

Car-to-Car Injury Risk Functions for Real-Time Emergency Decision-Making in Highly Automated Vehicles

Giulio Vichi

giulio.vichi1@unifi.it

University of Florence

Michelangelo-Santo Gulino

University of Florence

Anita Fiorentino

Stellantis Europe

Cyril Chauvel

Laboratory of Accidentology, Biomechanics and Driving Behaviour (LAB)

Laura Bigi

Laboratory of Accidentology, Biomechanics and Driving Behaviour (LAB)

Sophie Cuny

European Centre of Studies on Safety and Risk Analysis (CEESAR)

Dario Vangi

University of Florence

Article

Keywords:

Posted Date: September 15th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10588740/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: Competing interest reported. This study received financial support from Stellantis. A.F. is an employee of Stellantis Europe, C.C. and L.B. are employees of Stellantis Auto, and G.V. receives a PhD scholarship co-funded by Stellantis. Stellantis reviewed and approved the manuscript prior to submission. The authors declare that the funding and institutional affiliations did not

influence the study design, data analysis, interpretation of the results, or the decision to publish. The remaining authors declare no competing financial or non-financial interests.

Car-to-Car Injury Risk Functions for Real-Time Emergency Decision-Making in Highly Automated Vehicles

Giulio Vichi^{1,*}, Michelangelo-Santo Gulino¹, Anita Fiorentino², Cyril Chauvel³, Laura Bigi³, Sophie Cuny⁴, and Dario Vangi¹

¹University of Florence, Department of Industrial Engineering (DIEF), Florence, 50139, Italy

²Stellantis Europe, Product Development & Technology, Tech Research, Pomigliano d'Arco, 80038, Italy

³Laboratory of Accidentology, Biomechanics and Driving Behaviour (LAB), Nanterre, 92000, France

⁴European Centre of Studies on Safety and Risk Analysis (CEESAR), Nanterre, 92000, France

*giulio.vichi1@unifi.it

ABSTRACT

Traditional active safety systems focus on collision avoidance, but unavoidable impacts require real-time occupant injury risk (IR) mitigation. Existing car-to-car injury risk functions (IRFs) rely on crash-phase metrics like velocity change (ΔV) and Energy Equivalent Speed (*EES*), requiring computationally intensive simulations that violate millisecond on-board decision budgets. Furthermore, most are trained on single-region databases, limiting their geographical transferability. To bridge this gap, this study proposes a framework to derive omni-directional, continuous IRFs based exclusively on pre-crash variables known at the first instant of contact. Optimal predictors are identified through a two-stage feature selection process, and the resulting functions are cross-validated on three independent accident databases. Application to the IGLAD, NASS-CDS, and VOIESUR datasets reveals excellent cross-dataset discrimination for severe (MAIS 3+) and fatal (MAIS 6) injury thresholds, with Area Under the Curve (AUC) metrics consistently exceeding 0.80, whilst minor injury prediction remains limited by occupant-intrinsic parameters unknown to the vehicle control unit. Application to real-world car-to-car crash test data confirms straightforward model applicability and validates viability under imperfect data regimes. By omitting crash-phase simulations entirely, this approach permits purely kinematic evaluations of autonomous steering and braking interventions at each sensor update, selecting the manoeuvre that minimises occupant IR. Ultimately, this framework provides a concrete technological foundation for real-time injury mitigation in automated vehicles.

Introduction

The landscape of automotive safety has undergone a profound transformation, evolving from historical enhancements in structural crashworthiness to sophisticated Advanced Driver Assistance Systems (ADAS) that use exteroceptive sensor suites for environmental perception¹. Contemporary ADAS interventions, such as Autonomous Emergency Braking (AEB), Adaptive Cruise Control (ACC) and Lane Keeping Assist (LKA), actively control vehicle dynamics to prevent collisions or mitigate human error², which remains the primary cause for approximately 94% of road traffic accidents^{3,4} and 1.19 million deaths worldwide in 2021⁵. The rise of Automated Driving (AD) further promises to enhance transport safety by potentially eliminating performance deviations such as distraction or delayed reaction times^{3,6-8}. Despite these advancements, the predominant objective for both longitudinal and lateral emergency interventions in ADAS and AD systems has focused on collision avoidance or proactive prevention of high impact risk conditions as the primary metric of safety performance⁹⁻¹¹. Empirical evidence suggests that even with universal market penetration of current AEB technologies, a significant proportion of accidents would be residual^{12,13}; specifically, even under an ideal 100% penetration rate, at least 20% of intersection collisions would still be unavoidable¹⁴. Moreover, real-world conflict scenarios frequently transcend the functional limits of current avoidance-based algorithms^{6,9}, whether based on braking or steering, due to fundamental constraints including sensing depth, sensor update rates, and reaction time budgets. This is further exacerbated by sensor measurement uncertainties, which necessitate the adoption of a conservative 0.95 confidence level to define reliable safe avoidance regions, alongside the inherent dynamic limitations of the vehicle¹⁵. These factors result in the persistence of Inevitable Collision States (ICS) requiring a strategic paradigm shift: when avoidance is not possible, the safety objective must focus on the systematic minimisation of occupant injury risk (IR)¹⁶. In this context, Autonomous Emergency Steering (AES) and related AD emergency interventions offer a strategic dimension for IR mitigation^{9,17}. Key severity predictors in car-to-car conflicts, such as the velocity change experienced by the individual vehicle in the impact (ΔV) and the Energy Equivalent Speed (*EES*)¹⁸⁻²¹, are functions of both impact

speed and collision geometry (eccentricity)^{22,23}. Unlike braking, which primarily reduces collision and closing speed (V_r), steering fundamentally alters the impact configuration and intrusion patterns^{6,24}. By increasing impact eccentricity, steering interventions can effectively attenuate these predictors, thereby lowering occupant IR²⁵. Specifically, quantitative evaluations in critical lane-departure scenarios have shown that, even under ideal sensing and instantaneous X-by-wire activation, risk-oriented steering manoeuvres can reduce the average severe injury risk for ego-vehicle occupants from 25.7% to 0.9%²⁶. Under identical evaluation conditions, this significantly outperforms pure AEB interventions, which yield a higher average residual risk of 12.5%²⁶. However, manoeuvres focused solely on avoidance rather than IR mitigation may inadvertently transition occupants toward more hazardous impact configurations; for instance, shifting from a frontal to a near-side car impact entails a drastically higher probability of serious injury at equivalent or even lower ΔV levels^{20,27}. Leveraging autonomous steering thus necessitates advanced probabilistic IR assessment within the pre-crash environment, updating synchronously with exteroceptive sensors. By evaluating braking and steering trajectory combinations, the system may identify the optimal manoeuvres that converge on the minimum IR configuration when avoidance is unfeasible^{6,26}. Nevertheless, the real-time execution of such risk-based logic is currently hindered by existing IR estimation methodologies, which were not developed to support instantaneous on-board assessments.

Established injury risk functions (IRFs) are primarily formulated to support prospective evaluation tools, utilising crash-phase metrics that describe the outcome of an impact rather than its initial conditions. Velocity-based parameters (such as ΔV) and energy-based metrics (such as *EES*) represent the most well-established predictors for estimating injury severity in car-to-car collisions^{20,21,28–31}. These parameters, however, remain unmeasurable during the emergency pre-crash phase, necessitating either detailed retrospective accident reconstruction or computationally intensive virtual forward simulations (VFS) to estimate impact dynamics^{18,32,33}. Current VFS techniques for car-to-car collisions exhibit a prohibitive conflict between computational demand and predictive precision. At the lower end of the complexity spectrum, impulse-momentum models assume an infinitesimal crash duration to resolve momentum conservation equations³⁴. While efficient, these models inherently neglect structural energy absorption; consequently, they fail to provide energy-based parameters and yield non-negligible errors in ΔV across numerous impact configurations^{34,35}. Pure multibody (MBD) systems offer higher computational efficiency than high-fidelity structural simulations, yet they typically treat vehicles as rigid bodies, precluding deformation analysis without onerous supplementary modules³⁶. More balanced methodologies, such as Reduced Order Dynamic Models (RODM), ensure high accuracy for both velocity-based and energy-based metrics^{35,37}. However, although significantly faster than the prohibitive durations of non-linear Finite Element Method (FEM) analysis, RODM solution times—ranging from 1.5 seconds to 5 minutes—remain incompatible with the millisecond reaction budgets of active safety systems^{16,37}. To bridge this gap, recent hybrid frameworks have combined real-time MBD simulations with pre-compiled databases to achieve latencies as low as 80 ms³⁸. Nevertheless, such approaches necessitate detailed structural models for every potential collision partner, rendering them unfeasible for on-board deployment in the uncatalogued complexity of real-world environments.

Beyond computational barriers, the on-board suitability of current IRFs is further hindered by the regional nature of their training databases. Most established functions derive from single-region datasets reflecting specific vehicle fleets and localised reporting procedures^{13,14,21,25}. This reliance on singular data limits geographical generalisability, as significant discrepancies exist between countries regarding crashworthiness standards and vehicle distributions^{25,39,40}. Consequently, such models often lack the robustness required for reliable cross-regional deployment. This limitation is further compounded by the mathematical fragmentation of IRFs into discrete crash modes or discretized impulse directions, typically defined by specific Principal Directions of Force (PDOF) or the 12-category Collision Deformation Classification (CDC) system^{41–44}. Such angular segmentation introduces mathematical discontinuities when evaluating the continuous collision space required for automated trajectory planning⁴⁵. Moreover, even existing omni-directional frameworks frequently rely on occupant-specific covariates, such as age and sex, which lack immediate applicability for real-time on-board deployment⁴⁵.

