

**Pedestrian and Bicyclist Safety at Highway-Rail Grade Crossings –
Assessment of Crash Predictions (Year 2/Phase 2 Report)**

Aemal J. Khattak, Ph.D.

Professor and Director
Mid-America Transportation Center
Department of Civil and Environmental
Engineering
University of Nebraska-Lincoln
330E Prem Paul Research Center at Whittier
School
2200 Vine Street, Lincoln 68588 Nebraska
khattak@unl.edu, (402) 472-8126
ORIC ID: 0000-0002-6695-9842

Moomal Bukhari

Graduate Research Assistant
School of Computing
University of Nebraska-Lincoln
123 Avery Hall
1144 T Street, Lincoln, NE 68588
mbukhari2@unl.edu
ORICD ID: 0009-0007-5028-6456

Muhammad Naveed Aman, Ph.D.

Assistant Professor
School of Computing
University of Nebraska-Lincoln
263 Avery Hall
1144 T Street, Lincoln, NE 68588
ORIC ID: 0000-0002-4629-7589

A Report on Research Sponsored by

University Transportation Center for Railway Safety (UTCRS)

University of Nebraska Lincoln (UNL)

January 2026

Technical Report Documentation Page

1. Report No. UTCRS-UNL-O4CY24	2. Government Accession No.	3. Recipient's Catalog No.	
4. Title and Subtitle Pedestrian and Bicyclist Safety at Highway-Rail Grade Crossings – Assessment of Crash Predictions (Year 2 Report/Phase II Report)		5. Report Date January 31, 2026	
		6. Performing Organization Code UTCRS-UNL	
7. Author(s) Aemal J. Khattak; M. Naveed Aman; Moomal Bukhari		8. Performing Organization Report No. UTCRS-UNL-O4CY24	
9. Performing Organization Name and Address University Transportation Center for Railway Safety (UTCRS) University of Nebraska Lincoln (UNL) 1400 R Street Lincoln, NE United States 68588		10. Work Unit No. (TRAIS)	
		11. Contract or Grant No. 69A3552348340	
12. Sponsoring Agency Name and Address U.S. Department of Transportation (USDOT) University Transportation Centers Program 1200 New Jersey Ave. SE Washington, DC, 20590		13. Type of Report and Period Covered Project Report June 1, 2024 – January 31, 2026	
		14. Sponsoring Agency Code USDOT UTC Program	
15. Supplementary Notes			
16. Abstract This report presents results of a multi-state effort to improve pedestrian and bicyclist safety at highway-rail grade crossings (HRGCs) by explicitly incorporating non-motorist traffic exposure into crash prediction models. Building on an AI-based video pipeline from previous research, empirical non-motorist counts from 18 rail crossings in Lincoln, Nebraska, were used to calibrate a regression model that predicted daily pedestrian and bicycle traffic volumes at 13,672 public rail crossings in Nebraska, Kansas, Missouri, and Iowa. These traffic volume predictions were integrated into a modified Accident Prediction and Severity (APS) framework using a Zero-Inflated Negative Binomial (ZINB) model, showing an average of 5.5% increase in predicted crashes and much larger increases at crossings with substantial non-motorist traffic. A two-stage AI approach, Random Forest classification followed by Random Forest regression, further improved predictive performance relative to the ZINB and provided a practical tool for screening and prioritizing high-risk rail crossings. The findings demonstrate that non-motorist traffic is a critical but previously underrepresented factor in rail crossing safety assessment and support a scalable, data-driven framework for multimodal safety management.			
17. Key Words Railroad Grade Crossings, Bicycle Crossings, Crosswalks, Safety, Artificial Intelligence, Object Detection, Data Science		18. Distribution Statement	
19. Security Classification (of this report) None	20. Security Classification (of this page) None	21. No. of Pages 45	22. Price

Table of Contents

List of Figures	5
List of Tables	5
Disclaimer	6
Acknowledgements	6
Executive Summary	7
1. Introduction	9
1.1. Problem Statement	10
1.2. Objectives	10
1.3. Research Framework	11
1.4. Organization of the Report	12
2. Literature Review	13
2.1. HRGC Crash Prediction Models	13
2.2. Non-Motorist Safety at Grade Crossings	13
2.3. AI Applications in Transportation Safety	14
2.4. Gaps in Existing Research	15
3. Methodology	17
3.1. Phase I: Non-Motorist Volume Prediction	17
3.1.1. Data Collection	17
3.1.2. Object Detection via AI Models	17
3.1.3. Non-Motorist Traffic Volume Prediction Model Preparation	17
3.1.4. Non-Motorist Traffic Volume Prediction Model	18
3.2. FRA Accident Prediction Model	18
3.2.1. Data Processing and Analysis	18
3.2.2. Dataset Preparation (Form 57, Form 71, Exposure)	19
3.2.3. Data Cleaning and Feature Engineering	19
3.2.4. Implementation of FRA’s Crash Prediction Model	20
3.3. Phase II: Crash Prediction Model Development	21
3.3.1. Statistical Model	21
3.3.1.1. Data and Target Overview	21

3.3.1.2.	Univariate Screening: Relationship to Accidents	21
3.3.1.3.	Screening and Encoding (Univariate Evidence → Modeling Choices)	23
3.3.1.4.	Zero-Inflation Rationale and Inflation Features	25
3.3.1.5.	Zero-Inflated Negative Binomial (ZINB) specification	26
3.3.1.6.	Model evaluation and interpretation	27
3.3.2.	AI Model	27
3.3.2.1.	Initial AI Modeling Attempts	28
3.3.2.2.	Single-Stage Models	28
3.3.2.3.	Rationale for a Two-Stage Design	28
3.3.2.4.	Two-Stage AI Pipeline: RFC (with SMOTE) → RFR	29
3.3.2.4.1.	Stage 1: Rare-event classification with Random Forest + SMOTE	29
3.3.2.4.2.	Stage 2: Count modeling for predicted-positive sites (Random Forest Regressor) 30	
4.	Analysis and Results	32
4.1.	Phase I: Volume Prediction Results	32
4.2.	FRA APS Variants: AADT-Only vs. AADT+Non-Motorist Exposure	32
4.3.	Phase II: Crash Prediction Model	33
4.3.1.	Statistical Model – Results and Analysis	33
4.3.2.	AI Model – Results and Analysis	34
5.	Conclusions	36
5.1.	Limitations	36
	References	37
	Appendix A	41
	U. S. DOT Crossing Inventory Form 71	44

List of Figures

Figure 1. Research framework.....	11
Figure 2. Numeric predictors vs. accidents.....	22
Figure 3. Categorical predictors vs. accidents	23
Figure 4. Box plots showing the distribution of percent differences in crash prediction outputs (Note: Each subplot compares the AADT-only prediction against the AADT plus non-motorist prediction, highlighting the variation in model sensitivity to non-motorist volume across the four different states)...	33
Figure 5. Confusion matrix summarizing Stage 1 predictions vs. actual crash presence.	35

List of Tables

Table 1. Description of variables used in the FRA crash prediction and non-motorist traffic volume models.....	20
Table 2. Univariate screening summary	25
Table 3. Stage 1-RFC: Hyperparameters and their roles	30
Table 4. Stage 2-RFR: Hyperparameters and their roles	31
Table 5. Percentage differences in crash predictions after incorporating non-motorist exposure across all four states.....	32
Table 6. Stage 1 Random Forest performance metrics for crash detection using SMOTE-balanced data .	34
Table 7. Stage 1 Random Forest performance metrics for crash detection evaluated on the original imbalanced dataset.....	34
Table 8. End-to-end performance of the two-stage AI crash prediction pipeline (Stage 1 classification followed by Stage 2 regression), reporting MAE and RMSE on the test dataset	34

List of Abbreviations

Federal Highway Administration (FHWA)

Federal Railroad Administration (FRA)

Highway-Rail Grade Crossings (HRGCs)

Negative Binomial Model (NB)

Random Forest Classifier (RFC)

Random Forest Regressor (RFR)

Root-Mean-Square Error (RMSE)

You Only Look Once Version 8 (YOLOV8)

Zero-Inflated Negative Binomial (ZINB)

Disclaimer

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of information presented herein. This document is disseminated under the sponsorship of the U.S. Department of Transportation's University Transportation Centers Program, in the interest of information exchange. The U.S. Government assumes no liability for the contents or use thereof.

Acknowledgements

The authors acknowledge the University Transportation Center for Railway Safety (UTCRS), headquartered at the University of Texas at Rio Grande Valley, for funding this project under the USDOT UTC Program Grant No. 69A3552348340.

Executive Summary

This report presents findings of a multi-state research effort to improve pedestrian and bicyclist safety at highway-rail grade crossings (HRGCs) by explicitly incorporating non-motorist traffic exposure into a crash prediction framework. Building on prior work that developed an artificial intelligence (AI)-based methodology to estimate non-motorist traffic volumes from video at HRGCs, this research scales those estimates to four Midwestern states: Nebraska, Kansas, Missouri, and Iowa, and evaluates both statistical and AI models for predicting crashes at 13,672 public HRGCs using a 2020–2024 analysis window.

In the previous study (Phase I), video footage from 18 urban and suburban HRGCs in Lincoln, Nebraska was processed using a fine-tuned YOLOv8n object detection model to automatically detect pedestrians and bicyclists. These empirically observed non-motorist counts were converted to average daily volumes and used to train a log-link linear regression model with bells, gates, crossbucks, and annual average daily traffic (AADT) as predictors. The model was then applied to inventory data across the four states to estimate daily non-motorist exposure at each crossing. The model explained approximately 81.6% of the variance in observed daily non-motorist volumes ($R^2 \approx 0.816$), with MAE ≈ 14.1 users/day and RMSE ≈ 18.0 users/day, indicating that it captures the dominant spatial variation in non-motorist activity with operationally acceptable error.

This study (Phase 2) integrates non-motorist exposure estimates into the Federal Railroad Administration’s (FRA) Accident Prediction and Severity (APS) Model and develops alternative crash prediction models. Two exposure scenarios are evaluated: (i) the current APS practice using AADT-only and (ii) an enhanced exposure metric combining AADT with predicted pedestrian and bicyclist volumes. Using a Zero-Inflated Negative Binomial (ZINB) implementation of the APS framework, the study finds that incorporating non-motorist exposure increases predicted annual crashes by an average of 5.5% across the four states, with larger increases (up to ~60%) at crossings where non-motorist activity is substantial. This demonstrates that ignoring non-motorist traffic volumes underestimates crash risk at many HRGCs.

A dedicated ZINB crash model is then estimated on the multi-state dataset to characterize the effects of roadway and crossing characteristics on crash frequencies. The model employs a two-part structure with a logit inflation component (log(Exposure), flashing lights, and gates) to account for structural zeros, and a Negative Binomial count component with roadway class, surface type, bells, and maximum timetable speed as predictors. The selected specification achieves the best Akaike Information Criterion/Bayesian Information Criterion (AIC/BIC) among competing models while maintaining low hold-out error (mean absolute error, MAE ≈ 0.091 , root means squared error, RMSE ≈ 0.235 for five-year crash counts). Results show that higher roadway class, rougher or more complex surfaces, additional bell devices, and higher train speeds are all associated with increased expected crash counts, whereas higher exposure reduces the likelihood that a crossing is in the “always-zero crash” state.

To further enhance predictive performance, the study develops a two-stage AI pipeline. Stage 1 uses a Random Forest Classifier with synthetic minority over-sampling or SMOTE-based technique to detect whether a crossing experiences any crash in the five-year window; Stage 2 uses a Random Forest Regressor to estimate crash counts at crossings predicted as at-risk. On the imbalanced test data ($\approx 5.4\%$ crossings with at least one crash), Stage 1 attains recall ≈ 0.97 and F1 ≈ 0.667 , missing only 3 out of 100 crash-prone crossings while maintaining a precision (≈ 0.51) markedly higher than the base rate. When combined with

Stage 2, the end-to-end AI system achieves $MAE \approx 0.0668$ and $RMSE \approx 0.27$ crashes on the test set, outperforming the ZINB baseline in pure predictive accuracy while providing a practical screening tool for prioritizing safety investments.

This research demonstrates that (i) AI-based video analytics can reliably estimate non-motorist traffic volumes at HRGCs, (ii) incorporating these traffic volumes into exposure measures materially affects crash risk estimates, and (iii) a combination of interpretable statistical models (ZINB) and high-capacity AI models (two-stage Random Forest pipeline) offers a robust framework for identifying and ranking high-risk HRGCs for pedestrians and bicyclists across multiple states. The work highlights the need for broader non-motorist data collection and state-specific calibration but provides a transferable methodology for advancing multimodal safety at highway-rail crossings.

1. Introduction

Highway-rail grade crossings (HRGCs) are critical points of interaction between rail and roadway users with a significant risk of conflict due to the convergence of different transportation modes. While considerable progress has been made in reducing motor vehicle crashes through engineering, enforcement, and education programs, safety threats for pedestrians and bicycles remain a major problem [1,2,3]. Incidents at HRGCs involving non-motorists are more serious due to the vulnerability of those involved and the large impact forces associated with train collisions.

Non-motorist safety at HRGCs involves specific issues that are distinct from those encountered with motor vehicles. Limited visibility, insufficient crossing infrastructure, inconsistent warning systems, and noncompliance with crossing rules all contribute to high danger levels [1,2,3]. Many present safety technologies are largely geared for automobile use, with little emphasis on pedestrian and cycling behavior. Furthermore, non-motorist decision-making patterns, such as the inclination to cross immediately before a train comes or to avoid gates, add behaviors that are difficult to monitor and forecast in the absence of specialist monitoring systems [7].

