Cooperative Intelligence for Heavy Commercial Vehicles: An Edge–Cloud V2X Architecture for Predictive Risk Mitigation

Anand Kumar Vedantham
UST Global
Chicago, IL, USA

Kiran Kumar Sunkara
KPIT Technologies Ltd
Chicago, IL, USA

Nanna Ram Parza
LTIMindtree LLC, US
Chicago, IL, USA

***Abstract*—Heavy commercial vehicles (HCVs) present disproportionate safety risks due to their large mass, extended braking distances, limited maneuverability, and long operating hours. Existing safety mechanisms, including Advanced Driver Assistance Systems (ADAS) and fleet telematics, remain largely reactive or retrospective, limiting their ability to prevent hazardous events before escalation.**

This paper proposes a cooperative intelligence architecture that integrates onboard edge computing, Vehicle-to-Everything (V2X) communication, and cloud-based predictive analytics to enable proactive, fleet-wide risk mitigation. The architecture combines local sensor fusion and real-time risk scoring at the vehicle edge, cooperative hazard exchange through V2V/V2I/V2N channels, and cloud-based predictive analytics for fleet-level forecasting and policy refinement. Using simulated fleet scenarios, the proposed approach is compared against local-only ADAS and cloud-only telematics baselines. The evaluation focuses on comparative architectural response characteristics, including hazard detection time, mitigation latency, and near-collision events. Results show that the cooperative architecture achieves earlier hazard detection, reduced reaction latency, and fewer near-collision events than the baseline approaches. These findings indicate that heavy vehicles can evolve from isolated reactive entities into predictive, network-aware safety agents, while also motivating future real-world validation.

***Index Terms*—V2X, connected vehicles, edge computing, predictive safety, heavy commercial vehicles, cooperative intelligence, cloud analytics**

I. Introduction

Heavy commercial vehicles (HCVs) are essential for freight movement and public transportation, yet their physical characteristics and operating conditions make them disproportionately represented in severe crashes [1]. Their large mass, extended braking distances, limited maneuverability, wide blind spots, and long driving shifts amplify the consequences of driver error, adverse weather, road geometry, and traffic disturbances.

Advanced Driver Assistance Systems (ADAS), including forward collision warning, automatic emergency braking, lane departure warning, and stability control, have improved vehicle safety. However, these functions rely predominantly on local sensor data and are typically designed to respond only after a hazard enters the vehicle's immediate perception horizon. In parallel, fleet telematics platforms collect rich information on driver behavior, vehicle health, and route conditions, but they are used primarily for reporting, compliance, and retrospective analysis rather than time-critical safety intervention.

Many critical HCV hazards evolve over broader spatial and temporal scales than a single vehicle can directly observe. Examples include shockwave braking that propagates through traffic, localized low-friction road segments, and recurrent high-risk curves whose severity depends on vehicle load, speed, and environmental conditions. These hazards are inherently network-level phenomena and therefore benefit from cooperation among vehicles, roadside infrastructure, and cloud intelligence.

Vehicle-to-Everything (V2X) communication and vehicular edge computing provide the technical foundations for such cooperation [2], [3], [5]–[7]. When vehicles can exchange safety-relevant information and consume fleet-level risk forecasts, HCV safety can evolve from isolated reactive assistance toward cooperative and predictive risk mitigation.

Heavy commercial vehicle safety therefore requires more than either isolated onboard perception or delayed retrospective analytics. Unlike local-only ADAS systems, which are constrained by sensing range and line-of-sight limitations, and unlike cloud-only telematics systems, which are limited by communication and processing latency for immediate intervention, the framework proposed in this paper integrates low-latency onboard inference, cooperative V2X hazard exchange, and fleet-level cloud forecasting into a unified safety architecture. The primary contribution of this work lies in this cooperative architectural integration and in its comparative evaluation across alternative deployment strategies for predictive risk mitigation in HCV operations.

This paper makes the following contributions:

- **Integrated cooperative architecture for HCV safety:** We propose a three-layer edge–V2X–cloud architecture that combines onboard risk inference, cooperative hazard awareness, and fleet-level predictive intelligence for proactive safety support.
- **Multi-context risk modeling concept:** We define a risk-scoring approach that integrates vehicle state, environmental conditions, and cooperative context received through connected vehicle communication.
- **Comparative architectural evaluation:** We evaluate the proposed approach in simulated fleet scenarios against

local-only ADAS and cloud-only telematics baselines using hazard detection time, mitigation latency, and near-collision events as key metrics.
- **Deployment-oriented perspective:** We discuss industrial applicability, implementation considerations, and the relevance of cooperative intelligence for future software-defined vehicle ecosystems.

