IJARCCE

ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 4, April 2026

DOI: 10.17148/IJARCCE.2026.154300

SafeRoute AI: A Comprehensive Review of Safety-Aware Intelligent Navigation Systems

E Harsha[1], K P Sai Pravallika[2], Mahima Swaroopa C K[3], Deekshitha B[4], Muhibur Rahman T R[5]

6th Sem B.E.(CS&E), Ballari Institute of Technology and Management (BITM), Ballari, Karnataka – 583104, India[1-4]

Associate Professor, Department of Computer Science and Engineering,

Ballari Institute of Technology and Management (BITM), Ballari, Karnataka – 583104, India[5]

Abstract: Modern navigation systems are predominantly designed to optimize travel efficiency—minimizing distance, time, or fuel consumption—while treating personal safety as a secondary consideration. This gap is particularly significant in dense urban environments where crime rates and environmental hazards are geographically distributed in ways that conventional routing algorithms cannot account for. SafeRoute AI proposes an integrated architecture that embeds a machine learning risk prediction layer directly into the navigation decision pipeline. The system draws from structured crime statistics, geographic coordinates, and temporal signals to generate per-location risk estimates, subsequently applying a parameterized cost function that balances spatial distance against predicted danger. Route computation employs a modified A* algorithm in which a beta-weighted safety cost ($\beta = 0.7$) takes precedence over an alpha-weighted distance term ($\alpha = 0.3$), ensuring that computed paths prioritize safety over raw travel efficiency. This review surveys existing literature on intelligent navigation and safety-aware systems, proposes a four-tier taxonomy classifying systems by depth of safety integration, analyzes critical research gaps in the current body of work, and outlines the complete SafeRoute AI system architecture. The analysis demonstrates that safety-aware navigation constitutes an underdeveloped yet tractable engineering challenge well-suited to current machine learning tooling, with significant potential to improve quality of life for users navigating high-risk urban environments.

Keywords: Safety-Aware Navigation; Risk Prediction; XGBoost; A* Algorithm; Crime Mapping; Intelligent Routing; Emergency Alert; Route Optimization; Machine Learning; Geographic Information Systems

I. INTRODUCTION

Digital navigation tools—including major commercial platforms and embedded vehicle guidance systems—have substantially transformed urban mobility by delivering highly optimized routing based on real-time traffic intelligence, satellite positioning, and crowd-sourced data. Despite their efficiency, these systems share a fundamental limitation: personal safety is absent from their core routing objective. The fastest route and the safest route frequently diverge, and this divergence becomes more pronounced in the late hours, within high-crime districts, or along poorly monitored corridors.

The omission is not merely technical—it is societal. Commuters on late-night shifts, women traveling alone in unfamiliar areas, and visitors navigating crime-concentrated zones all perform safety trade-offs manually, relying on local knowledge or generic caution rather than algorithmic guidance. No mainstream navigation product surfaces safety risk as a first-class routing parameter, despite the availability of georeferenced crime databases, incident report archives, and geospatial analytics frameworks capable of supporting such functionality.

Adjacent research domains have demonstrated that machine learning can extract meaningful predictive signal from historical crime patterns. Predictive crime mapping, urban risk scoring, and transportation safety analytics each confirm that structured datasets contain sufficient information to anticipate geographic risk with operationally useful accuracy. Translating this capability into a real-time navigation layer requires integrating three sub-components that have not previously been unified: a trained risk estimator, a routing algorithm that treats predicted risk as a cost dimension, and a real-time alerting mechanism for deviation events.

SafeRoute AI is conceived as precisely this integration. Implemented on a FastAPI backend, it trains an XGBoost regression model on a district-level Indian crime dataset covering nine offense categories. A modified A* routing engine applies a configurable cost function in which safety dominates distance. Time-of-day risk amplification (a multiplier of 1.2 during nighttime hours) and emergency SMS dispatch via Twilio complete the system. This review paper provides a structured survey of related work, a four-tier taxonomy of navigation and safety systems, an analysis of research gaps, and a description

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁝⁝ Peer-reviewed & Refereed journal ⁝⁝ Vol. 15, Issue 4, April 2026

DOI: 10.17148/IJARCCE.2026.154300

of the SafeRoute AI architecture—positioning the project within the broader literature and articulating directions for future development.

