

Research article

Feature group tabular transformer: a novel approach to traffic crash modeling and causality analysis

Oscar Lares, Hao Zhen and Jidong J. Yang*

Smart Mobility and Infrastructure Lab, College of Engineering, University of Georgia, Athens, GA 30602, USA

* **Correspondence:** Email: Jidong.Yang@uga.edu.

Academic Editor: Chih-Cheng Hung

Abstract: Reliable and interpretable traffic crash modeling is essential for understanding causality and improving road safety. This study introduced a novel approach to predicting crash types by utilizing a comprehensive dataset fused from multiple sources, including weather data, crash reports, high-resolution traffic information, pavement geometry, and facility characteristics. An essential part of our proposed approach was a feature group tabular transformer (FGTT) model, which organizes disparate data into meaningful feature groups, represented as tokens. These group-based tokens serve as rich semantic components, enabling effective identification of collision patterns and interpretation of causal mechanisms. The FGTT model was compared with widely used tree ensemble models, including random forest, XGBoost, and CatBoost, demonstrating better predictive performance. Furthermore, the attention heatmaps from the FGTT model revealed key influential factor interactions, providing fresh insights into the underlying causality of distinct crash types.

Keywords: road safety; crash analysis and modeling; causality analysis; tabular transformer; tree ensemble; feature grouping; attention heatmap

1. Introduction

According to the national highway traffic safety administration (NHTSA), an estimated 40,990 people died in motor vehicle traffic crashes in 2023 [1], highlighting the urgent need for enhanced road safety measures and policies. The inherent complexity of traffic dynamics, coupled with evasive factors, such as driver and human behavior, vehicular, and environmental conditions, makes crash modeling a challenging endeavor. To develop effective safety interventions, it is essential to gain a deeper understanding of the intricate and multifaceted interactions among road users, roadway infrastructure, and weather conditions. These can include enhancements to geometric design and traffic

control devices, the implementation of new traffic regulations, and the integration of advanced vehicle and infrastructure technologies. Collectively, these measures have the potential to significantly reduce the frequency and severity of traffic crashes.

Existing approaches to crash analysis have two critical gaps. First, traditional approaches rely on narrowly defined datasets, excluding critical variables such as detailed traffic patterns, environmental conditions, pavement characteristics, and driver behavior due to data accessibility challenges. This feature exclusion hampers the ability to model the multifaceted interactions among diverse factors. Second, while machine learning (ML) and deep learning (DL) models have shown strong predictive capabilities, their interpretability remains limited, particularly in deriving actionable insights for policy formulation or engineering design.

To address these challenges, we present a novel approach that leverages an extensive, integrated dataset encompassing weather conditions, traffic data, road geometry, pavement conditions, and driver characteristics. By fusing data from multiple sources, we aim to capture the full spectrum of factors underlying traffic crashes. Building upon this rich dataset, we introduce a novel model, namely the feature group tabular transformer (FGTT), designed to enhance both predictive performance and interpretability in crash type modeling. The FGTT organizes original features into semantically meaningful units or groups, such as traffic, event, vehicle, driver, environmental, geometric, pavement, and contextual factors. These distinct groups are then represented as tokens and used within a transformer architecture [2]. This semantic grouping allows the model to capture interdependencies among different feature groups and facilitates a more nuanced understanding of the factors contributing to different crash types. By explicitly modeling the interactions among feature groups through the attention mechanism, the FGTT enhances causal analysis in crash modeling, leveraging attention heatmaps to uncover critical insights.

In summary, our contributions are twofold:

(1) We compile and utilize a multisource dataset that integrates critical variables often omitted in the existing literature. This comprehensive dataset enables a more holistic analysis of traffic crashes by accounting for a wide range of influencing factors.

(2) We introduce the FGTT model, which improves predictive accuracy for crash types while enhancing interpretability through the analysis of attention weights associated with semantic feature groups. This model not only outperforms state-of-the-art ensemble methods such as random forest, XGBoost, and CatBoost but also provides valuable insights into the relationships among feature groups, aiding in the understanding of potential confounding effects common in crash data. The attention heatmaps analysis highlights the dominant role of event-specific details (e.g., vehicle maneuvers and crash locations) and their interactions with the driver feature group, leading to different crash types. Additionally, it uncovers more complex feature group interactions specific to angle crashes.

This study combines a comprehensive dataset with an innovative modeling approach to advance traffic safety analysis, providing both methodological advancements and practical insights to inform more effective road safety countermeasures and strategies.

2. Crash modeling

Crash collision and severity modeling has been a widely investigated topic utilizing various kinds of data. Typically, the approaches focus on modeling crash severity to identify factors mostly correlated with traffic crashes to provide effective countermeasures. However, not much emphasis has been placed on modeling and predicting crash types, as they are typically used as independent features within those studies. Kim et al. [3] conducted a study to model and predict crash types at intersections in rural Georgia. Their results highlight that countermeasures targeting crash severity might only address a subset of crashes. Besides crash severity, more effective countermeasures could be developed by understanding contributing factors to different crash types. The authors also note that crash types are associated with road geometry, environment, and traffic variables in different ways than crash severities are, which serve as justification for modeling them separately. While this study focuses on predicting the crash types, it is imperative to acknowledge the interconnection of crash types and severity. Typically, the methods employed for predicting these outcomes (crash severity vs. crash type) often overlap, leveraging similar datasets and analytical frameworks. Many crash severity studies and modeling approaches directly use the crash type as an input feature for modeling severity and risk [4–6]. Moreover, advanced driver-assistance systems (ADAS) and other interventions can also reduce the likelihood of certain types of collisions occurring in the first place [7]. Therefore, while the primary objective of this study focuses on crash types, a comprehensive literature review pertaining to predicting crash severity is also included.

Statistical models have widely been applied to analyze crash collision and risk with respect to varying features, offering a robust framework, and identifying key risk factors. Statistical methods offer some advantages when it comes to crash modeling due to their elegant forms and ease in modeling and interpretation. However, they are limited in adequately modeling complex relationships present within crash data. For example, they often assume explicit or fixed interactions among variables, potentially oversimplifying crash complexity and disregarding the dynamic interactions of variables, such as driver behavior, the environment, and traffic.

Additionally, the quality and availability of data can also significantly impact the model's accuracy and reliability, with issues such as missing data and class imbalance posing notable challenges. Zeng et al. [6] investigated the effects of real-time weather conditions and roadway geometry on the severity of freeway crashes. They utilized an ordered logit model due to the ordered nature of crash severity (low, medium, high). The variables included the hourly wind speed, temperature, precipitation, visibility, and humidity alongside the horizontal curvature and grading of the roadway, and other crash- and driver- related features. Their results indicated that heavier precipitation contributed more to medium severity crashes, and that more severe crashes tended to occur on freeway segments with a small horizontal curve radius and higher vertical gradients. However, they failed to incorporate essential traffic-related features, such as speed and traffic volume, which play a crucial role in crash types and severity [8, 9]. They acknowledged the fact that higher-resolution weather data could give additional insights into the crash severity outcomes.

Truck traffic impacts on crash severity have also been studied, since crashes involving trucks tend to be more severe due to the weight and size of the trucks involved. Xu et al. [10] looked at the effect of truck traffic characteristics on crash severity, leveraging weigh-in-motion (WIM) sensor data over a 5-year period. They utilized both vehicle-specific and summary data, along with geometric

roadway features, to reveal a correlation between mean vehicle weight and annual average daily traffic (AADT), and an increased risk of injuries and fatalities. Dutta and Fontaine [11] investigated the use of continuous count station (CCS) sensor data, including traffic speed and volume metrics, to improve crash prediction modeling. Their findings highlighted the significance of three time scales of those features (15-minute, hourly, yearly traffic data) and their impacts on model accuracy. Other studies have explored the influence of various features on truck traffic crash severity, such as environmental conditions, vehicle and driver characteristics, roadway geometry, and crash-specific features [12–15].

Aside from classic statistical approaches, various machine learning techniques have been employed for crash severity analysis and prediction [14, 16]. While statistical models often struggle to capture the intricate and non-linear relationships in crash and traffic data due to their inherent constraints, machine learning approaches offer a viable alternative. These methods excel in modeling complex relationships and handling high-dimensional feature spaces. Assi [4] utilized a hybrid approach comprised of principal component analysis (PCA) and support vector machines (SVM) as well as a multilayer-perceptron (MLP) using crash, roadway, driver, vehicle, and environmental characteristics to predict the severity of crashes. An increase in model performance and accuracy was noted when incorporating PCA features into the model over the initial features due to the reduction in dimensionality, however this study did not consider any traffic-related variables (e.g., speed or volume) which are known to influence the crash severity outcome.

Morris and Yang (2020) [17] conducted a comprehensive study on crash collision patterns on interstates in Georgia, leveraging linear discriminate analysis and extreme gradient boosting (XGBoost) to classify multi-vehicle crash collisions. Their comprehensive dataset included a wide range of variables, encompassing roadway, traffic, weather, environmental, and driver-related features. However, the final dataset size was relatively limited (approximately 3700 crash instances) and did not consider any pavement characteristics. In a separate study, Morris and Yang [18] presented an approach for addressing imbalanced data and analyzing the outcomes using three different machine learning methods. The imbalance inherently present in crash data can pose an issue for modeling due to the biases imposed as the minority classes can easily be overpowered by the majority class, resulting in poor prediction for the minority classes. To cope with the class imbalance, they investigated various resampling methods to balance the dataset more evenly using techniques like the synthetic minority oversampling technique (SMOTE) and adaptive synthetic sampling approach for imbalanced learning, and leveraging ensemble methods for prediction (CatBoost, XGBoost, random forest). Their results demonstrated that resampling methods enhanced crash collision predictions across all the models evaluated.