The Federal Railroad Administration (FRA) and state Departments of Transportation (DOTs) have long recognized the importance of improving safety at HRGCs, and numerous programs have been launched to increase infrastructure, public awareness, and enforcement [8]. Despite these efforts, non-motorist offenses have not decreased as much as motor vehicle violations. This gap highlights the importance of focused approaches to non-motorist detection, exposure estimate, and crash risk modeling. FRA maintains a large collection of HRGC inventory data, which includes characteristics such as, highway traffic control devices, pavement markings, land use, number of tracks, crossing surface, and traffic volume. Using these datasets, the FRA created an Accident Prediction and Severity (APS) Model in 2020 [10,11] to anticipate the predicted frequency of crashes at HRGCs for more efficient resource allocation and risk management. However, this crash prediction model does not incorporate the non-motorist exposure, which constitutes a significant gap. Accounting for non-motorists' crash exposure is critical for more thorough safety evaluations and targeted treatments, especially since non-motorists related crashes frequently occur under different circumstances than motor-vehicle crashes.

Advances in technology such as computer vision, artificial intelligence (AI), Internet of Things (IoT) sensors, and machine learning (ML) algorithms have created new opportunities for tracking and predicting safety concerns at HRGCs. Video analytics can be used to detect non-motorist presence, and track their movement patterns, providing useful data for predictive models. Using real-world data on non-motorist (pedestrian and bicyclist) traffic at HRGCs gathered from video surveillance in Phase I, we build models to predict their traffic at HRGCs in four Midwestern states: Missouri (MO), Kansas (KS), Nebraska (NE), and Iowa (IA).

In this study, we extend the FRA framework by including non-motorist traffic volume estimates in the accident prediction procedure. We began by implementing the FRA's 2020 crash prediction model using a 5-year inventory dataset. A specific non-motorist traffic prediction model was used to estimate non-motorist volumes, which were then incorporated into the FRA model by adding non-motorist counts to the Annual Average Daily Traffic (AADT). The crash forecasts were then updated for HRGCs in four midwestern

states: Nebraska, Kansas, Missouri, and Iowa. We then created both a statistical and an AI-based crash prediction model to improve predictive accuracy and assess the integration of non-motorist traffic to AADT.

1.1. Problem Statement

Despite continued attempts to improve railroad crossing safety, non-motorists continue to face significant risks at HRGCs [13]. Existing safety programs, infrastructure improvements, and warning systems are mostly built with motorists in mind, which limits their usefulness for non-motorist users [12]. This gap is concerning, given the distinct behaviors, mobility patterns, and exposure risks associated with crossing activities of non-motorists.

Current crash prediction model employed by FRA focuses on motor vehicle occurrences and rely primarily on aggregate crash history [10,11]. This crash prediction model does not incorporate non-motorists' exposure. Moreover, available crash prediction models do not incorporate non-motorists' exposure into predictive analyses.

Without specific, data-driven models customized for non-motorists, agencies struggle to reliably identify high-risk crossings and prioritize interventions. As a result, investments in safety improvements may not result in optimal reductions in pedestrian and bicycling crashes. There is a significant need for methodologies that integrate statistical and machine learning approaches to solve these limitations and enable more effective safety planning for vulnerable road users near railroad crossings.

1.2. Objectives

The objectives for this research are as follows:

1. Review existing crash frequency prediction models used for several transportation modes. Summarize the statistical and AI-based crash prediction models highlighting their applicability, scalability, and limitations.
2. Collect and integrate inventory, crash, and non-motorist exposure data for assembling a multi-state HRGC data repository for 5-year period.
3. Perform Exploratory Data Analysis (EDA) to analyze spatial, temporal and exposure related crash patterns using statistical summaries and visualizations.
4. Implement the FRA APS model on a 5-year HRGC inventory dataset. Employ the non-motorist volume prediction model to estimate the non-motorists counts and integrate these with AADT values to get updated annual crash predictions using FRA's APS model.
5. Develop statistical and AI-based crash prediction models incorporating exposure for motorist and non-motorists at HRGCs using techniques such as linear regression, zero-inflated models, negative

binomial regression, decision trees, support vector machines (SVMs), random forests, and neural networks.

6. Compare and evaluate the model performance using performance metrics such as RMSE (Root Mean Square Error), MAE (Mean Absolute Error), R^2 (R-squared), and AIC (Akaike Information Criterion). Assessments to evaluate computational efficiency, and feature importance for interpretability.
7. Identify the most effective model based on quantitative performance and practical applicability and recommend final model for HRGC safety applications and future crash prediction efforts.

1.3. Research Framework

Figure 1 presents the research framework. As depicted, the first step of the process was to assemble a five-year HRGC inventory dataset and crash history for Nebraska, Kansas, Missouri, and Iowa. Second step was to estimate the non-motorist traffic volumes for the four states. Using this estimation, we refine the exposure metrics by combining the AADT and non-motorist volume for each crossing. These updates exposure estimates were then used in the FRA's APS model to generate the updated annual crash predictions for the four states. Using the prepared multi-state HRGC dataset, we developed statistical and AI-based crash prediction models. Selection of final proposed crash prediction model was done after comparison using a set of performance indicators.

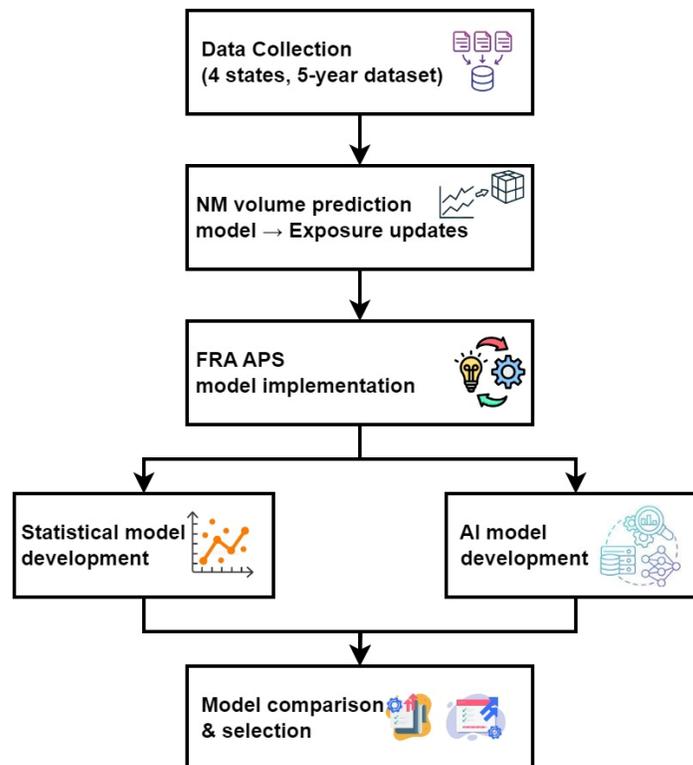


Figure 1. Research framework

1.4. Organization of the Report

This report consists of five chapters. Chapter 1 (this chapter) is the introduction followed by Chapter 2 that presents a review of pertinent literature. Chapter 3 presents the research methodology including development of non-motorist traffic volume predictions, the FRA accident prediction model, and the development of statistical and AI models. Chapter 4 provides the analysis and results while Chapter 5 presents the conclusions and research limitations.

2. Literature Review

HRGCs are critical locations of contact between road users and railroads, posing safety problems due to their multimodal nature [14]. Despite considerable improvements in infrastructure and warning mechanisms, HRGCs continue to experience crashes involving trains and a mix of motor vehicles, bicyclists, and pedestrians. According to the US Government Accountability Office (GAO), 89 pedestrians were killed at public HRGCs in 2023, which represented about forty-one percent of all HRGC fatalities that year. This was up from about thirty-one percent pedestrian fatalities at HRGCs in 2014 [15]. Pedestrians and bicyclists lack physical protection, which increases the likelihood of severe injury outcomes compared to motor vehicle users.

2.1. HRGC Crash Prediction Models

State transportation agencies ensure HRGC safety within their jurisdictions. They prioritize and select projects for improving HRGC safety by considering a variety of factors including output from crash prediction models [16]. Since the 1940's transportation agencies have relied on crash and hazard prediction models to quantify risk at HRGCs [17]. The two main ingredients of all models are measures of highway motor vehicle and rail traffic, i.e., Annual Average Daily Traffic (AADT) and daily train traffic. Together, these two account for crash exposure due to motor vehicle and train traffic, which represents the opportunities for occurrence of motor vehicle-train involved crashes. While AADT accounts for the portion of exposure arising from vehicular traffic, it is a subset of the total highway traffic. The complete exposure from highway traffic is the sum of motor vehicle and non-motor vehicle traffic (pedestrians, bicyclists, and other forms of non-motorized users). Thus, existing crash prediction models are predicting fewer crashes at HRGCs due to non-consideration of crash exposure from highway users such as pedestrians and bicyclists. However, the magnitude of the shortfall is unknown because non-motor vehicle traffic counts at HRGC are rarely available.

2.2. Non-Motorist Safety at Grade Crossings

With growing emphasis on multimodal mobility and active transportation, the vulnerability of non-motorists at HRGCs is becoming more important. An observational study found non-motorists to frequently indulge in unsafe behavior at HRGCs when trains were approaching [18]. This study reported no significant differences in the mean frequencies of pedestrian vs bicyclist gate-related violations at HRGCs; 1.27 gate-related violations were recorded per crossing event. Children of around 8 years of age or younger were involved in 25% more gate-related violations than older users. The study also reported increased gate violations with greater numbers of individuals present at the crossing, but the contribution from young children was disproportionately higher. Another study explored the use of population in the vicinity of HRGCs besides rail crossing attributes to predict HRGC crashes using a binary logit model [19] while another study [20] investigated factors influencing bicycle crashes at skewed HRGCs and reported approach angle as the most significant determinant of crashes with a critical traversing angle of 30 degrees.

A study of pedestrians and bicyclists with respect to their violations at HRGCs, distracted behaviors, and crossing speeds [21] reported 88.9 percent pedestrians and 93.3 percent bicyclists engaged in violations at HRGCs; 9.5 percent of pedestrians and 7.7 percent of bicyclists were distracted while crossing HRGCs. With respect to crossing speed, several pedestrian/bicyclist and HRGC/environmental characteristics impacted crossing speeds. Furthermore, the HRGC gate violation rates were higher than previous research [18].

A study of pedestrian and bicyclist warnings devices and signs at HRGCs [22] reported a lack of standards to analyze/quantify pedestrian risk at HRGCs and the use of signs and devices that were not compliant with the Manual of Uniform Traffic Control Devices (MUTCD). It further indicated that the effectiveness of any sign or device in reducing the risk of a collision between a pedestrian and a train was unknown. Along the same lines, a Volpe Center study [23] surveyed the state of pedestrian treatments and reported a lack of consistent evidence regarding the effectiveness of combinations of devices (e.g., pedestrian swing-gates, in-pavement lights, channelizing fences) in improving pedestrian safety. It recommended research to identify the efficacy of pedestrian treatments at HRGCs.

Overall, several studies have highlighted the increasing concern over pedestrian and bicyclist safety at HRGCs. Observational research has shown frequent gate-related violations by non-motorists, particularly by young children and in group settings, with no significant difference between pedestrian and bicyclist behavior. Factors such as distraction, approach angle, and crossing speed were found to influence crash risks and unsafe behaviors. Despite the use of various warning devices and signs at HRGCs, their effectiveness remains unclear due to non-standardization and limited supporting evidence, prompting calls for further research.

2.3. AI Applications in Transportation Safety

AI has increasingly emerged as a transformative tool in addressing transportation safety challenges, enabling more accurate predictions, enhanced risk assessment, and proactive safety interventions. In the context of HRGCs, AI applications build on advancements in data processing, machine learning algorithms, and pattern recognition to identify safety-critical factors that traditional statistical approaches may not fully capture.

The integration of AI into transportation safety has demonstrated transformative potential across multiple modes, including HRGCs, urban roadway systems, and maritime navigation. AI applications in this domain focus on analyzing driver and traveler behavior, enhancing control systems, predicting crash severity, and enabling proactive safety interventions.

At HRGCs, studies highlight the critical role of AI in mitigating risks. For instance, Ren et al. [24] analyzed traffic safety and injury severity at undivided two-way crossings, showing that non-compliance with traffic control devices remains a major contributor to accidents. They suggest AI-driven educational and enforcement strategies as effective countermeasures. Similarly, Yeh and Multer [25] reviewed traffic control devices and barrier systems, emphasizing the need for innovative intervention strategies to prevent driver violations of active warning signals. Environmental influences have also been investigated; Hao and Kamga [26] found that inadequate lighting correlates with increased crash severity at low-traffic crossings,

pointing to the potential for AI systems to incorporate real-time environmental data into predictive safety measures. Yeh et al. [27] further demonstrated how safety factors at crossings can be evaluated using signal detection theory, suggesting that AI could extend such analyses for real-time decision support.

Technological advancements have further expanded AI applications at HRGCs through connected vehicle systems and in-vehicle warnings. Wang et al. [28] proposed a connected vehicle-based warning system that proactively alerts drivers about approaching trains, while Nadri et al. [29] investigated in-vehicle auditory alerts to mitigate risks from driver inattention. Parallel to these efforts, Rana et al. [30] utilized machine learning techniques on crash data to identify hotspot locations, enabling targeted interventions, and Fan et al. [31] demonstrated the use of multinomial logit modeling for analyzing crash severity. Together, these efforts underscore AI's growing role in developing predictive models that guide safety countermeasures.

