ARTICLE INFO
Article ID: 12-09-03-0019

doi:10.4271/12-09-03-0019

Risk-Aware Navigation Framework for Autonomous and Human-Driven Vehicles: Integrating Crash Probability Data for Safer Mobility

Lars Skaug[1] and Mehrdad Nojoumian[2]

[1]*Florida Atlantic University, Computer Science, USA*
[2]*Florida Atlantic University, Electrical Engineering and Computer Science, USA*

Abstract

This article presents a system to incorporate crash risk into navigation routing algorithms, enabling safety-aware path optimization for autonomous and human-driven vehicles alike. Current navigation systems optimize travel time or distance, while our approach adds crash probability as a routing criterion, allowing users to balance efficiency with safety. We transform disparate data sources, including traffic counts, crash reports, and road network data, into standardized risk metrics. Because traffic volume data only exist for a small subset of road segments, we develop a solution to project average daily traffic estimates to an entire road inventory using machine learning, achieving sufficient coverage for practical implementation. The framework computes exposure-normalized crash rates weighted by severity and integrates these metrics into routing cost functions compatible with existing navigation algorithms. The key strength of our solution is its scalability. In addition to the mapping data required by the navigation system, it requires only two additional data sources commonly maintained by transportation authorities: geolocated crash reports and traffic counts, enabling deployment across diverse jurisdictions. For connected and automated vehicles, the framework provides quantitative risk assessment for path planning algorithms. For conventional vehicles, it enables drivers to make informed routing choices based on safety preferences. Our empirical validation demonstrates that risk-aware routing achieves substantial safety improvements while maintaining reasonable travel times. The methodology also serves transportation agencies by systematically identifying high-risk corridors and crash patterns across road networks. By establishing a standardized approach to safety-aware navigation, this work addresses an important gap in current routing systems and contributes to the development of safer transportation infrastructure.

History

Received: 19 May 2025
Revised: 15 Sep 2025
Accepted: 15 Dec 2025
e-Available: 09 Jan 2026

Keywords

Crash Prevention, Safe Navigation, Machine Learning, GIS, Spatial Analysis, Open-Source Software, Navigation, Routing

Citation

Skaug, L. and Nojoumian, M., “Risk-Aware Navigation Framework for Autonomous and Human-Driven Vehicles: Integrating Crash Probability Data for Safer Mobility,” *SAE Int. J. of CAV* 9(3):2026, doi:10.4271/12-09-03-0019.

ISSN: 2574-0741
e-ISSN: 2574-075X

1. Introduction

Navigation software generally optimizes for travel time, commonly with options to avoid toll roads or ferries. Our research introduces another feature to optimize for when navigating public roads: the risk of a crash. This is a novel idea, as current "research into risk-aware modeling for driving is almost entirely focused on hazardous materials transportation," according to Hu et al. [1]. Hazardous materials routing addresses catastrophic risk scenarios with mandatory regulatory constraints, focusing on protecting the public from potential chemical releases or explosions through prescribed route restrictions. Our solution, in contrast, is meant for general civilian navigation and addresses statistical crash probability based on historical traffic safety data, offering users the flexibility to balance personal safety preferences with travel efficiency rather than compliance with regulatory mandates. The advent of autonomous vehicles, with their requisite navigation systems, provides an obvious opportunity to include such preferences, but human drivers could easily incorporate safety preferences when using navigation software as well.

When it comes to road safety analysis itself, it has traditionally focused on crash hotspots and countermeasures through examination of historical crash patterns. Significant work has been done to develop statistical methodologies for crash prediction and safety assessment [2, 3], to understand where and how crashes happen. A framework for quantifying and modeling crash risk across different road types was developed by Ezra Hauer [4]. His empirical Bayes (EB) methodology has been incorporated into the U.S. Federal Highway Administration's Highway Safety Improvement Program (HSIP), as it allows unusual crash frequencies to be modified by Safety Performance Functions to compare segments based on common features and account for regressions to the mean. Moving beyond post-fact analysis, real-time crash risk calculations have been explored with concepts such as extreme value theory models and traffic conflict data [5], though these approaches often require detailed analyses of specific infrastructure segments, which limits their scalability. Alternative methodologies in crash risk assessment include EB approaches that adjust observed crash frequencies using Safety Performance Functions [4], probabilistic risk mapping based on detailed infrastructure characteristics, and traffic conflict-based models that analyze near-miss events to predict crash likelihood. While these approaches provide valuable insights for localized safety analysis, they typically require extensive data collection efforts, specialized sensing infrastructure, or detailed geometric and operational characteristics that are not uniformly available across large road networks. Our solution requires only two commonly available data sources and integrates directly with existing navigation systems, providing immediate scalability across diverse geographic regions.

We use state-of-the-art GIS and data engineering solutions (GDAL, Spark, Sedona, Overture Maps) to incorporate crash risk assessments directly into route planning algorithms for both autonomous and human-driven vehicles, with the explicit goal of global applicability. The framework provides safety intelligence for autonomous vehicles while simultaneously offering enhanced navigation options for human drivers. The core algorithm for integrating safety metrics into routing decisions is the subject of a U.S. patent application [6], representing a novel approach to navigation optimization. Our implementation requires only two widely available data sources: (1) crash reports from law enforcement officers (LEOs) that include geolocation or clear location descriptions, and (2) traffic counts for a subset of roads. These data requirements are met across large parts of the world, enabling safety-aware navigation beyond the constraints of previous infrastructure-specific approaches. To demonstrate our solution, we processed 3.3 million crashes spanning from January 2013 to February 2025 from the U.S. state of Ohio, to quantify relative crash risk for individual road segments based on historical patterns and traffic exposure metrics. Details on the implementation and its integration with navigation software are presented in the following sections.

Our solution is a natural fit for autonomous vehicles, as a navigation system is an integral part of their operation. For people driving, it allows for a balance between travel time and risk exposure, creating a new dimension in route optimization. This dual-purpose solution is not only appropriate for individual travelers but also for connected vehicle fleets, autonomous mobility services, and transportation agencies working to mitigate systemic safety challenges. By providing a standardized methodology for risk assessment that functions across diverse operational environments, our framework contributes to both immediate safety improvements and long-term CAV development objectives.

2. Methodology

The proposed methodology incorporates crash risk assessment into route optimization algorithms through a systematic approach combining data preparation, risk quantification, and integration into routing cost functions. To clarify the different uses of the term "cost" in this article, we use the following distinct notation:

- C_e = Economic cost of crashes on segment e (dollars)
- $cost_e$ = Routing cost function for optimization (dimensionless)
- r_e^c = Cost-weighted crash rate (dollars per DVMT)

2.1. Risk Estimation

Let $G = (V, E)$ represent the road network, where V is the set of vertices (intersections) and E is the set of edges (road segments). For each segment $e \in E$, we define the following key attributes:

$$l_e = \text{length of segment } e \left(\text{miles}\right) \tag{1}$$

$$\text{ADT}_e = \text{Average Daily Traffic on segment } e \left(\text{vehicles/day}\right) \tag{2}$$

$$c_e = \text{number of crashes on segment } e \text{ durang study period} \tag{3}$$

$$T = \text{duration of study period} \left(\text{days}\right) \tag{4}$$

To normalize for differential exposure across segments, we compute the daily vehicle miles traveled (DVMT) as:

$$\text{DVMT}_e = \text{ADT}_e \times l_e \tag{5}$$

The exposure-normalized crash rate r_e for each segment is then calculated as:

$$r_e = \frac{c_e}{\text{DVMT}_e \times T} \tag{6}$$

This provides the raw crash rate. While we use this raw rate in our calculations, it's worth noting that for reporting and comparison with standard literature, this rate can be expressed as:

$$r_e^{\text{reported}} = r_e \times 10^6 \tag{7}$$

This scaling expresses the rate in crashes per million vehicle miles traveled, consistent with standard practice used by FHWA and transportation safety literature [7].

2.2. Crash Cost by Severity

The severity of crashes matters greatly to the actual as well as perceived risk. Similar to analyses by the Federal Highway Administration (FHWA) [8], we account for the severity of crashes in our calculations by assigning economic costs to crashes based on their severity level using the KABCO scale: For each segment e, we compute the total crash cost C_e as:

$$C_e = \sum_{k \in K} n_{e,k} \times cost_k \tag{8}$$

where K is the set of severity levels, $n_{e,k}$ is the number of crashes of severity k on segment e, and $cost_k$ is the economic cost associated with severity level k. This allows us to calculate a cost-weighted crash rate r_e^c for each segment:

TABLE 1 Crash costs by KABCO severity classification. (Data taken from Ref. [8].)