II. Background and Related Work

A. Heavy Vehicle Safety and Reactive Systems

HCV crashes tend to be more severe than passenger-vehicle crashes and are strongly influenced by road alignment, visibility, and driver violations [1]. While ADAS functions mitigate some of these risks, they share three core limitations:

- **Local visibility:** They cannot see beyond line-of-sight or sensor range.
- **No network context:** They are unaware of upstream traffic dynamics or remote incidents.
- **Short prediction horizon:** Decisions are based on immediate kinematics rather than multi-minute risk evolution.

Consequently, ADAS is necessary but not sufficient for proactive HCV safety.

B. V2X and Cooperative Perception

V2X communication covers vehicle-to-vehicle (V2V), vehicle-to-infrastructure (V2I), vehicle-to-network (V2N), and vehicle-to-pedestrian (V2P) exchanges. Standards have evolved from IEEE 802.11p to cellular V2X (C-V2X) and are progressing toward 5G/6G-enabled Internet of Vehicles architectures [2], [3], [10]. V2X supports messages such as basic safety messages, emergency brake lights, hazard warnings, and signal phase and timing.

Cooperative perception extends this by sharing sensor-level or feature-level data among vehicles and infrastructure to overcome occlusions and limited line-of-sight [4]. Most work, however, targets passenger vehicles and focuses on message dissemination rather than integrated edge–cloud prediction for heavy fleets.

C. Vehicular Edge Computing and ML for Safety

Vehicular edge computing (VEC) and fog computing shift computation closer to vehicles, reducing latency and offloading the cloud [5]–[7]. Architectures typically assign fast, local tasks to onboard units or roadside nodes, and slower global analytics to the cloud. In parallel, machine learning models, including LSTM and CNN variants, have been applied to real-time crash risk prediction using traffic and environmental data [8], [9].

These strands show that:

- Edge compute is suitable for low-latency risk inference.
- Cloud platforms can learn longer-horizon risk patterns.

However, there is limited work on a single architecture that marries HCV-oriented edge intelligence, V2X cooperation, and cloud-based predictive safety.

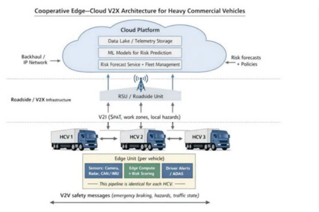

Fig. 1. Cooperative Edge–Cloud V2X Architecture. Vehicles perform local risk scoring, exchange hazards via V2X, and receive predictive insights and policy updates from the cloud.

III. Cooperative Edge–Cloud V2X Architecture

The proposed cooperative intelligence architecture comprises three layers:

1) Edge layer – onboard HCV intelligence,
2) V2X communication layer – cooperative exchange of safety information,
3) Cloud intelligence layer – fleet-level analytics and prediction.

A. Edge Layer: Onboard HCV Intelligence

Each HCV hosts an edge unit connected to:

- Environment sensors (cameras, radar, optionally lidar),
- GNSS and inertial sensors,
- CAN bus (speed, engine torque, RPM, brake pressure, steering, stability systems),
- Optional driver monitoring sensors.

The edge unit performs:

- **Sensor fusion:** Align heterogeneous streams into a local state.
- **Feature extraction:** Derive compact features such as headway, lateral acceleration, jerk, lane-keeping measures, braking patterns, and duty-cycle indicators.
- **Local risk scoring:** Evaluate a light-weight model at 5–10 Hz to obtain a local risk score $R_{\text{edge}} \in [0, 1]$.
- **Immediate mitigation:** If R_{edge} exceeds threshold τ_{edge}, issue alerts or adjust assistance systems within predefined safety envelopes.

Time-critical loops remain local, ensuring deterministic response even with intermittent connectivity.

B. V2X Communication Layer: Cooperative Context

The V2X layer extends awareness beyond the ego vehicle's sensor horizon. Representative message classes include:

- **Events:** Emergency braking, stability control activation, loss of traction, infrastructure-detected hazards, lane closures.
- **Context:** Local traffic density and speed, weather and friction indicators, temporary speed limits.