To address these limitations, this study proposes a novel methodological framework for deriving IRFs tailored for real-time, on-board emergency decision-making. The approach ensures geographical generalisability through cross-validation across diverse accident databases and achieves input (predictor) objectivization via a two-stage feature selection process incorporating model-agnostic feature ranking. This methodology is applied to develop three distinct, omni-directional continuous IRFs for minor, severe, and fatal injury outcomes based exclusively on pre-crash variables. Finally, the practical viability of the framework is demonstrated through its application to real-world car-to-car crash test data, validating the applicability of purely pre-crash predictors in real contexts. Serving as a fully resolved physical reference where all experimental parameters are known, this crash test defines a realistic baseline to achieve an initial quantification of the estimation errors propagated when a real-time, on-board decision unit lacks advanced and exhaustive information regarding the opponent vehicle. By restricting the predictive computation to pre-crash kinematics even when resolving unavoidable impact configurations, the developed methodology is designed to eliminate the need for crash-phase simulations, thereby offering a computationally efficient foundation for real-time mitigation strategies.

On-board real-time injury risk estimation strategy

To contextualise the strict operational constraints that the derived IRFs must satisfy within a vehicle-embedded environment, this section outlines the fundamental rationale governing the closed-loop control sequence designed for on-board implementation, as illustrated in Figure 1. Rather than formulating this control loop from the ground up, the proposed workflow adapts and extends the model-in-the-loop Adaptive Logic concept originally conceptualised by Vangi et al.³⁷, utilising it specifically as the host deployment environment to define the boundary requirements for the IRFs. Under this integrated framework, when a vehicle enters an active emergency phase, the environment is continuously scanned via an exteroceptive sensor suite to feed the imminent crash detection and trajectory generation modules. These modules process available physical escape paths in parallel, including emergency braking and steering and coordinated longitudinal and lateral X-by-wire interventions.

As outlined in the architecture, the control flow bifurcates at a critical decision node based on whether a collision-free trajectory is available. If the trajectory planner identifies escape paths that guarantee absolute collision avoidance while maintaining global vehicle stability, preventing secondary collisions or road departure, a collision avoidance strategy is executed. Under this regime, to maximize the safety margins, the control logic selects the specific trajectory that maximizes the minimum clearance distance from the opponent vehicle, thereby favouring the path that inherently provides the largest safety buffer^{16,26}. Conversely, if no physically feasible path can avert an impact under safe operating conditions, the vehicle enters an ICS, triggering the collision mitigation branch. Within this mitigation regime, the safety objective shifts from geometric clearance to the systematic minimization of injury probability for the occupants. For each candidate trajectory leading to an unavoidable impact, the potential consequences are estimated using pre-crash parameters, completely omitting any simulation of the collision phase. The vehicle then identifies the optimal emergency manoeuvre by selecting the path that minimizes the occupant IR, directly using the calculated IR value as an algorithmic cost function.




Figure 1. Functional block diagram of the proposed real-time, on-board emergency decision-making architecture.

To operationalise this mitigation logic, the vehicle control loop must execute the IR computations for a set of possible interventions within strict execution limits determined by perception hardware update rates. Once an emergency scenario is identified, the entire decision-making pipeline depicted in Figure 1 must execute iteratively immediately after a sensor scan update and before the subsequent scan begins. Depending on the vehicle hardware profile, these sensor update intervals usually range from 0.05 seconds for high-performance perception suites to 0.2 seconds for standard setups. Within these restricted temporal windows, it is computationally prohibitive to estimate traditional crash-phase parameters, such as ΔV or EES , which underpin established car-to-car IRFs. Resolving impact dynamics to extract these crash-phase parameters would require simulating the physical collision phase. While low-fidelity impulse-momentum formulations could mathematically fit within the available time budget, they inherently neglect structural energy absorption and exhibit significant errors in ΔV predictions, which become especially severe in complex and highly eccentric collision configurations. Such predictive inaccuracies would inevitably propagate into the subsequent IR calculation, severely compromising the reliability of the decision³⁵. Consequently, real-time on-board IR mitigation requires bypassing crash-phase or structural deformation simulations entirely, relying exclusively on pre-crash variables. The trajectory planner projects the vehicle pre-crash kinematics up to the virtual point of impact to predict the expected contact configuration. To successfully estimate IR using only pre-crash predictors, the architecture requires highly reliable IRFs. These functions must be geographically independent to ensure they can be utilized anywhere regardless of the specific training datasets, and they must be omni-directional to remain mathematically applicable regardless of the impact direction. Within this operational framework, IRFs predicting the probability of sustaining at least severe injuries (MAIS3+) are well suited for on-board deployment. They provide an objective, clinically relevant, and continuous metric of severe trauma, allowing the selection module to reliably differentiate between candidate crash trajectories and immediately command the optimal mitigation inputs before the physical impact occurs.

Materials and methods

To operationalise the on-board IR mitigation strategy outlined in the previous section, this study establishes a structured, reproducible workflow for deriving and validating pre-crash predictor-based IRFs. The methodological framework for IRF derivation and validation is established by detailing the general statistical procedures for data preparation and harmonisation, two-stage feature selection, and model training and validation. This framework is formulated to be database-independent, capable of being applied to any trio of geographically diverse in-depth accident datasets and adapted to any target injury threshold. To populate this statistical framework, the general methodology is applied to the three specific accident databases utilised in this work, targeting occupant injury severity levels focusing on minor, severe, and fatal outcomes.

Methodological framework for IRF derivation

The methodological framework for deriving and validating the IRFs is structured into four consecutive operational steps, as mapped out in the functional block diagram in Figure 2a, with the corresponding cross-validation matrix detailed in Figure 2b. The primary objective of this framework is to obtain IRFs that rely exclusively on pre-crash predictors and are robustly cross-validated across entirely independent accident databases, ultimately returning the probability of sustaining an injury of a designated severity level for the most severely injured occupant of an individual vehicle involved in a car-to-car collision.

To evaluate and demonstrate its practical performance, this general framework is applied in the present study to the IGLAD, NHTSA NASS-CDS, and VOIESUR databases.

Data preparation and accident databases

The initial data preparation and harmonisation phase involves isolating impacts from three geographically independent in-depth accident datasets based on objective inclusion criteria. The sample is restricted exclusively to collisions involving exactly two passenger cars, excluding secondary impacts, and requiring that at least one occupant within the crash sustained an injury. Cases involving the non-use or misuse of seatbelts are omitted due to their confounding effects on the baseline probability of injury²⁵. Rollover events are likewise excluded because of their complex, three-dimensional dynamics and the dominant role that subsequent secondary impacts play in causing injuries⁴⁶. Within this filtered sample, candidate variables are restricted to pre-crash parameters representing conditions known at the exact point of first contact between the two vehicles. Crucially, any pre-crash parameters that are not accessible by the vehicle control unit, such as occupant age, height, and sex, are excluded to ensure that the derived IRFs remain straightforwardly applicable using only information available on-board. The remaining predictors must either be encoded across all target databases or derivable through deterministic relationships from existing variables, provided the underlying parameters are recorded. To ensure data reliability, these remaining observations should be filtered to remove cataloguing errors⁴⁷ or internal inconsistencies⁴⁸. The framework handles both continuous variables, which undergo normalisation, and categorical variables, which are structured via sum coding.

These standardised screening procedures were systematically applied across the three distinct, geographically independent datasets used in this study. The first, the Initiative for the Global harmonization of Accident Data (IGLAD) database, is an




Figure 2. (a) Flowchart outlining the four sequential operational stages from initial data preparation and predictor selection to model building and validation strategy. (b) Schematic of the external cross-validation matrix, illustrating the parallel application of logit, probit, and complementary log-log link functions across three independent accident datasets evaluated via auto-testing (dashed arrows) and cross-testing (solid arrows) to establish model generalisability.

international repository collecting in-depth accident data across 13 different countries. The standardised data scheme comprises 127 variables organised into four relational tables: accidents, participants, occupants, and safety systems⁴⁹. From an initial extraction pool of 3,380 collisions involving exactly two passenger cars between 2013 and 2023, the data filtering process yielded 4,624 valid observations (from 2312 collisions). The National Automotive Sampling System-Crashworthiness Data System (NASS-CDS), curated by the US National Highway Traffic Safety Administration (NHTSA), serves as a primary benchmark for road safety studies. Individual crash impacts within this dataset are statistically sampled to ensure they are representative of real-world collisions occurring across the United States. A total of 13,452 collisions involving exactly two passenger cars and occurring between 2009 and 2015 were initially targeted. Upon implementing the procedural filters defined in the data preparation phase, the active sample was refined to 3,287 observations. The French in-depth accident database VOIESUR was utilised as the third independent data source. This dataset was established within the context of the VOIESUR project launched in 2012, funded in part by the French National Research Agency (ANR), and conducted through the collaboration of four major road safety research institutions in France: LAB, CEREMA, IFSTTAR, and CEESAR. The database was compiled from the comprehensive analysis of more than 8,500 French police reports from the year 2011, capturing 100% of all fatal accidents and a 5% sample of all injury-producing accidents nationwide. Consequently, the VOIESUR database is representative of national traffic injury patterns for that specific year upon the application of case-specific statistical weighting coefficients⁵⁰. Within this dataset, the initial sample of car-to-car conflicts comprised 825 collisions, which are representative of 8,292 collisions once the statistical weighting coefficients are applied. The data filtering process yielded 998 VOIESUR valid observations (from 499 collisions).

Regarding the target output variable for both predictor selection and model building purposes, the Maximum Abbreviated Injury Scale framework is utilised, where the notation MAISX+ defines a binary outcome indicating whether an injured occupant of a vehicle sustained an injury severity equal to or greater than the threshold level X. While injury severity can be evaluated for discrete body regions, such as the head, thorax, or pelvis, this framework considers the global MAIS indicator to capture overall trauma severity. The target injury severities selected for analysis encompass at least light (MAIS 1+), at least severe (MAIS 3+), and fatal (MAIS 6) outcomes, determined by the maximum injury severity sustained by the most severely injured occupant of each vehicle. The corresponding vehicle-level injury distributions for these active samples are illustrated in Figure 3.