Deep learning methods have also been applied to crash modeling, leveraging their innate ability to handle large, high-dimensional datasets and to learn complex features and intricate relationships. However, these models can be computationally intensive and are susceptible to overfitting without careful tuning and proper regularization. While these models can deliver exceptional results, they often face the “black box problem”, where the lack of interpretability make it unclear why certain predictions were made. This opacity hinders the ability to derive actionable insights from the results, limiting their practical applicability [19]. Dong et al. [20] developed a deep learning model for predicting crash severity counts, utilizing a two-module architecture. The first module encodes inputs into feature representations, which are then processed by the second module for fine-tuned encoding before being passed to a regression layer to generate final crash severity predictions. The researchers evaluated their

model with and without the regression layer in the final module and compared its performance to a traditional SVM model. The proposed model with the regression layer performed the best, followed by that without the regression layer, and finally the SVM. These results highlighted that even without the regression layer, the encoded feature representations significantly contributed to prediction accuracy. The study also investigated how the model grouped closely related features (e.g., AADT, speed limit, truck percentage) into shared representation nodes, which were labeled as different representation groups based on their constituent features. Their findings indicated that geometric, pavement, and traffic feature representations had the most direct impacts on major injury crashes, minor injury crashes, and property-damage-only crashes, respectively.

Sattar et al. [21] also conducted a study predicting crash severity and comparing three deep learning models: an MLP, an MLP with embedding and TabNet, a popular tabular deep learning model based on transformer architectures. Their evaluation considered not only accuracy, but also precision, recall, and F1 scores, emphasizing the fact that misclassifying a severe injury is more crucial than less severe ones. Additionally, the study examined training times, given the resource-intensive nature of training deep learning models. The results showed that all three models yielded similar results for the severity classification task. However, the authors highlighted the significant differences in training times. The TabNet architecture, being more complex than the MLP models, required substantially longer training time, approximately 700% more. TabNet has also been explored to predict the duration of traffic incidents using tabular data [22]. This method demonstrated better performance compared to previous machine learning methods, excelling not only in prediction accuracy but also in interpretability. Its feature importance mechanism, driven by attention weights within the transformer architecture, enables a more intricate understanding of the model's outputs. Building on this foundation, our proposed FGTT introduces an innovative feature grouping strategy. By encoding semantically related features into tokens, the FGTT enhances both interpretability and predictive performance, offering a more transparent and effective framework for modeling traffic collisions and improving decision-making in traffic safety.

3. Data collection and compilation

Multiple sources were fused together to create a comprehensive dataset. Traffic counts, pavement condition data, and crash data were sourced from the Georgia Department of Transportation (GDOT), while public weather data was readily available from online platforms. This fused dataset enabled a more nuanced understanding of crash collisions. The dataset is not only used to identify the strongly correlated features with crash types but also to provide a broader perspective on the events surrounding crash instances, thereby enhancing crash modeling in ways that other studies may not have considered. Furthermore, this study specifically focused on multi-vehicle (MV) crashes. Previous studies [23, 24] have noted differences in the factors affecting single-vehicle (SV) crashes compared to MV crashes. For instance, Wang and Feng [23] found that factors like traffic volume and speed variance had no significant influence on SV crashes, while the opposite was true of MV crashes.

High-resolution traffic data, such as traffic counts and speeds, were obtained at 5-minute intervals from various CCSs throughout the state of Georgia. Crash incidents from police report data were first matched to these CCS stations both temporally and spatially. After the sites and crashes were matched, the high-resolution traffic data was aggregated into hourly intervals for modeling purposes. Weather

data for the crash instances was obtained from *Weather Underground* [25], with weather variables extracted from weather stations that were spatially matched to the corresponding CCS sites. Finally, pavement condition data was matched to the appropriate CCS sites based on geographic coordinates, providing pavement condition details for each crash instance.

After validating and confirming the correct matches among all crashes, CCS sites, pavement features, and weather variables, the data was fused into a single combined dataset. Subsequently, feature selection and reduction were carried out. Since the crash data comes from statewide police reports, several features were purposely excluded due to their post-incident nature. Given that this study aimed to identify the predictive variables prior to crash occurrence, after-effect features were excluded. Additionally, other crash-related features were also removed, particularly those influenced by the subjective judgment of the police officer writing the report, such as the contributing factors to the crash and the specific harmful event elements pertaining to each crash. Other redundant crash features from the police reports such as the road characteristics (e.g., whether the road was straight or curved), were omitted in favor of more accurate, quantified road and pavement measurements that better portray the road geometry and pavement conditions.

The compiled dataset had missing values for variables such as driver ages, traffic speeds, and precipitation data. These missing values were imputed using averages from relevant groups to maintain consistency across similar conditions. For precipitation accumulation data, missing values were imputed by calculating the group mean based on the categorical variables *City* and *Date_element*. The imputation was performed using a group-wise approach to preserve local weather patterns. Similarly, missing values in the *Hourly_avg_speed* column were imputed by grouping the data based on *No_lanes*, *Day_of_week*, *Facility_type*, *Area_type*, and *Time_of_day*, and replacing the missing values with the mean of each respective group. This method ensured that imputations accounted for contextual variations in traffic conditions. After processing, the final dataset consisted of 6810 MV crash instances and 33 features. Table 1 lists the variables analyzed in this study.

The dataset had a total of 14 numerical features and 19 categorical features. Each of the numerical and categorical feature statistics is shown in Tables 2 and 3.

Table 1. Total dataset description.

Feature	Description	Unit
<i>City</i>	City limits where the crash took place	-
<i>Crash_type</i>	Manner of collision of the car crash	-
<i>Crash_location</i>	Location on the roadway where the crash took place	-
<i>Lighting</i>	The lighting conditions noted when the crash occurred	-
<i>Surface</i>	The surface roadway condition noted when the crash occurred	-
<i>Driver1_safety_equip</i>	The safety equipment in use by the primary (at-fault) driver	-
<i>Driver2_safety_equip</i>	The safety equipment in use by the secondary driver	-
<i>Veh1_type</i>	Type of vehicle of the primary (at-fault) driver	-
<i>Veh2_type</i>	Type of vehicle of the secondary driver	-
<i>Veh1_maneuver</i>	Maneuver that the primary driver was carrying out to cause the crash	-

Continued on next page

Feature	Description	Unit
<i>Road_composition</i>	Make-up composition of the roadway where the crash occurred	-
<i>Trafficway_layout</i>	The general roadway description of how traffic flow is controlled	-
<i>Wind_speed</i>	Measured wind speed at the hour closest to the crash	mph
<i>Gust</i>	Measured gust at the hour closest to the crash	mph
<i>Precip_rate</i>	The precipitation rate for the hour closest to the crash	in/hr
<i>Precip_accum</i>	The total precipitation accumulation for the day up until an hour before the crash	inches
<i>Hourly_truck_ratio</i>	Truck percentage (FHWA Classes 4–13) of total traffic in the crash direction during the preceding hour	-
<i>Hourly_volume</i>	The total volume of traffic that was measured in the same roadway direction as the crash for the hour prior to the crash occurrence	count
<i>Hourly_avg_speed</i>	The average speed of traffic measured in the same roadway direction as the crash for the hour prior to the crash occurrence	mph
<i>IRI_avg</i>	International Roughness Index average	in/mile
<i>Rut_avg</i>	Longitudinal surface depression in the asphalt pavement of the average two-wheel paths measured between the width limits of the lane	inches
<i>Faulting_avg_3d</i>	Faulting values that fall across the entire lane	in/mile
<i>Heading_angle</i>	Measures the beginning and ending of the horizontal curvature to determine the heading	degrees
<i>Percent_grade</i>	Measures the grade classification of pavement sections	percentage
<i>Cross_section_slope</i>	Critical design element for pavement cross-sections	percentage
<i>Crack_percentage</i>	Percentage of pavement surface exhibiting cracking	percentage
<i>Day_of_week</i>	The day of the week the crash occurred	-
<i>Driver1_agerange</i>	The age range of the primary (at-fault) driver	-
<i>Driver2_agerange</i>	The age range of the secondary driver	-
<i>Curvature</i>	The curvature classification of the roadway	-
<i>Facility_type</i>	The classification type of the roadway facility	-
<i>Area_type</i>	The area where the roadway facility is located, either urban or rural	-
<i>Num_lanes</i>	The number of lanes of the roadway (for one direction)	count
<i>Time_of_day</i>	The time of day that the crash occurred	-

Table 2. Numerical feature statistics.

Feature	Mean	Std	Min	Max
<i>Wind_speed</i>	1.22	1.89	0	79
<i>Gust</i>	2.32	3.07	0	79
<i>Precip_rate</i>	0.01	0.09	0	3
<i>Precip_accum</i>	0.10	0.33	0	4.52
<i>Hourly_truck_ratio</i>	0.08	0.10	0	0.88
<i>Hourly_volume</i>	4370	2884	10	10,688
<i>Hourly_avg_speed</i>	49.06	15.56	4.18	80.27
<i>IRI_avg</i>	61.43	33.10	25	251
<i>Rut_avg</i>	0.107	0.061	0	0.41
<i>Faulting_avg_3d</i>	0.005	0.018	0	0.39
<i>Heading</i>	205	120.76	0.9	359.9
<i>Percent_grade</i>	-1.29	1.91	-7.40	3.70
<i>Cross_section_slope</i>	0.361	1.691	-4.70	3.4
<i>Crack_percentage</i>	8.11	10.74	0	57

Table 3. Categorical feature distributions.