Logical channels:

- **V2V:** Low-latency hazard messages between vehicles.
- **V2I:** Vehicle–RSU interactions for intersection, corridor, or work-zone information.

- **V2N:** Telemetry and risk advisories via cellular links.

To maintain scalability, the system prioritizes safety-critical messages and filters data by spatial and temporal relevance. Standard security mechanisms (certificates, signatures, encryption) protect integrity and privacy.

C. Cloud Intelligence Layer: Fleet-Level Prediction

The cloud layer aggregates:

- Edge-generated features and events,
- V2X logs from vehicles and RSUs,
- External data (weather, traffic, maps, grade, historical incidents).

Key functions:

- **Model training:** Train ML models for crash and rollover risk, driver behavior profiles, and route risk mapping [2], [5], [8].
- **Risk forecasting:** Produce medium-horizon risk forecasts (e.g., 1–10 minutes / 1–10 km ahead) per vehicle and route segment, combining vehicle behavior, cooperative signals, and environment.
- **Policy generation:** Transform forecasts into concrete recommendations (speed envelopes, route changes, driver coaching).
- **Model distribution:** Deliver updated models or parameters to edge nodes via over-the-air updates.

The cloud operates on slower time scales than the edge but with broader context, making it suitable for predicting risk build-up rather than immediate threats.

IV. Risk Prediction Model

The architecture relies on a common risk scoring concept that can be instantiated differently at edge and cloud.

A. Feature Groups

We group features into:

- **Vehicle state** x_v: Speed, acceleration, jerk, brake pressure, lateral acceleration, steering variability, load proxies, stability system status, duty-cycle metrics.
- **Environment** x_e: Road type, curvature, grade, lane count, historical incident density, weather and visibility, road surface condition indicators.
- **Cooperative context** x_c: Recent emergency braking and stability events nearby, speed variance and headways in vicinity, infrastructure-reported anomalies, corridor-level congestion.

Edge models use a compressed subset of these features over short time windows; cloud models can exploit richer historical context.

B. Model Families and Deployment

We adopt a heterogeneous stack:

- **Edge:** Gradient boosting or compact neural networks that can run on embedded hardware with low latency and limited memory.
- **Cloud:** Sequence models such as LSTM or temporal convolutional networks for time-series risk prediction [8], [9], potentially enriched with graph-based representations of vehicles and road segments.

The generic risk function is:

$$R = f(x_v, x_e, x_c) \in [0, 1] \tag{1}$$

Edge and cloud implementations share this structure but differ in complexity and input horizon.

C. Model Training and Scope of Evaluation

For the purposes of this study, both edge and cloud models are trained using labeled hazard and non-hazard events generated within the simulation environment. The goal is not to propose a novel predictive learning algorithm, but to evaluate how different deployment architectures influence the timeliness, responsiveness, and usefulness of safety mitigation.

Accordingly, the models are treated as representative implementations with realistic computational footprints and operational constraints. This supports controlled comparison of hazard detection timing, reaction latency, and mitigation effectiveness across local-only, cloud-only, and cooperative configurations.

D. Thresholds and Actions

Two thresholds guide actions:

- τ_{edge} for immediate, local mitigation,
- τ_{cloud} for medium-horizon interventions (e.g., routing or policy changes).

Actions range from escalated driver alerts and adaptive cruise adjustments to speed recommendations on high-risk segments. Fleets can tune thresholds by use case (e.g., hazardous materials vs. empty trailer).

V. Experimental Setup

We evaluate the proposed cooperative intelligence architecture using a simulation environment designed to approximate realistic logistics operations involving heavy commercial vehicles (HCVs) and mixed passenger traffic. The simulation is intended to compare architectural behavior under representative heavy-vehicle hazard conditions rather than to reproduce a specific deployed fleet or roadway system.

A. Fleet, Network, and Scenarios

- **Fleet:** The simulated fleet consists of 50 HCVs equipped with onboard sensing, edge computation, and V2X capability. Vehicle dynamics include simplified longitudinal and lateral behavior, with load-dependent braking performance and rollover sensitivity [1], [7].
- **Road network:** The environment uses a hybrid urban–highway topology containing straight corridors, intersections, curves, ramps, and variable grades.
- **Traffic mix:** HCVs operate alongside passenger vehicles under heterogeneous driving patterns and varying traffic densities.
- **Scenarios:**

- Propagating braking waves in dense traffic,
- Localized low-friction areas during adverse weather,
- Night-time operation with elevated fatigue probability,
- Downhill segments combining grade, speed, and cargo load.