Beyond HRGCs, AI has been increasingly applied in broader transportation safety contexts. Intelligent Transport Systems (ITS) are leveraging AI for real-time data analysis, traffic management, and decision-making. Kalašová et al. [32] emphasized the importance of AI in facilitating driver interaction with ITS by collecting, processing, and disseminating traffic-related information. Zhao and Chen [33], as well as Zhao [34], highlighted the role of deep learning in managing spatial and temporal challenges in urban transportation systems, demonstrating how AI enhances both safety and efficiency. Similarly, Jaidev et al. [35] showed that AI can improve situational awareness of road conditions and driver behavior, further reducing accident risks.

AI's impact is not limited to road transport. In maritime contexts, early applications such as the Automatic Identification System (AIS) have been used for collision avoidance and situational awareness. Grabowski and Dhimi [36] highlighted the adoption of AIS to prevent collisions through enhanced navigation information, while Vestre et al. [37] documented its broader role in improving operational safety in high-traffic maritime channels. These studies demonstrate how AI methodologies for predictive modeling, data-driven interventions, and real-time monitoring extend across transportation domains.

In summary, AI applications in transportation safety exhibit a multifaceted landscape: at HRGCs, AI enables improved prediction, warning systems, and environmental awareness; in urban networks, AI strengthens ITS and crash prevention; and in maritime operations, AI-driven systems like AIS enhance collision avoidance. Collectively, these advancements illustrate the transformative role of AI in reshaping safety frameworks and creating more resilient transportation systems.

2.4. Gaps in Existing Research

The analysis of non-motorist exposure at HRGCs, particularly for pedestrians and bicyclists, remains a critical yet underexplored area of transportation safety. Research to date has emphasized injury severity, behavioral patterns, and environmental factors, but has not sufficiently addressed how exposure itself influences predictive crash modeling at HRGCs.

Eluru et al. [38] employed a mixed generalized ordered response model to investigate pedestrian and bicyclist injury severity in broader traffic crashes. Their work underscores the importance of capturing non-motorist behavior through advanced modeling, offering insights into more effective safety strategies.

Similarly, Dong et al. [39] analyzed injury severity at mid-block locations and highlighted pre-crash behaviors, noting that although non-motorists constitute a smaller share of traffic, they account for a disproportionately high percentage of fatal outcomes. These findings point to the heightened vulnerability of non-motorists and the necessity of targeted safety interventions.

At HRGCs specifically, Metaxatos and Sriraj [40] advocated standardized warning devices and improved engineering practices that explicitly address pedestrian and bicyclist risks. This is complemented by Kizawi's [41] analysis of conflicts between vehicles and non-motorists, which demonstrated how changes in traffic patterns exacerbate risks for vulnerable road users. Environmental conditions have also been shown to play a substantial role: Easa et al. [42] emphasized the importance of improving sight distance at grade crossings to mitigate risks associated with limited visibility, while Al-Mahameed et al. [43] used structural equation modeling to identify roadway design and socioeconomic factors as significant contributors to crash likelihood among non-motorists.

Collectively, these studies provide valuable insights into injury severity and risk factors, but they reveal two critical gaps. First, existing research has largely examined behavioral and environmental determinants of non-motorist crashes, without explicitly quantifying non-motorist exposure as a predictor within crash frequency models at HRGCs. Second, despite evidence that non-motorists face significant risks, no studies to date have evaluated how incorporating non-motorist exposure data alters crash prediction outcomes compared to models based solely on vehicular traffic (e.g., AADT). Addressing these gaps is essential for developing more comprehensive and equitable safety models that capture the realities of multimodal travel at HRGCs.

3. Methodology

This chapter outlines the methodological framework adopted in this study, which is divided into two main phases. Phase I focuses on non-motorist volume prediction, building on Year 1 efforts where video data from selected urban and suburban HRGCs was processed using AI-based object detection techniques to quantify pedestrian and bicyclist exposure. The outputs of this phase formed the basis of NM volume prediction models. Phase II centers on the development and application of crash prediction models, leveraging both statistical and AI-based approaches. By integrating NM exposure estimates into the FRA existing crash prediction methodology, this study provides an enhanced framework for predicting crash frequency and identifying safety risks at HRGCs across multiple states.

3.1. Phase I: Non-Motorist Volume Prediction

3.1.1. Data Collection

In Phase I, data collection was conducted using Miovision cameras and City of Lincoln’s existing surveillance systems deployed at 18 rail crossings in Lincoln, Nebraska. These crossings were chosen to capture both urban and suburban environments, ensuring diversity in traffic conditions. Video data was recorded over several days to account for variations in non-motorist behavior and temporal travel patterns. From these recordings, footage was systematically processed to extract information on pedestrian and bicyclist movements at crossings.

3.1.2. Object Detection via AI Models

To quantify NM exposure, the study implemented YOLOv8n, a state-of-the-art computer vision model trained for real-time object detection. The model was fine-tuned to recognize pedestrians and bicyclists, allowing accurate frame-by-frame detection from the collected videos. Post-processing techniques aggregated detections across time intervals, converting raw counts into standardized traffic measures. This AI-based approach significantly reduced manual counting efforts and provided consistent, scalable detection performance across varied environments.

3.1.3. Non-Motorist Traffic Volume Prediction Model Preparation

By applying a prediction model that we estimated using empirically obtained video data from sample HRGCs to a larger inventory dataset, we were able to estimate non-motorist traffic at the various HRGCs under consideration. Using the same four-state crossing inventory, a different dataset was created to quantify non-motorist crash exposure, concentrating on important characteristics linked to the presence of bicyclists and pedestrians. Based on their availability and pertinence, the following variables were chosen as features in the model: Gates, Bells, XBuck (crossbuck), and AADT.

Over the course of many days, deployed cameras at 18 HRGCs in Lincoln, NE recorded video footage of pedestrian and bicyclist traffic at those crossings. The cameras captured continuous video footage, which

was processed with the YOLOv8n object detection model. The extracted data were checked with manual observations of the video footage to verify accuracy of the pedestrian and bicyclist counts. The entire number of observed non-motorists was added together and normalized by the total number of recording hours for each crossing to get an average daily non-motorist traffic volume. The regression model used this estimated value as its response or target variable. Throughout data preprocessing and feature extraction procedures, we followed the FRA’s crash prediction guidelines to guarantee uniformity and comparability. As a result, the non-motorist exposure estimations were easily included in the pre-existing crash prediction framework. The final dataset, which included 13,672 HRGCs after the cleaning, merging, and transformation processes were finished for all four states, was prepared for comparative safety analysis and predictive modeling.

3.1.4. Non-Motorist Traffic Volume Prediction Model

We developed a regression model to estimate the average daily non-motorist traffic volume at HRGCs. After preparing the dataset for non-motorist traffic volume prediction (described earlier), we trained the model using observed non-motorist counts as the response or target variable. A standard linear regression model was fitted between the log-transformed counts and selected features: Bells, Gates, AADT, and XBuck (described in Table 1). The observed counts were aggregated and averaged to yield a representative daily non-motorist volume for each crossing. The model follows a Poisson-like approach using a log-link function, ensuring compatibility with count data. Final predictions were obtained by exponentiating the linear predictor to revert from the log scale to daily counts. The equations of the model are:

$$\log(\lambda_i) = \beta_0 + \beta_1 Bells_i + \beta_2 Gates_i + \beta_3 AADT_i + \beta_4 XBuck_i \quad (1)$$

$$\hat{V}_i = \exp(\log(\lambda_i)) \quad (2)$$

where: \hat{V}_i = predicted average daily non-motorist traffic volume at crossing i , Bells, Gates, Aadt, XBuck are regression predictors, and coefficients β_k are estimated via linear regression on the log-link scale. After training on NE data, the learned coefficients were applied to the processed inventory data from MO, KS, NE, and IA. This resulted in non-motorist traffic volume estimates across 13,672 HRGCs, consistent with FRA data standards.

3.2. FRA Accident Prediction Model

3.2.1. Data Processing and Analysis

To apply the FRA APS Model for crash prediction, we first created a single dataset by combining the accident/incident records for the four MO, KS, NE, and IA with the Grade Crossing Inventory File [45]. The USDOT Crossing Number was used as the primary key in the merge process. The FRA APS Model requires the use of five-year crash data. For this research, the 2020-2024 period was used. A feature often encountered in crash data modeling is an unbalanced and significantly zero-inflated response or target variable, which was the case in this research due to the presence of thousands of crossings with zero reported

crashes over the five-year period. The response or target variable for modeling was the total five-year crash count, and it is represented by:

$$y_i = \text{The number of reported crashes at crossing "i" between 2020 and 2024.}$$

In this case, no binary conversion was used because the goal was to anticipate count data rather than record merely crash occurrence. To standardize their distributions and stabilize variance, several characteristics were log-transformed to satisfy the FRA's APS Model requirements (Zero-Inflated Negative Binomial or ZINB). A log-linear structure is defined by the APS Model (GX APS2-A) model, in which the crash rate equation incorporates the natural log of specific features. Further details are in subsection "Implementation of FRA's Crash Prediction Model".

We used the variables and equations specified in the official specification to create the ZINB crash prediction model developed by the FRA. These variables included Exposure, AADT, MaxTtSpd, TotalTrains, D_2 , D_3 , HwyClassCD, and XSurfIDs [5]. Table 1 presents the details of these variables. For model compatibility, the warning device type, WDCODE, categorical variable was also converted into two binary features as given in the data dictionary [45]. One characteristic, D_2 , shows whether flashing lights are installed at the crossing, while another feature, D_3 , shows whether gates are present. Passive crossings were defined as those without gates or lights (i.e., where both D_2 and D_3 equal zero). The safety impact of the various warning device kinds in the model is captured by this encoding.

The dataset that resulted from combining crash records and inventory data using Crossing ID was unbalanced since it had many zeros, i.e., crossings without any reported collisions. The ZINB model, which uses a two-part structure to account for excess zeros, is justified by this feature.

3.2.2. Dataset Preparation (Form 57, Form 71, Exposure)

We assembled a cross-sectional, multi-state dataset by merging the FRA Highway-Rail Crossing Inventory (Form 57) with the FRA Highway-Rail Incident (Form 71) crash records on the common crossing identifier (USDOT Crossing ID). Form 57 supplied static and dynamic site characteristics (e.g., warning devices, approach speed limits, track configuration, train volumes), while Form 71 contributed the five-year crash counts needed for model training and validation. Exposure was represented in two scenarios: (i) AADT-only (current Accident Prediction System (APS) practice) and (ii) AADT + estimated daily non-motorist volume (pedestrians + bicyclists) obtained from Phase I and appended to AADT to reflect total highway user exposure. The final integrated dataset comprised 13,672 HRGCs across MO, KS, NE, and IA after cleaning, merging, and transformation, following FRA crash-prediction guidelines to ensure consistency of inputs across scenarios.

3.2.3. Data Cleaning and Feature Engineering

We removed rows with missing key predictors or invalid codes, harmonized categorical encodings (e.g., rural/urban, warning device categories), and verified ranges/units for traffic and train volumes. Following the FRA APS transformations, AADT, Total Trains, Exposure, and Maximum Timetable Speed were log-transformed with variable-specific scaling factors (α) so that the transformation is centered at the variable's median (i.e., the log-transform equals zero at the median). This centering improves numerical stability and interpretability of the log-links used in the count process. Because many crossings experience zero crashes

in a five-year window, the crash outcome distribution is zero-inflated and over-dispersed, motivating a two-part model structure (zero-inflation + negative binomial count), as used by FRA’s APS framework.

3.2.4. Implementation of FRA’s Crash Prediction Model

We used the ZINB model as outlined in the GX APS2-A report [44] to simulate the crash prediction framework of the FRA. To determine the anticipated number of accidents over a five-year period, the model consists of two parts: a count model, and a zero-inflation model. For AADT, Total Trains, Exposure, and Maximum Timetable Speed, natural log-transformation is needed as follows.

$$X = \ln(1 + \alpha x) \quad (3)$$

Where X is the natural log transformed variable, x is the input variable, α is a scaling factor selected so that for the median value of x , the following equation holds true:

$$\ln(1 + \alpha x_{median}) = \ln(x_{median}) \quad (4)$$

Using a variety of crossing features, the model’s accident count component forecasts the anticipated number of collisions at each crossing across a five-year range. This is modeled using the Negative Binomial regression’s log-link form, given as:

$$N_{count} = e^{\beta_0 + \beta_1 \ln(1 + \alpha_1 Exposure) + \beta_2 D_2 + \beta_3 D_3 + \beta_4 HwyClassCD + \beta_5 XSurfID + \beta_6 \ln(1 + \alpha_6 AADT) + \beta_7 \ln(1 + \alpha_7 MaxTtSpd)} \quad (5)$$

$$P_{InflatedZero} = \frac{z}{1+z}, \quad z = e^{\gamma_0 + \gamma_1 \ln(1 + \alpha TotalTrains)} \quad (6)$$

$$N_{pred} = N_{count}(1 - P_{InflatedZero}) \quad (7)$$

By down weighting the forecast for HRGCs with excess zeros, this formulation guarantees that the model maintains its sensitivity to HRGC variability in crash risks. As explained in the FRA’s approach, we applied the Empirical Bayes (EB) modification to increase prediction reliability, particularly at low-volume crossings. This involves using a weight “ w ” that is calculated from the relative variance of prediction to combine the predicted value with the observed crash count, $N_{observed}$.

$$N_{EB} = wN_{pred} + (1 - w)N_{observed} \quad (8)$$

$$w = \frac{1}{1 + \frac{V[N_{pred}]}{N_{pred}}} \quad (9)$$

Where, $V[N_{pred}]$ is the variance of the crossing’s predicted number of crashes. Use of the above equations enabled estimation of crash predictions for the four states under consideration in this research.