KABCO	Severity code	Cost per crash (2016 dollars)
K	1	$11,295,400
A	2	$655,000
B	3	$198,500
C	4	$125,600
O	5	$11,900

$$r_e^c = \frac{C_e}{\text{DVMT}_e \times T} \tag{9}$$

Weighting the crash risk by severity biases our algorithm to avoid severe crashes while allowing more fender benders. This approach is consistent with standard safety evaluation practices and reflects typical drivers' risk perception and preferences, as research indicates drivers are willing to accept longer travel times to avoid routes with a higher likelihood of severe crashes [9, 10]. However, the practical implementation of such severity-weighted risk metrics in navigation systems depends critically on whether drivers will actually modify their routing behavior in response to safety information—a behavioral assumption that merits empirical examination (Table 1).

2.3. Risk Percentile Transformation

In order for risk values to be easier to compare across different road types and locations, we transform the raw crash rates into percentiles. For each road segment e with type t (e.g., primary, secondary, residential):

$$p_e = \text{percentile}\left(r_e, \left\{r_{e'} | e' \in E_t\right\}\right) \tag{10}$$

where E_t represents all segments of type t. This transformation normalizes risk values to a scale between 0 and 1, making them easier to use in routing algorithms while still recognizing that different types of roads have different underlying risks. Such a percentile is also easier to interpret for analysts reviewing the data.

3. Implementation

We selected GraphHopper to demonstrate our solution for three key reasons: (1) it is open source, providing full access to modify and extend its capabilities; (2) it utilizes OpenStreetMap (OSM) data, which is also open source and forms the basis of our road network data; and (3) it offers extensive customization options for routing algorithms, enabling the integration of our crash risk metrics directly into route calculations.

For this research, we created a fork of GraphHopper [11] to implement risk-aware routing modifications. In our implementation, each OSM way in the road network is augmented with a risk metric encoded as a tag with value $r \in [0,1]$ with resolution $\delta = 10^{-3}$. The risk values are stored using a 10-bit representation, providing sufficient granularity for risk differentiation while maintaining computational efficiency.

Our GraphHopper implementation directly corresponds to the multiplicative cost adjustment framework established in the methodology section. The routing engine applies the theoretical cost function (Equation 11):

$$cost_e = \frac{cost_e^{base}}{1 - k \cdot p_e} \tag{11}$$

through GraphHopper's priority-based weight system, where $cost_e^{base}$ represents the standard routing cost and $k \cdot p_e$ is stored as the road risk tag for each segment.

The model uses a priority function that adjusts the weight of each road segment based on its risk value:

```
{
  "priority": [
    {
      "if": "true",
      "multiply _ by": "1-road _ risk"
    }
  ]
}
```

This configuration implements the theoretical framework by setting the priority multiplier to (1 − road risk) = $(1 - k \cdot p_e)$ from Equation 11, ensuring consistency between the methodology and implementation. For a road segment with risk value r, the priority multiplier becomes $(1 - r)$, meaning that segments with higher crash risk receive proportionally higher costs and are less likely to be selected in the optimal route. By disabling GraphHopper's optimization algorithms (contraction hierarchies and landmarks) through API parameters, we ensure that the custom model is fully applied during route calculation.

The resulting routes balance safety considerations with practical routing constraints, demonstrating how crash risk data can be integrated into real-world navigation systems.

Our approach introduces no computational overhead compared to standard routing algorithms. Risk values are pre-computed and embedded as static attributes in the road network data, enabling CAV routing systems to achieve identical performance to conventional navigation while incorporating safety optimization. Although our implementation requires disabling GraphHopper's Contraction Hierarchies and Landmarks optimizations to enable custom priority models, performance benchmarking validates that this does not impact routing performance: our risk-aware configuration achieved identical response times to standard GraphHopper configuration with optimizations enabled—both averaging 3.5 ms per request across 1120 route calculations. For the typical commute patterns between adjacent counties tested in this study, disabling these optimizations to enable safety-aware routing imposes no performance penalty compared to conventional navigation systems.

3.1. Data Preparation

The integration of crash risk metrics with a navigation system requires rigorous data preparation to ensure spatial accuracy, compatibility with routing software, and computational efficiency. For different navigation systems, there will be different considerations to make, but the key step is to project geolocated crash counts onto official road inventory data with average daily traffic (ADT) counts. The result is then projected onto mapping data that can be used by routing software. The road inventory would be maintained by regional or national transportation authorities—Departments of Transportation (DOT) in the United States. Crash counts are collected by local LEOs (police). In our case, we were working with crash reports that had already been geolocated. In many localities, the description of the location of the crash would have to be geolocated first.

Figure 1 outlines the data processing steps described in this section.

The Ohio crash dataset contained geocoded coordinates for all 3.3 million crash records spanning January 2013 to February 2025. No crashes were missing geolocation data, and no attempt was made to validate or correct the geolocation accuracy in this work, though we acknowledge this as a limitation discussed in Section 4.1.2. All crash data used in this study consists of de-identified records from the Ohio state authorities. Our analysis uses only aggregated crash counts, severity classifications, and geographic coordinates, with no personally identifiable information accessed or retained.

To demonstrate our solution, we obtained traffic counts (ADT) from the Ohio Department of Transportation's road inventory [12]. Although this dataset contains a total of 418,329 segments, only 127,252 segments (30.42%) contained ADT values. To address this limitation, we implemented a gradient boosted tree (GBT) regression model using the LightGBM framework to project ADT values onto the remaining 291,077 segments lacking empirical measurements. We selected GBT over alternatives such as CNNs primarily due to practical considerations: tree-based models demonstrate robust performance with our training dataset size (101,801 samples) without requiring extensive hyperparameter tuning, and they naturally handle the mixed categorical and continuous features present in road inventory data without complex preprocessing pipelines. The primary benefits of this approach include computational efficiency, robust handling of

FIGURE 1 Data processing pipeline for risk-aware navigation system.

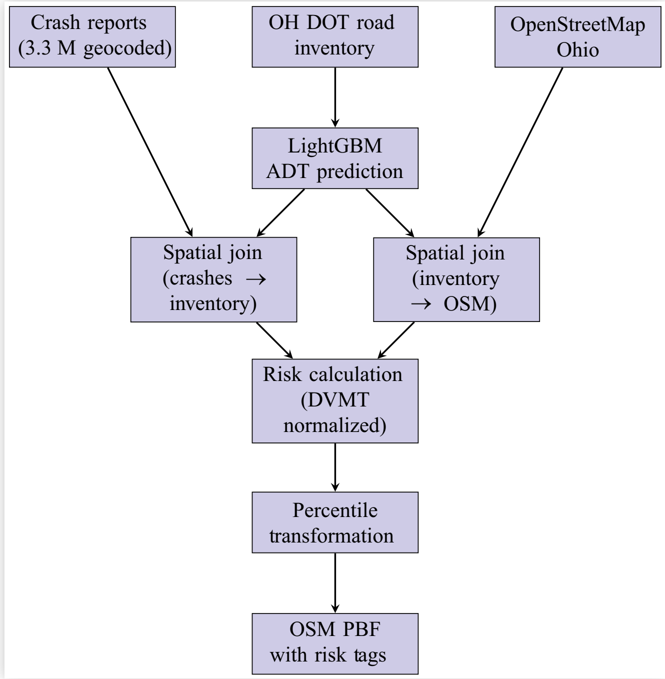

missing values, and interpretability for transportation domain experts. However, this choice also presents limitations: GBT models may not fully capture complex spatial relationships between neighboring road segments or leverage geographic proximity information that could improve predictions. Additionally, the model may have reduced generalizability to road types with limited representation in the training data.

3.1.1. Model Performance and Stratified Analysis Our GBT model achieved a high coefficient of determination (R^2) of 0.92 on the validation dataset, indicating robust predictive performance. An analysis of absolute errors shows a mean absolute error (MAE) of 1662.23, while the more robust median absolute deviation (MAD) was 863.03, reflecting the model's typical prediction error. While the mean absolute percentage error (MAPE) was elevated at 120.58% the symmetric mean absolute percentage error (sMAPE), a more stable measure, was 52.79%.

To further assess performance across different functional classes, the metrics were stratified by route type, as detailed in Table 2. The results show consistently strong performance on arterial routes (e.g., SR, US, IR) but varied accuracy on local roads and ramps. Notably, the model's overall R^2 of 0.92 is higher than that of any individual subgroup. This phenomenon, related to the Yule–Simpson effect, indicates that a substantial portion of the explained variance is attributable to the model's ability to capture the significant mean differences in traffic volume between distinct route types, in addition to explaining the variance within each type.

