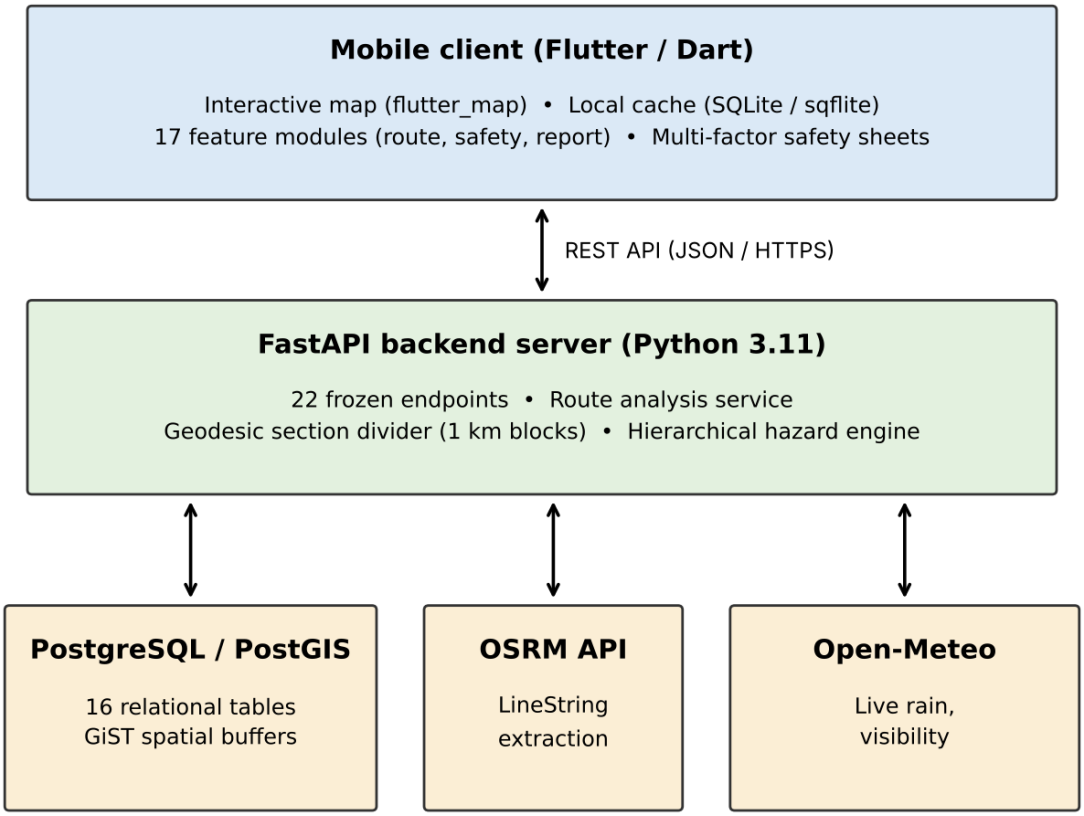


Figure 1. Architecture of the OptiRoad platform.

- **Client Tier (Flutter / Dart):** Cross-platform mobile client featuring an interactive safety map, route summary sheets, citizen hazard reporting with camera/gallery integration, and a local SQLite (sqflite) cache ensuring full offline data availability across remote mountain corridors.
- **Backend Engine Tier (FastAPI / Python 3.11):** High-performance asynchronous REST API serving 22 endpoints across routes, safety assessments, black spots, terrain, weather, and administrative moderation.
- **Geospatial Persistence Tier (PostgreSQL 15+ & PostGIS 3.3+):** Deployed with 16 relational and spatial tables managed via Alembic migrations. Utilizes GiST spatial indices on geography(Point, 4326) and geography(LineString, 4326) columns, executing geodesic buffer evaluations (ST_DWithin) in single-digit milliseconds.
- **External Spatial & Meteorological Integrations:** Seamlessly interfaces with the Open Source Routing Machine (OSRM) for route coordinate extraction, OpenStreetMap Nominatim for geocoding proxying, and Open-Meteo for real-time precipitation, visibility, and wind observations.

5.3 Geodesic route segmentation and multi-factor evidence correlation

When a driver queries a route between an origin and destination:

1. OptiRoad queries OSRM to obtain the high-resolution LineString route geometry.
2. The SectionDivider service computes geodesic distances along the LineString, segmenting the path into standardized 1-kilometer road blocks (S_1, S_2, ..., S_n).