B. Configurations

We compare three system configurations:

1) **Local-only (ADAS):** Vehicles rely exclusively on on-board perception and rule-based assistance without V2X or cloud input.
2) **Cloud-only (Telematics analytics):** Telemetry is transmitted to the cloud, where risk is computed centrally and advisories are returned to vehicles, without local edge inference or V2V/V2I exchange.
3) **Cooperative (Proposed):** Edge inference, V2X hazard dissemination, and cloud-generated risk forecasts operate together. Mitigation may be triggered either locally or through predictive cloud insight.

All three configurations are evaluated under the same traffic, hazard, sensing, and communication assumptions in order to isolate the effect of architectural distribution of intelligence.

C. Metrics

System performance is evaluated using the following metrics:

- **Hazard detection time:** The interval between ground-truth hazard emergence and system recognition.
- **Reaction latency:** The time from hazard detection to initiation of mitigation, such as driver warning or assistance adjustment.
- **Near-collision count:** Events in which time-to-collision falls below a critical threshold without resulting in physical impact.

These metrics were selected to capture both the speed of hazard awareness and its operational safety effect.

D. Simulation and Network Assumptions

The evaluation is performed using a discrete time-stepped traffic simulation with heterogeneous HCVs and passenger cars. Wireless communication is modeled through latency distributions representing a combination of direct V2V exchange and cellular connectivity. Unless otherwise noted, V2X penetration among HCVs is assumed to be 100%.

For V2V communication, one-hop broadcast latency is uniformly distributed between 20–60 ms with a packet delivery ratio of 95% and an effective communication range of approximately 300 m. V2N communication toward the cloud incorporates uplink transmission, core-network traversal, and processing delays, resulting in end-to-end latencies ranging from 400–900 ms.

Hazard ground truth is activated when the time-to-collision (TTC) falls below 2.5 s or when stability-related limits, such as excessive lateral acceleration, are violated. Edge inference executes at 10 Hz, while cloud-generated risk forecasts are refreshed every 5 s.

Sensor uncertainty is approximated using bounded Gaussian perturbations applied to speed, acceleration, and position signals. Each scenario is executed multiple times using different random seeds, and the reported metrics correspond to mean values together with their associated standard deviations.

The chosen parameters are aligned with ranges commonly reported in vehicular networking and traffic safety literature. These assumptions are reported explicitly to support reproducibility and to clarify that the present study evaluates comparative architectural behavior under controlled but realistic operating envelopes rather than claiming field-validated performance.

VI. Results

Table I summarizes the main results.

The cooperative configuration outperforms both baseline approaches across the primary safety metrics. Relative to the local-only baseline, it improves hazard detection time by approximately 61% and reduces near-collision occurrences by about 46%. These gains arise primarily from earlier awareness enabled by V2X hazard dissemination combined with predictive context obtained from the cloud.

In braking-wave scenarios, for example, upstream V2V notifications allow following vehicles to classify the situation as hazardous before local sensors alone observe severe deceleration. Because all configurations are evaluated under identical traffic, hazard, sensing, and communication assumptions, the observed differences primarily reflect the effect of where intelligence is placed in the system rather than differences in scenario difficulty.

Reaction latency remains low for both the local-only and cooperative configurations because mitigation is executed at the vehicle edge. By contrast, the cloud-only configuration exhibits higher latency due to communication and processing delays, making it less suitable for split-second intervention despite its value for broader fleet-level analysis.

The reduction in near-collision events in the cooperative configuration is particularly important from a safety perspective. Earlier hazard awareness increases the effective stopping distance and supports smoother deceleration, especially in dense traffic and on downhill grades where HCV dynamics are more difficult to control.

The cloud-only configuration is intended to approximate prevailing fleet telematics practice, in which centralized analytics provide advisory feedback but do not generally participate in sub-second control loops. The comparison therefore highlights the benefit of distributing intelligence closer to vehicles rather than implying that centralized analytics lack operational value.