Note that we employed a standard random cross-validation rather than spatial cross-validation techniques. This methodological choice reflects the specific nature of our task: intra-dataset interpolation of missing ADT values within a connected road network. Unlike predictive modeling scenarios where spatial independence is crucial for generalizability assessment, our interpolative task benefits from the model's ability to leverage spatial

TABLE 2 Stratified model performance by route type.

Code	Description	*N*	R^2	RMSE	MAE	sMAPE (%)
CR	County road	7942	0.82	2094.23	1244.61	54.12
SR	State route	5575	0.87	3137.92	1815.50	35.95
TR	Township road	4305	0.61	1178.40	554.94	81.33
MR	Municipal route	3822	0.67	3131.86	2081.91	61.08
US	US route	1859	0.89	3633.91	2419.76	24.18
RA	Ramp	1418	0.25	7003.15	2623.43	55.28
IR	Interstate route	530	0.88	12,426.76	7046.81	12.20

relationships between connected road segments. The resulting ADT estimates serve as input for exposure-normalized risk calculations rather than standalone predictions, making the practical performance within the study region the primary consideration.

The finalized model was subsequently applied to all road segments in the road inventory dataset lacking empirical ADT measurements. The combined actual and estimated ADT values served as critical inputs for the subsequent computation of risk metrics normalized by vehicle miles traveled. Note that although this set now had traffic counts for the inventory maintained by the state DOT, there are still segments unaccounted for, such as those maintained by federal or local authorities, as well as segments not designed for vehicular traffic. For reference, Table 3 shows each road (highway) type available in our dataset.

For the routing solution to work, we needed to project the segments from the state DOT inventory onto a standard mapping dataset. We chose to use OpenStreetMap (OSM) for three reasons: its widespread adoption in spatial navigation applications, its open licensing model that permits derivative works, and its direct compatibility with GraphHopper—our selected routing implementation platform.

The projection was done with a spatial join process that leveraged geometric relationships between the datasets. To perform this work, we used Apache Sedona, an open-source system extending Apache Spark with spatial data processing capabilities. Here are the steps involved:

First, we transformed both datasets into an appropriate local projection for Ohio (EPSG:3746) with units in meters to simplify distance calculations. For each OSM segment, we identified all potentially matching road inventory segments using the `ST_Intersects` predicate, while requiring a minimum overlap threshold of 25 m to exclude insignificant intersections. The overlap length was calculated using the `ST_Length` function applied to the intersection of geometries. To resolve cases where an OSM segment intersected multiple road inventory segments, we ranked each candidate match by the length of overlap and selected the road inventory segment with the maximum overlap value. This approach ensured that each OSM segment was associated with its most spatially congruent counterpart in the road inventory dataset.

Once the data had been matched to the OSM data, we calculated crash risk metrics for each segment based on the ratio of crashes to exposure (measured as DVMT). These metrics were normalized to a [0, 1] scale using percentile ranking, producing a risk score where higher values indicate greater relative crash risk. We then stored these risk values alongside their corresponding OSM segment identifiers, creating:

$$R = \left\{ \left(id_i, r_i \right) \mid id_i \in \text{OSM}, r_i \in \left[0,1 \right] \right\} \tag{12}$$

where R represents the set of OSM segment identifiers and their corresponding risk values. For segments lacking empirically measured traffic volumes (ADT) and therefore missing risk values (approximately 95% of the network), we implemented a hierarchical imputation strategy. Rather than employing a single global imputation value, we compute the median risk for each highway classification category:

$$\mu_h = \text{median}\left(\left\{ r_i \mid \left(id_i, r_i \right) \in R, \text{highway}\left(id_i \right) = h \right\} \right) \tag{13}$$

where μ_h represents the median risk for highway type h and highway(id_i) is a function that returns the highway classification of segment id_i. Table 3 shows each highway type available in the OSM dataset along with its calculated median crash risk.

While absolute crash rates vary substantially between road classes—ranging from under 400 to over 30,000 crashes per million VMT (see Appendix Table A.1)—our routing algorithm normalizes risk within each road class to percentiles. This ensures route selection accounts for relative safety differences within each road type rather than being dominated by baseline differences between road classifications. Alternative risk estimation approaches are discussed in Section 4.2.

Once the crash risk has been calculated for each segment, we need to make this new feature available for navigation. To this end, we traversed the OSM data structure to append a risk tag containing the appropriate risk value according to the following decision framework:

$$\text{risk}\left(id_i \right) = \begin{cases} r_i, & \text{if } \left(id_i, r_i \right) \in R \\ \mu_h, & \text{if highway}\left(id_i \right) = h \text{ and } \mu_h \text{ exists} \\ \mu_{\text{global}}, & \text{otherwise} \end{cases} \tag{14}$$

where μ_{global} represents the global median risk across all segments with empirically derived values.

This hierarchical approach ensures comprehensive coverage of the road network while maintaining the highest possible fidelity to observed risk patterns. The categorization by highway type provides valuable risk discrimination even in the absence of direct empirical measurements. Our analysis revealed substantive variations in median risk across different road classifications that would be obscured by a simplistic global median approach. As can be seen from Table 3, unclassified roads demonstrated a markedly elevated risk profile ($\mu_{\text{unclassified}}$ = 0.81) compared to service roads (μ_{service} = 0.33), despite both categories often being considered secondary in conventional road hierarchies.

We carefully chose the steps in our algorithm to achieve three objectives: First, it accommodates the inherent sparsity of direct risk measurements by implementing an imputation strategy that preserves meaningful differentiation across road types, rather than relying

TABLE 3 Median crash risk by highway type.

Highway type	Median risk	Count
Residential	0.48	82,746
Tertiary	0.70	25,825
Secondary	0.52	17,958
Primary	0.47	14,192
Footway	0.38	12,097
Motorway	0.56	9786
Service	0.33	9681
Unclassified	0.81	7234
Motorway link	0.35	6143
Trunk	0.57	4053
Cycleway	0.33	612
Trunk link	0.36	406
Primary link	0.14	323
Secondary link	0.06	184
Path	0.30	142
Tertiary link	0.09	115
Pedestrian	0.21	69
Track	0.76	49
Construction	0.41	40
Living street	0.22	18
Busway	0.49	12
Bridleway	0.70	10
Steps	0.37	3
Road	0.74	1

on global defaults. Second, it stays true to the original OSM data, allowing compatibility with solutions that depend on that. Third, it normalizes risk in a way that would be compatible with many algorithms while maintaining computational efficiency through selective data processing and appropriate file format transformations. The direct compatibility with GraphHopper demonstrates its usefulness in an established, high-performance routing engine capable of multi-criteria path optimization. Finally, the normalized risk expression also facilitates interpretation by users reviewing the data.

3.2. Risk-Weighted Route Evaluation

To evaluate the effectiveness of risk-aware routing, we developed a method to calculate the weighted average risk along a complete route. This metric provides a quantitative measure of overall risk exposure:

$$\text{Route Risk} = \frac{\sum_i \left(\text{distance}_i \times \text{risk}_i\right)}{\sum_i \text{distance}_i} \tag{15}$$

where distance_i represents the length of a road segment and risk_i its associated risk value. The implementation leverages GraphHopper's path details, which provide run-length–encoded information about various attributes along the route, including our custom risk metric and segment distances.

3.3. Performance Analysis

Our implementation enables systematic comparative analysis between traditional shortest-path routes and risk-optimized alternatives. The framework supports quantitative evaluation of the trade-offs between route efficiency and safety through multiple metrics captured simultaneously during route calculation.

The GraphHopper API configuration allows for the collection of comprehensive route characteristics through the details parameter:

```
"details": ["road_risk", "road_class", "distance"]
```

The three listed features are the normalized road risk, the highway type, and the distance for each segment traversed in a route. With these details, we can calculate the weighted risk exposure for an entire route.

By comparing risk-optimized routes with traditional shortest paths, the system provides a foundation for empirical assessment of the relationship between route efficiency and safety. We use this to assess the effectiveness with which our solution allows a user to prioritize safety versus travel time.

4. Methodological Considerations

The risk assessment framework presented in the previous section provides a mathematical foundation to integrate safety considerations into route optimization. Here, we discuss ways to improve some limitations of our approach.

4.1. Data Quality and Modeling Assumptions

While our methodology normalizes crash counts by exposure through the r_e calculation, it treats crashes as uniformly distributed along segments. In reality, risk often varies within segments, particularly near intersections or geometric features. The percentile transformation $p_e = F(r_e)$ mitigates some effects of outliers, but does not account entirely for heterogeneity within segments.