Table 1. Description of variables used in the FRA crash prediction and non-motorist traffic volume models.

Feature/variable Name	Description
-----------------------	-------------

AADT	Annual Average Daily Traffic
TotalTrains	Sum of total daylight through trains and total night through trains
Exposure	Product of AADT and Total Trains
MaxTtSpd	Maximum timetable speed
D ₂	Binary indicator for flashing lights (1 = lights present, 0 = no lights)
D ₃	Binary indicator for gates (1 = gates present, 0 = no gates)
XSurfIDs	Crossing surface
HwyClassCD	Crossing location type: Rural or Urban
XBuck	Number of crossbuck assemblies
Gates	Count of roadway gateway arms
Bells	Number of bells

3.3. Phase II: Crash Prediction Model Development

3.3.1. Statistical Model

3.3.1.1. *Data and Target Overview*

This section uses the same merged MO–KS–NE–IA dataset described in Data Processing and Analysis section, with the identical 2020–2024 five-year window, cleaning and merging steps (Form 57 + Form 71 via USDOT Crossing ID), and field definitions. The unit of analysis is a USDOT highway-rail grade crossing, and the response variable is the five-year crash count at each crossing; we model counts directly rather than binarizing occurrence. As established earlier, the outcome is highly zero-inflated and over-dispersed, motivating a Zero-Inflated Negative Binomial (ZINB) approach. From the prepared fields we employ exposure and operational measures (e.g., $\log(\text{Exposure})$, MaxTtSpd), device/configuration indicators (D₂ for lights, D₃ for gates, HwyClassCD, XSurfIDs, Bells, with rare level collapsing), and the same encodings and reference categories specified previously. Here we focus on model specification, screening, and results, while deferring all data-construction details to the prior section.

3.3.1.2. *Univariate Screening: Relationship to Accidents*

We began with univariate checks to understand how each candidate predictor relates to the five-year crash count at a crossing. For numeric variables (e.g., MaxTtSpd, AADT, Exposure) we inspected scatter/point plots of Accidents vs. the predictor with smoothed trends and binned means, and we computed rank correlations with the positive-count subset to avoid zero-inflation dominating the signal. For categorical variables (e.g., HwyClassCD, XSurfIDs, Bells, Gates), we compared groupwise means of Accidents and the proportion of zeros in each level. Collectively, these screens informed which variables enter the count model vs. the zero-inflation model and where level collapsing was warranted.

Categorical variables are encoded as indicator (dummy) variables using a reference (baseline) level; coefficients are interpreted relative to that baseline in the count model. We collapse rare levels (below a minimum count threshold) into “other” category to stabilize estimation.

Figure 2 summarizes numeric predictors (Exposure, MaxTtSpd, HwySpeed) by plotting, for equal-frequency bins, the probability that a crossing has ≥ 1 crash (blue line, left axis) alongside the average crash count conditional on ≥ 1 crash (gray bars, right axis). Exposure shows the clearest monotonic pattern: higher exposure is associated with a higher likelihood of any crash and slightly higher conditional means—consistent with exposure being the primary driver of risk. Speed variables show a milder, mostly increasing relationship with crash occurrence; conditional means remain near ~ 1 – 2 , indicating severity (given a crash) does not change much across speed bins. Figure 3 provides the categorical analogue: for each level of HwyClassCD, XSurfIDs, and Bells, the blue bar is the fraction of crossings with ≥ 1 crash and the orange bar is the mean count among those that crashed. Elevated probabilities appear for higher roadway classes (1, 3, 4), higher surface codes (3–4), and non-zero bell categories, which matches the directions seen in the ZINB coefficients. The “other” bell category shows high rates but is rare; those values should be interpreted cautiously. Note that the two y-axes differ in scale (probability vs. mean count).

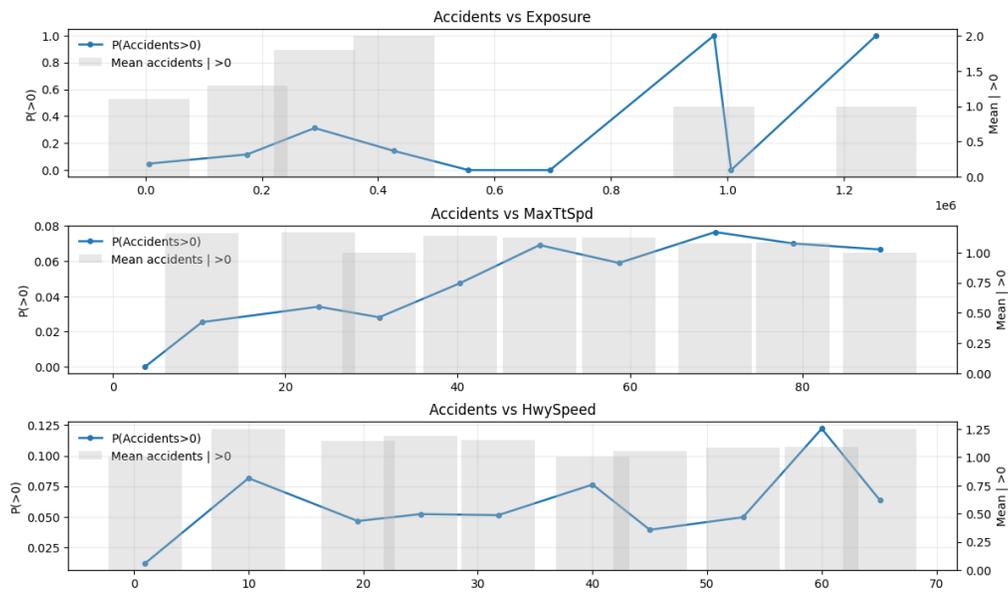


Figure 2. Numeric predictors vs. accidents

For Exposure, MaxTtSpd, and HwySpeed (binned), the blue line shows the fraction of crossings with ≥ 1 crash ($P[\text{Accidents} > 0]$; left axis) and gray bars show the average crash count conditional on ≥ 1 crash (right axis). Exposure exhibits the strongest positive association with crash occurrence; speed effects are weaker and mostly affect occurrence rather than conditional severity.

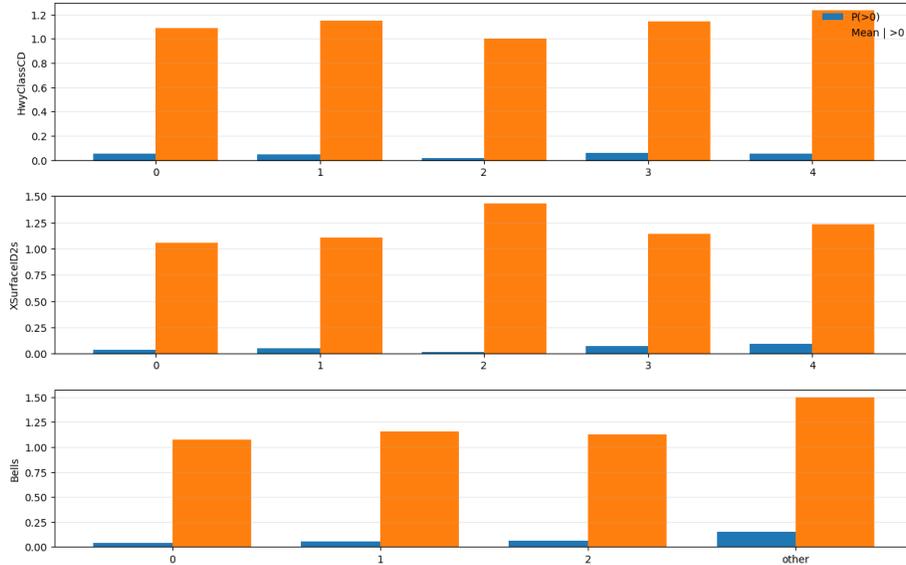


Figure 3. Categorical predictors vs. accidents

For each level of HwyClassCD, XSurfIDs, and Bells, blue bars show the fraction of crossings with ≥ 1 crash (left axis) and orange bars show the average crash count conditional on ≥ 1 crash (right axis). Higher roadway classes and surface codes are associated with higher crash occurrence; non-zero bell categories correlate with higher risk, while the rare “other” level should be interpreted with caution.

3.3.1.3. Screening and Encoding (Univariate Evidence \rightarrow Modeling Choices)

We began with univariate diagnostics to understand how each candidate feature relates to the five-year crash count at a crossing and to decide how predictors should enter the model. For numeric variables (e.g., MaxTtSpd, AADT, Exposure), we plotted *Accidents* against each predictor with smoothed trends and equal-frequency binning, and we computed rank correlations on the positive-count subset to avoid the zero-mass dominating the signal. For categorical variables (e.g., HwyClassCD, XSurfIDs, Bells, Gates), we compared (i) the fraction of crossings with ≥ 1 crash and (ii) the mean crash count conditional on ≥ 1 crash across levels. These views guided two linked decisions: (1) whether a predictor informs the count component (variation in non-zero rates) or the inflation component (propensity to be structurally zero), and (2) how to encode the predictor for stable estimation.

Key univariate findings (evidence \rightarrow model entry).

- “MaxTtSpd” shows a clear positive association with crash counts among crossings with $y > 0 \rightarrow$ kept in the count part.
- “XSurfIDs” displays strong level differences (levels 1, 3, 4 $>$ baseline in both non-zero probability and conditional mean) \rightarrow kept in the count part (with rare levels collapsed).
- “HwyClassCD” levels 1, 3, 4 exceed baseline; level 2 behaves similarly to baseline \rightarrow kept in the count part, with 2 \rightarrow baseline for stability.

- “Bells” has higher means for levels 1–2 vs baseline, but very sparse categories inflate variance → kept in the count part, with rare levels → “other.”
- “Gates” (counts) is weak/inconsistent once D_2/D_3 and other devices are considered → not prioritized for the count part.
- Exposure / AADT / TotalTrains. Exposure = AADT × TotalTrains tracks crash propensity but is collinear with its factors; we therefore use $\log(\text{Exposure})$ in the inflation equation to flag structural zeros (very low exposure → higher zero odds) and omit AADT/TotalTrains from the count side to avoid redundancy.
- “HwySpeed”, “XBuck”, “Whistban” add little incremental signal for positive counts and do not reduce zero proportions → not prioritized for the baseline model.

Encoding choices for stable estimation.

- Dummy (indicator) encoding is used for all categorical predictors with a designated reference (baseline) level; coefficients in the count model are interpreted relative to that baseline.
- Rare-level collapsing. To prevent unstable maximum-likelihood estimates from sparse cells, categories below a minimum frequency threshold are collapsed to “other” (e.g., Bells).
- Substantive merges. Where diagnostics showed a level empirically/operationally indistinguishable from baseline, we merge it into the baseline (e.g., HwyClassCD: 2.0 → baseline, XSurfIDs: 2 → baseline).
- Train/test alignment. We fix the same level set and reference categories across splits so that design matrices (dummy columns) match exactly during fitting and prediction.

Figure 2 (numeric) and Figure 3 (categorical) visualize the probability of any crash and the conditional mean count side-by-side, and Table 2 summarizes how these univariate patterns map to final modeling decisions (kept/dropped, count vs. inflation, and any collapsing/merging applied). See Appendix A for complete numeric and categorical univariate plots (probability of ≥ 1 crash and conditional mean counts), which underpin the variable inclusion, collapsing, and baseline selections summarized in Table 2.

Table 2. Univariate screening summary

Feature	Type	What the plot showed	Notes / interpretation	Used in model
Exposure (AADT × TotalTrains)	numeric	More positives at higher exposure; long right tail	Core exposure driver; variance grows with level	Inflation: log(Exposure)
AADT, TotalTrains	numeric	Weak individually; collinear with Exposure	Redundant with Exposure	Dropped
MaxTtSpd	numeric	Mild positive association with crash counts	Interpretable operating risk	Kept in count
HwySpeed	numeric	Very weak/irregular trend	Adds little beyond MaxTtSpd/Exposure	Dropped
HwyClassCD	categorical	Higher classes show more non-zeros than baseline	Traffic/context proxy	Kept in count (C(HwyClassCD_c), 2→0)
XSurfIDs	categorical	Levels 3–4 show more positives than 0/1	Road surface/condition proxy	Kept in count (C(XSurfIDs_c), 2→0)
Bells (WD sub-type)	categorical	Some configs (1,2) higher than baseline; rare levels noisy	Selection effect at riskier sites	Kept in count (rare→'other')
D ₂ (flashing lights)	binary	Mixed at raw level; separates structural zeros	Good in the zero model.	Kept in inflation
D ₃ (gates)	binary	Mixed at raw level; dominated by zeros	Same as D ₂	Kept in inflation
Gates (count)	categorical/ordinal	Sparse at higher counts; no clear monotone trend	Collinear with D ₃ /WD.	Excluded from count
XBuck (crossbucks)	ordinal	Weak/irregular; many zeros at all levels	Adds little beyond D ₂ /D ₃ /WD	Dropped
Whistban (whistle ban)	binary	Very sparse; no stable signal	Policy heterogeneity + sparsity	Dropped

3.3.1.4. Zero-Inflation Rationale and Inflation Features

The response distribution is extremely sparse ($\approx 95\%$ of crossings report zero crashes over five years), which is inconsistent with a single Negative Binomial process and indicates the presence of structural zeros—crossings effectively not at risk during the period (e.g., very low traffic or inactive operations). To address this, we estimate a Zero-Inflated Negative Binomial (ZINB) in which a logit inflation model first predicts whether a site falls into the always-zero state, and the NB count model governs crashes for the remaining

at-risk sites. Inflation-side predictors are selected to flag conditions associated with structural-zero behavior: $\log(\text{Exposure})$, D_2 (flashing lights), and D_3 (gates). In estimation, the coefficient on $\log(\text{Exposure})$ in the logit equation is negative, meaning that higher exposure is associated with lower odds of a site being structural zero. The device indicators D_2 and D_3 provide additional separation of zero-prone locations. This two-part specification avoids forcing excess zeros into the count process and improves inference for crossings where crashes are genuinely possible.