VII. Discussion and Limitations

The experiments assume full V2X penetration among participating HCVs in order to estimate an upper-bound cooperative benefit under favorable connectivity conditions. In real

TABLE I. COMPARATIVE PERFORMANCE (MEAN ± STD ACROSS RUNS)

Approach	Detection Time (s)	Latency (ms)	Near-Collisions
Local-Only	1.82 ± 0.41	95 ± 18	7.4 ± 1.9
Cloud-Only	4.63 ± 0.88	620 ± 140	4.9 ± 1.2
Cooperative	$\mathbf{0.71 \pm 0.22}$	110 ± 25	$\mathbf{4.0 \pm 1.0}$

deployments, partial adoption would reduce the availability of shared safety information; however, even moderate penetration levels are expected to preserve qualitative advantages because hazard information can still propagate earlier across connected vehicles than through isolated onboard sensing alone.

The results highlight the complementary roles of the three architectural layers:

- **Edge:** Provides fast, deterministic response through local sensing, local inference, and caching of cooperative data.
- **V2X:** Extends situational awareness by enabling early warning of upstream hazards and richer context from neighboring vehicles and infrastructure.
- **Cloud:** Learns from historical and external data to forecast conditions that are not yet locally visible and to support fleet-level policy refinement [2], [5], [8].

Together, these layers form a cooperative intelligence loop in which risk is assessed and mitigated using both local evidence and shared knowledge. The results should therefore be interpreted as evidence of architectural feasibility and comparative benefit under controlled simulation assumptions rather than as proof of deployment readiness.

There are several limitations to the present study:

- **Simulation-based evaluation:** The experimental validation is simulation-based and therefore does not fully capture the diversity of real-world driver behavior, sensing artifacts, traffic dynamics, or operational disruption conditions.
- **No real-world dataset validation:** The study does not include validation using public or proprietary fleet telemetry datasets, and therefore predictive behavior under operational data drift remains to be established.
- **Abstracted communication modeling:** Communication performance was modeled using realistic but abstract latency and delivery assumptions; actual V2X deployments may exhibit different reliability, congestion, and interference patterns.
- **No deployment or hardware-in-the-loop validation:** The evaluation is intended to compare architectural behavior rather than to demonstrate production readiness. Hardware-in-the-loop testing, pilot deployments, and field validation are required before operational conclusions can be drawn.
- **Fleet and regional adaptation:** Models trained on one fleet, geography, or regulatory context may require recalibration or adaptation before transfer to other operational environments.

Despite these limitations, the experiments indicate that integrating edge, V2X, and cloud intelligence is a promising direction for proactive HCV safety.

A. Deployment and Resource Considerations

The cooperative architecture is designed with practical deployment constraints in mind. Edge inference targets embedded compute platforms capable of executing compact models within tens of milliseconds, consistent with contemporary automotive-grade processors.

V2X communication prioritizes small, event-driven safety messages rather than continuous raw sensor streaming, thereby limiting bandwidth consumption and improving scalability. Cloud workloads focus on aggregation, medium-horizon forecasting, and policy refinement, where somewhat higher latency is acceptable.

Although the present study abstracts hardware diversity, the observed timing relationships remain representative of realistic operational envelopes. The architecture is therefore intended to support incremental deployment within existing fleet and vehicle ecosystems rather than requiring wholesale replacement of current onboard safety or telematics systems.

VIII. INDUSTRIAL APPLICABILITY

The architecture can be applied incrementally in:

- **Logistics and freight fleets:** Enhancing safety on long-haul routes and high-risk corridors.
- **Public transportation:** Improving safety and incident response in dense urban traffic.
- **Hazardous materials transport:** Supporting stricter safety envelopes and compliance reporting.

Deployment can start as advisory-only, integrated with existing telematics and ADAS platforms, and evolve toward tighter integration with vehicle actuation as regulations and technology mature. Security, privacy, and interoperability must be addressed using standard V2X security frameworks and careful data governance.

IX. IMPLICATIONS FOR SOFTWARE-DEFINED VEHICLES

The ongoing transition toward software-defined vehicles (SDVs) is changing how functionality is delivered and maintained throughout the vehicle lifecycle. Capabilities that were once implemented as fixed hardware features are increasingly realized as software services that can be updated, tuned, and validated after vehicles are in operation.

From this perspective, the cooperative architecture explored in this paper fits naturally into emerging platform strategies, although its realization will depend on integration maturity across OEM, supplier, and infrastructure ecosystems. Edge-resident risk logic can in principle be updated over-the-air, and cloud systems can continuously refine models using fleet observations. In practice, however, deployment cadence, certification requirements, and organizational readiness will influence how quickly such feedback loops can be established.