To populate the statistical framework, a homogeneous vector of pre-crash predictors common to all three databases was established, either through the direct retrieval of encoded variables or via deterministic physical relationships. These shared attributes encompass the *impact configuration* (classified as frontal, rear, near-side, or far-side⁵¹), *vehicle kerb mass*, *collision speed*, *closing speed* (V_r) at the beginning of the collision, *occupant seating row*, and *vehicle segment*. Furthermore, the common vector incorporates the Crash Momentum Index (CMI)²³ and the Crash Severity Index (CSI)²².

The specific operational derivation of these parameters when not encoded varies according to the parameters available within

SEVERITY DISTRIBUTION OF THE DATASETS




Figure 3. Vehicle-level injury severity distribution in car-to-car collisions across the IGLAD, NASS-CDS, and VOIESUR samples. Percentages represent at least light, at least severe, and fatal outcomes, determined by the maximum injury severity of the occupants for each vehicle. The share of vehicles with no injured occupants is the complement of the at least light injury distribution.

each data source. For the IGLAD dataset, the initial closing speed, V_r , was reconstructed utilising the encoded magnitudes of the collision speeds of both vehicles and their corresponding principal directions of force (PDOF)⁵². This reconstruction relies on the physical assumption that the initial closing speed vector aligns with the PDOF, which applies to all cases where no significant sliding occurs, representing the most severe crash configurations²³. The corresponding CMI and CSI metrics were then determined from their governing relationships²², where the former is expressed by equation (1)

$$CMI = \frac{\Delta V}{V_r} \quad (1)$$

and the latter is formulated via equation (2)

$$CSI = \frac{EES}{V_r} \quad (2)$$

As concerns the NASS-CDS dataset, the values of V_r were computed analytically for each case based on the respective vehicle masses and the structural energy absorbed during the deformation phase²⁵, which is encoded within the database as the *EES*. The corresponding *CMi* and *CSi* predictors were subsequently derived using the identical mathematical formulations (equations 1 and 2) applied to the IGLAD dataset. Finally, the VOIESUR dataset natively provides the baseline V_r values, allowing the kinematic predictors *CMi* and *CSi* to be obtained directly using the same physical definitions (equations 1 and 2).

Predictor selection

Predictor selection is conducted through a two-stage filter-wrapper approach, where the initial screening phase is executed independently of the underlying risk model structure. In the first stage, the model-agnostic ReliefF algorithm, which is a widely recognised variant within the Relief class of algorithms originally proposed by Kononenko⁵³, is implemented to rank the relative importance of candidate pre-crash predictors. Introduced to the accident analysis field in 2021²⁵, this feature scorer filters the dataset by evaluating each variable's capacity to differentiate between distinct outcomes. Mathematically, for a given instance R_i , the algorithm iteratively updates the weight W_f of a feature f by penalising the distance to its K nearest hits (H) of the same class, and proportionately rewarding the distance to its K nearest misses (M) from each different class C . In this framework, K was set to 10 nearest neighbours, yielding the formulation expressed by equation (3):

$$W_f \leftarrow W_f - \frac{\sum_{j=1}^K \text{diff}(f, R_i, H_j)}{m \cdot K} + \sum_{C \neq \text{class}(R_i)} \left[\frac{P(C)}{1 - P(\text{class}(R_i))} \frac{\sum_{j=1}^K \text{diff}(f, R_i, M_j(C))}{m \cdot K} \right] \quad (3)$$

where m is the total number of instances in the training sample, $P(C)$ represents the prior class probability, and $\text{diff}(\cdot)$ denotes the specific feature distance metric. This mechanism assigns importance weights accordingly to isolate a relevant subset of predictors prior to regression.

Following this initial filtering, only predictors exhibiting positive ReliefF weights are retained for the second selection stage. A greedy stepwise selection procedure is then applied to these remaining positive-weight variables. This greedy mechanism is specifically adopted to systematically enforce model parsimony, yielding compact yet mathematically robust IRFs. Restricting the model to a low-dimensional input vector is critical for ensuring straightforward algorithmic execution within real-time vehicle control loops and high model interpretability.

The stepwise procedure iteratively evaluates the inclusion or exclusion of predictors based on the Bayesian Information Criterion (BIC). The BIC operates as a strict statistical selection metric that balances model goodness-of-fit against parameter complexity, penalising the objective function based on the total sample size n and the number of estimated parameters p , comprising the intercept and all fixed-effect regression coefficients, as defined by equation (4):

$$\text{BIC} = -2\ln(\hat{L}) + p\ln(n) \quad (4)$$

where $\hat{L}$ is the maximised likelihood of the model. To ensure robust convergence, the criterion is combined with a Jeffreys prior, $|\mathcal{J}(\beta)|^{1/2}$, which is proportional to the square root of the determinant of the Fisher information matrix $\mathcal{J}(\beta)$, evaluated at the vector of fixed-effect parameters β . The integration of this prior yields a penalised log-likelihood function expressed by equation (5):

$$l^*(\beta) = l(\beta) + \frac{1}{2} \ln |\mathcal{J}(\beta)| \quad (5)$$

This structural penalty stabilises parameter estimates and prevents computational divergence in scenarios characterised by quasi-complete data separation.

Model Building

For each accident dataset, generalised linear models are trained using three distinct link functions: logit, probit, and complementary log-log (cloglog). Because the underlying distribution of biomechanical injury thresholds across the population is unknown a priori, evaluating different link functions allows for testing distinct mathematical assumptions regarding tail behaviour^{54,55}. The logit link, which assumes a symmetric and heavier-tailed risk escalation, is traditionally employed for car-to-car IRFs. The probit link, based on a symmetric standard cumulative normal distribution, assumes normally distributed thresholds and is utilised somewhat less frequently in this specific domain. Conversely, the cloglog link introduces an asymmetric transformation where risk approaches unity more steeply than zero. Although the cloglog link has traditionally been employed for pedestrian impacts⁵⁶, a literature review shows that it has never been utilised for car-to-car collisions. Following recommendations in international guidelines such as ISO/TS 18506:2014, all three candidate link functions are evaluated empirically rather than

assumed *a priori*, with the subsequent validation phase establishing which specific link function ultimately yields the single most robust and transferable IRF.

To account for the hierarchical structure inherent to accident data, generalised linear mixed-effects models (GLMMs) are implemented. In car-to-car collision databases, observations are naturally clustered because the two vehicles involved in the same specific crash event share unobserved environmental factors, structural alignment characteristics, and global crash dynamics. This nesting violates the fundamental statistical assumption of independent observations. To resolve this dependency, a random intercept $u_j \sim \mathcal{N}(0, \sigma^2)$ is specified at the unique crash event level (j), where the tilde symbol $\sim$ denotes the distributional operator and $\mathcal{N}(0, \sigma^2)$ represents a normal distribution characterised by a mean of zero and a between-crash variance of σ^2 . Conditioned on this random effect, the binary injury outcome Y_{ij} for the i -th occupant within the j -th crash follows a Bernoulli distribution, $Y_{ij} | \mathbf{x}_{ij}, u_j \sim \text{Bernoulli}(\pi_{ij})$, where the vertical bar " $|$ " denotes statistical conditioning. Within this notation, π_{ij} represents the conditional probability of sustaining an injury equal to or greater than the target threshold, and $\mathbf{x}_{ij}$ represents the corresponding vector of harmonised pre-crash predictors. In this mixed-effects framework, the mathematical formulation of the linear predictor η combines the fixed effects of the standard vector of covariates and the crash-specific random intercept, as formalised in Equation 6:

$$\eta = \beta_0 + \sum_k \beta_k \cdot \frac{x_k - \mu_k}{\sigma_k} + u_j \quad (6)$$

where β_0 represents the fixed global intercept, $\beta = \{\beta_k\}$ denotes the vector of estimated fixed-effect regression coefficients for each predictor x_k , while μ_k and σ_k represent the corresponding mean and standard deviation values established during the training data normalisation phase.

The selected link function dictates how this linear predictor maps directly to the predicted IR. $\pi_{i,j}$ is equivalently expressed as IR in the following equations, denoting the probability of sustaining an injury of at least the target severity. When evaluating the logit link function, the probability of sustaining an injury is modeled via a symmetric logistic cumulative distribution function, mapping the response as shown in Equation 7:

$$IR = \frac{1}{1 + e^{-\eta}} \quad (7)$$

Alternatively, the probit link function maps the response probability using the standard cumulative normal distribution function denoted by Φ , defined as $\Phi(\eta) = \int_{-\infty}^{\eta} \frac{1}{\sqrt{2\pi}} e^{-t^2/2} dt$, assuming normally distributed biomechanical tolerance thresholds across the target sample population, as detailed in Equation 8:

$$IR = \Phi(\eta) \quad (8)$$

Finally, the complementary log-log link function introduces an asymmetric mathematical transformation where the hazard curve scales more rapidly as the predictor values increase, establishing the IR definition expressed in Equation 9:

$$IR = 1 - e^{-e^{\eta}} \quad (9)$$

This structured model configuration establishes the candidate mathematical formulations whose predictive performance and cross-dataset stability are quantified in the subsequent validation phase.