3.3.1.5. *Zero-Inflated Negative Binomial (ZINB) specification*

We estimate a ZINB to handle two defining features of the data: (i) extreme zero inflation and (ii) over-dispersion among the non-zero counts. Let Y_i denote the five-year crash count at crossing i . We assume a two-component mixture:

$$Y_i \sim \begin{cases} 0, & \text{with probability } \pi_i \text{ (structural zero)} \\ \text{NB}(\mu_i, \alpha), & \text{with probability } (1 - \pi_i) \text{ (at-risk process)} \end{cases} \quad (10)$$

Where $\mu_i > 0$ is the mean of the negative binomial count component and $\alpha > 0$ is the dispersion (negative binomial variance inflation) parameter. Under the mean-dispersion parameterization:

$$E[Y_i | \text{at-risk}] = \mu_i, \text{Var}[Y_i | \text{at-risk}] = \mu_i + \alpha \mu_i^2 \quad (11)$$

so $\alpha > 0$ allows over-dispersion relative to Poisson.

To operationalize the two components, we specify the link functions and linear predictors, mapping the structural-zero probability through a logit link and the at-risk mean count through a log link as follows.

- Inflation (logit) part — probability of a structural zero:

$$\Pr(Y_i \text{ is structural zero}) = \pi_i, \quad \text{logit}(\pi_i) = z_i^T \gamma, \quad (12)$$

with

$$z_i = [\log(\text{Exposure}_i), D_2, D_3], \quad (13)$$

so higher $\log(\text{Exposure})$ typically reduces π_i (fewer structural zeros), while device indicators D_2, D_3 capture configurations associated with crossings that rarely experience crashes.

- Count (NB) part — expected crashes for at-risk crossings:

$$E[Y_i | \text{at-risk}] = \mu_i, \log(\mu_i) = x_i^T \beta, \quad (14)$$

with $x_i = [\text{dummies}(\text{HwyClassCD}_c), \text{dummies}(\text{XSurfIDs}_c), \text{dummies}(\text{Bells}_i), \text{MaxTtSpd}_i]$. Categorical covariates HwyClassCD_c , XSurfIDs_c , and Bells are the collapsed/cleaned factors used to stabilize estimation; they are encoded as indicator (dummy) variables with a designated baseline for each factor.

Given these links, the data-generating process is a mixture of a structural zero and a negative binomial count, which yields the following observation-level probabilities and likelihood. Let $f_{NB}(y|\mu_i, \alpha)$ denote

the negative binomial pmf under the mean–dispersion parameterization. The observation-level probabilities are:

$$\Pr(Y_i = 0) = \pi_i + (1 - \pi_i)f_{NB}(0|\mu_i, \alpha), \Pr(Y_i = y > 0) = (1 - \pi_i)f_{NB}(y|\mu_i, \alpha). \quad (15)$$

Under this parameterization, coefficients from the two components admit complementary interpretations. In the count (NB) part, a one-unit increase in a scalar covariate x_{ik} multiplies the expected crash rate μ_i by $\exp(\beta_k)$; thus $\exp(\beta_k)$ is the incident-rate ratio (IRR). For categorical predictors, IRRs compare each level’s rate to the designated baseline level. In the inflation (logit) part, a one-unit increase in a scalar covariate z_{ij} multiplies the odds that a crossing is a structural zero by $\exp(\gamma_j)$; hence $\exp(\gamma_j)$ is the odds ratio for being in the always-zero state.

Because the crash counts display both an extreme concentration of zeros and clear over-dispersion among the non-zero observations, a ZINB is preferred over simpler Poisson models. Empirically, many crossings report no crashes across the five-year window (excess zeros), while those with crashes have variance that exceeds the mean ($\text{Var} > E$). The ZINB explicitly handles these features: the inflation logit captures a structural-zero process—parameterized by $\log(\text{Exposure})$, D_2 , and D_3 —that separates crossings effectively not at risk, and the NB count component with $\alpha > 0$ accommodates extra-Poisson variability. This specification also preserves familiar incident-rate ratio (IRR) interpretation in the count part for roadway class (HwyClassCD_c), surface condition (XSurfIDs_c), bell devices, and maximum train speed, linking those factors to expected crash rates at crossings deemed at risk.

3.3.1.6. *Model evaluation and interpretation*

Candidate ZINB specifications were compared on a common train/test split, and the baseline was selected by the lowest information criteria (AIC/BIC) relative to fuller variants (e.g., with Gates, without category collapsing). Final performance is summarized on the held-out test set using MAE and RMSE for count predictions. Because an operational question is often “any crash vs. none,” we also assess the model’s ranking ability via ROC-AUC and report a confusion matrix at a tuned probability threshold for the minority (non-zero) class. For interpretability, count-part coefficients are presented as incident-rate ratios ($\text{IRR} = e^\beta$), expressing multiplicative changes in the expected crash rate relative to reference levels. The selected model and all metrics are reported in Section 4.3.

3.3.2. **AI Model**

AI models are data-driven learners that map inputs (features) to outcomes (targets) by finding patterns that generalize beyond the training data. In transportation safety, they are valuable for screening (flagging sites at risk), ranking (prioritizing limited resources), and forecasting (estimating expected crash counts). Common families include tree ensembles (Random Forests, Gradient Boosting/XGBoost) that capture non-linear interactions with minimal feature engineering; linear and generalized linear models (e.g., logistic/Poisson/negative binomial) that provide interpretable coefficients and uncertainty; support-vector machines for margin-based classification; and neural networks for high-capacity pattern learning when very large, labeled datasets are available.

3.3.2.1. *Initial AI Modeling Attempts*

We began by testing flexible, tree-based learning algorithms because they can capture nonlinear effects and interactions with little feature engineering. In practice, we first used XGBoost to predict crash counts directly (regression), and then as a binary classifier to predict “any crash vs. none.” We also tried Random Forest classification, adding class-imbalance remedies such as class weights, SMOTE oversampling, and post-hoc threshold tuning. These choices were motivated by the structure of the data: many engineered and categorical features with potential interactions, coupled with a very sparse outcome where roughly five percent of crossings record any crash over five years.

3.3.2.2. *Single-Stage Models*

The early results made clear why a single-stage model was not enough. For XGBoost regression, fit on the raw counts, the model tended to predict values very close to zero across the board; this minimized squared error in a dataset dominated by zeros, but it failed to learn the positive tail, yielding poor generalization (negative R^2 and small but misleading MAE/RMSE). Switching to binary classification improved interpretability, yet accuracy appeared deceptively high because a model can be “right” most of the time by predicting the majority class (no crash). More telling metrics—precision, recall, F1, ROC-AUC, and especially PR-AUC—showed the limitation. Precision measures how many predicted crashes were truly crashes; recall measures how many true crashes we caught; F1 is their harmonic mean, punishing lopsided trade-offs. ROC-AUC summarizes how well the model ranks positives above negatives across thresholds, but it can look acceptable even when the positive class is rare. PR-AUC, by contrast, focuses on the quality of positive detections only and therefore reflects the reality of rare-event screening; our PR-AUC remained low and recall for the crash class was weak.

We pushed the classifiers further with imbalance treatments. Using ‘scale_pos_weight’ in XGBoost and “balanced” class weights in Random Forest put more cost on missing a crash, nudging recall upward but often at the expense of precision and overall calibration. SMOTE generated synthetic minority examples to enrich the decision boundary, which sometimes helped ranking but also introduced variance and mild overfitting to synthetic structure. Threshold tuning on the classifier scores traded false positives for higher crash recall and improved F1 at a chosen operating point, but gains were modest because the underlying score distributions for crash vs. non-crash still overlapped substantially. In short, each remedy helped a bit, yet none overcame the fundamental tension of trying to learn both “if any crash happens” and “how many crashes” in one step or one loss function.

3.3.2.3. *Rationale for a Two-Stage Design*

The initial modeling attempts led us to the two-stage design we adopted. Stage 1 treats “any crash vs. none” as a focused detection problem, where we can evaluate and tune directly with recall, precision, F1, and PR-AUC to meet screening goals. Stage 2 conditions on crossings predicted as at-risk and models the positive counts with a separate regressor, which is evaluated with MAE, RMSE, and R^2 on the positive subset. This separation matches the operational question flow (first identify at-risk crossings, then size the risk), reduces the extreme class imbalance seen by the regressor, and allows each stage to be optimized for the metric that matters for its task.

3.3.2.4. *Two-Stage AI Pipeline: RFC (with SMOTE) → RFR*

3.3.2.4.1. *Stage 1: Rare-event classification with Random Forest + SMOTE*

Objective of this stage is to decide whether a crossing is likely to have any crash in the five-year window. Because the crash label is highly imbalanced ($\approx 5\%$ positives), models trained naively tend to default to the majority class (“no crash”), yielding high apparent accuracy but poor minority detection. To address this, we use SMOTE (Synthetic Minority Over-sampling Technique), which synthesizes additional positive examples by interpolating between nearest-neighbor minority samples in feature space. This augments the training set with realistic, diverse positives, allowing the classifier to learn meaningful decision boundaries rather than collapsing to an “all zeros” solution.

Prior to modeling, we removed administrative fields (CrossingID, State, Whistban) and dropped rows with missing values. To stabilize heavy-tailed scales, we applied $\log(1+x)$ to selected numerics (Aadt, Exposure, MaxTtSpd).

RFC with hyperparameter tuning via *GridSearchCV* (5-fold CV, scoring='f1', n_jobs=-1, random_state=42). The grid targeted model capacity (n_estimators, max_depth), regularization (min_samples_split), and imbalance handling (class_weight), enabling the RFC to capture non-linear interactions while controlling overfitting and weighting errors on the minority class. The selected model (from the grid search) was refit on the training data and evaluated on the held-out test fold at the default probability threshold 0.50, with accuracy, precision, recall, F1, confusion matrix, and a full classification report. In our runs, the cross-validated best configuration was: n_estimators=100, max_depth=15, min_samples_split=2, class_weight='balanced_subsample'. Table 3 summarizes the hyperparameters used for the Stage 1 RFC, including the search grid, the selected values from GridSearchCV, and the purpose of each parameter.

Table 3. Stage 1-RFC: Hyperparameters and their roles

Parameter	Values explored	Selected	Role / effect
n_estimators	100	100	Number of trees; higher can reduce variance and improve stability at extra compute cost.
max_depth	5, 10, 15	15	Maximum tree depth; governs model complexity. Larger depth captures more interactions but risks overfitting without enough regularization.
min_samples_split	2, 5, 10	2	Minimum samples to split an internal node; higher values regularize by preventing overly specific splits.
class_weight	balanced, balanced_subsample	balanced_subsample	Reweights classes inversely to frequency (globally or per bootstrap sample) so minority errors carry more penalty. Aids recall on rare crashes.
criterion	(default) gini	gini (default)	Node impurity measure; gini is standard for speed and accuracy; not varied here.
max_features	(default) sqrt	sqrt (default)	Features considered per split; promotes tree diversity and reduces correlation between trees.
bootstrap	(default) True	True (default)	Sample with replacement for each tree; foundational to Random Forest bagging.
random_state	42	42	Ensures reproducibility of splits/bootstraps.

3.3.2.4.2. Stage 2: Count modeling for predicted-positive sites (Random Forest Regressor)

Stage 2 estimates the magnitude of five-year crash counts for crossings that are at risk. Whereas Stage 1 answers “*is there any crash?*”, Stage 2 answers “*how many?*” so that prioritized sites can be triaged by expected burden. Consistent with the two-stage design, the regression model is trained and evaluated on the positive subset of crossings with at least one observed crash over 2020–2024 (i.e., Accidents > 0). This conditioning isolates the count-sizing problem from the detection problem, avoiding domination by zeros

and aligning the learning target with how Stage 2 is used operationally (only after a site is flagged by Stage 1).

Prior to modeling, selected skewed numeric variables are log-transformed upstream (e.g., Aadt, MaxTtSpd) to reduce scale dominance; then apply a StandardScaler to the regression feature matrix. Although tree ensembles are invariant to monotone transformations and do not require scaling, standardization is benign here and keeps feature ranges comparable for any downstream diagnostics. Table 4 summarizes the configuration used for the Stage 2 Random Forest Regressor (RFR), including the fixed values in this implementation and the purpose of each parameter.

Table 4. Stage 2-RFR: Hyperparameters and their roles

Parameter	Value	Role / effect
n_estimators	100	Number of trees; averaging reduces variance and improves stability.
max_depth	10	Controls tree complexity; deeper trees capture richer interactions but risk overfitting.
max_features	sqrt (default)	Random subset of features per split; decorrelates trees and improves generalization.
bootstrap	True (default)	Sample with replacement per tree; foundational to bagging/variance reduction.
min_samples_split	default	Regularization via minimum samples to split an internal node; prevents overspecific partitions.
min_samples_leaf	default	Minimum samples per leaf; guards against leaves with very few observations.
random_state	42	Reproducibility of ensemble construction and results.