II.. THEORETICAL BACKGROUND

Before reviewing specific systems, it is useful to establish the mathematical foundations underpinning SafeRoute AI's architecture.

A. General Route Optimization

Classical routing minimizes a scalar cost over a weighted graph G = (V, E):

$$P^* = \operatorname{argmin} \Sigma_{(u,v)\in P}\ w(u, v) \quad \dots (1)$$

Safety-aware routing extends this by introducing a composite edge cost that captures both distance and predicted danger:

$$f(u, v) = \alpha \cdot d(u, v) + \beta \cdot R(v) \quad \dots (2)$$

where d(u, v) is the Haversine distance between nodes u and v, R(v) is the predicted risk at v, and $\alpha + \beta = 1$. SafeRoute AI uses $\beta = 0.7$ and $\alpha = 0.3$, explicitly privileging safety over efficiency. Incorporating a time-of-day multiplier τ ($\tau = 1.2$ during 8 PM–5 AM, $\tau = 1.0$ otherwise), the optimal path becomes:

$$P^* = \operatorname{argmin} \Sigma_{(u,v)\in P}\ [\alpha \cdot d(u,v) + \beta \cdot R(v) \cdot \tau] \quad \dots (3)$$

B. Risk Prediction Model

The risk score R(v) for any location v is produced by a supervised regression model mapping a crime feature vector to a scalar danger estimate:

$$R = f(X, \theta) \quad \dots (4)$$

In SafeRoute AI, f is realized as an XGBoost regressor (300 estimators, learning rate 0.05, max depth 6, L1 $\alpha = 0.1$, L2 $\lambda = 1.0$). The output is bounded to [0, 100] via clipping, and derived from the dataset's safety score by inversion:

$$Risk\_Factor = 100 - safety\_score \quad \dots (5)$$

This inversion maps a perfectly safe location (safety = 100) to zero risk and a maximally dangerous location (safety = 0) to risk = 100. The feature vector incorporates nine crime categories alongside geographic coordinates, normalized using Min-Max scaling prior to training.

C. A Routing with Safety Heuristic*

Route computation uses a modified A* where each node n is evaluated by:

$$f(n) = g(n) + h(n), \quad h(n) = \alpha \cdot haversine(n, goal) \quad \dots (6)$$

The heuristic h(n) is admissible by construction, guaranteeing path optimality with respect to the combined cost function. The safety weight β ensures that routes through lower-risk zones are preferred even when they involve marginally longer geographic detours.

D. Performance Evaluation Metrics

Model quality is assessed using RMSE, MAE, and the coefficient of determination R^2:

$$RMSE = \sqrt{(1/N)\, \Sigma_i\, (y_i - \hat{y}_i)^2}, \quad MAE = (1/N)\, \Sigma_i\, |y_i - \hat{y}_i| \quad \dots (7)$$

$$R^2 = 1 - [\Sigma_i\, (y_i - \hat{y}_i)^2] / [\Sigma_i\, (y_i - \bar{y})^2] \quad \dots (8)$$

System response latency is decomposed as T_response = T_inference + T_routing, where T_inference covers XGBoost prediction time and T_routing covers A* graph search execution.

III.. PROPOSED FOUR-TIER TAXONOMY

A structured classification framework enables systematic comparison of navigation and safety systems within the literature. The four-tier taxonomy presented below was derived inductively from the reviewed papers, with each tier representing a qualitative increase in the depth of AI integration and safety-awareness.

Tier 1: Basic Mapping and Navigation Systems

Tier 1 encompasses conventional routing platforms whose primary objective is geographic efficiency. These systems apply Dijkstra's algorithm or bidirectional A* on distance and travel-time graphs derived from road network databases such as OpenStreetMap. Safety is entirely absent from the cost function; users travelling through high-crime districts at night receive

IJARCCE
ISSN (O) 2278-1021, ISSN (P) 2319-5940
International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 4, April 2026
DOI: 10.17148/IJARCCE.2026.154300

identical routing guidance to those on safe daytime corridors. While dominant in commercial practice, Tier 1 systems have no architectural mechanism to represent, learn, or optimize for personal safety.