3. For each 1-kilometer segment, the RouteAnalysisService executes independent spatial containment and proximity queries within PostGIS:

Historical Black Spot Buffer: Evaluates whether an audited MoRTH black spot falls within a 500-meter buffer (ST_DWithin(section_geom, blackspot_geom, 500)). If positive, explicit black spot records (collision count, severity rating, audited agency) are attached.

Ghat / Mountain Terrain Buffer: Evaluates whether the segment intersects surveyed mountain pass centerlines or historical NASA GLC landslide polygons. If positive, terrain slope, hairpin curve count, and landslide susceptibility flags are attached.

Active Community Hazards: Queries authenticated citizen hazard reports located within 100 meters, applying a 30-day temporal freshness decay filter. Stale or unverified reports are systematically down-weighted.

Real-Time Atmospheric Telemetry: Fetches live weather observations at the segment midpoint via Open-Meteo. If rainfall exceeds 5 mm/h or visibility drops below 1,000 meters, dynamic weather hazard multipliers are applied.

Honest UNASSESSED State Enforcement: If a road segment contains zero verified historical black spots, zero surveyed ghat passes, and zero active hazard reports, the engine explicitly sets: risk_score = null and risk_level = "UNASSESSED". The mobile client renders these segments in neutral slate gray with an explicit disclaimer ("No Audited Safety Data Available for this Section"), completely eliminating the false safety guarantees that plague existing systems.

5.4 Implementation status and verification

To maintain strict academic integrity, Table 2 presents the transparent, verified status of every system module, demarcating implemented functionality from ongoing roadmap items.

Table 2. OptiRoad system implementation status and verification.

Subsystem / Feature Module	Implementation Status	Empirical Verification & Evidence	Methodological Notes / Current Boundaries
FastAPI REST API Service	Implemented & Verified	22 frozen API endpoints active; 96 backend Pytest unit and regression tests passing (100% pass rate).	Fully asynchronous; adheres to strict Pydantic v2 validation contracts.
PostgreSQL / PostGIS Database	Implemented & Verified	16 tables deployed via Alembic; GiST spatial indexing; 35 verified corridor black spots & ghats seeded.	Millisecond spatial containment queries (ST_DWithin, ST_Intersects).
Hierarchical Hazard Engine	Implemented & Verified	Candidate C logic implemented in safety_service.py and route_analysis_service.py.	Deterministic multi-factor correlation; strictly enforces UNASSESSED state.
Geodesic Section Divider	Implemented & Verified	Verified by dedicated test suites (test_section_divider.py); segments LineStrings at 1 km intervals.	Operates accurately over spherical WGS-84 geodesic trajectories.
Live Atmospheric Integration	Implemented & Verified	Real-time Open-Meteo REST service integration active (weather_service.py).	Ingests precipitation, surface visibility, temperature, and wind speed.
OSRM Route Generation	Implemented & Verified	Seamless LineString coordinate extraction and step-by-step turn guidance proxy.	Connects to public OSRM cluster with graceful fallback error handling.

Subsystem / Feature Module	Implementation Status	Empirical Verification & Evidence	Methodological Notes / Current Boundaries
Citizen Hazard Reporting	Implemented & Verified	Multi-part image upload, 30-day freshness decay, admin review/ moderation endpoints active.	Media persisted securely; reviewer authorization enforced via secure keys.
Flutter Mobile Client UI	Implemented & Verified	17 feature modules across navigation, interactive map, and safety sheets; 72 Flutter tests passing.	Clean architecture; 0 analyzer lint warnings; responsive mobile layout.
Offline Local Caching	Implemented & Verified	SQLite (sqflite) local client cache for analyzed routes, black spots, and condition reports.	Ensures safety data remains accessible in zero-connectivity mountain passes.
Turn-by-Turn GPS Tracking	Partially Implemented	Mobile map tracking and static simulated route navigation operational (60% milestone stage).	Background continuous GPS telemetry and hardware compass heading in progress.
Push / Audio Safety Warnings	Partially Implemented	Proximity detection logic structured; local foreground notification hooks implemented.	Background OS geofencing and hands-free text-to-speech audio alerts in progress.
Supervised ML Classification	Planned (Deferred)	Candidate A (XGBoost) deliberately evaluated and deferred in methodology.py.	Requires multi-corridor AADT exposure denominators and iRAD incident logs.
Empirical Bayes Rate Modeling	Planned (Deferred)	Candidate B evaluated and documented in methodology.py.	Awaiting integration of state highway traffic census and volume tallies.