Feature	Distribution of Categories
<i>City</i>	<i>Atlanta</i> : 2375 (40.16%) <i>Metro Area Outside of Atlanta</i> : 2202 (32.33%) <i>Unincorporated</i> : 1873 (27.5%)
<i>Crash_location</i>	<i>On Roadway - Non-Intersection</i> : 5513 (80.95%) <i>On Roadway - Crossing/Intersection/Crosswalk/Roundabout</i> : 717 (10.53%) <i>Entrance/Exit Ramp</i> : 359 (5.27%) <i>Private Property/Off Roadway</i> : 133 (1.95%) <i>Shoulder/Median/Gore</i> : 88 (1.29%)
<i>Lighting</i>	<i>Daylight</i> : 5066 (74.39%) <i>Dark-Lighted</i> : 1028 (15.09%) <i>Dark-Not Lighted</i> : 574 (8.43%) <i>Dawn</i> : 77 (1.13%) <i>Dusk</i> : 65 (0.95%)
<i>Surface</i>	<i>Dry</i> : 5465 (80.25%) <i>Wet/Snow/Ice</i> : 1345 (19.75%)
<i>Driver1_safety_equip</i>	<i>Lap/Shoulder Belt/Helmet Used</i> : 4911 (72.11%) <i>Unknown</i> : 1,787 (26.24%) <i>None Used</i> : 112 (1.64%)
<i>Driver2_safety_equip</i>	<i>Lap/Shoulder Belt/Helmet Used</i> : 5541 (81.37%) <i>Unknown</i> : 1208 (17.74%) <i>None Used</i> : 61 (0.90%)

Continued on next page

Feature	Distribution of Categories
<i>Veh1_type</i>	<i>Passenger Car/Pickup/Van/SUV</i> : 6060 (88.99%) <i>Truck/Trailer</i> : 457 (6.71%) <i>Unknown</i> : 206 (3.02%) <i>Other</i> : 55 (0.81%) <i>Motorcycle/Bicycle/ATV</i> : 32 (0.47%)
<i>Veh2_type</i>	<i>Passenger Car/Pickup/Van/SUV</i> : 6191 (90.91%) <i>Truck/Trailer</i> : 413 (6.06%) <i>Unknown</i> : 142 (2.09%) <i>Other</i> : 47 (0.69%) <i>Motorcycle/Bicycle/ATV</i> : 17 (0.25%)
<i>Veh1_maneuver</i>	<i>Straight</i> : 3809 (55.93%) <i>Changing Lanes/Passing</i> : 1886 (27.69%) <i>Negotiating a Curve</i> : 369 (5.42%) <i>Turning (left, right, u-turn)</i> : 285 (4.19%) <i>Other</i> : 251 (3.69%) <i>Backing</i> : 89 (1.31%) <i>Stopped/Parked</i> : 85 (1.25%) <i>Entering/Leaving Parking/Driveway</i> : 36 (0.53%)
<i>Road_composition</i>	<i>Black Top</i> : 5921 (86.95%) <i>Concrete/Other</i> : 889 (13.05%)
<i>Trafficway_layout</i>	<i>Two-Way Trafficway With A Physical Barrier/Separation</i> : 3762 (55.28%) <i>One-Way Trafficway</i> : 1275 (18.72%) <i>Two-Way Trafficway With No Physical Barrier/Separation</i> : 944 (13.86%) <i>Continuous Turning Lane</i> : 32 (0.47%)
<i>Day_of_week</i>	<i>Monday</i> : 963 (14.14%) <i>Tuesday</i> : 1096 (16.09%) <i>Wednesday</i> : 935 (13.73%) <i>Thursday</i> : 1135 (16.67%) <i>Friday</i> : 1282 (18.83%) <i>Saturday</i> : 784 (11.51%) <i>Sunday</i> : 615 (9.03%)
<i>Driver1_agerange</i>	<i>Under 25</i> : 1758 (25.81%) <i>25–34</i> : 1620 (23.79%) <i>35–44</i> : 1856 (27.25%) <i>45–54</i> : 723 (10.62%) <i>55 and up</i> : 853 (12.53%)
<i>Driver2_agerange</i>	<i>Under 25</i> : 1244 (18.27%) <i>25–34</i> : 1869 (27.44%) <i>35–44</i> : 1520 (22.32%) <i>45–54</i> : 1112 (16.33%) <i>55 and up</i> : 1065 (15.64%)

Continued on next page

Feature	Distribution of Categories
<i>Curvature</i>	<i>A</i> : 5869 (86.18%) <i>B</i> : 869 (12.76%) <i>C or more</i> : 72 (1.06%)
<i>Facility_type</i>	<i>Interstate</i> : 5342 (78.44%) <i>Principal Arterial - Other</i> : 751 (11.03%) <i>Minor Arterial</i> : 393 (5.77%) <i>Principal Arterial - Other Freeways and Expressways</i> : 324 (4.76%)
<i>Area_type</i>	<i>Urban</i> : 6,691 (98.25%) <i>Rural</i> : 119 (1.75%)
<i>Num_lanes</i>	2: 1746 (25.64%) 3: 618 (9.07%) 4: 585 (8.59%) 5: 954 (14.01%) 6: 709 (10.41%) 7: 2198 (32.28%)
<i>Time_of_day</i>	<i>Early morning</i> : 428 (6.28%) <i>Peak morning</i> : 1185 (17.40%) <i>Midday</i> : 1315 (19.31%) <i>Peak afternoon</i> : 2857 (41.95%) <i>Late evening</i> : 1025 (15.05%)

4. Crash type classification

In this study, we frame crash type inference as a multi-class classification problem. To enhance both performance and interpretability in tabular data modeling, we propose the FGTT, leveraging the transformer architecture [2]. The FGTT is evaluated against three tree ensemble baselines, including random forest (RF), extreme gradient boosting (XGBoost), and categorical boosting (CatBoost). These tree ensemble methods are well-established and widely recognized for their robust performance in supervised learning tasks involving tabular data. This section provides a brief overview of each ensemble method, followed by a detailed description of our proposed FGTT.

4.1. Baselines

Ensemble methods, which combine multiple decision trees to make predictions, have gained significant popularity due to their flexibility and robustness in handling large tabular datasets. Among the most widely used ensemble methods are random forest (RF) and gradient boosting. RF is a parallel ensemble method, while gradient boosting is a sequential ensemble method. Within gradient boosting, two state-of-the-art approaches are XGBoost and CatBoost.

4.1.1. Random forest

Random forest is a versatile machine learning algorithm. It operates by constructing a multitude of decision trees at the training time and outputting the most common prediction (for classification) or the

average prediction (for regression) from the individual trees. This helps reduce prediction variance and prevent overfitting, a common issue with single decision trees. The algorithm randomly selects subsets of both data and features at each split point, making it more robust. For classification, each tree votes for a class, and the class with the most votes becomes the final prediction. RF is particularly well-suited for predicting crash types due to its ability to handle complex, non-linear relationships between features like speed, weather conditions, and road conditions. For example, Khanum et al. [26] and Morris and Yang [17] both implemented RF in their approaches in modeling crash types and severity predictions, respectively. Given its effectiveness, RF was chosen as a competitive baseline for comparison in this study.

4.1.2. XGBoost

XGBoost is an advanced implementation of gradient boosting algorithms [27]. It has gained popularity in the machine learning community for its speed, performance, and versatility. Different from the conventional first-order tree ensemble, XGBoost is a second-order method with an objective function expressed in Eq (1).

$$\mathcal{L}(\phi) = \sum_{i=1}^n l(y_i, \hat{y}_i) + \sum_k \Omega(f_k), \quad (1)$$

where $l(y_i, \hat{y}_i)$ is the loss function measuring the discrepancy between the predicted \hat{y}_i and the actual target y_i , which is approximated by a second-order Taylor expansion; $\Omega(f_k)$ denotes the regularization term, which penalizes the complexity of tree f_k , typically by the number of nodes and the magnitude of leaf weights. XGBoost recursively chooses a feature split that maximizes the gain (or a reduction in loss). In the context of crash type classification, Yang et al. [28] utilized XGBoost for predicting crash severity and employed the Shapley additive explanation (SHAP) for model interpretation.

4.1.3. CatBoost

CatBoost, short for categorical boosting, is a state-of-the-art, open-source, gradient boosting library developed by Yandex [29]. It is specifically designed to work well with categorical data and is known for its performance, accuracy, and ease of use. CatBoost can efficiently handle categorical features through ordered target statistics while addressing the target leakage issue. Another distinguishing feature of CatBoost is ordered boosting, a permutation-based enhancement to traditional gradient boosting. This approach addresses bias and prediction shift by utilizing independent datasets at each gradient step, a process formally described by Eq (2).

$$\hat{y}_i^{(t)} = \hat{y}_i^{(t-1)} + \alpha \cdot h_t(x_i, \mathcal{D}_{\text{train}}^{<i}), \quad (2)$$

where

- $\hat{y}_i^{(t)}$ is the prediction for the i -th instance at iteration t ,
- $h_t(x_i, \mathcal{D}_{\text{train}}^{<i})$ is the model built on the training subset $\mathcal{D}_{\text{train}}^{<i}$, which includes only data points that precede x_i in the current permutation,
- α is the learning rate, controlling the step size of the boosting algorithm.

In addition to the ordered boosting, CatBoost uses oblivious trees as base predictors, which are computationally efficient and lead to faster predictions while reducing the risk of overfitting. As a

state-of-the-art model for tabular data with categorical features, CatBoost naturally serves as a strong baseline for crash type prediction in this study. Previous research has successfully applied CatBoost to crash modeling, as demonstrated by Hasan et al. [30], Li et al. [31], and Morris and Yang [18]. For interpreting tree ensemble models, the SHAP framework, developed by Lundberg et al. [32], has been widely adopted [18, 30, 31]. SHAP values quantify the contribution of each feature to a specific prediction by comparing the model’s output with and without the inclusion of that feature, offering insights into the model’s decision-making process.