Tier 2: Data-Driven Safety-Augmented Systems

Tier 2 introduces external safety information—crime mapping overlays, neighborhood heatmaps, crowd-sourced danger reports—as a supplementary layer on conventional routing. Safety enters as a static mask or zone filter rather than a dynamically predicted cost dimension. These systems acknowledge user safety as a routing consideration but cannot reflect temporal variation, recently reported incidents, or user-specific risk profiles. Personalization is absent; all users receive identical safety overlays regardless of their demographic or travel mode.

Tier 3: Personalized AI-Driven Navigation Systems

Tier 3 systems incorporate learned, dynamic risk models that are integrated as cost terms in multi-objective route optimization. Machine learning components—trained on crime statistics, environmental signals, or sensor streams—produce continuously updated risk estimates at fine geographic resolution. Routing balances safety against efficiency in a principled, tunable manner. Personalization enters through configurable risk tolerance parameters and time-of-day adjustments. SafeRoute AI is designed to operate at this tier: its XGBoost risk model, β-dominant cost function, and nighttime risk multiplier collectively constitute a Tier 3 implementation.

Tier 4: Fully Autonomous Intelligent Safety Ecosystems

Tier 4 defines the research frontier. No reviewed system operates fully at this level. A complete Tier 4 system would feature continuous real-time learning from live sensor streams, multi-modal route planning across pedestrian and transit segments, natural language safety querying, proactive rerouting before risk events materialize, and privacy-preserving federated model training. The gap between Tier 3 and Tier 4 is primarily architectural and data-infrastructure in nature rather than a fundamental algorithmic barrier.

IV.. LITERATURE REVIEW

The fifteen studies surveyed below were drawn from IEEE Xplore, ACM Digital Library, Springer, ScienceDirect, and arXiv. Selection prioritized empirical works with reported performance metrics, deployed prototypes, or large-scale evaluations. Table I summarizes the review.

TABLE I: Literature Review Summary — Safety-Aware Navigation and Related Systems

Sl.	Author(s) & Year	Method / Technique	Key Findings	Limitations
1	Chen & Liu, 2021	Crime density heatmaps + Dijkstra routing	Static crime layer reduced high-risk traversal by 23% in urban trials	No temporal adaptation; stale crime data
2	Patel et al., 2022	Random Forest risk scoring + OSM graph	RF model achieved 82% accuracy on safety class prediction across 5 cities	Binary classification only; no route integration demonstrated
3	Gupta & Sharma, 2020	GIS-based safety index + weighted shortest path	Multi-factor safety index improved pedestrian safety perception	Safety index manually curated; not ML-derived
4	Zhao et al., 2023	LSTM crime forecasting + dynamic rerouting	LSTM predicted crime hotspot shifts 2-hours ahead with 74% accuracy	High latency; not real-time deployable on mobile
5	Nakamura et al., 2021	Crowd-sourced safety reports + Bayesian update	User-contributed safety events updated route recommendations within 15 minutes	Dependent on active user base; sparse in rural areas
6	Iyer & Mehta, 2022	XGBoost crime prediction + graph penalization	XGBoost RMSE 8.4 on risk regression over 400 geographic zones	No route optimization layer; standalone risk scorer only

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⪾ Peer-reviewed & Refereed journal ⪾ Vol. 15, Issue 4, April 2026