Our study period (January 2013 to February 2025) includes the COVID-19 pandemic years (2020–2021), during which traffic patterns and volumes differed substantially from normal conditions. While our exposure-normalized approach using DVMT should partially account for reduced traffic volumes, the fundamental changes in travel behavior—including shifts toward local trips, changes in trip timing, and altered route preferences—during this period may still influence our risk calculations. This temporal bias could be addressed in future work through time-weighted modeling that de-emphasizes pandemic years or separate risk models for different temporal regimes.

Our use of a decade-long study period ($T \approx 3650$ days) stabilizes estimates by reducing short-term variability. This approach assumes that crash patterns remain relatively stable over time. When significant roadway modifications occur within the study period, the calculated r_e values may not accurately reflect current risk conditions. An improvement could incorporate time-weighted averages that emphasize more recent data:

$$r_e = \frac{\sum_{y=1}^{Y} w_y \cdot c_{e,y}}{\sum_{y=1}^{Y} w_y \cdot \mathrm{DVMT}_{e,y}} \tag{16}$$

where $c_{e,y}$ and $\mathrm{DVMT}_{e,y}$ represent crashes and exposure in year y and w_y are temporal weights that increase for more recent years. Ohio's population remained relatively stable during our study period, supporting the assumption of consistent crash patterns over time. This methodological consideration would be more critical in regions experiencing significant demographic growth or decline, as such changes typically drive infrastructure expansion and modify traffic patterns.

4.1.1. Injury Severity Assessment Accuracy Crash costs in this study are anchored to the KABCO scale, which classifies injuries on the basis of an officer's on-scene judgement. Yet, as Burdett et al. [13] demonstrated, those judgements can diverge sharply from medical findings. Reviewing records from 520 Wisconsin agencies (2010–2019), they found that injuries coded as "A" (serious) were inflated in 45%–90% of cases, depending on the agency; statewide, nearly two-thirds (65%) of "A" designations were later shown to be less severe. Because our cost model assumes that the reported KABCO levels are correct, this upward bias in severity coding can propagate through the economic estimates and our risk severity estimates, causing bias in the risk calculations.

While the magnitude of this bias is concerning, the extent to which KABCO misclassifications affect route selection in our framework is unclear. If severity inflation varies significantly by jurisdiction or road type—as suggested by the 45%–90% variation between agencies reported by Burdett et al.—this could systematically bias risk calculations in ways that alter route recommendations. Our percentile-based ranking approach may mitigate some effects if bias is relatively uniform within road type categories, but this assumption requires validation. A comprehensive sensitivity analysis examining how various misclassification scenarios affect route selection represents a critical validation step for future implementations of this framework.

4.1.2. Crash Location Accuracy Our framework consists of binding every recorded crash to the exact road segment where it took place. Unfortunately, the source data often misses that mark—a wide-ranging audit by Ahmed et al. [14] revealed that almost 27% of the police crash reports were incorrectly geocoded, shifting some events onto neighboring links or even parallel roads. These misplacements distort segment-level risk tallies and, in turn, any routing choices that rely on them. Before analysis, a rigorous round of spatial validation—and, where needed, targeted correction—should therefore be standard practice. Concrete routines for detecting and amending such positional errors are laid out by Skaug and Nojoumian [15].

4.2. Statistical Refinements

The current formulation assumes independence between the road segments to allow for straightforward risk calculations and optimization. This simplification ignores potential spatial correlations. For example, the crash risk of an interchange can affect multiple connected segments, violating the independence assumption underlying our percentile-based approach.

The substantial variation in absolute crash rates across road classes (see Appendix Table A.1 for details) illustrates the complexity of risk patterns that could benefit from more sophisticated statistical modeling approaches.

An EB approach could enhance the stability of our risk estimates, particularly for segments with small sample sizes. Using this method, the estimated crash rate would become:

$$r_e^{EB} = w_e \cdot r_e + \left(1 - w_e\right) \cdot r_e^{SPF} \tag{17}$$

where r_e^{SPF} is the expected crash rate from a calibrated Safety Performance Function based on segment characteristics and $w \in \left[0,1\right]$ is a weight determined by the reliability of the observed data. This approach would be particularly valuable for refining the median substitution we currently employ for segments with missing data (as defined in Equation 10).

4.3. Implementation Challenges

The risk assignment framework defined in Equation 11 modifies standard routing costs by incorporating crash risk through a multiplicative penalty. A key implementation challenge involves calibrating the risk parameter k in this formulation. A high k value could divert traffic to paths that, while statistically safer, might introduce practical issues such as significantly longer travel times or routing through inappropriate corridors. A low k value, conversely, might render the safety adjustment ineffective.

The current formulation applies a linear relationship between the risk percentile p_e and the cost penalty. Alternative formulations could consider nonlinear transformations that apply greater penalties to the highest-risk segments:

$$cost_e = \frac{cost_e^{base}}{1 - k \cdot f\left(p_e\right)} \tag{18}$$

where $f(p_e)$ could be p_e^2 or $e_e^p - 1$ to create steeper penalties for high-risk segments.

A further challenge is that traffic counts (ADT), a necessary component of our system, are not available for all segments. To account for this, we apply the median crash rate by road type:

$$\begin{aligned} r_e &= \text{median}\left(\left\{r_i | i \in E_t, \text{data available for } i\right\}\right) \\ &\text{if data unavailable for } e \end{aligned} \tag{19}$$

where E_t represents all segments of highway type t.

Assigning the median risk to segments with missing data (Equation 10) provides a practical solution, which could be improved by imputation methods that consider spatial and network context. A Bayesian hierarchical model, for instance, could potentially provide more nuanced estimates for these segments by leveraging information from similar roads within the network.

4.4. Driver Acceptance and Behavioral Considerations

The practical effectiveness of risk-aware routing depends on the driver willingness to modify route choices in response to safety information. Empirical evidence from transportation choice experiments provides support for the assumption that safety considerations influence route selection behavior. Flügel et al. [16] conducted discrete choice experiments examining safety attributes in transportation decisions, finding that safety was "the most crucial determinant of the choices made" among alternative routes. Approximately 37% of respondents chose lexicographically with respect to safety attributes (i.e., selecting the safest option regardless of other factors), indicating that a substantial portion of drivers prioritize safety above other route characteristics. The research further demonstrates loss aversion for safety deterioration, with participants valuing safety losses 8%–19% higher than equivalent safety gains. Additional research confirms drivers' willingness-to-accept longer travel times to avoid routes with a higher likelihood of severe crashes [9, 10]. Meta-analytic evidence indicates a 20% premium in willingness-to-accept versus willingness-to-pay for safety changes across transportation contexts [16], suggesting that safety considerations represent a measurable component of route choice utility functions. While individual risk preferences exhibit heterogeneity, these findings provide empirical foundation for the assumption that risk-aware routing options may find acceptance among a portion of the driving population. For connected and automated vehicles, such behavioral evidence supports the integration of safety preferences into routing algorithms as a feature that aligns with demonstrated human preferences rather than imposing artificial constraints.

5. Empirical Validation

To validate the proposed risk-informed navigation approach, we used commute data from the Census Transportation Planning Products (CTPP [17]) for the 15 Metropolitan Statistical Areas (MSAs) in Ohio. For each MSA, we identified the most populous county as the *destination county*, under the assumption that densely populated counties likely have both high residential density and significant employment opportunities. Subsequently, we selected the most populous adjacent county, which is not already part of another MSA, as the *origin county*. This final selection criterion ensures clear delineation of origin–destination commute pairs that are meaningful in terms of regional transportation dynamics (Table 4).

Once we have the set of county pairs, we query the CTPP API to extract census tract-level commuting data, focusing on tract-to-tract commuter counts to identify high-volume commute corridors. We found that census tracts had sufficient data while providing a compact geographical area. Note that a census tract is a relatively small geographical area, and even with the robust population sampling employed by the U.S. Census Bureau, the margin of error is significant. In our study, we chose to filter out origin and destination pairs where the margin of error exceeded 75% of the total number of workers in any given census tract. This selection process yielded 56 high-demand routes distributed across Ohio's urban and suburban landscape, providing robust test cases for evaluating our risk-informed navigation approach in balancing travel efficiency and crash risk minimization.

To refine these locations into precise routing points, we employed Overture Maps [18], an open-source mapping dataset with comprehensive building information. For each high-volume tract pair, we identified the tallest buildings within both the origin and destination tracts, assuming that taller buildings in residential areas typically represent high-density housing, while in commercial districts they generally indicate major employment centers. This approach establishes precise geographic coordinates for the starting and ending points of each

TABLE 4 County pairs selected for origin–destination commute analysis.