4.2. Feature group tabular transformer

This section introduces the proposed method, the feature group tabular transformer (FGTT). We begin with a brief overview of transformers and the multi-head self-attention mechanism, followed by an explanation of our feature grouping approach.

4.2.1. Transformer encoder

Transformers [2] initially designed for sequential data, have become the architecture of choice for a wide range of natural language processing (NLP) and computer vision tasks. The core innovation of the transformer is the self-attention mechanism which enables the model to focus on relevant input tokens based on their relationships to one another. This mechanism, referred to as scaled dot-product attention, shown in Eq (3), computes the attention weight by extracting information from other tokens in the input sequence. This allows the model to effectively capture dynamic context, semantics, and relationships within the data.

$$\text{Attention}(Q, K, V) = \text{Softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V, \quad (3)$$

where Q and K are the query and key vectors, with the embedding dimension of d_k . V is the value vector. The dot products of the queries and keys are computed and divided by $\sqrt{d_k}$, followed by a softmax function to obtain the weights. This process is repeated across multiple heads n -times and results are concatenated together.

In this study, only the transformer encoder is employed, as its feature encoding is expected to enhance classification performance. The prediction approach is similar to that of Dong et al. [20], where a linear regression head was appended to the encoder model. However, in this work, an MLP classifier was used instead of a linear regression model to make final crash type predictions [21]. The transformer encoder architecture utilized in this study is illustrated in Figure 1.

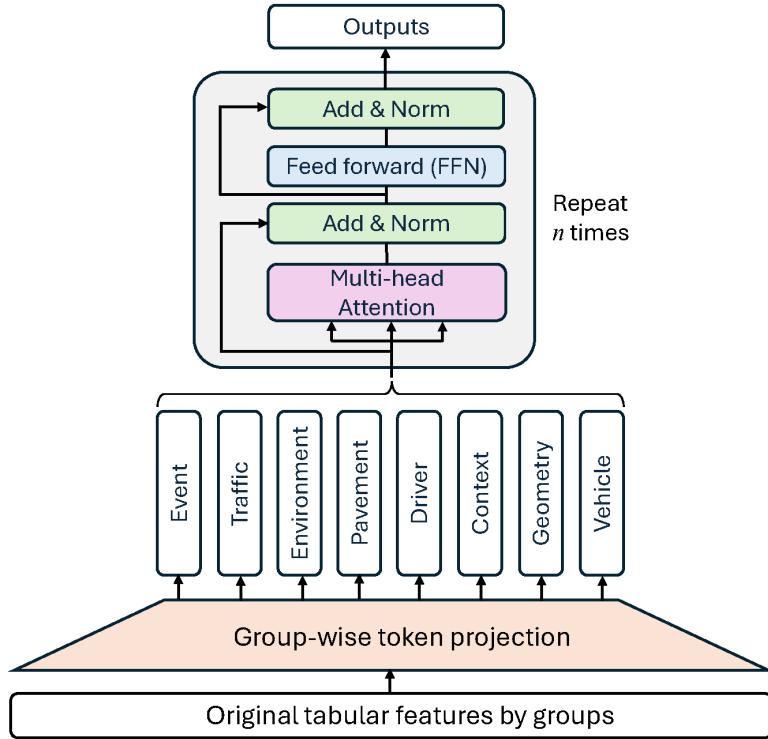


Figure 1. Transformer encoder.

4.2.2. Feature group tokens

Transformers are typically used in the NLP domain, where they process sequential words (or subwords) as vectorized tokens to learn context-sensitive semantic relationships between words to process sentences. Inspired by this concept, the proposed FGTT approach adapts this framework for tabular data in this study, where a *crash sequence* is created, consisting of distinct tokens that encapsulate various semantic aspects of a crash event. By assigning unique semantics to each token, this approach differs from tree-based models, which treat each original feature independently.

Instead of tokenizing each individual feature in the dataset, an additional step was introduced to encode features by meaningful groups that represent distinct semantic aspects of a crash event. To achieve this, related features were grouped based on their shared characteristics and roles in determining crash outcomes. This grouping reflects different aspects of the events surrounding a crash instance while addressing the overlap between certain features. Each feature group was then encoded into a token vector using an MLP projector, where the resulting token represented the specific semantic contribution of that feature group to the crash instance. This process produced a sequence of tokens, with each encapsulating a meaningful aspect of the crash. By organizing features into these groups, the approach aimed to provide a richer semantic representation of crashes compared to treating features independently. This richer representation was expected to better leverage the learning capability of the transformer encoder and deliver improved results over traditional ensemble methods. Table 4 shows the eight feature groups defined for this study. These groups were derived from features most commonly identified in the literature as influential in determining crash outcomes.

Table 4. Feature groups and corresponding features.

Feature group	Features
Event features	<i>Crash_location</i> <i>Veh1_maneuver</i>
Traffic features	<i>Hourly_truck_ratio</i> <i>Hourly_volume</i> <i>Hourly_avg_speed</i>
Environment features	<i>Gust</i> <i>Wind_speed</i> <i>Precip_rate</i> <i>Precip_accum</i> <i>Lighting</i>
Pavement features	<i>IRI_avg</i> <i>Rut_avg</i> <i>Faulting_avg_3d</i> <i>Percent_grade</i> <i>Cross_section_slope</i> <i>Crack_percentage</i> <i>Road_composition</i> <i>Surface</i>
Driver features	<i>Driver1_AgeRange</i> <i>Driver2_AgeRange</i> <i>Driver1_safety_equip</i> <i>Driver2_safety_equip</i>
Contextual features	<i>Day_of_week</i> <i>Time_of_day</i> <i>City</i>
Geometric features	<i>Heading_angle</i> <i>Curvature</i> <i>Trafficway_layout</i> <i>Number_of_lanes</i> <i>Facility_type</i> <i>Area_type</i>
Vehicle features	<i>Veh1_type</i> <i>Veh2_type</i>

4.2.3. Proposed FGTT model

For the proposed FGTT model, features from the dataset are first organized into distinct feature groups, each representing a specific aspect of the crash data. These groups, varying in size and dimensionality due to the differing number of features they contain, are then processed through a *Feature Group Token Projection* layer. This step ensures that all feature group tokens are projected to a uniform dimension which is a necessary condition for input into the transformer encoder. The

projection process not only standardizes the input but also transforms the feature tokens into a shared representation space, enabling the model to better capture and understand the relationships and interactions among different feature groups. For this purpose, an MLP model was utilized as the feature projector. Once the tokens are projected to a common dimension, they are passed through the transformer encoder block. The encoded outputs are then fed to an MLP classifier, followed by a softmax layer to predict probabilities of crash types. The proposed FGTT model is shown in Figure 2.

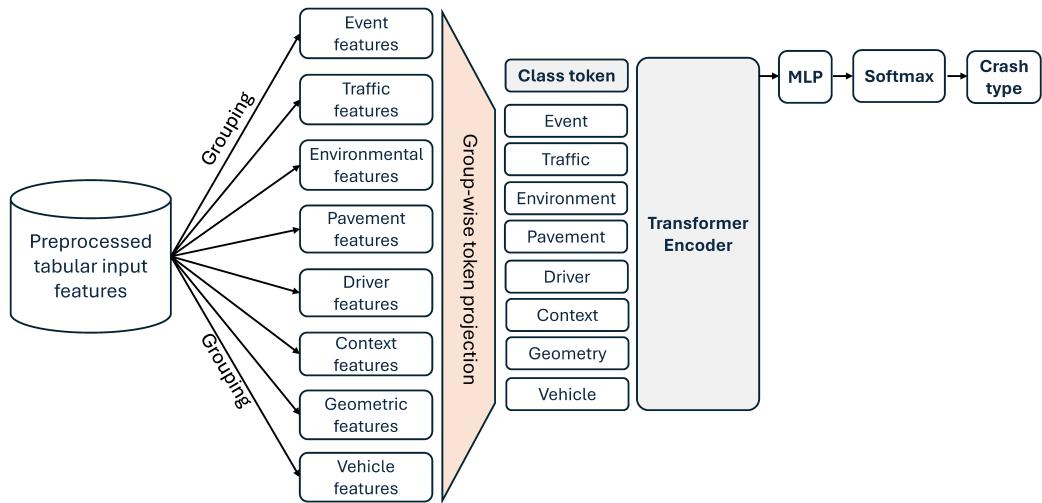


Figure 2. Illustration of the FGTT.

5. Experimental results and discussion

This section outlines the experiments conducted, including the implementation of each model discussed in Section 4 and their interpretation to better understand the possible crash causality chains. Specifically, SHAP value plots are used to explain the ensemble models, while attention weight heatmaps are generated for the FGTT model to provide deeper insights into the impact and interactions of the feature groups on crash type predictions.

Data partition and normalization are crucial steps in machine learning workflows. The original dataset, as described previously, was divided into three subsets: training, validation, and testing. Instead of using random data splitting, stratified splitting was adopted to ensure that each subset (i.e., training, validation, and testing) retains the same class distribution as the original dataset (13% of angle crashes; 29% of sideswipe crashes; and 58% of rear-end crashes). This is particularly important in scenarios like this study, where certain class labels, such as angle and sideswipe crashes, are underrepresented. As a result, the original dataset of 6810 multi-vehicle crashes was partitioned using stratified sampling based on the class label *Crash type*, leading to the following splits.

- Training Set: 6026 samples (88.5%);
- Validation Set: 392 samples (5.75%);
- Testing Set: 392 samples (5.75%).