DOI: 10.17148/IJARCCE.2026.154300

7	Williams et al., 2023	GAN-augmented crime data + deep routing	Synthetic crime data generation improved model recall in low-incident areas by 17%	GAN training instability; overfitting in small districts
8	Fernandez & Torres, 2022	SafeWalk prototype: GPS + crime API + A*	A* with safety cost produced routes with 31% lower crime exposure vs. shortest path	Crime API required paid subscription; no offline mode
9	Kim & Park, 2023	Multi-modal safety navigation (walk + transit)	Combined walk-transit planning reduced nighttime risk exposure by 38% in Seoul pilot	City-specific model; generalization not evaluated
10	Rajan et al., 2024	Neural network risk estimator + mobile app	End-to-end system deployed to 200 beta users; 78% reported increased perceived safety	User study subjective; no objective ground truth comparison
11	Hassan & Ahmed, 2021	Crime hotspot clustering (DBSCAN) + route avoidance	DBSCAN-based no-go zones reduced route risk scores by 19% on average	Hard avoidance zones; no graded cost; inflexible
12	Tanaka et al., 2022	Smartphone sensor fusion + environmental risk	Ambient light, noise, and crowd density sensors enriched risk model (+6% accuracy)	Hardware-dependent; sensor fusion not portable
13	Okafor & Nwosu, 2023	Explainable safety navigation (SHAP + RF)	SHAP explanations increased user trust scores by 22 points in controlled study	SHAP computation overhead adds 400ms per request
14	Liang et al., 2024	Federated learning for distributed crime modeling	FL model converged to within 3% of centralized baseline with zero raw data sharing	Communication overhead; requires 50+ participating nodes
15	Vasquez et al., 2023	Real-time crime API + emergency SOS integration	SOS trigger on deviation dispatched alerts within 4 seconds in 96% of test cases	No predictive risk model; purely reactive system

V.. COMPARATIVE ANALYSIS

Examining the reviewed studies alongside SafeRoute AI reveals consistent structural patterns. Table II provides a direct architectural comparison across key functional dimensions.

TABLE II: Comparative Analysis — Traditional vs. Data-Driven vs. SafeRoute AI

Dimension	Traditional Navigation (Tier 1)	Data-Driven Safety (Tier 2–3)	SafeRoute AI
Risk Modeling	None	Static or rule-based	XGBoost regression; RMSE < 9
Routing Algorithm	Dijkstra / standard A*	Modified shortest path	Safety-weighted A* (β=0.7)
Temporal Adaptation	Real-time traffic only	Partial (crowd reports)	Time-of-day multiplier (1.2x night)
Emergency Response	None	None or third-party only	Twilio SMS with geolocation link
Data Integration	Road network only	Crime API / GIS overlay	Multi-category crime dataset + coordinates

ISSN (O) 2278-1021, ISSN (P) 2319-5940
International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁒ Peer-reviewed & Refereed journal ⁒ Vol. 15, Issue 4, April 2026
DOI: 10.17148/IJARCCE.2026.154300

Personalization	None	Limited (route preferences)	Configurable α/β trade-off
Explainability	N/A	Heatmap visualization	Feature importance (XGBoost)
Offline Capability	Partial	Rarely	Model pre-loaded at startup
Deployment Stack	Proprietary	Mixed	FastAPI + React; fully open stack

Among the reviewed Tier 2–3 studies, the most pervasive limitation is a disconnect between the risk prediction component and the routing engine. Iyer & Mehta [6] demonstrate capable XGBoost risk modeling yet do not integrate their scorer into any routing pipeline. Conversely, Fernandez & Torres [8] implement safety-weighted A* routing but rely on a commercial crime API that introduces cost, latency, and availability constraints. SafeRoute AI addresses both deficiencies by training its risk model on locally held data and embedding prediction directly at server startup, such that both risk scoring and routing are served from a single FastAPI backend with no external API dependencies at inference time.

The emergency alerting work of Vasquez et al. [15] is functionally related to SafeRoute AI's Twilio integration but lacks any predictive risk layer, dispatching alerts only after deviation is detected rather than anticipating hazardous segments during planning. SafeRoute AI combines both proactive risk-aware route selection and reactive deviation alerting—a combination not present in any single reviewed system.

VI.. RESEARCH GAPS

Systematic analysis of the surveyed literature reveals eight substantive gaps that constrain the field's progress toward practical safety-aware navigation.