A key implication is that safety improvement becomes partly a data and lifecycle management problem. Mechanisms for monitoring model behavior, validating updates, and providing safe rollback paths are as important as prediction accuracy itself. Moreover, interoperability across vendors and jurisdictions remains a non-trivial barrier that will require alignment beyond technical architecture.

Nevertheless, the ability to blend local perception with shared and learned context offers a pragmatic direction for incrementally enhancing safety capabilities. Even partial adoption can provide benefits in high-risk corridors or specific operational profiles, creating a pathway toward broader cooperative functionality over time.

X. Conclusion and Future Work

This study examined how combining onboard edge inference, V2X communication, and cloud analytics can improve the timeliness of hazard recognition for heavy commercial vehicles. Across the simulated scenarios, the cooperative configuration consistently identified risk earlier than local-only ADAS and cloud-only telematics baselines while preserving rapid local execution of mitigation actions at the vehicle edge.

The results suggest that distributing intelligence across architectural layers provides practical safety advantages, particularly when critical signals originate beyond the immediate sensing range of an individual vehicle. Access to upstream cooperative knowledge, even when partial or imperfect, can extend the decision horizon and support earlier, smoother, and more context-aware responses.

At the same time, the present findings should be interpreted within the scope of simulation-based evaluation. Moving from simulation to operational deployment introduces additional challenges, including communication reliability, fleet heterogeneity, model calibration, regulatory approval, and integration with existing operational processes. The proposed architecture should therefore be viewed as a reference framework for cooperative predictive safety rather than as a turnkey deployment solution.

Future work will extend the study beyond simulation through pilot implementations, hardware-in-the-loop experiments, and evaluation under reduced V2X penetration levels. Additional validation using real fleet telemetry, where available, will be important to assess robustness under operational data drift and cross-fleet variability. Further research is also needed to understand how cooperative predictions interact with driver behavior, fleet maintenance practices, and existing driver assistance and fleet management workflows.

Overall, the evidence indicates that structured cooperation among vehicles, infrastructure, and cloud services can form a meaningful step toward more anticipatory, network-aware safety operations in heavy transport.

References

[1] Z. Tian, F. Chen, S. Ma, and M. Guo, “Analysis of the severity of heavy truck traffic accidents under different road conditions,” *Applied Sciences*, vol. 14, no. 22, Art. no. 10751, Nov. 2024.

[2] H. Zhou, W. Xu, J. Chen, and W. Wang, “Evolutionary V2X technologies toward the Internet of Vehicles: Challenges and opportunities,” *Proceedings of the IEEE*, vol. 108, no. 2, pp. 308–323, Feb. 2020.

[3] X. Wang, M. X. Gong, and S. Mao, “An overview of 3GPP cellular vehicle-to-everything standards,” *GetMobile: Mobile Computing and Communications*, vol. 21, no. 3, pp. 19–25, Sep. 2017.

[4] T. Huang *et al.*, “Vehicle-to-everything cooperative perception for autonomous driving,” *Proceedings of the IEEE*, vol. 113, no. 5, pp. 443–477, May 2025.

[5] L. Liu, C. Chen, Q. Pei, S. Maharjan, and Y. Zhang, “Vehicular edge computing and networking: A survey,” *Mobile Networks and Applications*, vol. 26, pp. 1145–1168, 2021.

[6] S. Raza, S. Wang, M. Ahmed, and M. R. Anwar, “A survey on vehicular edge computing: Architecture, applications, technical issues, and future directions,” *Wireless Communications and Mobile Computing*, vol. 2019, Art. ID 3159762, 2019.

[7] T. S. J. Darwish and K. A. Bakar, “Fog based intelligent transportation big data analytics in the Internet of Vehicles environment: Motivations, architecture, challenges, and critical issues,” *IEEE Access*, vol. 6, pp. 15679–15701, 2018.

[8] P. Li, M. Abdel-Aty, and J. Yuan, “Real-time crash risk prediction on arterials based on LSTM-CNN,” *Accident Analysis and Prevention*, vol. 135, Feb. 2020, Art. no. 105371.

[9] Y. Yuan, L. Zhang, and X. Qin, “Real-time crash risk prediction using long short-term memory recurrent neural network,” *Transportation Research Record*, vol. 2673, no. 6, pp. 314–326, 2019.

[10] K. N. Qureshi, S. Din, G. Jeon, and F. Piccialli, “Internet of vehicles: Key technologies, network model, solutions and challenges with future aspects,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 3, pp. 1777–1786, Mar. 2021.