Many of the features in this study are categorical. Given the small number of categories for each categorical variable, one-hot encoding is employed. To expedite and stabilize training, numeric features are standardized by subtracting the mean and divided by the standard deviation. To prevent information

leakage, the mean and standard deviation are computed training dataset and then applied consistently across all three subsets.

This standardization process ensures that each numeric feature contributes equally to the distance computations, which is important for metric-oriented learning algorithms, such as transformers. Standardizing the numeric data in this study is particularly important in this study, as some numeric features, such as the *Hourly_volume*, span a large range of values.

5.1. Ensemble methods

A GridSearch cross-validation (CV) strategy was employed for hyperparameter tuning of the ensemble models, including random forest, XGBoost, and CatBoost. CV is a commonly used technique in machine learning for hyperparameter tuning. It involves partitioning the training data into subsets, training the model on some subsets, and validating the model on the remaining subsets. This process is repeated multiple times, cycling through the subsets, to ensure that each subset serves as both training and validation data. The results are averaged to provide a final performance evaluation. To ensure proper class label distribution, the data subsets are stratified. GridSearch systematically explores the combinations of hyperparameters, cross-validating each combination to determine the best-performing set. The final models were trained with the optimal hyperparameters identified.

All tree ensemble models were trained using a consistent methodology, which included hyperparameter tuning through 5-fold GridSearch CV and leveraging a validation set to implement early stopping to mitigate the risk of overfitting. The hyperparameters for the 5-fold CV process for XGBoost, random forest, and CatBoost are listed in Table 5.

Table 5. Cross-validation hyperparameter tuning for XGBoost, random forest, and CatBoost (**bold** indicates the selected value).

Model	Cross-validation Parameters
Random Forest	η :{0.01, 0.05 , 0.1, 0.3}; <i>n_estimators</i> :{100, 200 , 500}; <i>max_depth</i> :{None, 10, 20, 30 }; <i>min_samples_split</i> :{2, 5 , 10}
XGBoost	η :{0.01, 0.05 , 0.1, 0.3}; <i>n_estimators</i> :{ 100 , 200, 500}; <i>max_depth</i> :{3, 5, 7 , 9}
CatBoost	η :{0.01, 0.05 , 0.1, 0.3}; <i>iterations</i> :{100, 200, 500, 1000, 2000 , 4000}; <i>depth</i> :{3, 5, 7 , 9}

5.2. FGTT

A different strategy was employed for tuning the parameters of the FGTT model. Bayesian optimization (BO) was used to efficiently find the optimal parameter settings. Unlike GridSearch, which exhaustively evaluates every combination of parameters, BO leverages a Gaussian process to construct a posterior distribution of the objective function based on prior evaluations. It then uses an acquisition function to decide where to sample next. This is particularly advantageous for tuning deep learning models, where the parameter space can be extremely large, and training can be computationally intensive and time-consuming. By intelligently choosing the next set of parameters to evaluate, BO significantly reduces the time and resources required as compared to the GridSearch method. This method has proven to be efficient and effective in tuning the hyperparameters of complex models [33, 34]. The Optuna package [35] was used for BO implementation.

The FGTT model was trained using focal loss, which was initially implemented for object detection tasks to address the challenge associated with class imbalance [36]. It achieves this by down-weighting the loss contribution from well-classified examples, typically belonging to majority classes, and focusing more on misclassified examples, often from minority classes. This adjustment enables the model to prioritize learning from underrepresented but significant classes, improving robustness and performance on imbalanced datasets where traditional loss functions may struggle. Focal loss has been successfully applied in various domains [37–39], demonstrating improved prediction outcomes. The parameter settings for the FGTT are listed in Table 6.

Table 6. Parameter settings for the FGTT.

Parameter	Range of Values Evaluated	Optimal Value
<i>learning rate</i>	[0.001, 0.1]	0.017
<i>optimizer</i>	<i>Adam, SGD, RMSProp</i>	<i>SGD</i>
<i>FFN dimension</i>	16, 24, 32, 64	64
<i>hidden dimension</i>	16, 24, 32, 64	64
<i>MLP dropout rate</i>	0.1, 0.2, 0.3, 0.4	0.2
<i>number of heads</i>	2, 3, 4, 6	4
<i>number of layers</i>	2, 3, 4, 5, 6	3

5.3. Evaluation metrics

Evaluating machine learning models on imbalanced datasets requires metrics that capture performance beyond simple accuracy. Although overall accuracy, defined as the ratio of correct predictions to total predictions, offers a general assessment, it falls short in scenarios with uneven class distributions. In such cases, metrics like precision, recall, and their harmonic mean, the F1 score, offer a more nuanced and comprehensive evaluation.

Precision measures the proportion of correctly predicted positive instances among all predicted positives:

$$\text{Precision} = \frac{TP}{TP + FP}. \quad (9)$$

Recall quantifies the ability to identify actual positive instances among all true positives:

$$\text{Recall} = \frac{TP}{TP + FN}. \quad (10)$$

Since Precision and Recall often exhibit a trade-off, balancing them is critical. The F1 score achieves this balance by calculating their harmonic mean:

$$F1 = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}. \quad (11)$$

Unlike the arithmetic mean, the F1 score penalizes extreme imbalances between Precision and Recall, making it suitable for tasks involving critical minority classes. In this study, where rare crash types, such as sideswipe and angle crashes, are of primary interest, the F1 score serves as a more reliable and informative metric than Accuracy.

6. Results

The model performance metrics are summarized in Table 7.

Table 7. Comparison of model performance.

Modeling Method	Label	Acc.	Prec.	Rec.	F1	Rear-end	Sideswipe	Angle
Random forest	Rear-end	93.4%	0.845	0.934	0.887	93.0%	6.6%	0.4%
	Sideswipe	77.2%	0.740	0.772	0.756	17.5%	77.2%	5.3%
	Angle	29.4%	0.682	0.294	0.411	37.3%	33.3%	29.4%
	<i>Weighted Avg</i>	80.4%	0.793	0.804	0.787	–	–	–
XGBoost	Rear-end	92.5%	0.850	0.925	0.886	92.5%	5.3%	2.2%
	Sideswipe	79.8%	0.752	0.798	0.775	15.8%	79.8%	4.4%
	Angle	27.5%	0.583	0.275	0.373	37.3%	35.3%	27.5%
	<i>Weighted Avg</i>	80.4%	0.787	0.804	0.787	–	–	–
CatBoost	Rear-end	91.6%	0.863	0.916	0.889	91.6%	5.7%	2.6%
	Sideswipe	74.6%	0.752	0.746	0.749	15.8%	74.6%	9.6%
	Angle	41.2%	0.553	0.412	0.472	29.4%	29.4%	41.2%
	<i>Weighted Avg</i>	80.1%	0.790	0.801	0.794	–	–	–
FGTT	Rear-end	92.1%	0.867	0.921	0.893	92.1%	5.7%	2.2%
	Sideswipe	77.2%	0.733	0.772	0.752	17.5%	77.2%	5.3%
	Angle	39.2%	0.645	0.392	0.488	23.5%	37.3%	39.2%
	<i>Weighted Avg</i>	80.9%	0.799	0.809	0.799	–	–	–

The FGTT model outperformed all ensemble models, achieving the highest overall weighted F1 score of 0.799, the highest precision of 0.799, the highest recall of 0.809, and the highest accuracy of 80.9%. This demonstrates the FGTT’s ability to deliver a more balanced performance across all crash types while maintaining strong overall predictive capability.

The CatBoost model also showed strong performance, with a weighted F1 score of 0.794, only slightly lower than the FGTT. Notably, it excelled in predicting the minority class, Angle, achieving the highest accuracy (41.2%) and recall (0.412) for this category. This highlights CatBoost’s effectiveness in handling imbalanced datasets and identifying the minority crash type. Additionally, CatBoost exhibited high precision (0.863) and a high F1 score (0.889) for the rear-end category, showcasing its reliability for majority class predictions.

The XGBoost model, while slightly trailing in overall performance, achieved a competitive weighted F1 score of 0.787. It displayed balanced predictive performance across categories but struggled in the Angle class, with the lowest accuracy (27.5%) and F1 score (0.373).

The random forest model excelled in predicting the majority class (Rear-end), achieving the highest accuracy (93.4%) and F1 score (0.887) among all models for this category. However, its performance declined for the minority class, Angle, with an accuracy of 29.4% and an F1 score of 0.411. Despite this, the RF model maintained a competitive overall weighted F1 score of 0.787 matching XGBoost.

In summary, the FGTT delivered the best overall performance, while CatBoost demonstrated better capabilities for minority class prediction. Both models showcased the advantages of ensemble learning techniques in achieving balanced and accurate predictions across crash types.

In addition to comparing overall model performance, further analysis and interpretation of model

predictions were performed using SHAP values for the ensemble methods and attention weight heatmaps for the proposed FGTT model.

SHAP values provided insights into feature importance by quantifying each feature's contribution to the prediction, enabling a deeper understanding of how individual features influenced the model's decisions for different crash types. Specifically, the SHAP value for a given feature measures its contribution to the final model prediction, weighted relative to the contributions of all other features. In this study, the SHAP values for the CatBoost model were averaged across all classes (rear-end, sideswipe, angle) since CatBoost has the best performance among the ensemble models. This analysis highlights the magnitude of each feature's impact on the model's prediction.

6.1. Feature importance by SHAP values for CatBoost

Figure 3 shows the top fifteen features based on SHAP values for the CatBoost model. The analysis highlights that the event features, such as the maneuver of vehicle 1 (*Veh1_maneuver*) and crash location, and traffic features are consistently the most significant features in inferring crash types, as evidenced by their high SHAP values. For the event features, the maneuver of vehicle 1 such as changing lanes or passing and keeping straight all have the most impact with the highest SHAP values for the three crash types. Specifically, the maneuver executed by the at-fault vehicle, such as changing lanes, traveling straight, or performing other actions, directly affects the positioning and interaction between vehicles, thereby influencing the type of crash. The crash location also plays an important role with higher SHAP values of *Crash_location_On roadway - Non-intersection* for rear-end and sideswipe crashes than angle crash.