4. Analysis and Results

4.1. Phase I: Volume Prediction Results

The model’s predictive accuracy was assessed with several standard indicators. On the 17 hold-out locations, the coefficient of determination was $R^2 = 0.816$, indicating that the model explains about 81.6% of the variability in observed daily non-motorist volumes. The mean absolute error (MAE) was 14.1 users/day and the root-mean-squared error (RMSE) was 18.0 users/day, reflecting typical absolute deviations on the order of a few tens of users per day across a range of 0–175. As a practical check, we also report a tolerance accuracy: 82.4% of predictions fall within ± 15 users/day of ground truth. Taken together, these metrics indicate a strong, well-calibrated fit on the validation set, capturing the bulk of site-to-site variation while keeping errors at magnitudes that are operationally acceptable for downstream planning and screening.

4.2. FRA APS Variants: AADT-Only vs. AADT+Non-Motorist Exposure

The linear regression model was used to predict non-motorist traffic volumes at HRGCs. Evaluation of the two scenarios utilized the FRA’s APS model: the first scenario with only AADT and the second with AADT plus the predicted daily non-motorist traffic volume. All other variable values between the two scenarios were similar. Table 5 presents summary statistics for the percentage difference in crash predictions between the two scenarios. The main finding is that on average predicted yearly crashes at HRGCs increased by 5.51 percent when non-motorist traffic was considered in the FRA APS Model. The median value indicates that there are significant deviations at the distribution’s tails. The IQR shows moderate dispersion among mid-range predictions; the maximum increase in crash prediction was 60.59 percent while the minimum was negligible.

Table 5. Percentage differences in crash predictions after incorporating non-motorist exposure across all four states

Mean Percent Difference	Median Percent Difference	Interquartile Range (IQR)	Maximum Percent Difference	Minimum Percent Difference
5.51%	1.28%	7.73%	60.59%	0.0017%

The box plots in Figure 4 show the percentage differences in crash predictions across the four states when moving from AADT-only exposure to exposure that includes both AADT and non-motorist traffic volume. The distributions for MO, KS, and IA show notable variation, with broader interquartile ranges and higher maximum values, demonstrating that non-motorist traffic volume has a significant impact on predicted HRGC crashes. Missouri has the greatest distribution and outliers, indicating considerable non-motorist traffic prevalence at HRGCs. Overall, the box plots visually support the conclusion that non-motorist traffic at HRGCs varies by state, with stronger effects where predicted non-motorist activity is more prevalent.

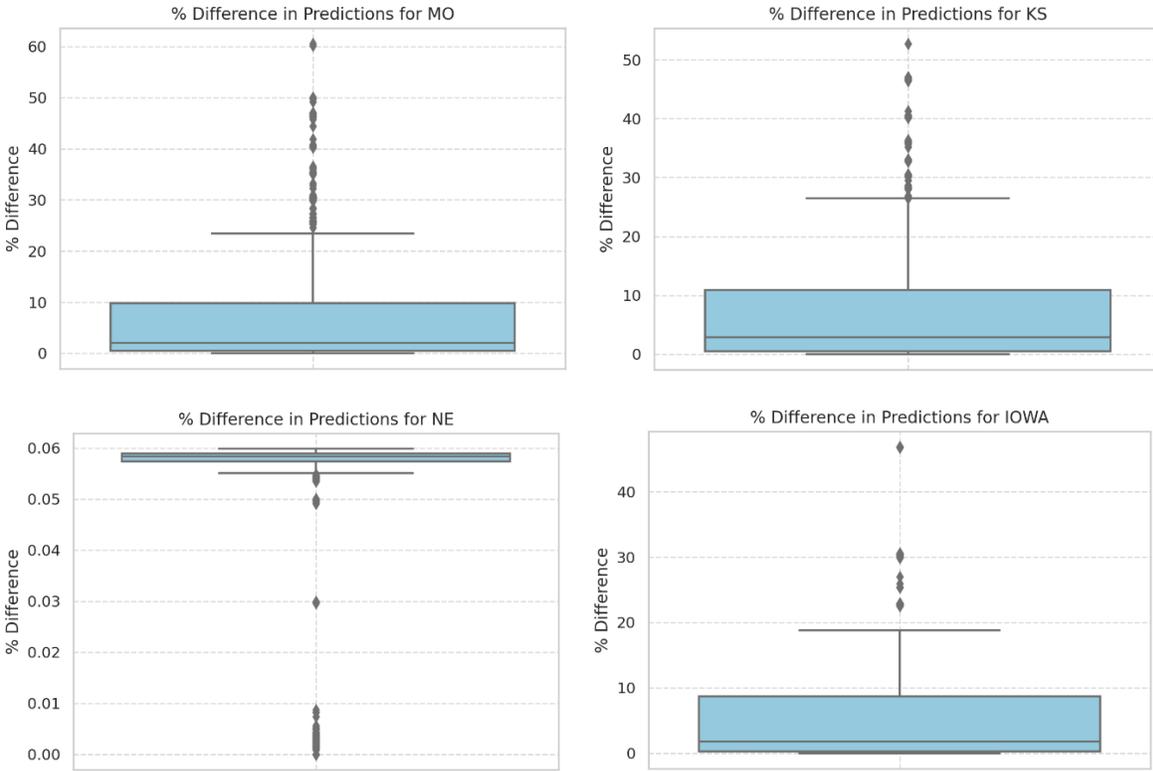


Figure 4. Box plots showing the distribution of percent differences in crash prediction outputs (Note: Each subplot compares the AADT-only prediction against the AADT plus non-motorist prediction, highlighting the variation in model sensitivity to non-motorist volume across the four different states)

4.3. Phase II: Crash Prediction Model

4.3.1. Statistical Model – Results and Analysis

We evaluated several ZINB specifications and building directly on the screening, encoding, and design-matrix decisions in subsection 3.3.1.3 (Screening and Encoding) and the model formulation in subsection 3.3.1.5 (ZINB Specification), selected the configuration with a logit inflation on $\log(\text{Exposure}) + D2 + D3$ and a log-count on $C(\text{HwyClassCD}_c) + C(\text{XSurfaceID2s}_c) + C(\text{Bells}) + \text{MaxTtSpd}$. This model achieved the best information criteria on the common sample (lowest AIC/BIC) while keeping hold-out error unchanged relative to fuller variants ($\text{MAE} \approx 0.091$, $\text{RMSE} \approx 0.235$ on five-year counts). Excluding D2/D3 from the inflation equation degraded fit substantially ($\Delta\text{AIC} > 40$), confirming their importance. In the count part, HwyClassCD_c (levels 1, 3, 4), XSurfaceID2s_c (levels 1, 3, 4), Bells (levels 1, 2, “other”), and MaxTtSpd show positive, statistically significant associations with expected crash counts relative to baselines; the estimated over-dispersion ($\alpha > 0$) supports Negative Binomial over Poisson. In the inflation part, $\log(\text{Exposure})$ is negative and highly significant—higher $\text{AADT} \times \text{TotalTrains}$ lowers the odds that a site is “always-zero”, while D2 and D3 are positive and significant, identifying locations prone to structural zeros. Overall, the chosen ZINB reconciles the heavy zero mass and over-dispersion described in subsection

3.3.1.1 (Data and Target), delivers competitive predictive accuracy, and yields interpretable effects aligned with the Chapter 3 setup.

4.3.2. AI Model – Results and Analysis

We implemented a two-stage learning pipeline to (i) classify whether a crossing experiences at least one crash in the observation window and, conditional on a positive classification, (ii) predict the crash count. Prior to modeling, numeric skew was reduced with a \log_{1p} transform applied to Exposure, and MaxTtSpd using equation (16), consistent with the preprocessing shown in the code.

$$\log(1 + x) \tag{16}$$

Stage 1 used an RFC with SMOTE to handle class imbalance. The final feature set was: \log_{1p} of Exposure, \log_{1p} of MaxTtSpd, NightThru, HwySpeed, and Gates. Stage 2 feature set was: HwySpeed, \log_{1p} of MaxTtSpd, \log_{1p} of Aadt, Bells, TotalTrains, and NightThru. A Random Forest Regressor was trained and evaluated on the corresponding positive subset. Table 6 and Table 7 summarize Stage 1 classification performance under SMOTE-balanced and original imbalanced test distributions, respectively; Table 8 reports end-to-end performance (MAE, RMSE) of the full two-stage pipeline.

Table 6. Stage 1 Random Forest performance metrics for crash detection using SMOTE-balanced data

Metric	Value
Accuracy	0.9628
Precision	0.9505
Recall	0.9763
F1	0.9633

Table 7. Stage 1 Random Forest performance metrics for crash detection evaluated on the original imbalanced dataset

Metric	Value
Accuracy	0.9477
Precision	0.508
Recall	0.97
F1	0.667

Table 8. End-to-end performance of the two-stage AI crash prediction pipeline (Stage 1 classification followed by Stage 2 regression), reporting MAE and RMSE on the test dataset

Metric	Value
MAE	0.0668
RMSE	0.27

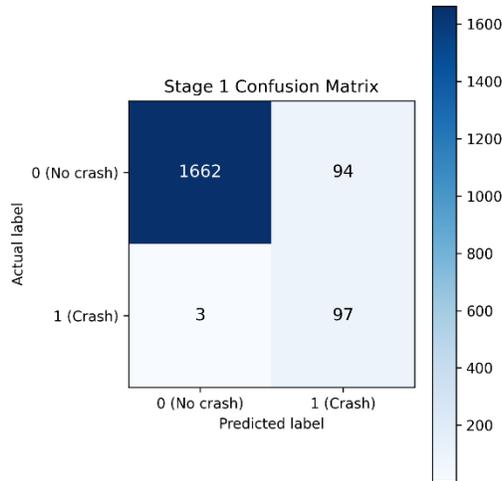


Figure 5. Confusion matrix summarizing Stage 1 predictions vs. actual crash presence.

On a SMOTE-balanced test sample, the stage 1 classifier achieves accuracy = 0.9628, precision = 0.9505, recall = 0.9763, and F1 = 0.9633. Table 6 presents these balanced sample results, which reflect the model’s capacity when classes are artificially equalized and serve as a diagnostic upper bound on what the learner can capture under balanced conditions.

Evaluated on the original (imbalanced) test distribution (positives \approx 5.4%), performance is accuracy = 0.9477, precision = 0.508, recall = 0.970, and F1 = 0.667, with confusion matrix TN = 1662, FP = 94, FN = 3, TP = 97 (Figure 5). Table 7 presents these real-world metrics. The very high recall (0.97) indicates the screen misses only 3 of 100 crash-prone crossings, appropriate for a safety-first objective, while precision (\sim 0.51) is roughly nine times the base rate, showing meaningful discrimination despite class imbalance. The confusion matrix shows a dominant TN cell and a very small FN cell: most negatives are correctly rejected, and nearly all positives are detected. Although the TP count is smaller in absolute terms (due to the low base rate), the recall is 0.97. False positives (94) are moderate relative to 1,756 negatives (specificity \approx 0.95).

Combining Stage 1 decisions with Stage 2 predictions to produce counts for all test crossings yields MAE = 0.0668 and RMSE = 0.27 crashes, indicating that predicted totals closely track observed totals and provide a reliable ranking of flagged sites for resource allocation and mitigation planning. Table 8 presents these end-to-end results for the AI model (Stage 1 + Stage 2 combined).

5. Conclusions

Incorporating modeled non-motorist exposure (adding pedestrian and bicyclist volumes to AADT) materially improves crash risk estimation at HRGCs; APS predictions increase by an average ~5–6% with larger uplifts at high-activity sites, correcting systematic under-prediction when non-motorists are ignored. The statistical crash model, ZINB, captures the data’s structure ($\approx 95\%$ zeros and over-dispersion) and yields interpretable effects: higher HwyClassCD_c, XSurfaceID2s_c, and MaxTtSpd increase expected crashes at at-risk sites, while higher $\log(\text{Exposure})$ lowers the odds of structural zeros, with D2/D3 improving zero/at-risk separation. Complementing this, our two-stage AI pipeline (Stage-1 RFC + SMOTE for “any crash,” Stage-2 RFR for counts) achieved lower held-out MAE/RMSE than the ZINB baseline, indicating better predictive accuracy on this dataset; the gain likely comes from the AI model’s ability to learn nonlinearities and interactions (e.g., exposure \times device configuration) that a log-linear count model cannot fully express. In practice, ZINB remains valuable for policy-facing interpretation (IRRs, device/class effects), while the AI model provides a higher-accuracy screen for prioritization and near-term mitigation targeting.

5.1. Limitations

This study has several limitations that should be considered when interpreting the results. First, the non-motorist volume model was trained on observations collected only in Lincoln, Nebraska, then applied to crossings in Kansas, Missouri, and Iowa. That geographic mismatch can introduce domain-shift bias if patterns of pedestrian/bicyclist activity, crossing design, or enforcement differ across states. Second, temporal coverage was limited—daily non-motorist counts were not seasonally normalized—so predicted exposure may not fully capture month-to-month variation. Third, our FRA APS “with-non-motorists” scenario operationalizes exposure by adding predicted non-motorist counts to AADT, implicitly assuming their marginal effect on crash risk parallels motor-vehicle exposure; this simplifying assumption may under- or over-state true sensitivity. Fourth, inventory data quality (e.g., device coding, speed limits) and cross-state reporting practices can affect inputs, and while we cleaned and screened variables, residual measurement error may persist. Finally, sample sizes for some validations were modest (e.g., the 17-site ground-truth set for non-motorist volumes), which constrains precision of out-of-sample error estimates. Future work will expand geographic/seasonal coverage for non-motorist counts, re-estimate exposure response directly (rather than equating it to AADT), and test calibration with state-specific hold-outs to quantify transferability.