6 Conclusion and future scope

6.1 Conclusion

This literature review has comprehensively examined the state of the art in road accident hotspot identification, highway crash prediction, environmental hazard modeling, and safety-aware navigation. By evaluating twenty-four genuine academic publications and national engineering guidelines, this study has exposed critical methodological vulnerabilities in current traffic safety literature:

- Supervised machine-learning classifiers applied to sparse crash points suffer from severe spatial autocorrelation and coordinate memorization, producing uncalibrated, illusory risk scores when denied traffic exposure (AADT) denominators.
- Mainstream digital navigation platforms operate with a singular focus on travel time and congestion velocity, remaining dangerously oblivious to underlying roadway black spots and terrain risks.
- Conventional safety maps commit the fallacy of negative proof, falsely designating uninspected and data-sparse highway segments as "safe" and providing drivers with deceptive assurances.
- High-risk mountain passes ("ghats") subject to monsoon-triggered landslides and severe geometric hazards are completely neglected in digital route planning.
- Crowdsourced incident logs are vulnerable to severe spatial reporting biases and rapidly decay into stale noise without strict temporal decay filtering and administrative verification.

In response to these findings, the OptiRoad platform establishes a technically sound, evidence-based alternative. By deploying Candidate C (Hierarchical Spatial Hazard Engine) over a robust PostgreSQL/

PostGIS and FastAPI backend, OptiRoad correlates live route LineStrings against verified MoRTH black spots, surveyed mountain passes, NASA GLC landslide records, and real-time Open-Meteo weather observations across standardized 1-kilometer geodesic blocks. Most importantly, OptiRoad introduces complete epistemic honesty into digital navigation by strictly enforcing the UNASSESSED status for unmonitored road segments, establishing an ethical benchmark for public safety informatics.

6.2 Future scope

Building upon the verified software foundation established in OptiRoad, several promising research and engineering directions emerge:

- **Integration with the Integrated Road Accident Database (iRAD):** Collaborating with state transport and police departments to establish secure API pipelines with India's national iRAD portal, expanding historical accident micro-data across nationwide highway corridors.
- **Highway Traffic Exposure (AADT) Ingestion:** Integrating National Highways Authority of India (NHAI) electronic toll collection (FASTag) volume data to establish empirical vehicle exposure baselines, enabling the defensible deployment of Candidate B (Empirical Bayes rate estimation).
- **Corridor-Level Spatial Block Cross-Validation:** Upon acquiring multi-corridor datasets, training supervised tree ensembles (XGBoost / LightGBM) utilizing rigorous spatial block cross-validation to ensure true generalizability across unseen geographic regions.
- **Autonomous Pavement Defect Telemetry:** Expanding the mobile client to incorporate smartphone inertial sensors (accelerometers and gyroscopes) to automate road roughness and pothole detection during vehicular transit.
- **Acoustic and Heads-Up Safety Alerts:** Finalizing background OS geofencing and text-to-speech voice warnings to deliver hands-free audio safety prompts as motorists approach audited black spots or entering severe ghat descents.

References

[1] World Health Organization. *Global status report on road safety 2023*. Geneva: World Health Organization; 2023. https://www.who.int/publications/i/item/9789240086517

[2] Ministry of Road Transport and Highways, Transport Research Wing. *Road accidents in India 2022*. New Delhi: Government of India; 2023. https://morth.nic.in/road-accidents-in-india

[3] Ministry of Road Transport and Highways. *Protocol for identification and rectification of road accident black spots on National Highways*. Circular RW/NH-15017/109/2006-P&M. New Delhi: Government of India; 2018.

[4] Indian Roads Congress. *Guidelines for the assessment and safety improvement of blackspots*. IRC:131-2022. New Delhi: Indian Roads Congress; 2022.