MSA	Destination	FIPS	Origin	FIPS
Cleveland–Elyria	Cuyahoga	39035	Ashland	39005
Columbus	Franklin	39049	Knox	39083
Cincinnati	Hamilton	39061	Adams	39001
Dayton–Kettering–Beavercreek	Montgomery	39113	Darke	39037
Akron	Summit	39153	Holmes	39075
Toledo	Lucas	39095	Sandusky	39143
Youngstown–Warren	Mahoning	39099	Columbiana	39029
Canton–Massillon	Stark	39151	Tuscarawas	39157
Springfield	Clark	39023	Champaign	39021
Lima	Allen	39003	Hardin	39065
Mansfield	Richland	39139	Crawford	39033
Weirton–Steubenville	Jefferson	39081	Harrison	39067
Wheeling	Belmont	39013	Monroe	39111
Huntington–Ashland	Lawrence	39087	Gallia	39053
Sandusky	Erie	39043	Huron	39077

commuting route, creating realistic test cases that reflect actual commuting patterns.

5.1. Results of Empirical Testing

Our analysis of 56 high-volume commute corridors reveals significant heterogeneity in the safety-efficiency tradeoffs achieved by risk-aware routing (Table 5).

The right-skewed distributions indicate that while most routes (median) achieve modest safety improvements (10.8% risk reduction) with minimal time penalties (21 s), some routes demonstrate exceptional safety gains exceeding 70% risk reduction.

Notably, 19 routes (33.9%) achieved safety improvements while simultaneously reducing travel distance, with the most efficient alternative reducing distance by over 4 km. This suggests that conventional shortest-path algorithms may overlook inherently safer corridors that are also more efficient.

The negative minimum values for distance and time changes indicate that 25% of routes achieved safety improvements while simultaneously reducing travel time or distance. This counterintuitive finding observed in GraphHopper's routing decisions suggests that even state-of-the-art routing algorithms may select suboptimal paths when safety considerations are excluded, potentially due to outdated traffic assumptions or insufficient consideration of road hierarchy in route scoring.

The 75th percentile values (45% risk reduction, 536 m distance increase, 319 s time increase) indicate that substantial safety improvements are achievable for most commute patterns, with the majority of routes requiring less than 9 additional minutes of travel time. These findings convincingly demonstrate that risk-informed routing can achieve significant safety improvements while imposing minimal penalties on travel efficiency. While this demonstration using Ohio commuting data establishes the feasibility of integrating crash risk into navigation systems, the methodology is designed for broader applicability, and the precision of risk assessments can be enhanced through the improvements discussed in Section 6.

Visual route comparisons are provided in Appendix B, demonstrating the practical application of our GraphHopper implementation for representative origin–destination pairs.

5.2. Implications for Transportation Planning

Our empirical analysis of commuting data in Ohio suggests that substantial crash risk reductions can be achieved with acceptable travel time penalties. Transportation planners can leverage such analysis to strategically prioritize infrastructure investments. Specifically, high-risk segments along frequently traversed

TABLE 5 Distribution of route optimization outcomes across 56 origin–destination pairs.

Metric	Mean	Median	Std. dev.	Min	Max	25th %ile	75th %ile
Risk reduction (%)	22.5	10.8	25.8	–1.7	73.3	0.0	45.0
Distance change (m)	369	35	1652	–4269	6466	0.0	536
Time change (s)	142	21	222	–85	915	0.0	319

commuter corridors between major employment centers (such as those identified in our MSA county-pair analysis) represent prime targets for safety interventions. The spatial distribution of these high-leverage segments varies significantly across the state's urban and rural contexts.

Notably, we identified several corridors where safer routing alternatives coincide with shorter distances. Such instances may be worth a review by local road engineers. We identified significant regional heterogeneity in risk-reduction potential, which suggests that reviews need to be localized. By quantitatively integrating crash risk into analyses, regional and national transportation engineers can target interventions where they are most needed.

6. Future Work

The risk assessment and routing optimization framework presented in this article establishes a foundation for data-driven safety-aware navigation. In this section, we will discuss possible future developments that would increase the accuracy and precision of our framework.

A limitation of our solution is that it treats crash risk as a static function of travel volumes. An improvement to this approach would be to run iterative models to determine the equilibrium state of risk distribution, similar to established practices in regional travel demand modeling. Recent research by Li et al. [19] has demonstrated the feasibility of incorporating safety considerations into network equilibrium frameworks. As drivers alter their routing decisions based on risk information, the dynamics of the entire transportation network may evolve, potentially causing formerly safe segments to become riskier. This feedback mechanism can be modeled as:

$$r_e^{(k+1)} = f\left(r_e^{(k)}, \Delta V_e^{(k)}\right) \tag{20}$$

where $r_e^{(k)}$ represents the risk on segment e at iteration k and $\Delta V_e^{(k)}$ denotes the change in traffic volume resulting from route choice adjustments. Building on Li et al.'s framework [19], which classifies drivers into groups based on acceptable risk thresholds, our approach would need to account for heterogeneous driver populations with varying safety sensitivities. Such an equilibrium-seeking approach would provide more realistic estimates of long-term risk distributions.

6.1. Behavioral Response and Risk Perception

The effectiveness of safety-aware routing for human drivers fundamentally depends on each driver's willingness to modify behavior in response to risk information. Crashes are inherently rare events, and psychological research indicates that most drivers exhibit optimism bias, meaning that they consider themselves to be less vulnerable than the average driver [20]. Future research could include stated preference surveys and behavioral experiments to understand which demographic segments are most responsive to crash risk information and how trip urgency affects risk tolerance (e.g., commuting versus leisure travel). Such behavioral research could inform the development of utility functions that better capture actual driver decision-making, such as:

$$U_{i,j} = -\alpha_i \cdot t_j - \beta_i\left(\text{urgency}\right) \cdot \psi_i\left(p_j\right) \tag{21}$$

where α_i and β_i are individual-specific parameters and $\psi_i(\cdot)$ represents driver i's subjective perception of risk percentile p_j.

Perceptions and behavior likely vary by age and experience, as young and inexperienced drivers are known to exhibit elevated crash risks [21].

6.2. Temporal Variations in Risk

The current implementation treats crash risk as temporally invariant, yet risk profiles exhibit significant variation by time of day and day of week. While GraphHopper, our chosen routing engine, currently lacks support for time-dependent edge weights, future implementations should prioritize this capability. The risk function would be extended to:

$$r_e\left(t,d\right) = \frac{c_e\left(t,d\right)}{\text{DVMT}_e\left(t,d\right) \times T} \times 10^6 \tag{22}$$

where t represents time of day and d denotes day of week. This enhancement would enable routing recommendations that account for rush hour congestion, nighttime visibility limitations, weekend recreational traffic patterns, and other temporal factors affecting crash probability.

6.3. Traffic Volume and Congestion Effects

Although it may seem intuitive that crash severity goes down as congestion increases due to slower traffic, research on the topic is mixed. For instance, Quddus et al. [22] did not find any correlation between levels of congestion and crash severity on highway M25 in London, but follow-up studies that expanded the area of study to a wider region found that severity increased with higher congestion [23]. It was hypothesized that congestion may have caused greater speed differences between vehicles and that this difference may have caused severe incidents. Any correlation between levels of congestion and crash severity is likely nonlinear and would need to be modeled carefully. Future work could incorporate speed-dependent severity models:

$$s_e\left(v_e\right) = \begin{cases} s_{\text{low}} & \text{if } v_e < v_{\text{threshold}} \\ s_{\text{base}} \cdot g\left(v_e\right) & \text{otherwise} \end{cases} \tag{23}$$

where s_e represents expected crash severity, v_e is the average speed on segment e, and $g(\cdot)$ is a function capturing the speed-severity relationship.

6.4. Sensors and Vehicle Collaboration

For CAVs, system-specific factors, including sensor configurations, perception algorithms, and the operational design domain (ODD), can affect safety calculations. For example, the choice and integration of sensor technologies—such as LiDAR, radar, and advanced camera systems—along with associated data processing and machine learning algorithms, directly influence the vehicle's ability to perceive and react to road conditions and hazards [24].

Additionally, the robustness of traffic control devices and the underlying infrastructure further affects operational risk and safety performance [25]. Integrating such factors into risk calculations would likely result in more detailed and accurate assessments of risk.