References

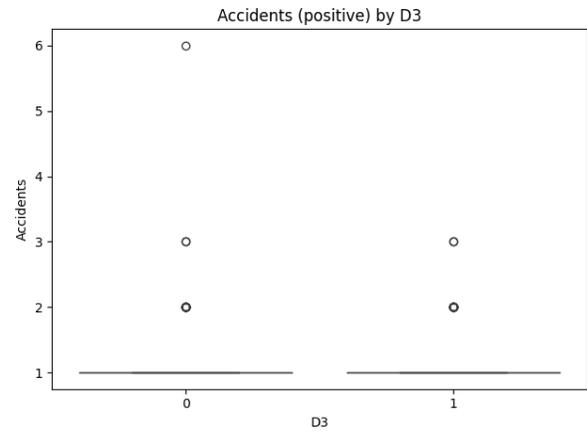
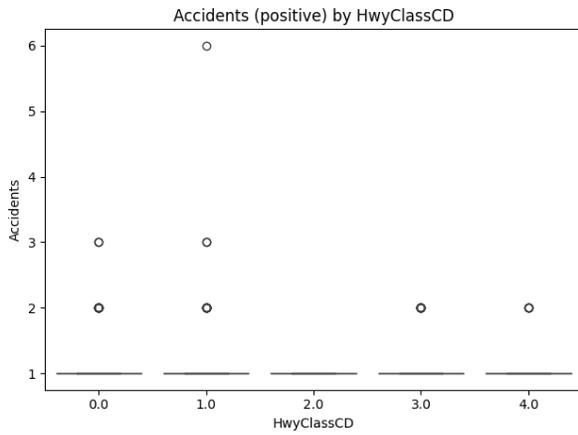
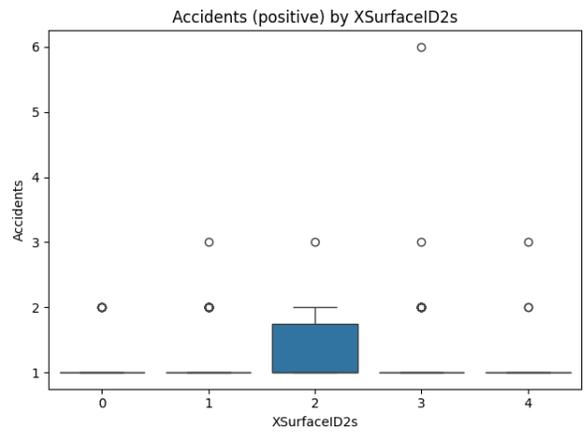
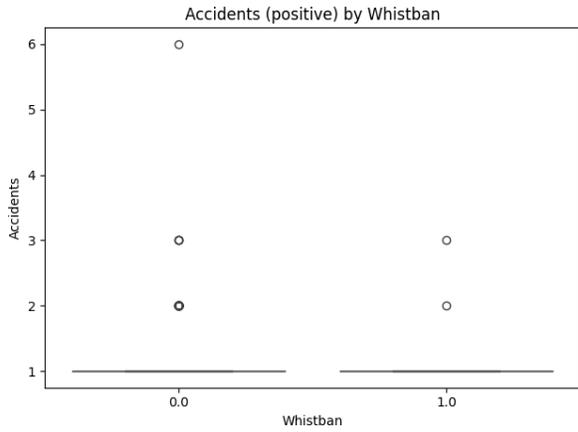
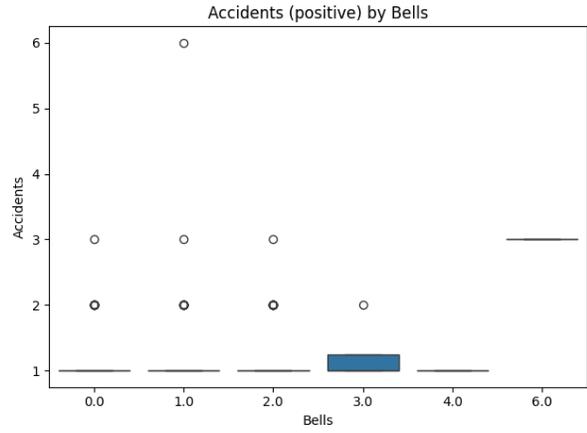
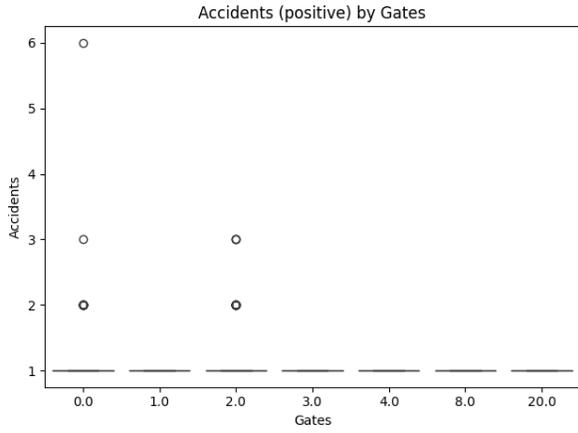
- [1] DaSilva, M., and Ngamdung, T. (2014). *Trespass Prevention Research Study-West Palm Beach, FL* (No. DOT-VNTSC-FRA-14-02). United States. Federal Railroad Administration. Office of Research and Development.
- [2] daSilva, M. P. (2011). *Railroad Infrastructure Trespass Detection Performance Guidelines* (No. DOT/FRA/ORD-11/01). United States. Federal Railroad Administration. Office of Research and Development.
- [3] Baillargeon, J., and Doran, J. (2021). *Automated Video Inspection System for Grade Crossing Safety [Research Results]* (No. RR 21-05). United States. Department of Transportation. Federal Railroad Administration.
- [4] Farooq, M. U., and A. J. Khattak. Investigating highway-rail grade crossing inventory data quality's role in crash model estimation and crash prediction. *Applied Sciences*, 2023. Vol. 13, No. 20, 11537.
- [5] Khattak, A. J., Farooq, M. U., and Farhan, A. (2024). Motor vehicle drivers' knowledge of safely traversing highway-rail grade crossings. *Transportation research record*, 2678(7), 604-621.
- [6] Zhao, L., Farooq, M. U., and Khattak, A. J. (2024). Data accuracy matters: Improving highway-rail grade crossings crash predictions through inventory verification. *Transportation Research Record*, 03611981241270179.
- [7] Khattak, A. J., Farooq, M. U., Aman, M. N., and Bukhari, M. (2024). *Pedestrian and Bicyclist Safety at Highway-Rail Grade Crossings (Year 1 Report)* (No. UTCRS-UNL-O3CY23). University Transportation Center for Railway Safety (UTCRS) Tier-1 University Transportation Center (UTC).
- [8] Federal Railroad Administration (FRA). Safety data and reporting. Available online: <https://railroads.dot.gov/safety-data> (accessed 3 May 2024).
- [9] Jacobini, F. B., and Ngamdung, T. (2022). *Railroad Trespass Detection Using Deep Learning-Based Computer Vision [Research Results]* (No. RR 22-04). United States. Department of Transportation. Federal Railroad Administration.
- [10] Brod, D., Gillen, D., and Decisiontek, L. L. C. (2020). *A new model for highway-rail grade crossing accident prediction and severity* (No. DOT/FRA/ORD-20/40). United States. Department of Transportation. Federal Railroad Administration.
- [11] Federal Railroad Administration. (n.d.). *Safety data*. U.S. Department of Transportation. Retrieved August 13, 2025, from <https://safetydata.fra.dot.gov/gxaps-app/#/>
- [12] Lautala, P. T., Jeon, M., Nelson, D., Landry, S., and Dean, A. (2020). *Driver behavior at highway-rail grade crossings using NDS and driving simulators* (No. DOT/FRA/ORD-20/47). United States. Department of Transportation. Federal Railroad Administration.

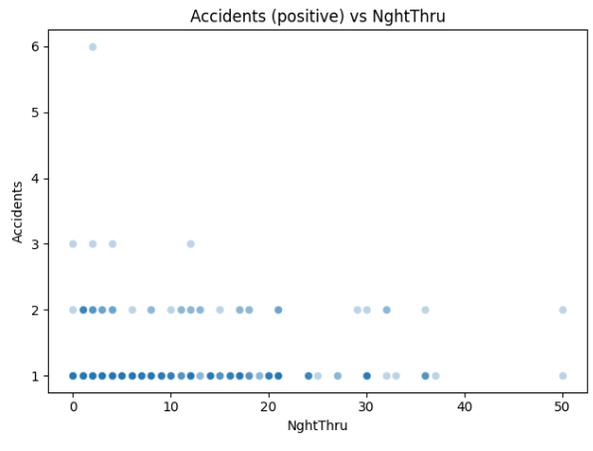
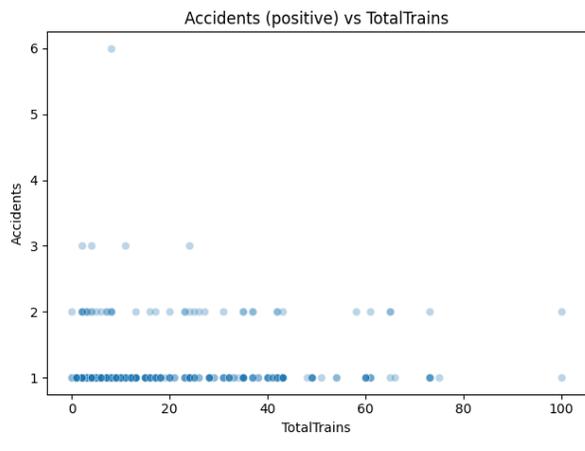
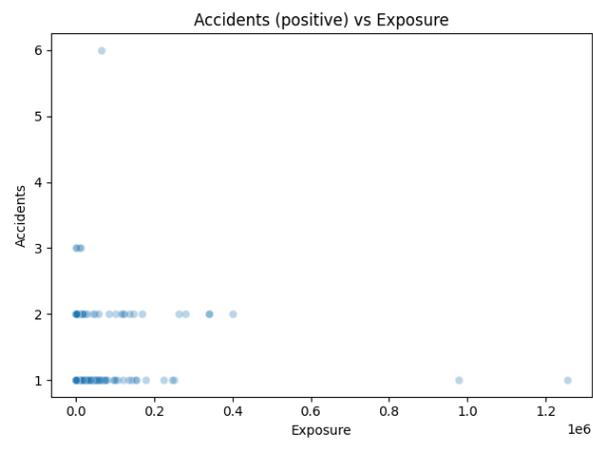
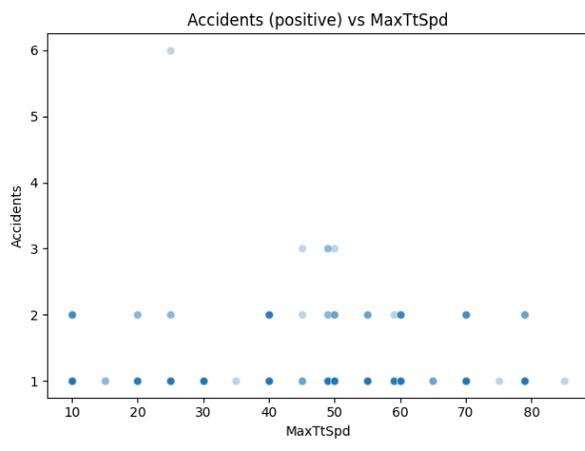
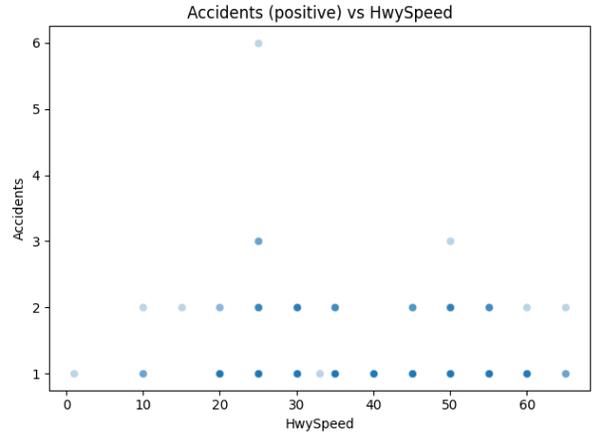
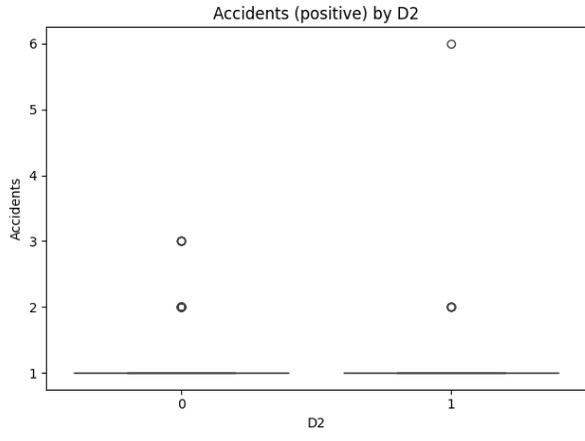
- [13] Gabree, S. H., Calabrese, C., Young, J., Mejia, B., and Prescott, J. (2019). *Review of Highway-Rail Grade Crossing Incidents 1986–2015: Trends and Potential Safety Gaps* (No. DOT/FRA/ORD-19/30). United States. Federal Railroad Administration. Office of Railroad Policy and Development.
- [14] Goodchild, Anne V., Edward McCormack, Anna Bovbjerg, and Manali Sheth. *Multi-Modal Intersections: Resolving Conflicts Between Trains, Motor Vehicles, Bicyclists and Pedestrians*. No. FHWA-OR-18-04. Oregon. Dept. of Transportation, 2017.
- [15] GAO, Improvements Needed to Federal Technical Assistance About Pedestrian Projects Related to Trespassing, United States Government Accountability Office Report GAO-25-107115, Washington, D.C., April 2025.
- [16] GAO, Grade-Crossing Safety: DOT Should Evaluate Whether Program Provides States Flexibility to Address Ongoing Challenges, United States Government Accountability Office Report GAO-19-80, Washington, D.C., Nov. 8, 2018.
- [17] Liu, H., M. Lee, and A. J. Khattak. Updating Annual Average Daily Traffic Estimates at Highway-Rail Grade Crossings with Geographically Weighted Poisson Regression. *Transportation Research Record: Journal of the Transportation Research Board*, 2019, Vol. 2673(10) 105–117; DOI: 10.1177/0361198119844976.
- [18] Khattak, A. J., and Z. Luo. Pedestrian and Bicyclist Violations at Highway-Rail Grade Crossings. *Transportation Research Record: Journal of the Transportation Research Board*, 2250, No. 1, 2011; 76-82.
- [19] Khan, I. U, E. Lee, and M. A. Khan. Developing a Highway Rail Grade Crossing Accident Probability Prediction Model: A North Dakota Case Study. *Safety* 4, No. 2, 2018; 22.
- [20] Ling, Z., C. R. Cherry, and N. Dhakal. Factors Influencing Single-Bicycle Crashes at Skewed Railroad Grade Crossings. *Journal of Transport and Health*, 7, 2017; 54-63.
- [21] Russo, B. J., E. James, T. Erdmann, and E. Smaglik. Pedestrian and Bicyclist Behavior at Highway-Rail grade Crossings: An Observational Study of Factors Associated with Violations, Distraction, and Crossing Speeds During Train Crossing Events. *Journal of Transportation Safety and Security*, 13, No. 11, 2021; 1263-1281.
- [22] Metaxatos, P., and P. S. Sriraj. Pedestrian/Bicyclist Warning Devices and Signs at Highway-Rail and Pathway-Rail Grade Crossings. Report FHWA-ICT-13-013, Urban Transportation Center, University of Illinois at Chicago, 2013.
- [23] Thompson, A., and B. J. Kennedy. Engineering Design for Pedestrian Safety at Highway-Rail Grade Crossings. USDOT FRA Report DOT/FRA/ORD-16/24, 2016.
- [24] Ren, Q., M. Xu, B. Zhou, and S. Chung. 2024. Traffic safety assessment and injury severity analysis for undivided two-way highway-rail grade crossings. *Mathematics*, 12(4): 519. <https://doi.org/10.3390/math12040519>