[5] Montella A. A comparative analysis of hotspot identification methods. *Accid Anal Prev*. 2010;42(2):571–581. https://doi.org/10.1016/j.aap.2009.09.025

[6] Luxen D, Vetter C. Real-time routing with OpenStreetMap data. In: *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (GIS '11)*; 2011 Nov; Chicago, IL, USA. p. 513–516. https://doi.org/10.1145/2093973.2094062

[7] Elvik R. A survey of operational definitions of hazardous road locations in some European countries. *Accid Anal Prev*. 2008;40(6):1830–1835. https://doi.org/10.1016/j.aap.2008.08.001

[8] Lord D, Mannering F. The statistical analysis of crash-frequency data: A review and assessment of methodological alternatives. *Transp Res Part A Policy Pract*. 2010;44(5):291–305. https://doi.org/10.1016/j.tra.2010.02.001

[9] Mannering FL, Bhat CR. Analytic methods in accident research: Methodological frontier and future directions. *Anal Methods Accid Res*. 2014;1:1–22. https://doi.org/10.1016/j.amar.2013.09.001

[10] Anderson TK. Kernel density estimation and K-means clustering to profile road accident hotspots. *Accid Anal Prev*. 2009;41(3):359–364. https://doi.org/10.1016/j.aap.2008.12.014

[11] Xie Z, Yan J. Kernel density estimation of traffic accidents in a network space. *Comput Environ Urban Syst*. 2008;32(5):396–406. https://doi.org/10.1016/j.compenvurbsys.2008.05.001

[12] Bíl M, Andrášik R, Janoška Z. Identification of hazardous road locations of traffic accidents by means of kernel density estimation and cluster significance evaluation. *Accid Anal Prev*. 2013;55:265–273. https://doi.org/10.1016/j.aap.2013.03.003

[13] Theofilatos A, Yannis G. A review of the effect of traffic and weather characteristics on road safety. *Accid Anal Prev*. 2014;72:244–256. https://doi.org/10.1016/j.aap.2014.06.017

[14] Theofilatos A. Incorporating real-time traffic and weather data to explore road accident likelihood and severity in urban arterials. *J Safety Res*. 2017;61:9–21. https://doi.org/10.1016/j.jsr.2017.02.003

[15] Parsa AB, Movahedi A, Taghipour H, Derrible S, Mohammadian AK. Toward safer highways, application of XGBoost and SHAP for real-time accident detection and feature analysis. *Accid Anal Prev*. 2020;136:105405. https://doi.org/10.1016/j.aap.2019.105405

[16] Bao J, Liu P, Ukkusuri SV. A spatiotemporal deep learning approach for citywide short-term crash risk prediction with multi-source data. *Accid Anal Prev*. 2019;122:239–254. https://doi.org/10.1016/j.aap.2018.10.015

[17] Galbrun E, Pelechrinis K, Terzi E. Urban navigation beyond shortest route: The case of safe paths. *Inf Syst*. 2016;57:160–171. https://doi.org/10.1016/j.is.2015.10.005

[18] Jiang S, Jafari M, Kharbeche M, Jalayer M, Al-Khalifa KN. Safe route mapping of roadways using multiple sourced data. *IEEE Trans Intell Transp Syst*. 2022;23(4):3169–3179. https://doi.org/10.1109/TITS.2020.3032643

[19] Goodchild MF. Citizens as sensors: The world of volunteered geography. *GeoJournal*. 2007;69(4):211–221. https://doi.org/10.1007/s10708-007-9111-y

[20] Amin-Naseri M, Chakraborty P, Sharma A, Gilbert SB, Hong M. Evaluating the reliability, coverage, and added value of crowdsourced traffic incident reports from Waze. *Transp Res Rec*. 2018;2672(43):34–43. https://doi.org/10.1177/0361198118790619

[21] Eriksson J, Girod L, Hull B, Newton R, Madden S, Balakrishnan H. The pothole patrol: Using a mobile sensor network for road surface monitoring. In: *Proceedings of the 6th International Conference on Mobile Systems, Applications, and Services (MobiSys '08)*; 2008 Jun; Breckenridge, CO, USA. p. 29–39. https://doi.org/10.1145/1378600.1378605

[22] Roberts DR, Bahn V, Ciuti S, Boyce MS, Elith J, Guillera-Arroita G, et al. Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. *Ecography*. 2017;40(8):913–929. https://doi.org/10.1111/ecog.02881

[23] Kirschbaum D, Stanley T, Zhou Y. Spatial and temporal analysis of a global landslide catalog. *Geomorphology*. 2015;249:4–15. https://doi.org/10.1016/j.geomorph.2015.03.016

[24] Bhagya SB, Sumi AS, Balaji S, Danumah JH, Costache R, Rajaneesh A, et al. Landslide susceptibility assessment of a part of the Western Ghats (India) employing the AHP and F-AHP models and comparison with existing susceptibility maps. *Land*. 2023;12(2):468. https://doi.org/10.3390/land12020468