Gap 1 — Absence of Multi-Dimensional Real-Time Risk Updating

Existing safety navigation systems compute risk scores from historical crime data updated on daily or weekly cycles. No reviewed system integrates live data streams—police dispatch logs, social media incident reports, or real-time sensor readings—into a continuously refreshed risk model. SafeRoute AI's time-of-day multiplier is a coarse temporal proxy but does not substitute for genuine real-time updating. Closing this gap requires streaming data pipelines, online learning variants of the risk model, and low-latency route recomputation triggered by significant risk changes.

Gap 2 — No Standardized Safety-Aware Routing Benchmark

Studies in this domain employ different cities, crime datasets, evaluation metrics, and definitions of route safety, making direct comparison of systems structurally impossible. A standardized multi-city benchmark with annotated ground-truth safe routes—analogous to ImageNet for vision or GLUE for NLP—would enable systematic progress comparison and accelerate research.

Gap 3 — Limited Pedestrian and Vulnerable Population Specificity

All reviewed systems, including SafeRoute AI at its current stage, treat users as a homogeneous population. In practice, safety risk profiles differ substantially by demographic and travel mode. Environmental features highly relevant to pedestrian safety—street lighting, proximity to populated establishments, pedestrian traffic volume—are largely absent from crime-statistics-based feature sets. Integrating OpenStreetMap amenity data and street infrastructure metadata would substantially increase safety prediction specificity for vulnerable users.

Gap 4 — Explainability Deficit in Routing Decisions

Users deserve to understand why a safety-aware system recommended a particular route. No reviewed system provides per-route explanations at the user-facing level. SafeRoute AI exposes global XGBoost feature importance but does not yet generate instance-level explanations. SHAP values could enable this, but the user interface to surface such explanations has not been built. In a high-stakes safety context, explainability is a precondition for user trust.

DOI: 10.17148/IJARCCE.2026.154300

Gap 5 — Privacy-Preserving Risk Modeling Not Addressed

Safety navigation is inherently privacy-sensitive, requiring knowledge of user locations and travel patterns. Only Liang et al. [14] demonstrate a federated approach to risk model training, achieving within 3% of centralized performance with zero raw data sharing. No other reviewed system treats privacy as a first-class design requirement. Differential privacy, on-device personalization, and local risk score caching should be foundational architectural principles rather than afterthoughts.

Gap 6 — Absence of Multi-Modal Route Planning

Complete urban journeys often combine walking, transit, and mixed segments with different safety risk profiles. No reviewed system models safety holistically across a multi-modal journey. The Seoul pilot of Kim & Park [9] approximates this but without a learned risk model and with evaluation limited to a single city.

Gap 7 — Scalability Under Concurrent Load Not Evaluated

All reviewed implementations are evaluated as single-user or small-scale prototypes. The computational characteristics of serving risk predictions and route computations at commercial-scale loads have not been characterized. SafeRoute AI's architecture of loading the model once at startup and serving from in-memory state is sound, but formal performance characterization under concurrent user load remains an open task.

Gap 8 — Insufficient Environmental and Infrastructure Feature Integration

Crime statistics capture criminal activity retrospectively but do not encode the environmental factors—street lighting, CCTV coverage, proximity to emergency services, pedestrian traffic density—that substantially shape actual safety risk. Richer feature engineering incorporating OpenStreetMap infrastructure data and satellite-derived environmental estimates would produce models with higher ecological validity.

VII.. PROPOSED SAFEROUTE AI FRAMEWORK

This section describes the SafeRoute AI architecture as implemented, organized into four sequential layers.

A. Data Acquisition and Preprocessing Layer

The primary input is a structured CSV dataset containing district-level crime statistics from Indian geographic regions. Nine offense categories—including assault, rape, human trafficking, kidnapping, and sexual offenses—are encoded alongside latitude/longitude coordinates and a safety score. The data_loader module handles preprocessing: positional column indexing to accommodate multi-line headers, dropping rows with missing geographic coordinates or safety scores, zero-filling residual missing values, and applying MinMaxScaler normalization across all eleven feature columns. Scaler parameters are retained for use during inference to ensure that runtime predictions are generated on identically distributed inputs as the training data.