Validation strategy

The predictive capacity and generalisation performance of the candidate models are evaluated using the Area Under the Curve (AUC) of the Receiver Operating Characteristic (ROC). This metric is specifically selected due to its mathematical robustness in configurations characterised by highly unbalanced data, which is a typical feature of injury severity datasets where critical or fatal outcomes are statistically rare⁵⁷. The underlying curve is established by plotting the true positive rate against the false positive rate across various decision thresholds, and the subsequent calculation of the area under this trajectory provides an indicator of model discrimination capacity. The validation protocol executes the cross-testing matrix illustrated in Figure 2b. For each of the three link functions trained on a given dataset, an auto-testing AUC is computed by evaluating the model on that same training data. Concurrently, two cross-testing AUC values are calculated by applying the pre-trained, unadjusted model equations directly to the remaining two independent datasets. This procedure generates an evaluation matrix encompassing the candidate functions derived across the databases. To identify the function that exhibits the highest transferability, a global

performance score is calculated for each candidate model as the arithmetic mean of its highest and lowest obtained AUC values across both the auto-testing and cross-testing evaluations, as formulated in equation (10):

$$\text{score} = \frac{\text{AUC}_{\text{highest}} + \text{AUC}_{\text{lowest}}}{2} \quad (10)$$

The candidate mathematical formulation that achieves the maximum global score is selected as the final model, establishing the framework that minimises performance degradation when deployed across geographically distinct datasets.

Results

This section presents the results of applying the proposed methodological framework across the three geographically diverse accident datasets. To establish a comprehensive point of reference, Table 1 provides an initial summary matrix of the predictive performance metrics obtained across the entire processing pipeline, serving as a global overview before entering into the detailed analysis of each specific target threshold. Each dataset was utilised to train three distinct candidate IRFs using logit, probit, and complementary log-log (cloglog) link functions, respectively. This model building pipeline was executed systematically for each of the three target injury severity thresholds: minor (MAIS 1+), severe (MAIS 3+), and fatal (MAIS 6) outcomes.

Table 1 outlines the AUC values calculated for each trained function, evaluated both via internal assessment on the respective training dataset (auto-testing) and via external evaluation on the remaining two unencountered datasets (cross-testing). AUC entries of IRFs compromised by severe overfitting, identified in this study by an absolute drop in AUC greater than 0.25 between the internal auto-testing baseline and any of the corresponding cross-testing evaluations, were omitted due to their lack of predictive validity and are represented by a backslash (\). In accordance with the validation strategy, each candidate function received a global score derived as the arithmetic mean of its maximum and minimum AUC values; functions rejected due to overfitting were assigned a score of 0. For each target variable, the formulation that maximised the global score—representing the optimal balance between predictive capacity and cross-dataset stability—is indicated with an asterisk (*) next to its score value.

TRAINING DATASET	TESTING DATASET	LIGHT (MAIS 1+)			SEVERE (MAIS 3+)			FATAL (MAIS 6)		
		<i>Logit</i>	<i>Probit</i>	<i>Cloglog</i>	<i>Logit</i>	<i>Probit</i>	<i>Cloglog</i>	<i>Logit</i>	<i>Probit</i>	<i>Cloglog</i>
IGLAD	IGLAD	0.6654	0.6641	0.6431	0.8297	0.8293	0.834	0.8769	0.8769	0.8769
	NASS-CDS	0.6833	0.6825	0.6826	0.8544	0.8557	0.8555	0.894	0.894	0.894
	VOIESUR	0.8023	0.8018	0.7683	0.8138	0.8079	0.7997	0.6438	0.6438	0.6438
	Score	0.7339*	0.73295	0.7057	0.8341	0.8318	0.8276	0.7689	0.7689	0.7689
NASS-CDS	IGLAD	0.5509	0.6045	0.6083	0.8174	0.8217	0.8171	0.9021	0.9018	0.9065
	NASS-CDS	0.6226	0.6669	0.6275	0.875	0.8757	0.8719	0.9032	0.9032	0.9032
	VOIESUR	0.6636	0.7631	0.8143	0.8336	0.8315	0.8304	0.6665	0.6661	0.6778
	Score	0.60725	0.6838	0.7113	0.8462	0.8487*	0.8445	0.7846	0.7847	0.7922
VOIESUR	IGLAD	\	\	\	0.7217	0.7422	0.7134	0.8342	0.8448	0.8340
	NASS-CDS	\	\	\	0.8005	0.8104	0.7952	0.8723	0.8769	0.8731
	VOIESUR	\	\	\	0.9503	0.9473	0.9696	0.9401	0.93	0.9451
	Score	0	0	0	0.8360	0.8448	0.8415	0.8872	0.8874	0.8896*

Table 1. Predictive validation and cross-dataset transferability matrix of the derived IRFs. Omitted entries (\) and zero-valued scores indicate functions rejected due to severe overfitting, whereas asterisks (*) identify the highest global transferability score selected for each target variable.

The subsequent subsections detail the mathematical composition and predictive behaviour of the optimal IRF selected for each of the three target injury thresholds.

Optimal IRF for at least light injury prediction (MAIS 1+)

The model that achieved the highest global score for predicting at least light injuries (0.7339) of the most injured car occupant is the logit function trained on the IGLAD dataset. Starting from an initial sample of 4,264 observations available in the filtered IGLAD dataset, the application of the ReliefF algorithm identified the following predictors with positive importance weights, listed in descending order: *CSI*, *CMI*, *vehicle kerb mass* and *closing speed*. Subsequent execution of the stepwise selection procedure identified the optimal predictor combination as a triplet comprising *CMI*, *vehicle kerb mass*, and *closing speed*, yielding a BIC value of 2493.84. Restricting the dataset to these three selected predictors and modelling outcomes equal to or greater than MAIS 1+, the number of complete observations containing non-missing values across all parameters

settled at 2,274. Table 2 reports the total number of available observations for each dataset at the time of model training and cross-validation, alongside their corresponding injury severity distributions.

DATASET	TOTAL	MAIS 1+	MAIS 3+	MAIS 6
IGLAD	2274	68.12%	18.03%	3.83%
VOIESUR	373	78.10%	20.67%	3.55%
NASS-CDS	3284	63.7%	16.75%	4.72%

Table 2. Sample size and cumulative injury severity distribution across the training (IGLAD) and testing (VOIESUR, NASS-CDS) datasets after restricting observations to the optimal predictor triplet for MAIS 1+ IR modelling.

Equations (6) and (7) present the resulting mathematical formulation, whose estimated coefficients are detailed in Table 3.

PREDICTOR	Variable	β_k	μ_k	σ_k
CMI	x_1	1.1224	0.4410	0.3988
Vehicle mass	x_2	-0.26007	1526.5183	381.3459
Closing speed	x_3	0.70748	52.4805	34.7342
\	<i>intercept</i>	0.82854	\	\

Table 3. Coefficients of the optimal logit model derived for predicting at least light injury (MAIS 1+) outcomes. β_0 is indicated as the *intercept*.

The derived IRF was evaluated both internally on the training dataset and externally on the VOIESUR and NASS-CDS datasets. Figure 4 illustrates the ROC curves and corresponding AUC metrics for internal auto-testing on IGLAD (Figure 4a), alongside external cross-validation via cross-testing performance evaluations across the other two independent datasets (Figure 4b and Figure 4c). The predictive performance metrics plotted in Figure 4 demonstrate that the selected optimal function exhibits a relatively modest discriminatory capacity. The internal auto-testing evaluation on the IGLAD development dataset yields an AUC of 0.6654, representing a low-level baseline fit that reflects a constrained predictive power within the training data itself. This limited classification performance is mirrored in the external validation across the larger cross-testing sample, namely the NASS-CDS database, which returns an AUC of 0.6833. Conversely, the function exhibits a localised performance spike when applied to the VOIESUR dataset, achieving an AUC of 0.8023. Despite this specific increase which mathematically elevates the final global score, the combined empirical results confirm that the derived IRF possesses a generally weak capacity to differentiate between uninjured states and at least light injury outcomes, thereby qualifying as a low-performance predictive model (AUC<0.70) within the developed framework⁵⁸.




Figure 4. Performance evaluation metrics for the optimal at least light injury (MAIS 1+) risk function trained on the IGLAD database. The three panels display the ROC curves and corresponding AUC values for internal validation via auto-testing on IGLAD (a), and external cross-validation via cross-testing on the VOIESUR (b) and NASS-CDS datasets (c).

Optimal IRF for at least severe injury prediction (MAIS 3+)

The model that achieved the highest global score for predicting at least severe injuries (0.8487) of the most injured car occupant is the probit function trained on the NASS-CDS dataset. Starting from an initial sample of 3,287 observations available in the filtered NASS-CDS dataset, the application of the ReliefF algorithm identified the following predictors with positive importance weights, listed in descending order: *Closing speed*, *CMI*, *CSI* and *impact configuration*. Subsequent execution of the stepwise selection procedure identified the optimal predictor combination as a triplet comprising *Closing speed*, *CMI* and *CSI*, yielding a BIC value of 1271.75. Restricting the dataset to these three selected predictors and modelling outcomes equal to or greater than MAIS 3+, the number of complete observations containing non-missing values across all parameters settled at 3,284. Table 4 reports the total number of available observations for each dataset at the time of model training and cross-validation, alongside their corresponding injury severity distributions.

DATASET	TOTAL	MAIS 1+	MAIS 3+	MAIS 6
NASS-CDS	3284	63.70%	16.75%	4.72%
VOIESUR	389	76.30%	18.96%	3.53%
IGLAD	2204	68.83%	17.06%	3.58%

Table 4. Sample size and cumulative injury severity distribution across the training (NASS-CDS) and testing (VOIESUR, IGLAD) datasets after restricting observations to the optimal predictor triplet for MAIS 3+ IR modelling.

Equations (6) and (8) present the resulting probit formulation, whose estimated coefficients are detailed in Table 5.

PREDICTOR	Variable	β_k	μ_k	σ_k
Closing speed	x_1	0.62955	51.76212	25.2365
CMI	x_2	0.323	0.4472	0.0816
CSI	x_3	0.19567	0.4916	0.1204
\	intercept	-1.8592	\	\

Table 5. Coefficients of the optimal probit model derived for predicting at least severe injury (MAIS 3+) outcomes. β_0 is indicated as the *intercept*.