The availability of vehicle telematics is a promising development both for human drivers and autonomous vehicles. Modern in-vehicle monitoring systems capture granular behavioral metrics—including acceleration patterns, braking intensity, and lane-keeping consistency—that serve as proximate indicators of crash risk [26]. For CAVs, these behavioral metrics could be expanded to include perception confidence scores, disengagement frequencies in specific environments, and sensor degradation patterns under various weather conditions [27]. Such solutions would likely be particularly valuable for CAV fleet operators preparing for full SAE Level 4 or 5 automation. Commercial platforms such as Arity [28] and Cambridge Mobile Telematics [29] have demonstrated the feasibility of large-scale telematics data collection and analysis for human drivers, while analogous platforms are emerging for CAV operational monitoring, offering potential integration points for our framework.

The risk modeling approach could further be enhanced through environmental sensing data from CAV perception systems. Real-time detection of road hazards, visibility conditions [30], and infrastructure degradation captured by autonomous vehicle sensor suites could be aggregated to augment historical crash data with dynamic environmental risk factors.

With a combination of sensor data and telematics, risk maps could be updated continuously to reflect current conditions rather than rely on historical patterns. This capability would be particularly valuable for CAVs operating in edge cases or unusual environmental conditions where historical data alone may prove insufficient.

It is important to acknowledge that the stochastic nature of crashes, including complex interactions between contributing risk factors, creates significant challenges for predictive accuracy. For CAVs specifically, the challenge is compounded by limited real-world crash data and the rapidly evolving nature of autonomous technology [31]. Recent work by Sarigiannis et al. [32] highlighted this difficulty: even sophisticated models with extensive feature engineering achieved high recall (above 90%) only at the expense of precision, with about 70% of high-risk predictions representing false positives.

The V2X (vehicle-to-everything) communication capabilities of CAVs present additional opportunities for enhanced risk modeling. By aggregating near-miss incidents, traffic density information, and infrastructure status updates communicated through V2V (vehicle-to-vehicle) and V2I (vehicle-to-infrastructure) channels, our framework could incorporate network-level intelligence beyond what is possible with isolated vehicle data [33]. Such collaborative solutions would likely be particularly useful in driving situations not covered well by traditional testing methods [34].

6.5. Model Refinements and Technical Improvements

The current ADT prediction model using gradient boosted trees achieves satisfactory performance (R2 = 0.92) but offers opportunities for enhancement in production deployments. Our evaluation of log-transformation of the target variable yielded significant improvements, reducing MAPE from 121% to 65%. This finding suggests that systematic investigation of alternative transformations—including Box–Cox and variance-stabilizing approaches—could further improve prediction stability. Spatial regression models that explicitly incorporate road network connectivity and graph neural networks designed for transportation applications warrant consideration for future implementations.

Technical refinements could also enhance other framework components. The current spatial matching process between road inventory and OpenStreetMap data employs geometric overlap thresholds, but approaches incorporating network topology and attribute similarity could improve matching accuracy. The percentile-based risk normalization, while effective, could benefit from hierarchical modeling that accounts for regional variations in baseline risk levels, potentially providing more refined risk assessments across diverse geographic contexts.

Given that ADT prediction serves as an intermediate step in the broader risk assessment methodology, these enhancements should be evaluated based on their contribution to overall route selection quality rather than isolated prediction performance metrics.

6.6. Severity Classification Validation and Bias Correction

The documented inflation of KABCO severity classifications—with "A" designations overstated in up to 65% of cases according to Burdett et al. [13]—introduces systematic bias into cost-weighted risk calculations that could propagate through route selection algorithms. The economic cost differential between severity levels spans two orders of magnitude (from $11,900 for "O" crashes to $11.3 million for "K" crashes), meaning that misclassification directly affects the relative weighting of road segments in routing decisions.

Future research should investigate the sensitivity of route selection to severity misclassification through controlled simulation studies. By systematically adjusting severity distributions according to documented bias patterns—such as downgrading a percentage of "A" classifications to "B" or "C" levels—researchers can quantify how misclassification affects both segment-level risk rankings and actual route recommendations. This analysis would determine whether the percentile-based transformation provides sufficient robustness or whether bias correction mechanisms are necessary for reliable implementation.

Where jurisdictions maintain linked crash–medical records, comparative studies could validate field severity assessments against clinical injury documentation. Such validation would enable the development of region-specific correction factors that account for local reporting practices. The research would also inform whether KABCO bias exhibits systematic patterns by road type, crash circumstances, or geographic factors that could be incorporated into improved risk assessment models.

6.7. Autonomous Vehicle-Specific Risk Modeling

The current framework's reliance on historical crash data from human-driven vehicles represents a fundamental limitation that future research must address. As CAV market penetration increases and crash data specific to autonomous systems becomes available, the development of differentiated risk models will become both feasible and necessary. Such models would account for:

- Systematic behavioral differences between algorithmic and human decision-making
- CAV-specific failure modes distinct from human error patterns
- Interaction effects between CAVs and human drivers in mixed traffic
- Sensor-dependent risk factors under various environmental conditions

The transition from human-derived to CAV-specific risk metrics will likely require probabilistic frameworks that can accommodate the deterministic nature of autonomous systems while capturing edge case behaviors and system limitations [31].

6.8. Network-Level Effects and Traffic Distribution

Our framework demonstrates a complete, implementable solution for integrating crash risk into navigation decisions. While our primary objective was to establish the feasibility and effectiveness of risk-aware routing, we recognize that large-scale deployment would benefit from additional constraints to manage traffic distribution across the network.

In developing this proof-of-concept, we focused on enabling individual route optimization based on empirically derived risk metrics. This approach successfully demonstrates significant safety improvements (as shown in Section 5) while maintaining computational efficiency and compatibility with existing navigation systems. Notably, our empirical analysis reveals that safety-optimized routes differ only modestly from standard routes—averaging just 369 m longer—suggesting that the algorithm naturally avoids dramatic diversions through residential neighborhoods or other inappropriate corridors. However, system-wide implementation would naturally evolve to incorporate explicit network-level considerations.

Specifically, future deployments could extend our framework to include:

- Time-varying restrictions for sensitive areas (e.g., school zones during arrival/dismissal times)
- Volume thresholds based on functional classification and design capacity
- Land use buffers to preserve neighborhood character

These enhancements would build upon our routing cost function through additional constraint terms:

$$\text{Cost}_e = \alpha \cdot t_e + \beta \cdot p_e \cdot l_e + \gamma \cdot \max\left(0, v_e - C_e\right) \quad (24)$$

where the third term penalizes routes exceeding capacity thresholds C_e. Such extensions align with the equilibrium modeling approaches discussed in this section, where iterative optimization would naturally balance individual route safety with system-wide traffic distribution. Transportation agencies implementing our framework have the flexibility to incorporate local priorities and constraints while leveraging the core risk assessment methodology we have established.

6.9. Commercialization Challenges

The core risk-aware navigation framework operates without requiring personally identifiable information (PII) or real-time user tracking. Risk assessments are computed from anonymized historical crash data and traffic counts, with routing decisions made locally within navigation software. This design minimizes privacy concerns for the fundamental solution.

However, real-time enhancements incorporating live traffic and weather data would introduce additional data requirements and cybersecurity considerations requiring secure API frameworks. Such enhancements would face distinct regulatory challenges across global markets. In the European Union, these enhanced versions would require compliance with GDPR for location tracking, while the EU Data Act mandates fair access to transportation authority datasets. The EU Digital Markets Act may impose interoperability requirements for large-scale deployments, potentially affecting proprietary risk algorithms.

Deployment in emerging markets across the Global South, ASEAN, and Africa faces different challenges primarily related to data availability and reliability. The core requirement for comprehensive crash reporting systems and regular traffic counting may not exist or may lack consistent update intervals necessary for reliable risk assessment. Our framework was primarily designed with the robust data infrastructure of the United States in mind, where crash reports and traffic inventories are systematically maintained and regularly updated.

Across all markets, automotive functional safety standards (ISO 26262) require systematic validation of safety claims, while product liability frameworks remain undefined for navigation services, making explicit safety recommendations.

7. Conclusion

This article introduces a novel way to account for crash risk in navigation software. We showed how to use road inventory data combined with crash reports by LEOs to develop risk assessments that we then used to minimize crash risk when selecting a route for travel. By formalizing the relationship between exposure-normalized crash rates and routing decisions, we have demonstrated the feasibility of navigation that balances traditional metrics of efficiency with quantifiable safety considerations—a critical advancement for connected and automated vehicles as well as manually driven vehicles.