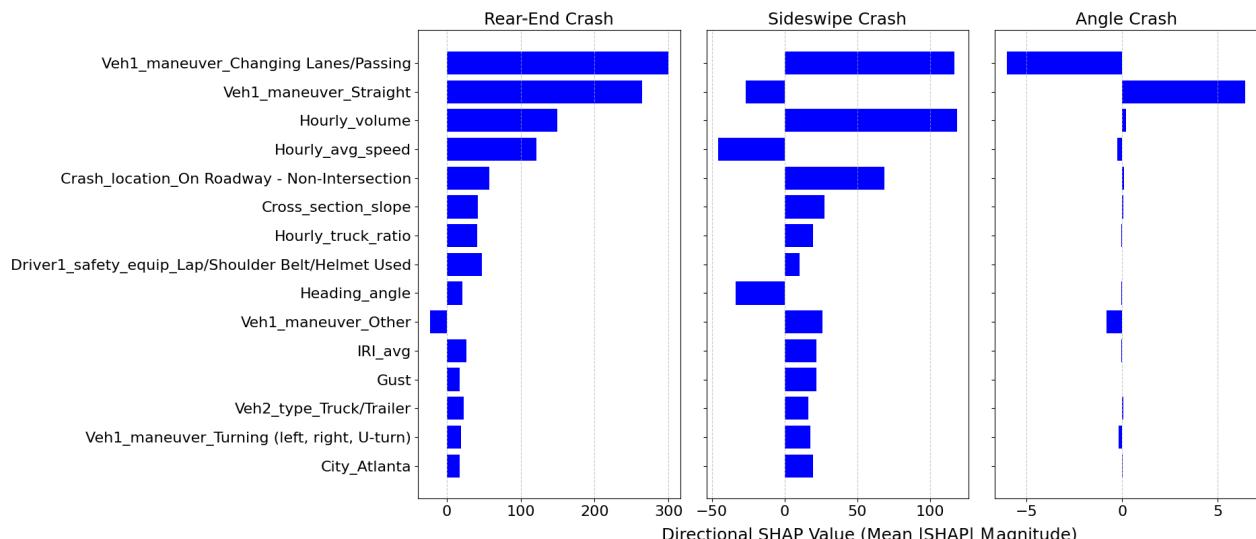


Figure 3. Directional SHAP values for CatBoost.

The traffic features, such as the hourly volume (*Hourly_volume*), hourly average speed (*Hourly_avg_speed*), and truck percentage (*Hourly_truck_ratio*), see detailed explanations in Table 1, are listed among the top seven most important features across crash types. These results underscore the critical role of vehicle behavior and traffic flow in affecting crash types. Similarly, hourly volume, speed, and truck ratio reflect traffic density, status, and exposure, which directly correlate

with crash risks. The hourly average speed emerges as a critical feature across crash types, with notable contributions in rear-end and sideswipe crashes. Its significance can be attributed to the direct relationship between speed and the likelihood of these crashes. High average speeds reduce the reaction time available to drivers, increasing the risk of rear-end collisions, especially in high-traffic scenarios. For sideswipe crashes, elevated speeds may amplify the difficulty of safely changing lanes or maintaining control during complex driving maneuvers. The *Hourly_truck_ratio*, which is unique in our dataset, emerges as a notable feature in the SHAP analysis, contributing to the understanding of crash type dynamics, particularly for rear-end and sideswipe collisions. This feature represents the proportion of trucks in the traffic flow during a given hour and is closely tied to crash likelihood due to the unique characteristics and limitations of trucks. The SHAP values indicate that a higher *Hourly_truck_ratio* generally increases the probability of sideswipe crashes, likely because trucks occupy more space on the roadway, have larger blind spots, and require greater distances to maneuver safely. These factors can create conditions that make lane changes and overtaking more challenging for other vehicles, leading to sideswipe incidents. In addition, *Hourly_truck_ratio* plays a role in rear-end crashes, albeit with a smaller magnitude of contribution. The large size, heavy weight, and slower acceleration of trucks can create speed differentials in mixed traffic, particularly during peak hours, where traffic density is high. Such conditions may increase the likelihood of rear-end collisions when vehicles following trucks fail to adjust their speed or maintain a safe following distance. These findings underscore the need for targeted interventions, such as lane restrictions for trucks during peak traffic hours, better signage to alert drivers of truck-heavy zones, and enhanced training for drivers on navigating safely around trucks.

Additionally, other factors, including pavement features (*cross-section slope*, *IRI_avg*), geometric features (*Heading_angle*), and weather-related features (*gust*), provide further insights into the multifaceted nature of crash dynamics.

6.2. FGTT attention heatmaps

To better understand feature importance for the FGTT model, attention weights were extracted from the last attention layer and examined at different levels. In the context of attention mechanisms, there are two key elements, the query and key tokens, that play a pivotal role. These tokens are transformed representations of the input data. Specifically, query tokens are used to probe the input data, while key tokens are matched against these queries. The model evaluates the similarity between each query token and all key tokens to compute attention scores. These scores quantify the relevance of each part of the input sequence to the query. Using these scores, the model generates a weighted sum of values, an amalgamation of contextually embedded information that highlights the features deemed most relevant for making predictions. In this study, attention weights for the FGTT model were extracted for three distinct crash types from the test set and passed through the final trained model to identify the features most associated with each crash type. The attention weights were obtained from the last layer of the transformer encoder, which is the layer that produces the final transformer output used for predictions.

Figure 4 illustrates the class token attention scores toward each feature group for crash type inference. These attention scores signify the importance of each feature group in predicting three crash types: rear-end, sideswipe, and angle.

The results show that the event feature group consistently receives the highest attention scores across all crash types from the class token. This indicates the dominant role of event-specific details,

such as vehicle maneuvers and crash locations, affecting the risks of different crash types. Notably, sideswipe crashes exhibit higher attention to the event feature group compared to angle crashes and rear-end crashes, likely reflecting the critical influence of dynamic actions, such as lane changes, which are commonly associated with these types of collisions. In contrast, rear-end and angle crashes may depend less on immediate events and more on environmental, contextual, and pavement factors.

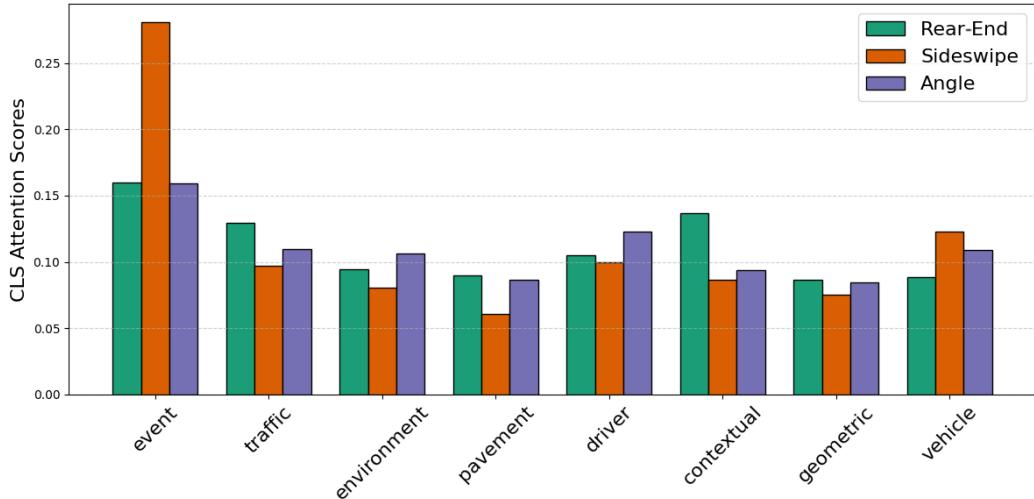


Figure 4. Class token (CLS) attention scores by feature groups.

The traffic, environment, pavement, driver, and contextual feature groups demonstrate moderate attention scores, suggesting their relevance in providing additional context for crash type prediction. For example, angle crashes exhibit relatively higher attention to environmental features, which may highlight the role of conditions like *gust*, *wind speed*, *precipitation*, and *lighting* in angle crash scenarios, particularly on the driver's driving perception and pavement conditions. Similarly, the driver and contextual feature groups receive comparable attention across crash types, suggesting the importance of driver characteristics, potentially driver behavior indicated by the driver characteristics, and contextual factors, such as traffic patterns, reflected by the time of day and the day of the week. These features likely serve as secondary factors that modulate the primary event-driven dynamics of crashes. This finding aligns with the understanding that immediate crash events and driver interactions often have a stronger impact on outcomes than static roadway or vehicle characteristics.

Figures 5–7 present the attention heatmaps depicting the interactions between different feature groups for three crash types: rear-end, sideswipe, and angle. The intensity of the attention weights in these heatmaps reflects the importance of the relationship for each feature group pair in the model's predictions.

Analyzing the heatmaps reveals distinct patterns of feature group interactions for each crash type. Comparing the heatmaps across crash types reveals that certain feature group interactions are universally influential, while others are crash-type specific. One notable finding across the three heatmaps is that the event feature group shows more influence than other feature groups across the three crash types. For the rear-end crash type, there are prominent interactions between feature groups, including *event* and *geometric*, *traffic* and *driver*, *event* and *vehicle*, *event* and *contextual*, *event* and *driver* feature groups. For sideswipe crash type, the heatmap highlights substantial interactions of the

event feature group with the *driver*, *contextual*, *vehicle* feature groups.