The derived IRF was evaluated both internally on the training dataset and externally on the VOIESUR and IGLAD datasets. Figure 5 illustrates the ROC curves and corresponding AUC metrics for internal auto-testing on NASS-CDS (Figure 5a), alongside external cross-validation via cross-testing performance evaluations across the other two independent datasets (Figure 5b and Figure 5c). The predictive performance metrics presented in Figure 5 demonstrate a highly robust discriminatory capacity, marking a substantial improvement over the minor injury configuration. The internal auto-testing evaluation on the NASS-CDS development dataset yields an AUC of 0.8757, indicating an excellent baseline fit ($0.8 \leq \text{AUC} < 0.9$) and strong classification power⁵⁸. Crucially, this high predictive accuracy is systematically maintained during external cross-validation across the remaining independent datasets. The model achieves an AUC of 0.8315 when tested on the VOIESUR database and an AUC of 0.8217 on the IGLAD dataset. The marginal performance degradation observed between the internal fit and external evaluations proves the structural absence of overfitting, confirming the significant geographical generalisability and numerical stability of the derived probit formulation when predicting at least severe injury outcomes.




Figure 5. Performance evaluation metrics for the optimal at least severe injury (MAIS 3+) risk function trained on the NASS-CDS database. The three panels display the ROC curves and corresponding AUC values for internal validation via auto-testing on NASS-CDS (a), and external cross-validation via cross-testing on the VOIESUR (b) and IGLAD datasets (c).

Optimal IRF for fatality prediction (MAIS 6)

The model that achieved the highest global score for predicting fatal injuries (0.8896) of the most injured car occupant is the complementary log-log function trained on the VOIESUR dataset. Starting from an initial sample of 998 observations available in the filtered VOIESUR dataset, the application of the ReliefF algorithm identified the following predictors with positive importance weights, listed in descending order: *vehicle kerb mass*, *closing speed*, *CSI* and *CMI*. Subsequent execution of the stepwise selection procedure identified the optimal predictor combination as a triplet comprising *vehicle kerb mass*, *closing speed* and *CSI*, yielding a BIC value of 896.34. Restricting the dataset to these three selected predictors and modelling outcomes equal to MAIS 6, the number of complete observations containing non-missing values across all parameters settled at 399. Table 6 reports the total number of available observations for each dataset at the time of model training and cross-validation, alongside their corresponding injury severity distributions.

Equations (6) and (9) present the resulting cloglog formulation, whose estimated coefficients are detailed in Table 7.

The derived IRF was evaluated both internally on the training dataset and externally on the IGLAD and NASS-CDS datasets. Figure 6 illustrates the ROC curves and corresponding AUC metrics for internal auto-testing on VOIESUR (Figure 6a), alongside external cross-validation via cross-testing performance evaluations across the other two independent datasets (Figure 6b and Figure 6c). The predictive performance metrics presented in Figure 6 demonstrate a high overall discriminatory capacity, although the cross-testing ROC trajectories exhibit a distinct step-like profile. This marked discontinuity is a direct mathematical consequence of sample sparsity, given that fatal outcomes constitute rare events representing only 3.00% of the 399 development

DATASET	TOTAL	MAIS 1+	MAIS 3+	MAIS 6
VOIESUR	399	82.82%	17.21%	3.00%
IGLAD	2184	68.64%	17.63%	3.75%
NASS-CDS	3284	63.70%	16.75%	4.72%

Table 6. Sample size and cumulative injury severity distribution across the training (VOIESUR) and testing (IGLAD, NASS-CDS) datasets after restricting observations to the optimal predictor triplet for MAIS 6 IR modelling.

PREDICTOR	Variable	β_k	μ_k	σ_k
Vehicle kerb mass	x_1	-1.5797	1177.2927	268.1789
Closing speed	x_2	2.1709	104.6921	37.2217
CSI	x_3	1.6573	0.4697	0.1118
\	<i>intercept</i>	-2.0244	\	\

Table 7. Coefficients of the optimal cloglog model derived for predicting fatal injury (MAIS 6) outcomes. β_0 is indicated as the *intercept*.

observations in VOIESUR, 3.75% of the 2,184 observations in IGLAD, and 4.72% of the 3,284 observations in NASS-CDS. The internal auto-testing evaluation on the VOIESUR dataset yields an elevated AUC of 0.9451, reflecting a strong classification power despite the low absolute number of positive instances. This predictive performance translates effectively to external cross-validation, where the model achieves an AUC of 0.8340 on the IGLAD database and an AUC of 0.8731 on the NHTSA NASS-CDS dataset. While these high cross-testing AUC metrics confirm the geographical generalisability of the derived cloglog formulation, the jagged nature of the curves across all three panels emphasises the underlying statistical limits imposed by the rare nature of MAIS 6 events.




Figure 6. Performance evaluation metrics for the optimal fatal injury (MAIS 6) risk function trained on the VOIESUR database. The three panels display the ROC curves and corresponding AUC values for internal validation via auto-testing on VOIESUR (a), and external cross-validation via cross-testing on the IGLAD (b) and NASS-CDS datasets (c).

Application Case Study: Evaluating pre-crash based IRF applicability under realistic conditions

The optimal IRF derived for predicting at least severe injuries (equations 6 and 8, Table 5) is highly suited for embedded on-board implementation within real-time emergency decision-making architectures (Figure 1) to identify the manoeuvre

yielding the lowest injury consequences for vehicle occupants during an ICS. Notably, the core strategy for reducing road traffic accidents endorsed by the European Commission via the Vision Zero Programme aims to reduce road fatalities and serious injuries by 50% by 2030, and to achieve zero road fatalities and serious injuries by 2050. Within this context, the deployment of MAIS 3+ IRFs aligns seamlessly with these major international road safety initiatives.

Evaluating the probability of sustaining at least severe injuries for the most critically injured occupant requires three fundamental inputs: the initial closing speed between the conflicting vehicles at the point of first contact, the CMI, and the CSI of the target vehicle (ego). The initial closing speed can be extracted in real time from an on-board kinematic simulation fed by exteroceptive sensors that track the opponent vehicle's kinematics (such as velocity and acceleration) and geometry. Similarly, both CMI and CSI can be estimated via purely pre-crash kinematic tracking once the expected impact configuration is identified, utilising the prospective analytical formulations proposed by Vangi²² defined in equations 11 and 12:

$$CMI = \frac{\gamma_1 \gamma_2 (1 + \varepsilon)}{(\gamma_2 + \gamma_1 R_m)} \quad (11)$$

$$CSI = \sqrt{\frac{\gamma_1 \gamma_2 (1 - \varepsilon^2)}{(1 + R_k)(\gamma_2 + \gamma_1 R_m)}} \quad (12)$$

where the mass reduction factor γ_i takes the mathematical form defined in Equation 13, ε is the coefficient of restitution evaluated at the vehicle first point of impact, R_m is the vehicle mass ratio ($R_m = m_{\text{ego}}/m_{\text{opp}}$), and R_k represents the structural stiffness ratio between the conflicting vehicles ($R_k = k_{\text{ego}}/k_{\text{opp}}$).

$$\gamma_i = \frac{k_i^2}{k_i^2 + h_i^2} \quad (13)$$

where k_i is the radius of gyration of the i -th vehicle and h_i denotes the lever arm of the resultant impact force vector.

The parameter γ_i can be calculated from the pre-crash kinematic projection at the initial instant of contact by estimating k_i and h_i , considering that the PDOF aligns with the relative speed vector. The coefficient of restitution ε , which increases as the closing speed decreases, can be estimated via empirical correlations⁵⁹; however, for relative impact speed exceeding 36 km/h, it moves rapidly from 0.1 to almost null values. Conversely, evaluating the ratios R_m and R_k introduces additional complexity, as it requires the ego vehicle to possess prior knowledge of the opponent vehicle's mass and structural stiffness, alongside its own properties. Consequently, the on-board control unit must have access to a comprehensive vehicular fleet database containing the mass and stiffness properties of the active vehicle fleet, which can be queried once the opponent vehicle model is identified via computer vision techniques. Alternatively, a future and ideal operational architecture would rely on Vehicle-to-Vehicle (V2V) communication networks, enabling vehicles to exchange their exact inertial and structural parameters in the pre-crash phase. However, in the absence of such cooperative driving infrastructure, obtaining these data remains notoriously difficult, as they are typically proprietary manufacturer information or must be reverse-engineered from standardised crash test data.

To preliminarily evaluate these operational challenges, a real-world crash test involving a collision between a Mitsubishi Lancer A 170 and a Ford Sierra 2.0 GL—conducted at Wildhaus in June 1999 by the AREC Group—is presented. This case study serves a dual purpose. First, it demonstrates the straightforward applicability of the proposed MAIS 3+ IRF when the ego vehicle has access to exact structural parameters, and thus precise values for R_m and R_k . Second, it provides a quantitative assessment of the predictive error propagated into the IR estimation when the exact stiffness of the opponent vehicle is unknown (representing the most realistic current operating scenario) and must be approximated via correlations with commercially available vehicle descriptors. A historical crash test involving discontinued vehicle models was intentionally selected to ensure the analysis serves as a methodological validation without favouring or comparing the commercial safety performance of contemporary vehicles. Table 8 summarises the primary kinematic and structural parameters of interest from the crash test (measured or derived through reconstruction²²), and the specific front-to-side impact configuration is illustrated in Figure 7a.

	Configuration	Mass (kg)	EES (km/h)	ΔV (km/h)	γ	Closing speed (km/h)	ε	R_k
Mitsubishi	Frontal	1076	35	27.5	0.73	81	0.1	0.41
Ford	Lateral	1142	27	26.0	0.61			

Table 8. Primary kinematic and structural parameters from the real-world crash test conducted at Wildhaus in June 1999 by the AREC Group, involving a front-to-side collision between a Mitsubishi Lancer A 170 and a Ford Sierra 2.0 GL.