The methodology presents several key contributions to the field of connected and automated vehicle technology and transportation safety:

First, our percentile-based risk ranking approach provides a standardized, interpretable measure of relative safety across heterogeneous road networks that can be seamlessly integrated into CAV decision-making systems. This transformation mitigates the challenges posed by outliers and data sparsity while maintaining an intuitive scale that can be effectively communicated to both autonomous systems and human users through appropriate interfaces. For CAVs specifically, this standardized risk metric can serve as a validation parameter for ODD assessment and safety assurance frameworks.

Second, the routing cost function (Equation 11) offers a flexible mechanism for integrating safety considerations into existing CAV path planning frameworks without sacrificing computational efficiency. The parameterization through α and β coefficients enables customization based on operational needs, safety requirements, and vehicle-specific performance characteristics. The adaptability is useful for CAV fleet operators seeking to optimize routing systematically but also for drivers with different risk profiles and preferences.

Third, our case study using Ohio's road network and crash data demonstrates a real-world implementation of the solution. The methodology is generalizable to any region with available crash and traffic data, requiring only minimal adaptation to local data structures and road classification systems. The scalability of our solution is a key feature that sets it apart from other studies constrained to small areas or particular types of infrastructure.

The implementation using the open-source GraphHopper routing library with OpenStreetMap data illustrates how the theoretical framework can be translated into a functioning navigation system compatible with existing CAV software architectures. This implementation pathway could be adopted by both commercial CAV developers and transportation agencies seeking to enhance both autonomous and human-driven navigation safety. By providing risk-aware navigation capabilities, the system supports functional safety requirements for SAE Level 3+ automation while simultaneously offering benefits for conventional vehicles.

As autonomous vehicles transition from testing to widespread deployment and connected vehicle technologies become standard across the transportation ecosystem, the incorporation of safety metrics into route planning represents a significant opportunity for crash reduction and enhanced CAV reliability. For autonomous vehicles specifically, risk-aware navigation can serve as an additional safety layer beyond immediate perception-based obstacle avoidance, potentially addressing edge cases and long-tail risks that challenge current validation approaches. Future refinements to this approach—particularly those addressing the spatiotemporal dynamics of risk, V2X information integration, and system-specific risk profiles—could further enhance its impact on autonomous vehicle safety verification and validation.

The goal of this study is to contribute to the broader goal of creating safer automated and connected transportation systems while developing solutions that can be used broadly including by human drivers. By bridging current and future mobility paradigms, we provide a practical pathway toward the incremental safety improvements necessary for public acceptance and regulatory compliance of autonomous vehicle technology.

Declaration of AI-Assisted Technologies

During the preparation of this work the author(s) used Claude 3.7 Sonnet to assist with LaTeX document preparation. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the publication.

Acknowledgements

The authors thank the Ohio Department of Transportation for providing access to crash data.

Contact Information

Lars Skaug, corresponding author
lskaug2022@fau.edu

References

1. Hu, Q., Mehdizadeh, A., Vinel, A., Cai, M. et al., "Shortest Path Problems with a Crash Risk Objective," *Transportation Research Record* 2678, no. 6 (2023): 284-300, doi:https://doi.org/10.1177/03611981231195053.
2. Mannering, F.L. and Bhat, C.R., "Analytic Methods in Accident Research: Methodological Frontier and Future Directions," *Analytic Methods in Accident Research* 1 (2014): 1-22, doi:https://doi.org/10.1016/j.amar.2013.09.001.
3. Skaug, L., Nojoumian, M., Dang, N., and Yap, A., "Road Crash Analysis and Modeling: A Systematic Review of Methods, Data, and Emerging Technologies," *Applied Sciences* 15, no. 13 (2025), doi:https://doi.org/10.3390/app15137115.
4. Hauer, E., "Empirical Bayes Approach to the Estimation of "Unsafety": The Multivariate Regression Method," *Accident Analysis & Prevention* 24, no. 5 (1992): 457-477.
5. Ghoul, T., Sayed, T., and Fu, C., "Real-Time Safest Route Identification: Examining the Trade-Off between Safest and Fastest Routes," *Analytic Methods in Accident Research* 39 (2023): 100277, doi:https://doi.org/10.1016/j.amar.2023.100277.
6. Nojoumian, M. and Skaug, L., Road-risk awareness system (ras) in semi or fully autonomous vehicles. U.S. Patent No. 19/016,240, 2025.
7. American Association of State Highway and Transportation Officials, *Highway Safety Manual* (Washington, DC: American Association of State Highway and Transportation Officials, 2010). https://www.highwaysafetymanual.org/.
8. Harmon, T., Bahar, G., and Gross, F., "Crash Costs for Highway Safety Analysis," Technical Report FHWA-SA-17-071, Federal Highway Administration, Washington, DC, 2018, https://highways.dot.gov/sites/fhwa.dot.gov/files/2022-09/fhwasa17071.pdf.
9. Huang, H., Han, C., Xu, G., Jiang, M. et al., "Incorporating Safety Reliability into Route Choice Model: Heterogeneous Crash Risk Aversions," *Analytic Methods in Accident Research* 25 (2020): 100112, doi:https://doi.org/10.1016/j.amar.2020.100112.
10. Han, C., Xu, G., Pervez, A., Gao, F. et al., "Modeling Traveler's Speed-Route Joint Choice Behavior with Heterogeneous Safety Concern," *Analytic Methods in Accident Research* 37 (2023): 100253, doi:https://doi.org/10.1016/j.amar.2022.100253.
11. Skaug, L., "Graphhopper: Custom Graphhopper Fork," 2025, accessed April 21, 2025, https://github.com/larsskaug/graphhopper.
12. Ohio Department of Transportation, "Road Inventory Dataset," accessed May 1, 2025, https://tims.dot.state.oh.us/tims/data-glossary/dataset/8999c0f518cc46b99b454c9fa51ce409.
13. Burdett, B., Bill, A., and Noyce, D., "Evaluation of Law Enforcement Agency Injury Severity Assessments," *Transportation Research Record* 2676, no. 9 (2022): 246-255, doi:https://doi.org/10.1177/03611981221086628.
14. Ahmed, A., Sadullah, A.F.M., and Yahya, A.S., "Errors in Accident Data, Its Types, Causes and Methods of Rectification-Analysis of the Literature," *Accident Analysis & Prevention* 130 (2019): 3-21, doi:https://doi.org/10.1016/j.aap.2017.07.018.
15. Skaug, L. and Nojoumian, M., "Llm-Assisted GIS Validation: A Computational Method for Improving Traffic Crash Location Data, Manuscript in Preparation," 2025.
16. Flügel, S., Elvik, R., Veisten, K., Rizzi, L.I. et al., "Asymmetric Preferences for Road Safety: Evidence from a Stated Choice Experiment among Car Drivers," *Transportation Research Part F: Traffic Psychology and Behaviour* 31 (2015): 112-123, doi:https://doi.org/10.1016/j.trf.2015.04.001.
17. American Association of State Highway and Transportation Officials, "Census Transportation Planning Products (CTPP) Datasets," accessed April 6, 2025, https://transportation.org/ctpp/datasets/.
18. Overture Maps Foundation, "Overture Provides Free and Open Map Data," accessed April 6, 2025, https://overturemaps.org/.
19. Li, M.-S. and Huang, H.-L., "Road-Safety Recognition and Network Equilibrium with Perceived Route-Choice Sets," *Transportation Safety and Environment* 1, no. 2 (2019): 126-134, doi:https://doi.org/10.1093/tse/tdz009.
20. Măirean, C., Havârneanu, G.M., Barić, D., and Havârneanu, C., "Cognitive Biases, Risk Perception, and Risky Driving Behaviour," *Sustainability* 14, no. 1 (2022): 77, doi:https://doi.org/10.3390/su14010077.
21. Williams, A.F., "Young Driver Risk Factors: Successful and Unsuccessful Approaches for Dealing with Them and an

Agenda for the Future," *Injury Prevention* 12, no. S1 (2006): i4-i8, doi:https://doi.org/10.1136/ip.2006.011783.