Moreover, the *event* feature group has consistent strong attention toward the *driver* feature group across all three heatmaps. The persistent interaction underscores the fundamental role of *driver-related* features, or potential latent factors, in the occurrence of different crash types.

In the angle crash heatmap, there is more pronounced emphasis on overall feature group interactions compared to rear-end and sideswipe crashes. This may suggest a more complex combined effect of feature group interactions on the model's inference for angle crashes. Notably, the strong attention associated with *driver-event* feature group pairs further highlights their significant influence in predicting the angle crash type.

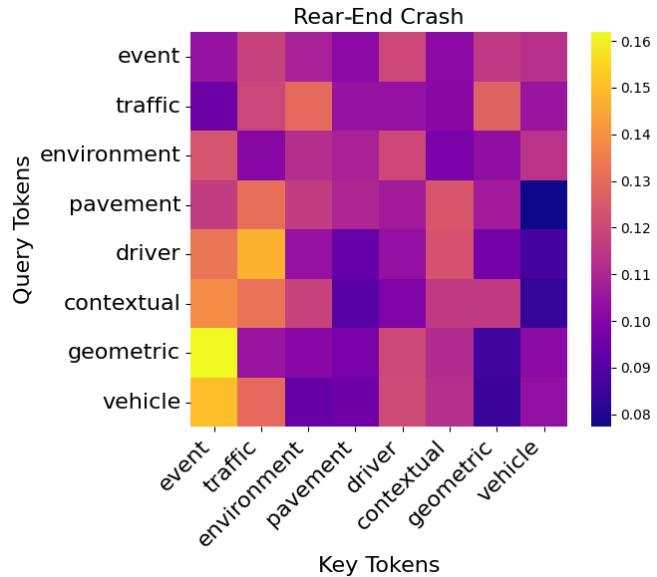


Figure 5. Feature groups attention heatmap for rear-end crashes.

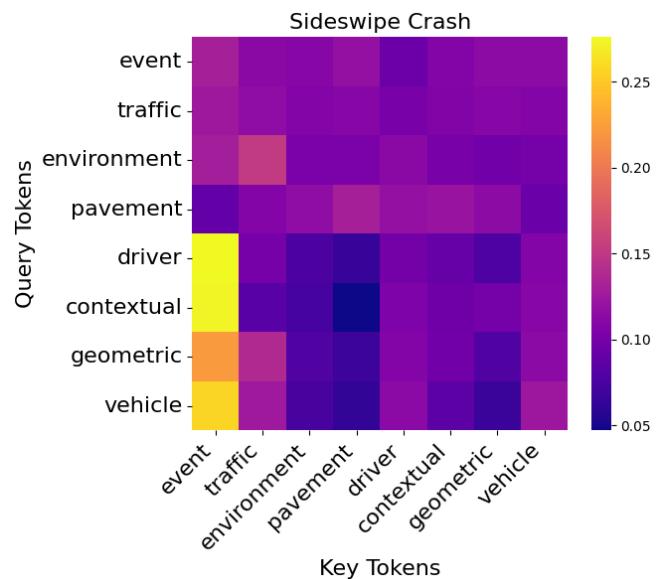


Figure 6. Feature groups attention heatmap for sideswipe crashes.

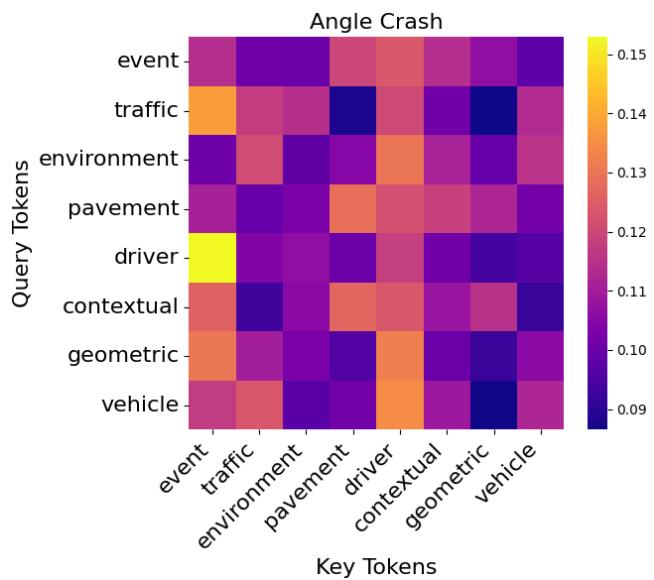


Figure 7. Feature groups attention heatmap for angle crashes.

These findings demonstrate the model’s ability to capture the intricate relationships among different feature groups in inferring crash types. The highlighted feature group interactions specific to each crash type reveal the potential for developing targeted intervention strategies. For instance, mitigating rear-end crashes could involve prioritizing driver assistance technologies and improving pavement conditions. Reducing sideswipe crashes might require redesigning roadway geometries and managing traffic flows. Addressing angle crashes could benefit from enhancing intersection designs and implementing environmental hazard warning systems.

In summary, the attention heatmaps offer valuable insights into how the transformer-based model utilizes feature group interactions to distinguish between crash types. Understanding these interactions not only enhances the interpretability of the model but also guides the development of targeted safety measures to mitigate specific crash types effectively.

7. Conclusions and future directions

Drawing from the collective insights gathered throughout the research, it becomes evident that predictive modeling of traffic crash outcomes benefits significantly from the integration of diverse data types and sources. The primary contributions of this study lie in the use of a comprehensive and highly descriptive dataset and the introduction of the FGTT, which efficiently processes and encodes data into group-specific tokens, enabling it to uncover complex relationships between feature groups through attention heatmaps. These semantically rich tokens enhance the model’s ability to capture nuanced interactions, resulting in both improved predictive performance and greater interpretability. While the FGTT model demonstrated better performance in our experiments compared to popular ensemble methods, opportunities for further refinement remain. In this study, the token feature groups were constructed using carefully selected features identified as semantically similar based on domain knowledge. Future research could explore more flexible grouping schemes that allow the model to dynamically learn optimal groupings to enhance the semantic understanding of feature interactions.

and further improve the model’s ability to infer crash outcomes.

Further insights were gained by examining the attention weights extracted from the FGTT model, providing a deeper understanding of the factors influencing different crash types. In particular, for angle crashes, distinct attention patterns emerged, with driver and environment features receiving greater weight compared to the other two crash types. For sideswipe and rear-end crashes, event-related features, such as crash location and vehicle maneuvers, and contextual-related features are more prominent. Moreover, the event feature group has relatively strong attention toward the driver feature group across all three crash types, underscoring the importance of driver features. Considering the growing emphasis on explainable machine learning, enhancing the interpretability of complex models like the FGTT is critical. Developing methods to make the model’s decision-making more comprehensible will not only bolster trust but also support its adoption in policy-making and real-world applications, where transparency and interpretability are essential.

Furthermore, the comparative evaluation of different modeling methods underscored the effectiveness of the CatBoost model in handling the minority classes, evidenced by its leading performance in predicting the angle crash category. Future research should explore SMOTE or generative models, such as generative adversarial networks or variational autoencoders, to synthetically augment minority classes with high-fidelity samples [40]. These approaches could help create a more balanced training dataset, ultimately enhancing the model’s predictive power across classes. Additionally, the exploration of alternative data augmentation techniques, innovative feature engineering strategies, and the integration of physics-guided or domain-specific embeddings could also be invaluable.

From a data perspective, in addition to CCS data utilized in this study, weigh-in-motion (WIM) data is another valuable data source than could enhance crash outcome prediction. Previous studies have identified correlations between truck weights and crash outcomes [10, 41], suggesting that incorporating WIM data could improve model performance. However, a challenge in utilizing WIM data, however, is its sparse distribution, which would limit the amount of crash instances available for analysis at these WIM sites. While this study considered hourly summaries, further insights can be gained through the exploration and usage of finer time resolutions (e.g., 5 minutes or 15 minutes), as Dutta and Fontaine [11] found that diverse ways to aggregate the traffic data had impacts on their modeling results.

The growth of the electric vehicle (EV) and connected and autonomous vehicles (CAV) present new opportunities for future research. Many new EVs on the market are equipped with an array of cameras and sensors that can provide a richer data source, offering deeper insights into driver behavior, actions, and vehicle maneuvers. In conclusion, this research represents a significant step forward in enhancing road safety analytics. By leveraging the FGTT’s ability to uncover complex relationships within the data through semantically grouped features, it lays the groundwork for more informed and effective safety strategies. Additionally, the comprehensive dataset utilized in this study underscores the importance of considering diverse features that impact crash outcomes. It is the convergence of data-driven insights and technological innovation that will drive progress toward safer roadways and more effective policy making for a safer future.

Use of AI tools declaration

The authors declare they have not used Artificial Intelligence (AI) tools in the creation of this article.

Conflict of interest

Jidong J. Yang is an editorial board member for Applied Computing and Intelligence and was not involved in the editorial review and the decision to publish this article.