Utilising the experimental crash data reported in Table 8, the initial closing speed is established at 81 km/h, which yields a *CMi* of 0.39 and a *CSI* of 0.49 for the Mitsubishi (designated as the ego vehicle). In an ideal on-board tracking scenario, where the ego vehicle's sensors perfectly capture the environment, its own frontal stiffness and mass are predefined, and the opponent's mass and lateral stiffness are fully known, either through exhaustive on-board vehicle databases or prospective V2V communication networks, the execution of the optimal probit model (equations 6 and 8) using the coefficients in Table 5 yields a baseline IR of $IR_{MAIS3+} = 8.7\%$.

Conversely, under realistic operating conditions where the exact structural stiffness of the opponent vehicle is unavailable, this value must be approximated using available macroscopic indicators. For instance, by implementing the empirical relationships developed by the NHTSA^{60,61} to estimate the opponent's lateral stiffness as a function of its vehicle segment and wheelbase, a stiffness ratio of $R_k = 0.70$ is obtained. Substituting this estimated ratio into Equation 12 yields a predicted *CSI* of 0.45, corresponding to a relative error of -8.8% (Figure 7b) compared to the true experimental crash test baseline. This structural misestimation propagates into the IRF, resulting in a predicted IR of 7.6%, which translates to a final relative error of -12.1% (Figure 7c). More broadly, even in the absence of literature-derived statistical correlations, it is highly reasonable to approximate that an opponent vehicle's lateral structural stiffness is roughly 20% ($R_k = 0.80$) to 60% ($R_k = 0.40$) lower than the frontal structural stiffness of the ego vehicle. Sweeping across this operational approximation range restricts the resulting relative error to between -11.3% and 0.5% for the *CSI*, and between -15.4% and 0.7% for the final IR prediction, as shown in Figure 7b and c.




Figure 7. Initial front-to-side impact configuration of the conflicting vehicles at the point of first contact (a). Relative error propagated into the Crash Severity Index (*CSI*) prediction (b) and the corresponding at least severe injury risk (IR_{MAIS3+}) estimation (c) evaluated across a realistic range of the estimated vehicle stiffness ratio (R_k).

Discussion

The finding that geographically transferable IRFs can be successfully established using exclusively pre-crash variables carries key interpretive value for embedded active safety architectures. Specifically, the high predictive capacity demonstrated by macro-level pre-crash parameters, such as closing speed, *CMi*, and *CSI*, indicates that sufficient physical information is accessible prior to contact to guide real-time emergency manoeuvres, thereby bypassing the need for structural crash-phase simulations. The stable MAIS 3+ predictive performance and the high theoretical potential of the MAIS 6 model observed during cross-testing demonstrate that the proposed parsimonious method yields mathematically robust IRFs featuring few pre-crash predictors, allowing their reliable transferability to vehicle fleets and crash scenarios across different geographical regions. However, a comparative analysis of the three IRFs across different injury severity levels highlights a clear divergence in model performance. The IRF selected for at least light injuries (MAIS 1+), a threshold of limited practical interest for AD/ADAS implementation, shows modest discriminatory capacity. This low performance finds a clear physical explanation in minor injury mechanisms, as light trauma is largely governed by microscopic, occupant-intrinsic factors such as age and sex⁶². Because these parameters cannot be captured by exteroceptive vehicle sensors and are rarely available to the vehicle control unit, they were intentionally excluded from the framework, leaving the MAIS 1+ IRF with an inherent predictive limitation. Although the corresponding auto-testing result of 0.6654 and cross-test result of 0.6833 sit immediately below the traditional acceptability threshold

established in literature⁵⁸, the safety-critical constraints of automotive active safety dictate a more stringent interpretation. Within this specific domain, these findings qualify the model as a low-performance predictive tool within the developed framework, as an AUC below 0.70 presents substantial reliability challenges. Consequently, practical on-board applications demanding real-time injury risk estimation inherently require more robust and higher-performing classifiers to ensure safe and dependable deployment. Conversely, the substantial performance increase observed for at least severe injuries (MAIS 3+), where AUC values remain consistently above 0.80 even across completely independent validation datasets, confirms that severe trauma probability can be accurately modelled using only pre-crash predictors, independently of occupant characteristics. This outcome aligns with the current literature, where studies deploying machine learning techniques²⁴ or explainable AI frameworks such as Shapley Additive Explanations (SHAP)¹⁸ similarly demonstrate the prominent ranking of pre-crash variables, including closing speed, *CMI*, and *CSI*, among the primary predictors of car-to-car injury severity. Regarding fatal injury prediction (MAIS 6), the distinct step-like profile of the *ROC* curves reflects the mathematical consequence of severe sample sparsity in real-world databases, where fatality frequencies remain systematically below 5%. Although the complementary log-log (*cloglog*) link function accommodated this statistical asymmetry better than traditional logistic approaches, these results must be interpreted with strict caution. Training empirical models on rare events introduces a high risk of over-sensitivity to localized data anomalies or unique database reporting biases. When deployed within real-time, safety-critical vehicle control loops, such functions could manifest unpredictable mathematical volatility or generate dangerous misestimations when encountering unmodelled edge-case scenarios. Consequently, while the cross-dataset transferability metrics, peaking at a cross-testing AUC of 0.87, suggest that this continuous predictor combination holds theoretical merit, the MAIS 6 model should currently be treated as a preliminary, exploratory baseline rather than a production-ready control metric. Robust verification against significantly larger, consolidated fatal crash repositories remains mandatory before assessing its operational readiness.

Regarding practical on-board implementation, applying the MAIS 3+ model to the real-world crash test case study provides valuable, albeit preliminary, insight into the operational robustness of the architecture under realistic conditions. In current real traffic scenarios, the ego vehicle control unit operates with partial information because the exact structural stiffness of the opponent vehicle, and consequently the stiffness ratio (R_k) required to calculate the prospective pre-crash *CSI*, is generally unknown. The case study demonstrates that the proposed IRF is straightforwardly applicable; however, while error propagation happens to be attenuated in this specific impact configuration, these findings raise critical issues concerning global error propagation when detailed structural parameters are missing. Given that an absence of exhaustive fleet data available on-board represents the most probable present-day operating scenario, this aspect demands rigorous attention. Although prospective V2V communication networks could eventually resolve this limitation by enabling the exchange of exact structural parameters between vehicles prior to contact, current technology must still rely on imperfect data regimes. Consequently, future research prior to real-world deployment must systematically analyse the sensitivity and cumulative effects of potential estimation errors across all kinematic and structural parameters that constitute the *a priori* formulations of both *CMI* and *CSI*, as well as any other input predictors deployed by the implemented IRF. Nevertheless, the preliminary results of this case study remain highly promising, showcasing that the framework can successfully function under such data constraints and reliably identify the optimal kinematic manoeuvre to minimise occupant IR during an ICS.

Conclusions

This study has developed and cross-validated three omni-directional, continuous injury risk functions (IRFs) based exclusively on pre-crash variables. This framework enables a fundamental paradigm shift from conventional collision avoidance and sole emergency braking during Inevitable Collision States (ICS) toward real-time occupant injury risk (IR) mitigation strategies for assisted and automated driving systems. By restricting calculations to variables known at the exact instant of first contact, prior to the actual car-to-car crash phase, the proposed methodology establishes the conceptual feasibility of implementing emergency steering interventions in unavoidable impacts. This approach eliminates the need for computationally intensive crash-phase virtual forward simulations to select the safest manoeuvre, which are otherwise prohibitive for real-time, on-board deployment. Reliable pre-crash based IRFs permit purely kinematic simulations, thereby matching the millisecond reaction budgets required by embedded safety architectures to actively expand upon traditional pure braking. The underlying development process integrates a two-stage feature selection technique and hierarchical mixed-effects modelling, ensuring that the derived functions are robustly transferable, a characteristic demonstrated through cross-validation across geographically diverse and independent databases. The empirical validation across these independent accident databases reveals a clear divergence in model capabilities based on trauma severity when exclusively utilising pre-crash predictors. While the at least light injury (MAIS 1+) function exhibits limited discriminatory capacity due to the exclusion of occupant-intrinsic parameters that are typically unavailable to the vehicle control unit executing the emergency manoeuvre, the at least severe injury (MAIS 3+) and fatal injury (MAIS 6) models demonstrate high geographical transferability and robust predictive performance using only macro-level pre-crash kinematic and structural predictors. Application to real-world car-to-car crash test data further confirms the straightforward applicability of these functions and validates their robust performance under imperfect data regimes, while highlighting the

need for caution when opponent structural parameters are not fully resolved. Although the IRFs for at least severe injuries are particularly promising for real-time, on-board risk estimation aimed at minimising injury risk in ICS, full deployment within production vehicle control loops requires addressing operational challenges regarding parameter uncertainties under realistic traffic conditions. Accurate injury risk estimation depends heavily on the reliability of external sensors and the capacity of computer vision tracking systems to properly identify and localise the opponent vehicle. To resolve uncertainties regarding the opponent's structural and inertial properties, host systems must currently rely on exhaustive on-board vehicle fleet databases. Alternatively, prospective Vehicle-to-Vehicle (V2V) communication networks represent the ideal future architecture, enabling vehicles to exchange their inertial, structural, and geometrical properties alongside dynamically updated kinematic data, such as speed and acceleration, prior to contact. Ultimately, this methodology offers the possibility to derive IRFs that permit algorithmically viable execution to actively guide combined braking and autonomous steering interventions on-board, in real time, and during emergency scenarios. This provides a concrete technological foundation to support major global road safety initiatives aimed at mitigating the consequences of road traffic accidents, such as the Vision Zero programme.

Data availability

The NHTSA NASS-CDS dataset analysed during the current study is publicly available from the National Highway Traffic Safety Administration repository. Conversely, the IGLAD and VOIESUR datasets are not publicly available due to proprietary licensing and third-party restrictions imposed by their respective organisations, meaning that access to these restricted data remains subject to approval from the corresponding third-party providers. Furthermore, the data in their elaborated form cannot be provided or shared by the authors due to industrial confidentiality constraints.