22. Quddus, M.A., Wang, C., and Ison, S.G., "Road Traffic Congestion and Crash Severity: Econometric Analysis Using Ordered Response Models," *Journal of Transportation Engineering* 136, no. 5 (2010): 424-435, doi:https://doi.org/10.1061/(ASCE)TE.1943-5436.0000044.
23. Wang, C., Quddus, M., and Ison, S., "A Spatio-Temporal Analysis of the Impact of Congestion on Traffic Safety on Major Roads in the UK," *Transportmetrica A: Transport Science* 9, no. 2 (2013): 124-148, doi:https://doi.org/10.1080/18128602.2010.538871.
24. Khanmohamadi, M. and Guerrieri, M., "Advanced Sensor Technologies in CAVS for Traditional and Smart Road Condition Monitoring: A Review," *Sustainability* 16, no. 19 (2024): 8336, doi:https://doi.org/10.3390/su16198336.
25. Gopalakrishna, D., Carlson, P., Sweatman, P., Raghunathan, D. et al., "Impacts of Automated Vehicles on Highway Infrastructure," Technical Report FHWA-HRT-21-015, Federal Highway Administration, U.S. Department of Transportation, Washington, DC, 2021, https://www.fhwa.dot.gov/publications/research/operations/21015/21015.pdf.
26. Bell, J.L., Taylor, M.A., Chen, G.X., Kirk, R.D. et al., "Evaluation of an in-Vehicle Monitoring System (IVMS) to Reduce Risky Driving Behaviors in Commercial Drivers: Comparison of In-Cab Warning Lights and Supervisory Coaching with Videos of Driving Behavior," *Journal of Safety Research* 60 (2017): 125-136, doi:https://doi.org/10.1016/j.jsr.2016.12.008.
27. Haseltalab, A. and Negenborn, R.R., "Sensing, Modeling and Decision Making for Autonomous Vehicles: How Much Information Is Enough?" *Annual Reviews in Control* 49 (2020): 223-236, doi:https://doi.org/10.1016/j.arcontrol.2020.09.001.
28. Arity, "Arity: Driving Change with Driving Data," accessed April 6, 2025, https://arity.com/.
29. Cambridge Mobile Telematics, "Cambridge Mobile Telematics—The World Leader in Telematics," accessed April 6, 2025, https://www.cmtelematics.com/.
30. Guo, Y., Li, Z., Wu, Y., and Xu, C., "Study on the Influence of Sun Glare on Driving Safety," *Building and Environment* 228 (2023): 109902, doi:https://doi.org/10.1016/j.buildenv.2022.109902.
31. Kalra, N. and Paddock, S.M., "Driving to Safety: How Many Miles of Driving Would It Take to Demonstrate Autonomous Vehicle Reliability?" *Transportation Research Part A: Policy and Practice* 94 (2016): 182-193, doi:https://doi.org/10.1016/j.tra.2016.09.010.
32. Sarigiannis, D., Atzemi, M., Oke, J., Christofa, E. et al., "Feature Engineering and Decision Trees for Predicting High Crash-Risk Locations Using Roadway Indicators," *Transportation Research Record* 2678, no. 8 (2024): 535-548, doi:https://doi.org/10.1177/03611981231217497.
33. Jung, C., Lee, D., Lee, S., and Shim, D.H., "V2x-Communication-Aided Autonomous Driving: System Design and Experimental Validation," *Sensors* 20, no. 10 (2020): 2903, doi:https://doi.org/10.3390/s20102903.
34. Rahmani, S., Rieder, S., de Gelder, E., Sonntag, M. et al., "A Systematic Review of Edge Case Detection in Automated Driving: Methods, Challenges and Future Directions," arXiv Preprint arXiv:2410.08491, 2024, doi:https://doi.org/10.48550/arXiv.2410.08491.

Appendix A: Supplementary Crash Risk Data

This appendix provides detailed metrics for absolute crash risk across all surveyed road classes with more than 100 segments.

TABLE A.1 Absolute crash risk metrics by road class.

Road Class	Segment count	Mean	Median	P25	P75
Residential	84,299	11,638.43	562.01	171.62	2683.02
Tertiary	26,174	25,691.63	1917.10	313.11	12,321.43
Secondary	17,738	18,465.81	750.95	152.78	4797.85
Primary	13,512	6305.76	530.71	124.15	2386.09
Footway	12,214	1074.40	332.84	142.29	879.46
Motorway	9831	5286.28	884.15	197.24	3903.90
Service	9663	1577.30	248.09	91.39	786.54
Unclassified	6771	31,260.41	6461.54	915.01	27,269.66
Motorway link	6277	661.12	278.64	110.19	691.07
Trunk	3959	7496.57	963.96	211.64	4657.65
Cycleway	629	1176.93	260.96	88.94	935.04
Trunk link	393	1101.63	304.14	101.38	1079.12
Primary link	309	351.31	87.11	28.45	291.59
Secondary link	174	374.84	46.94	17.30	196.56
Tertiary link	134	2057.73	59.37	14.56	308.45

Crash rates per million vehicle miles traveled (VMT).
Only road classes with >100 segments shown.

Appendix B: Visual Route Comparison

Figure B.1 demonstrates the practical implementation of risk-aware routing using our GraphHopper modification. The route pair shows a representative trade-off: the safer route reduces the risk score from 0.433 to 0.281 (a drop of 15 percentile points) at the cost of an additional 1.3 km distance and 89-second time penalty.

FIGURE B.1 Route comparison showing different path selections between risk-aware and default routing algorithms for the same origin–destination pair.

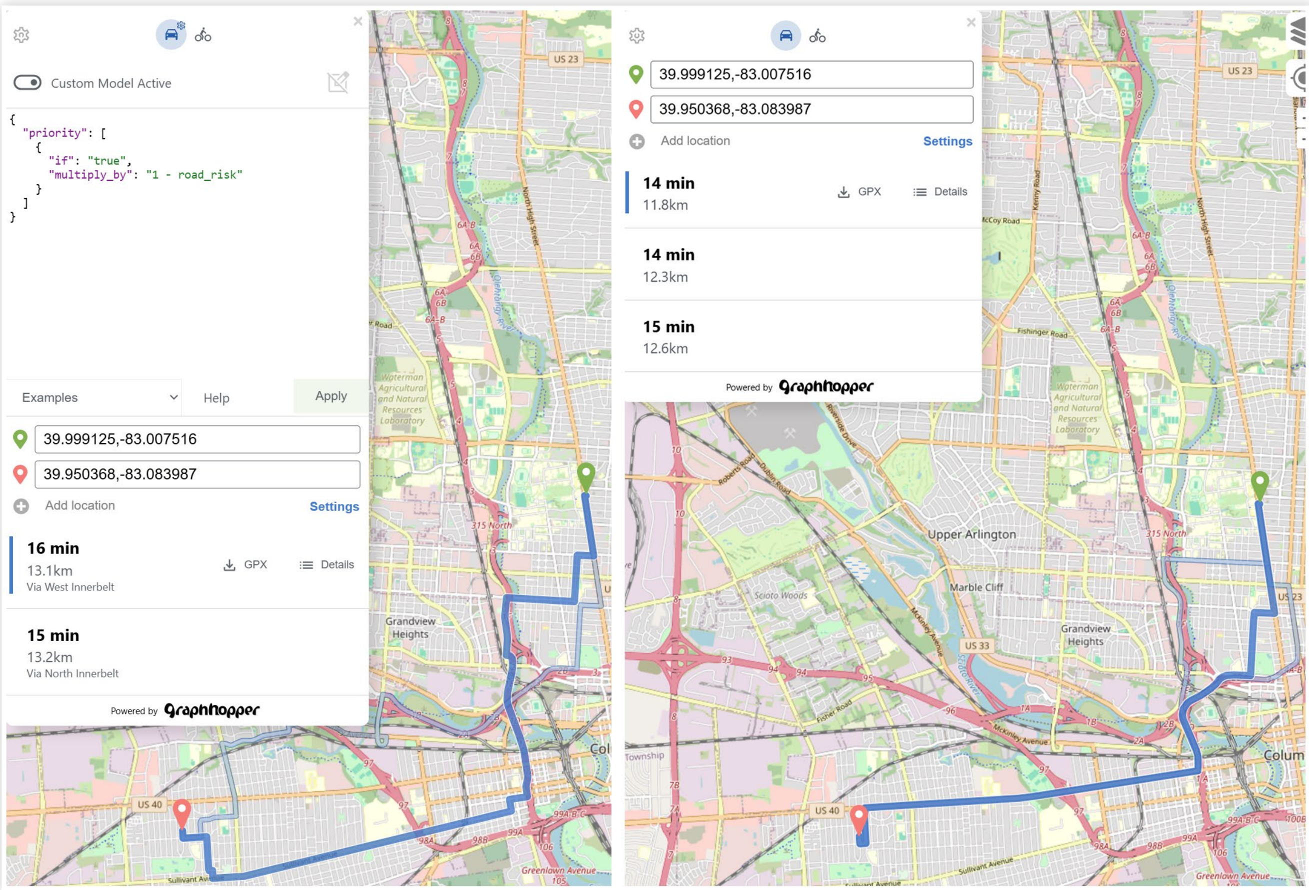

© SAE International

Positions and opinions advanced in this work are those of the author(s) and not necessarily those of SAE International. Responsibility for the content of the work lies solely with the author(s).