References

1. *National Highway Traffic Safety Administration, Early estimate of motor vehicle traffic fatalities in 2023*, U.S. Department of Transportation, 2023. Available from: <https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/813561>.
2. A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. Gomez, et al., Attention is all you need, *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017, 6000–6010.
3. D. Kim, S. Washington, J. Oh, Modeling crash types: new insights into the effects of covariates on crashes at rural intersections, *J. Transp. Eng.*, **132** (2006), 282–292. [https://doi.org/10.1061/\(ASCE\)0733-947X\(2006\)132:4\(282\)](https://doi.org/10.1061/(ASCE)0733-947X(2006)132:4(282))
4. K. Assi, Traffic crash severity prediction—a synergy by hybrid principal component analysis and machine learning models, *Int. J. Environ. Res. Public Health*, **17** (2020), 7598. <https://doi.org/10.3390/ijerph17207598>
5. G. Kadilar, Effect of driver, roadway, collision, and vehicle characteristics on crash severity: a conditional logistic regression approach, *Int. J. Inj. Control Sa.*, **23** (2016), 135–144. <https://doi.org/10.1080/17457300.2014.942323>
6. Q. Zeng, W. Hao, J. Lee, F. Chen, Investigating the impacts of real-time weather conditions on freeway crash severity: a bayesian spatial analysis, *Int. J. Environ. Res. Public Health*, **17** (2020), 2768. <https://doi.org/10.3390/ijerph17082768>
7. R. Spicer, A. Vahabaghiae, D. Murakhovsky, G. Bahouth, B. Drayer, S. Lawrence, Effectiveness of advanced driver assistance systems in preventing system-relevant crashes, *SAE Int. J. Adv. Curr. Prac. in Mobility*, **3** (2021), 1697–1701. <https://doi.org/10.4271/2021-01-0869>
8. M. Ahmed, M. Abdel-Aty, The viability of using automatic vehicle identification data for real-time crash prediction, *IEEE Trans. Intell. Transp.*, **13** (2011), 459–468. <https://doi.org/10.1109/TITS.2011.2171052>
9. A. Høye, I. Hesjevoll, Traffic volume and crashes and how crash and road characteristics affect their relationship—a meta-analysis, *Accident Anal. Prev.*, **145** (2020), 105668. <https://doi.org/10.1016/j.aap.2020.105668>
10. C. Xu, K. Ozbay, H. Liu, K. Xie, D. Yang, Exploring the impact of truck traffic on road segment-based severe crash proportion using extensive weigh-in-motion data, *Safety Sci.*, **166** (2023), 106261. <https://doi.org/10.1016/j.ssci.2023.106261>

11. N. Dutta, M. Fontaine, Improving freeway segment crash prediction models by including disaggregate speed data from different sources, *Accident Anal. Prev.*, **132** (2019), 105253. <https://doi.org/10.1016/j.aap.2019.07.029>
12. R. Avelar, K. Dixon, S. Ashraf, A comparative analysis on performance of severe crash prediction methods, *Transport. Res. Rec.*, **2672** (2018), 109–119. <https://doi.org/10.1177/0361198118794052>
13. M. Bedard, G. Guyatt, M. Stones, J. Hirdes, The independent contribution of driver, crash, and vehicle characteristics to driver fatalities, *Accident Anal. Prev.*, **34** (2002), 717–727. [https://doi.org/10.1016/S0001-4575\(01\)00072-0](https://doi.org/10.1016/S0001-4575(01)00072-0)
14. A. Iranitalab, A. Khattak, Comparison of four statistical and machine learning methods for crash severity prediction, *Accident Anal. Prev.*, **108** (2017), 27–36. <https://doi.org/10.1016/j.aap.2017.08.008>
15. B. Naik, L. Tung, S. Zhao, A. Khattak, Weather impacts on single-vehicle truck crash injury severity, *J. Safety Res.*, **58** (2016), 57–65. <https://doi.org/10.1016/j.jsr.2016.06.005>
16. K. Santos, J. Dias, C. Amado, A literature review of machine learning algorithms for crash injury severity prediction, *J. Safety Res.*, **80** (2022), 254–269. <https://doi.org/10.1016/j.jsr.2021.12.007>
17. C. Morris, J. Yang, Understanding multi-vehicle collision patterns on freeways—a machine learning approach, *Infrastructures*, **5** (2020), 62. <https://doi.org/10.3390/infrastructures5080062>
18. C. Morris, J. Yang, Effectiveness of resampling methods in coping with imbalanced crash data: crash type analysis and predictive modeling, *Accident Anal. Prev.*, **159** (2021), 106240. <https://doi.org/10.1016/j.aap.2021.106240>
19. A. Géron, *Hands-on machine learning with Scikit-Learn, Keras, and TensorFlow*, Sebastopol: O'Reilly Media, Inc., 2022.
20. C. Dong, C. Shao, J. Li, Z. Xiong, An improved deep learning model for traffic crash prediction, *J. Adv. Transport.*, **2018** (2018), 3869106. <https://doi.org/10.1155/2018/3869106>
21. K. Sattar, F. Chikh Oughali, K. Assi, N. Ratrou, A. Jamal, S. Masiur Rahman, Transparent deep machine learning framework for predicting traffic crash severity, *Neural Comput. Applic.*, **35** (2023), 1535–1547. <https://doi.org/10.1007/s00521-022-07769-2>
22. H. Li, Y. Li, A novel explanatory tabular neural network to predicting traffic incident duration using traffic safety big data, *Mathematics*, **11** (2023), 2915. <https://doi.org/10.3390/math11132915>
23. X. Wang, M. Feng, Freeway single and multi-vehicle crash safety analysis: influencing factors and hotspots, *Accident Anal. Prev.*, **132** (2019), 105268. <https://doi.org/10.1016/j.aap.2019.105268>
24. S. Geedipally, D. Lord, Investigating the effect of modeling single-vehicle and multi-vehicle crashes separately on confidence intervals of poisson—gamma models, *Accident Anal. Prev.*, **42** (2010), 1273–1282. <https://doi.org/10.1016/j.aap.2010.02.004>
25. *Weather underground, Wundermap weather information*, The Weather Company, LLC., 2023. Available from: <https://www.wunderground.com/wundermap?lat=33.751&lon=-84.39>.

26. H. Khanum, A. Garg, M. Faheem, Accident severity prediction modeling for road safety using random forest algorithm: an analysis of indian highways, *F1000Res.*, **12** (2023), 494. <https://doi.org/10.12688/f1000research.133594.2>

27. T. Chen, C. Guestrin, Xgboost: a scalable tree boosting system, *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, 785–794. <https://doi.org/10.1145/2939672.2939785>

28. Y. Yang, K. Wang, Z. Yuan, D. Liu, Predicting freeway traffic crash severity using xgboost-bayesian network model with consideration of features interaction, *J. Adv. Transport.*, **2022** (2022), 4257865. <https://doi.org/10.1155/2022/4257865>

29. L. Prokhorenkova, G. Gusev, A. Vorobev, A. Dorogush, A. Gulin, Catboost: unbiased boosting with categorical features, *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, 2018, 6639–6649.

30. A. Hasan, M. Jalayer, S. Das, M. Kabir, Application of machine learning models and shap to examine crashes involving young drivers in new jersey, *International Journal of Transportation Science and Technology*, **14** (2024), 156–170. <https://doi.org/10.1016/j.ijtst.2023.04.005>

31. J. Li, X. Wang, X. Yang, Q. Zhang, H. Pan, Analyzing freeway safety influencing factors using the catboost model and interpretable machine-learning framework, shap, *Transport. Res. Rec.*, **2678** (2024), 563–574. <https://doi.org/10.1177/03611981231208903>

32. S. Lundberg, G. Erion, S. Lee, Consistent individualized feature attribution for tree ensembles, arXiv: 1802.03888. <https://doi.org/10.48550/arXiv.1802.03888>

33. R. Turner, D. Eriksson, M. McCourt, J. Kiili, E. Laaksonen, Z. Xu, I. Guyon, Bayesian optimization is superior to random search for machine learning hyperparameter tuning: analysis of the black-box optimization challenge 2020, *Proceedings of the NeurIPS 2020 Competition and Demonstration Track*, 2021, 3–26.

34. A. Victoria, G. Maragatham, Automatic tuning of hyperparameters using bayesian optimization, *Evol. Syst.*, **12** (2021), 217–223. <https://doi.org/10.1007/s12530-020-09345-2>

35. T. Akiba, S. Sano, T. Yanase, T. Ohta, M. Koyama, Optuna: a next-generation hyperparameter optimization framework, *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2019, 2623–2631. <https://doi.org/10.1145/3292500.3330701>

36. T. Lin, P. Goyal, R. Girshick, K. He, P. Dollár, Focal loss for dense object detection. *Proceedings of IEEE International Conference on Computer Vision (ICCV)*, 2017, 2980–2988. <https://doi.org/10.1109/ICCV.2017.324>

37. J. Tian, P. Tsai, K. Zhang, X. Cai, H. Xiao, K. Yu, et al., Synergetic focal loss for imbalanced classification in federated xgboost, *IEEE Transactions on Artificial Intelligence*, **5** (2023), 647–660. <https://doi.org/10.1109/TAI.2023.3254519>

38. X. Liu, J. Lu, X. Chen, Y. Fong, X. Ma, F. Zhang, Attention based spatio-temporal graph convolutional network with focal loss for crash risk evaluation on urban road traffic network based on multi-source risks, *Accident Anal. Prev.*, **192** (2023), 107262. <https://doi.org/10.1016/j.aap.2023.107262>

- 39. R. Yu, Y. Wang, Z. Zou, L. Wang, Convolutional neural networks with refined loss functions for the real-time crash risk analysis, *Transport. Res. C-Emer.*, **119** (2020), 102740. <https://doi.org/10.1016/j.trc.2020.102740>
- 40. M. Hernandez, G. Epelde, A. Alberdi, R. Cilla, D. Rankin, Synthetic data generation for tabular health records: a systematic review, *Neurocomputing*, **493** (2022), 28–45. <https://doi.org/10.1016/j.neucom.2022.04.053>
- 41. Y. Jo, C. Oh, S. Kim, Estimation of heavy vehicle-involved rear-end crash potential using wim data, *Accident Anal. Prev.*, **128** (2019), 103–113. <https://doi.org/10.1016/j.aap.2019.04.005>



AIMS Press

© 2025 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0>)