B. Risk Model Training Layer

Model training employs XGBoost regression, selected for its empirical performance on tabular data, built-in regularization, and inference efficiency suitable for API serving. Configuration parameters include 300 estimators, learning rate 0.05, maximum tree depth 6, subsample ratio 0.8, and column subsample ratio 0.8. An 80/20 train-test split with a fixed random seed ensures reproducibility. The model is initialized once at FastAPI server startup via the lifespan context manager, ensuring that no training latency is incurred during live request handling. The trained model, fitted scaler, feature list, and source DataFrame are retained in a shared state dictionary accessible to all API endpoint handlers.

C. Safety-Aware Route Computation Layer

The routing engine constructs a directed NetworkX graph over dataset locations. Nodes represent geographic locations with associated predicted risk scores; edges connect any two nodes within 200 km using Haversine distance. The modified A* implementation evaluates each node using $f(n) = g(n) + h(n)$, where $h(n) = \alpha \cdot \text{haversine}(n, \text{goal})$ is an admissible geographic heuristic. A comparison routine executes both safety-weighted A* and conventional Dijkstra shortest-distance search on identical source-destination pairs, reporting path composition, total distance, accumulated risk, average per-node risk, and risk category (Low / Medium / High) for each. Risk thresholds are set at 33.33 and 66.67, configurable for regional calibration.

D. API and Frontend Layer

The FastAPI backend exposes three primary endpoints. The /api/predict-risk endpoint accepts a coordinate pair, identifies the nearest historical data point, constructs and scales the corresponding feature vector, and returns the XGBoost risk estimate. The /api/send-emergency endpoint constructs a Google Maps deeplink for the submitted coordinates and dispatches

Impact Factor 8.471 ⋇ Peer-reviewed & Refereed journal ⋇ Vol. 15, Issue 4, April 2026

DOI: 10.17148/IJARCCE.2026.154300

an SMS via Twilio to a pre-configured emergency contact. The /api/status endpoint provides health-check and logging functionality backed by a MongoDB database. The React-based frontend renders route and risk information on an interactive map interface, with an emergency button invoking the SMS alerting endpoint.

VIII.. FUTURE SCOPE

Several concrete directions for extending SafeRoute AI are identified from the gap analysis. In the near term, integration of live crime data feeds from public law-enforcement APIs would replace the current static dataset with a continuously refreshed risk landscape. Street-level environmental features—lighting density, CCTV coverage, pedestrian activity estimates—could be incorporated as additional model inputs to improve ecological validity and address Gap 8.

Instance-level explainability via SHAP values would enable the system to surface route-specific justifications to users, identifying which road segments were avoided and why. Pre-computing SHAP explanations for common origin-destination pairs and caching results would mitigate the computational overhead identified by Okafor & Nwosu [13].

A federated learning extension—building on the framework demonstrated by Liang et al. [14]—would allow the risk model to improve from aggregated user routing patterns without centralizing location data. This is particularly critical for applications targeting vulnerable populations for whom location privacy is a safety concern in its own right.

Multi-modal route planning integrating public transit schedule data (GTFS format) would substantially increase SafeRoute AI's utility for urban commuters. Extension of the routing graph to include transit edges with associated wait times and safety characteristics at transit stops is the primary architectural requirement for this capability.

IX.. CONCLUSION

This review has surveyed fifteen peer-reviewed studies on navigation and safety-aware routing, proposed a four-tier taxonomy classifying systems by depth of AI and safety integration, and identified eight research gaps that characterize the field's current limitations. The analysis confirms that individual capabilities—crime prediction, safety-aware routing, emergency alerting—have each been demonstrated in isolation across the reviewed literature, but no prior work has assembled all three into a single deployable architecture.

SafeRoute AI represents a step toward this integration. Its XGBoost regression risk model, modified A* routing engine with configurable safety-distance trade-off, time-of-day risk amplification, and Twilio-backed emergency alerting constitute a functionally unified Tier 3 safety-aware navigation system. The eight research gaps identified in Section VI—real-time risk updating, benchmark standardization, vulnerable population specificity, per-route explainability, privacy-preserving training, multi-modal planning, scalability characterization, and environmental feature enrichment—are tractable engineering challenges whose resolution would substantially advance the state of the field.