References

1. Bengler, K. *et al.* Three decades of driver assistance systems: Review and future perspectives. *IEEE Intell. Transp. Syst. Mag.* **6**, 6–22, DOI: [10.1109/mits.2014.2336271](https://doi.org/10.1109/mits.2014.2336271) (2014).
2. Neumann, T. Analysis of advanced driver-assistance systems for safe and comfortable driving of motor vehicles. *Sensors (Basel, Switzerland)* **24**, DOI: [10.3390/s24196223](https://doi.org/10.3390/s24196223) (2024).
3. Mueller, A. S., Cicchino, J. B. & Zuby, D. What humanlike errors do autonomous vehicles need to avoid to maximize safety? *J. safety research* **75**, 310–318, DOI: [10.1016/j.jsr.2020.10.005](https://doi.org/10.1016/j.jsr.2020.10.005) (2020).
4. Singh, S. Critical reasons for crashes investigated in the National Motor Vehicle Crash Causation Survey. Traffic Safety Facts Crash Stats DOT HS 812 506, National Highway Traffic Safety Administration, Washington, DC (2018).
5. World Health Organization. Global status report on road safety 2023. Report, World Health Organization, Geneva, Switzerland (2023).
6. Chougule, A., Chamola, V., Sam, A., Yu, F. & Sikdar, B. A comprehensive review on limitations of autonomous driving and its impact on accidents and collisions. *IEEE Open J. Veh. Technol.* **5**, 142–161, DOI: [10.1109/ojvt.2023.3335180](https://doi.org/10.1109/ojvt.2023.3335180) (2024).
7. Dobberstein, J., Lich, T. & Schmidt, D. Accident data analysis-remaining accidents and crash configurations of automated vehicles in mixed traffic. *OSCCAR Proj. Deliv. D 1*, 1 (2019).
8. Scanlon, J. *et al.* Waymo simulated driving behavior in reconstructed fatal crashes within an autonomous vehicle operating domain. *Accid. analysis prevention* **163**, 106454, DOI: [10.1016/j.aap.2021.106454](https://doi.org/10.1016/j.aap.2021.106454) (2021).
9. Li, H., Li, Y., Zheng, T. & Ye, Q. Emergency collision avoidance strategy for autonomous vehicles based on steering and differential braking. *Sci. Reports* **12**, DOI: [10.1038/s41598-022-27296-3](https://doi.org/10.1038/s41598-022-27296-3) (2022).
10. Liu, Z., Wen, G., Liu, W., Tan, X. & Wu, G. Research on automatic emergency steering collision avoidance and stability control of intelligent driving vehicle. *Front. Robotics AI* **10**, DOI: [10.3389/frobt.2023.1120658](https://doi.org/10.3389/frobt.2023.1120658) (2023).
11. Guo, W., Luo, J., Song, X. & Li, J. A crash-injury-aware driving safety field for real-time risk assessment and its application in autonomous vehicle motion planning. *Accid. analysis prevention* **222**, 108224, DOI: [10.1016/j.aap.2025.108224](https://doi.org/10.1016/j.aap.2025.108224) (2025).
12. Kaempchen, N., Schiele, B. & Dietmayer, K. Situation assessment of an autonomous emergency brake for arbitrary vehicle-to-vehicle collision scenarios. *IEEE Transactions on Intell. Transp. Syst.* **10**, 678–687, DOI: [10.1109/tits.2009.2026452](https://doi.org/10.1109/tits.2009.2026452) (2009).
13. Sander, U., Lubbe, N. & Pietzsch, S. Intersection aeb implementation strategies for left turn across path crashes. *Traffic Inj. Prev.* **20**, S119 – S125, DOI: [10.1080/15389588.2019.1602728](https://doi.org/10.1080/15389588.2019.1602728) (2019).

14. Sander, U. & Lubbe, N. Market penetration of intersection aeb: Characterizing avoided and residual straight crossing path accidents. *Accid. analysis prevention* **115**, 178–188, DOI: [10.1016/j.aap.2018.03.025](https://doi.org/10.1016/j.aap.2018.03.025) (2018).
15. Park, J., Kim, D. & Huh, K. Emergency collision avoidance by steering in critical situations. *Int. J. Automot. Technol.* **22**, 173 – 184, DOI: [10.1007/s12239-021-0018-2](https://doi.org/10.1007/s12239-021-0018-2) (2021).
16. Vangi, D., Virga, A. & Gulino, M. Adaptive intervention logic for automated driving systems based on injury risk minimization. *Proc. Inst. Mech. Eng. Part D: J. Automob. Eng.* **234**, 2975 – 2987, DOI: [10.1177/0954407020931228](https://doi.org/10.1177/0954407020931228) (2020).
17. Gulino, M., Vichi, G., Cecchetto, F., Lillo, L. D. & Vangi, D. A combined comfort and safety-based approach to assess the performance of advanced driver assistance functions. *Eur. Transp. Res. Rev.* **17**, DOI: [10.1186/s12544-024-00696-4](https://doi.org/10.1186/s12544-024-00696-4) (2025).
18. Gu, C., Xu, J., Li, S., Gao, C. & Ma, Y. Injury risk assessment and interpretation for roadway crashes based on pre-crash indicators and machine learning methods. *Appl. Sci.* DOI: [10.3390/app13126983](https://doi.org/10.3390/app13126983) (2023).
19. Kononen, D. W., Flannagan, C. & Wang, S. C. Identification and validation of a logistic regression model for predicting serious injuries associated with motor vehicle crashes. *Accid. analysis prevention* **43** 1, 112–22, DOI: [10.1016/j.aap.2010.07.018](https://doi.org/10.1016/j.aap.2010.07.018) (2010).
20. Jurewicz, C., Sobhani, A., Woolley, J., Dutschke, J. & Corben, B. Exploration of vehicle impact speed–injury severity relationships for application in safer road design. *Transp. research procedia* **14**, 4247–4256, DOI: [10.1016/j.trpro.2016.05.396](https://doi.org/10.1016/j.trpro.2016.05.396) (2016).
21. Andricevic, N., Junge, M. & Krampe, J. Injury risk functions for frontal oblique collisions. *Traffic Inj. Prev.* **19**, 518 – 522, DOI: [10.1080/15389588.2018.1442926](https://doi.org/10.1080/15389588.2018.1442926) (2018).
22. Vangi, D. Impact severity assessment in vehicle accidents. *Int. J. Crashworthiness* **19**, 576 – 587, DOI: [10.1080/13588265.2014.921974](https://doi.org/10.1080/13588265.2014.921974) (2014).
23. Gulino, M., Fiorentino, A. & Vangi, D. Prospective and retrospective performance assessment of advanced driver assistance systems in imminent collision scenarios: the cmi-vr approach. *Eur. Transp. Res. Rev.* **14**, 1–14, DOI: [10.1186/s12544-022-00527-4](https://doi.org/10.1186/s12544-022-00527-4) (2022).
24. Gulino, M., Sortino, A. & Vangi, D. The effect of pre-crash variables on occupants' injury risk in car impacts: assessment by feature ranking, regression, and machine learning for increasing injury scales. *Eur. Transp. Res. Rev.* **17**, DOI: [10.1186/s12544-025-00736-7](https://doi.org/10.1186/s12544-025-00736-7) (2025).
25. Gulino, M., Gangi, L. D., Sortino, A. & Vangi, D. Injury risk assessment based on pre-crash variables: The role of closing velocity and impact eccentricity. *Accid. analysis prevention* **150**, 105864, DOI: [10.1016/j.aap.2020.105864](https://doi.org/10.1016/j.aap.2020.105864) (2020).
26. Gulino, M., Damaziak, K., Fiorentino, A. & Vangi, D. Handling inevitable collision states by advanced driver assistance systems functions: software-in-the-loop performance assessment of an injury risk-based logic in a “lane departure” scenario. *J. Intell. Transp. Syst.* **28**, 1011 – 1031, DOI: [10.1080/15472450.2023.2277713](https://doi.org/10.1080/15472450.2023.2277713) (2023).
27. Bahouth, G., Graygo, J., Digges, K., Schulman, C. & Baur, P. The benefits and tradeoffs for varied high-severity injury risk thresholds for advanced automatic crash notification systems. *Traffic Inj. Prev.* **15**, S134 – S140, DOI: [10.1080/15389588.2014.936011](https://doi.org/10.1080/15389588.2014.936011) (2014).
28. Schmitt, K.-U., Cronin, D. S., Morrison III, B., Callaghan, J. P. & Muser, M. H. *Trauma biomechanics: an introduction to injury biomechanics* (Springer Nature, 2025).
29. Struble, D. E. & Struble, J. D. *Automotive accident reconstruction: practices and principles* (CRC Press, 2020).
30. Prochowski, L., Ziubiński, M., Dziewiecki, K. & Szwajkowski, P. Impact energy and the risk of injury to motorcar occupants in the front-to-side vehicle collision. *Nonlinear Dyn.* **110**, 3333 – 3354, DOI: [10.1007/s11071-022-07779-8](https://doi.org/10.1007/s11071-022-07779-8) (2022).
31. Ji, A. & Levinson, D. M. An energy loss-based vehicular injury severity model. *Accid. analysis prevention* **146**, 105730, DOI: [10.1016/j.aap.2020.105730](https://doi.org/10.1016/j.aap.2020.105730) (2020).
32. Alvarez, S. *et al.* Prospective effectiveness assessment of adas and active safety systems via virtual simulation: a review of the current practices. In *25th International Technical Conference on the Enhanced Safety of Vehicles (ESV) National Highway Traffic Safety Administration* (2017).
33. Smit, S. *et al.* Evaluation of a momentum based impact model in frontal car collisions for the prospective assessment of adas. *Eur. Transp. Res. Rev.* **11**, 1–9, DOI: [10.1186/s12544-018-0343-3](https://doi.org/10.1186/s12544-018-0343-3) (2019).

34. Kolk, H., Tomasch, E., Bakker, J., Dobberstein, J. & Sinz, W. Evaluation of a momentum based impact model and application in an effectivity study considering junction accidents. In *International Conference" ESAR-Expert Symposium on Accident Research"* (2016).
35. Gulino, M.-S., Vichi, G. & Vangi, D. Validation of a virtual forward simulation tool for adas assessment. *Eur. Transp. Res. Rev.* **17**, 25, DOI: [10.1186/s12544-025-00724-x](https://doi.org/10.1186/s12544-025-00724-x) (2025).
36. Sousa, L., Veríssimo, P. & Ambrósio, J. Development of generic multibody road vehicle models for crashworthiness. *Multibody Syst. Dyn.* **19**, 133–158, DOI: [10.1007/s11044-007-9093-z](https://doi.org/10.1007/s11044-007-9093-z) (2008).
37. Vangi, D., Begani, F., Spitzhüttl, F. & Gulino, M. Vehicle accident reconstruction by a reduced order impact model. *Forensic science international* **298**, 426.e1–426.e11, DOI: [10.1016/j.forsciint.2019.02.042](https://doi.org/10.1016/j.forsciint.2019.02.042) (2019).
38. Bosia, L. *et al.* Real-time collision mitigation strategies for autonomous vehicles. *IEEE Transactions on Intell. Transp. Syst.* **25**, 13483–13493, DOI: [10.1109/tits.2024.3400680](https://doi.org/10.1109/tits.2024.3400680) (2024).
39. Gomes, M. J. T. L., Cunto, F. & da Silva, A. R. Geographically weighted negative binomial regression applied to zonal level safety performance models. *Accid. analysis prevention* **106**, 254–261, DOI: [10.1016/j.aap.2017.06.011](https://doi.org/10.1016/j.aap.2017.06.011) (2017).
40. Soroori, E., Moghaddam, A. M. & Salehi, M. Modeling spatial nonstationary and overdispersed crash data: Development and comparative analysis of global and geographically weighted regression models applied to macrolevel injury crash data. *J. Transp. Saf. & Secur.* **13**, 1000 – 1024, DOI: [10.1080/19439962.2020.1712671](https://doi.org/10.1080/19439962.2020.1712671) (2020).
41. Funk, J. R., Cormier, J. M. & Gabler, H. C. Effect of delta-v errors in nass on frontal crash risk calculations. In *Annals of Advances in Automotive Medicine/Annual Scientific Conference*, vol. 52, 155 (2008).
42. Sunnevång, C., Rosén, E. & Bostrom, O. Real-life fatal outcome in car-to-car near-side impacts—implications for improved protection considering age and crash severity. *Traffic Inj. Prev.* **10**, 194 – 203, DOI: [10.1080/15389580802594127](https://doi.org/10.1080/15389580802594127) (2009).
43. Gabler, H. C., Digges, K., Fildes, B. N. & Sparke, L. Side impact injury risk for belted far side passenger vehicle occupants. Tech. Rep., SAE Technical Paper (2005). DOI: [10.4271/2005-01-0287](https://doi.org/10.4271/2005-01-0287).
44. Viano, D. C. & Parenteau, C. S. Serious injury in very-low and very-high speed rear impacts. In *SAE World Congress & Exhibition* (SAE Technical Paper, 2008).
45. McMurry, T., Cormier, J., Daniel, T., Scanlon, J. & Crandall, J. An omni-directional model of injury risk in planar crashes with application for autonomous vehicles. *Traffic Inj. Prev.* **22**, S122 – S127, DOI: [10.1080/15389588.2021.1955108](https://doi.org/10.1080/15389588.2021.1955108) (2021).
46. Liu, S., Fan, W. & Li, Y. Injury severity analysis of rollover crashes for passenger cars and light trucks considering temporal stability: A random parameters logit approach with heterogeneity in mean and variance. *J. safety research* **78**, 276–291, DOI: [10.1016/j.jsr.2021.06.013](https://doi.org/10.1016/j.jsr.2021.06.013) (2021).
47. Schlögl, M. & Stütz, R. Methodological considerations with data uncertainty in road safety analysis. *Accid. analysis prevention* DOI: [10.1016/j.aap.2017.02.001](https://doi.org/10.1016/j.aap.2017.02.001) (2017).
48. Vangi, D., Gulino, M. & Cialdai, C. Coherence assessment of accident database kinematic data. *Accid. analysis prevention* **123**, 356–364, DOI: [10.1016/j.aap.2018.12.004](https://doi.org/10.1016/j.aap.2018.12.004) (2019).
49. Bakker, J. *et al.* Iglad-international harmonized in-depth accident data. In *in Proceedings of the 25th International Conference on Enhanced Safety of Vehicles (ESV), Detroit, MI* (2017).
50. Amoros, E., Lardy, A., Martin, J., Wu, D. & Viallon, V. Méthodologie redressement et extrapolation. I3 deliverable voiesur project. anr-11-vppt-0007. https://www.researchgate.net/publication/282663658_Methodologie_redressement_et_extrapolation (2015).
51. Laberge-Nadeau, C., Bellavance, F., Messier, S., Vézina, L. & Pichette, F. Occupant injury severity from lateral collisions: a literature review. *J. safety research* **40** 6, 427–35, DOI: [10.1016/j.jsr.2009.08.005](https://doi.org/10.1016/j.jsr.2009.08.005) (2009).
52. Gulino, M.-S., Vichi, G., Kullgren, A. & Vangi, D. Assessing future road safety with the advent of l-category quadricycles: Simulation insights from current crash patterns. *Accid. Analysis & Prev.* **233**, 108561, DOI: <https://doi.org/10.1016/j.aap.2026.108561> (2026).
53. Kononenko, I. Estimating attributes: Analysis and extensions of relief. 171–182, DOI: [10.1007/3-540-57868-4_57](https://doi.org/10.1007/3-540-57868-4_57) (1994).
54. McMurry, T. & Poplin, G. Statistical considerations in the development of injury risk functions. *Traffic Inj. Prev.* **16**, 618 – 626, DOI: [10.1080/15389588.2014.991820](https://doi.org/10.1080/15389588.2014.991820) (2015).
55. Petitjean, A. *et al.* Injury risk curves for the worldsid 50th male dummy. *Stapp car crash journal* **53**, 443–76, DOI: [10.4271/2009-22-0016](https://doi.org/10.4271/2009-22-0016) (2009).

56. Martin, J.-L. & Wu, D. Pedestrian fatality and impact speed squared: Cloglog modeling from french national data. *Traffic Inj. Prev.* **19**, 101 – 94, DOI: [10.1080/15389588.2017.1332408](https://doi.org/10.1080/15389588.2017.1332408) (2018).
57. Boughorbel, S., Jarray, F. & El-Anbari, M. Optimal classifier for imbalanced data using matthews correlation coefficient metric. *PLoS ONE* **12**, DOI: [10.1371/journal.pone.0177678](https://doi.org/10.1371/journal.pone.0177678) (2017).
58. Mandrekar, J. Receiver operating characteristic curve in diagnostic test assessment. *J. thoracic oncology : official publication Int. Assoc. for Study Lung Cancer* **5** **9**, 1315–6, DOI: [10.1097/jto.0b013e3181ec173d](https://doi.org/10.1097/jto.0b013e3181ec173d) (2010).
59. Antonetti, V. W. Estimating the coefficient of restitution of vehicle-to-vehicle bumper impacts. Tech. Rep., SAE Technical Paper (1998). DOI: [10.4271/980552](https://doi.org/10.4271/980552).
60. Willke, D. T. & Monk, M. W. Crash iii model improvements: Vehicle categorization for stiffness parameters, volume i. Tech. Rep. DOT HS 807 336, U.S. Department of Transportation, National Highway Traffic Safety Administration, Washington, D.C. (1987).
61. Willke, D. T. & Monk, M. W. Crash iii model improvements: Derivation of new side stiffness parameters from crash tests, volume 2. Tech. Rep. DOT HS 807 353, U.S. Department of Transportation, National Highway Traffic Safety Administration, Washington, D.C. (1987).
62. Aziz, K. *et al.* An interpretable dynamic ensemble selection multiclass imbalance approach with ensemble imbalance learning for predicting road crash injury severity. *Sci. Reports* **15**, DOI: [10.1038/s41598-025-08935-x](https://doi.org/10.1038/s41598-025-08935-x) (2025).

Author contributions

Conceptualization: G.V., M.-S.G., C.C., D.V. Methodology: G.V., M.-S.G., C.C. Validation: G.V., M.-S.G., D.V. Software: G.V., M.-S.G. Formal analysis: G.V., M.-S.G., L.B., S.C. Investigation: G.V., M.-S.G., A.F., L.B. Writing - original draft: G.V., M.-S.G. Writing - review and editing: S.C., L.B., C.C. Visualization: G.V., D.V. Supervision and Project administration: M.-S.G., A.F., C.C., D.V. All authors have read and agreed to the published version of the manuscript.

Additional informations

Competing interests

This study received financial support from Stellantis. A.F. is an employee of Stellantis Europe, C.C. and L.B. are employees of Stellantis Auto, and G.V. receives a PhD scholarship co-funded by Stellantis. Stellantis reviewed and approved the manuscript prior to submission. The authors declare that the funding and institutional affiliations did not influence the study design, data analysis, interpretation of the results, or the decision to publish. The remaining authors declare no competing financial or non-financial interests.

Acknowledgements

This work was supported by Stellantis in collaboration with the Department of Industrial Engineering of the University of Florence, under the Italian National Recovery and Resilience Plan (PNRR). The authors acknowledge the contribution of LAB (Laboratoire d'Accidentologie, de Biomecanique et d'étude du comportement conducteur) and CEESAR (Centre Européen d'Etudes de Sécurité et d'Analyse des Risques) to the activities supporting this work.