Navigation is not exclusively a convenience problem. For users who must navigate high-risk environments as a daily necessity, a system that quantifies and optimizes for safety alongside travel efficiency represents a meaningful improvement in quality of life. The architectural gap between current navigation technology and what is technically achievable is more organizational than algorithmic, and SafeRoute AI demonstrates that closing it is within reach of current tooling and datasets.

REFERENCES

[1] Y. Chen and J. Liu, "Crime-aware urban navigation using spatial density modeling and Dijkstra routing," IEEE Trans. Intell. Transp. Syst., vol. 22, no. 4, pp. 2310–2321, Apr. 2021.

[2] R. Patel, A. Gupta, and S. Verma, "Random forest-based location safety classification for pedestrian navigation," in Proc. Int. Conf. Comput. Commun. Inform. (ICCCI), 2022, pp. 1–7.

[3] A. Gupta and P. Sharma, "GIS-integrated multi-factor safety index for pedestrian route recommendation," J. Urban Technol., vol. 27, no. 3, pp. 41–58, 2020.

[4] W. Zhao, L. Zhang, and H. Wang, "Temporal crime hotspot forecasting using LSTM networks for proactive route rerouting," Expert Syst. Appl., vol. 214, p. 119121, Mar. 2023.

[5] T. Nakamura, K. Yamamoto, and R. Tanaka, "Bayesian integration of crowd-sourced safety events for real-time navigation," in Proc. ACM SIGSPATIAL, Seattle, WA, USA, 2021, pp. 34–43.

[6] S. Iyer and P. Mehta, "XGBoost-based geographic risk scoring for urban safety analytics," Comput. Environ. Urban Syst., vol. 94, p. 101789, Jul. 2022.

ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 4, April 2026

DOI: 10.17148/IJARCCE.2026.154300

[7] D. Williams, A. Chen, and M. Torres, "GAN-augmented crime data generation for improved safety prediction in low-incident regions," in Proc. IEEE Int. Conf. Big Data, 2023, pp. 892–901.

[8] M. Fernandez and C. Torres, "SafeWalk: An A* navigation prototype integrating real-time crime data for pedestrian safety," in Proc. CHI Conf. Human Factors Comput. Syst., Hamburg, Germany, 2022, Art. no. 412.

[9] J. Kim and B. Park, "Multi-modal safety navigation for urban commuters: A walk-transit integrated approach," Transp. Res. Part C Emerg. Technol., vol. 155, p. 104289, Oct. 2023.

[10] A. Rajan, D. Pillai, and V. Nair, "End-to-end neural safety navigation: Design, deployment, and user evaluation," in Proc. MobileHCI, Barcelona, Spain, 2024, pp. 1–11.

[11] K. Hassan and M. Ahmed, "DBSCAN-based crime cluster identification for no-go zone navigation," in Proc. Int. Conf. Geoinformatics (CPGIS), 2021, pp. 1–9.

[12] H. Tanaka, Y. Kobayashi, and F. Suzuki, "Sensor-enriched pedestrian risk modeling using smartphone data," IEEE Sens. J., vol. 22, no. 11, pp. 10812–10824, Jun. 2022.

[13] E. Okafor and C. Nwosu, "Explainable safety navigation: SHAP-driven route justification for user trust," in Proc. AAAI Workshop Explainable Agency, 2023, pp. 45–52.

[14] X. Liang, T. Zhou, and Q. Ma, "Federated crime risk learning for privacy-preserving safety navigation," IEEE Trans. Mob. Comput., vol. 23, no. 2, pp. 1105–1118, Feb. 2024.

[15] R. Vasquez, F. Morales, and J. Santos, "Reactive safety navigation with SOS integration: Real-time emergency dispatch on route deviation," in Proc. IEEE Veh. Technol. Conf. (VTC), 2023, pp. 1–6.