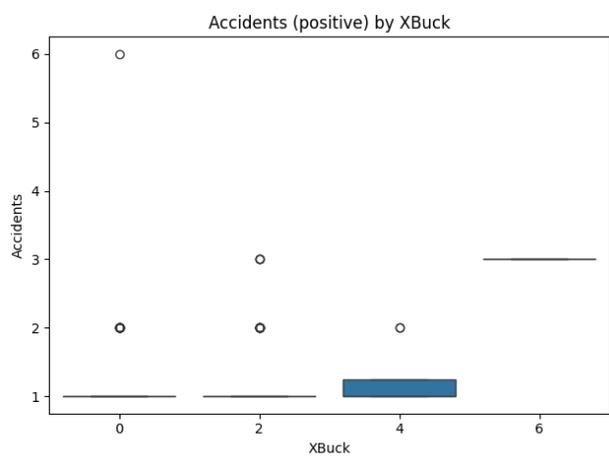
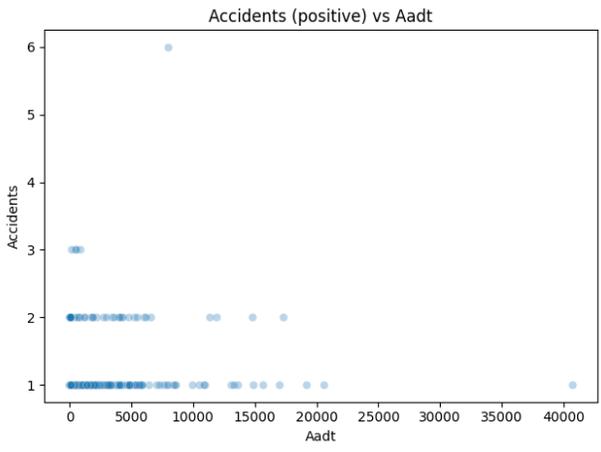
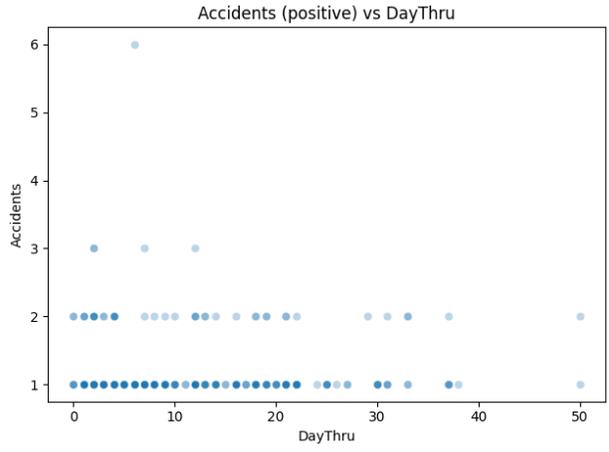
- [25] Yeh, M., and J. Multer. 2007. Traffic control devices and barrier systems at grade crossings: Literature review. *Transportation Research Record: Journal of the Transportation Research Board*, 2030(1): 69–75. <https://doi.org/10.3141/2030-10>
- [26] Hao, W., and C. Kamga. 2015. The effects of lighting on driver's injury severity at highway-rail grade crossings. *Journal of Advanced Transportation*, 50(4): 446–458. <https://doi.org/10.1002/atr.1353>
- [27] Yeh, M., T. Raslear, and J. Multer. 2012. Evaluating the impact of grade crossing safety factors through signal detection theory. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 56(1): 2226–2230. <https://doi.org/10.1177/1071181312561469>
- [28] Wang, X., J. Li, C. Zhang, and T. Qiu. 2019. Active warning system for highway-rail grade crossings using connected vehicle technologies. *Journal of Advanced Transportation*, 2019: 1–11. <https://doi.org/10.1155/2019/3219387>
- [29] Nadri, C., S. Lee, S. Kekal, Y. Li, X. Li, P. Lautala, *et al.* 2021. Effects of auditory display types and acoustic variables on subjective driver assessment in a rail crossing context. *Transportation Research Record: Journal of the Transportation Research Board*, 2675(9): 1457–1468. <https://doi.org/10.1177/03611981211007838>
- [30] Rana, P., F. Sattari, L. Lefsrud, and M. Hendry. 2023. Machine learning approach to enhance highway railroad grade crossing safety by analyzing crash data and identifying hotspot crash locations. *Transportation Research Record: Journal of the Transportation Research Board*, 2678(7): 1055–1071. <https://doi.org/10.1177/03611981231212162>
- [31] Fan, W., M. Kane, and E. Haile. 2015. Analyzing severity of vehicle crashes at highway-rail grade crossings: Multinomial logit modeling. *Journal of the Transportation Research Forum*, 54(2). <https://doi.org/10.5399/osu/jtrf.54.2.4282>
- [32] Kalašová, A., K. Čulík, and S. Skřivánek-Kubíková. 2022. The interaction between a driver and intelligent transport systems. *Acta Polytechnica CTU Proceedings*, 39: 16–22. <https://doi.org/10.14311/app.2022.39.0016>
- [33] Zhao, Z., and J. Chen. 2025. Application of artificial intelligence technology in the economic development of urban intelligent transportation system. *PeerJ Computer Science*, 11: e2728. <https://doi.org/10.7717/peerj-cs.2728>
- [34] Zhao, Y. 2024. The application of artificial intelligence in intelligent transportation system. *Highlights in Science, Engineering and Technology*, 83: 209–216. <https://doi.org/10.54097/5vy3ar35>
- [35] Jaidev, B., S. Garg, and S. Makkar. 2019. Artificial intelligence to prevent road accidents. *International Journal of Machine Learning and Networked Collaborative Engineering*, 3(1): 35–45. <https://doi.org/10.30991/ijmlnce.2019v03i01.003>
- [36] Grabowski, M., and H. Dhimi. 2005. Early adoption technology performance impact: AIS on the St. Lawrence Seaway. *Journal of Navigation*, 58(1): 17–30. <https://doi.org/10.1017/S0373463304003030>

- [37] Vestre, A., A. Bakdi, E. Vanem, and Ø. Engelhardtzen. 2021. AIS-based near-collision database generation and analysis of real collision avoidance manoeuvres. *Journal of Navigation*, 74(5): 985–1008. <https://doi.org/10.1017/S0373463321000357>
- [38] Eluru, N., C. Bhat, and D. Hensher. 2008. A mixed generalized ordered response model for examining pedestrian and bicyclist injury severity level in traffic crashes. *Accident Analysis and Prevention*, 40(3): 1033–1054. <https://doi.org/10.1016/j.aap.2007.11.010>
- [39] Dong, B., X. Ma, and F. Chen. 2018. Analyzing the injury severity sustained by non-motorists at mid-blocks considering non-motorists' pre-crash behavior. *Transportation Research Record: Journal of the Transportation Research Board*, 2672(38): 138–148. <https://doi.org/10.1177/0361198118777354>
- [40] Metaxatos, P., and P. Sriraj. 2015. Pedestrian safety at rail grade crossings: Focus areas for research and intervention. *Urban Rail Transit*, 1(4): 238–248. <https://doi.org/10.1007/s40864-016-0030-4>
- [41] Kizawi, A. 2021. Analysis of vehicle–pedestrian and bicyclist conflicts in Győr, Hungary using Swedish conflict technique. *Acta Technica Jaurinensis*, 14(4): 377–405. <https://doi.org/10.14513/actatechjaur.00605>
- [42] Easa, S., X. Qu, and E. Dabbour. 2017. Improved pedestrian sight-distance needs at railroad–highway grade crossings. *Journal of Transportation Engineering, Part A: Systems*, 143(7): 04017034. <https://doi.org/10.1061/JTEPBS.0000047>
- [43] Al-Mahameed, F., X. Qin, R. Schneider, and M. Shaon. 2019. Analyzing pedestrian and bicyclist crashes at the corridor level: Structural equation modeling approach. *Transportation Research Record: Journal of the Transportation Research Board*, 2673(7): 308–318. <https://doi.org/10.1177/0361198119845353>
- [44] Brod, D. and D. Gillen. *A New Model for Highway-Rail Grade Crossing Accident Prediction and Severity*. No. DOT/FRA/ORD-20/40. US Department of Transportation, Federal Railroad Administration, 2020.
- [45] U.S. Department of Transportation, Federal Railroad Administration. *FRA Data Dictionary for External Use–Grade Crossing*. Washington, DC, 2023. Accessed January 1, 2025. <https://railroads.dot.gov/safety-data/forms-guides-publications/guides/fra-data-dictionary-external-use-grade-crossing>.

Appendix A







U. S. DOT Crossing Inventory Form 71

U. S. DOT CROSSING INVENTORY FORM					
DEPARTMENT OF TRANSPORTATION FEDERAL RAILROAD ADMINISTRATION				OMB No. 2130-0017	
<p>Instructions for the initial reporting of the following types of new or previously unreported crossings: For public highway-rail grade crossings, complete the entire inventory Form. For private highway-rail grade crossings, complete the Header, Parts I and II, and the Submission Information section. For public pathway grade crossings (including pedestrian station grade crossings), complete the Header, Parts I and II, and the Submission Information section. For Private pathway grade crossings, complete the Header, Parts I and II, and the Submission Information section. For grade-separated highway-rail or pathway crossings (including pedestrian station crossings), complete the Header, Part I, and the Submission Information section. For changes to existing data, complete the Header, Part I Items 1-3, and the Submission Information section, in addition to the updated data fields. Note: For private crossings only, Part I Item 20 and Part III Item 2.K. are required unless otherwise noted. An asterisk * denotes an optional field.</p>					
A. Revision Date (MM/DD/YYYY) ____/____/____	B. Reporting Agency <input type="checkbox"/> Railroad <input type="checkbox"/> Transit <input type="checkbox"/> State <input type="checkbox"/> Other		C. Reason for Update (Select only one) <input type="checkbox"/> Change in Data <input type="checkbox"/> New Crossing <input type="checkbox"/> Closed <input type="checkbox"/> No Train Traffic <input type="checkbox"/> Re-Open <input type="checkbox"/> Date <input type="checkbox"/> Change in Primary Operating RR <input type="checkbox"/> Quiet Zone Update <input type="checkbox"/> Admin. Correction		D. DOT Crossing Inventory Number _____
Part I: Location and Classification Information					
1. Primary Operating Railroad		2. State		3. County	
4. City / Municipality <input type="checkbox"/> In <input type="checkbox"/> Near		5. Street/Road Name & Block Number _____ (Street/Road Name) * (Block Number)		6. Highway Type & No.	
7. Do Other Railroads Operate a Separate Track at Crossing? <input type="checkbox"/> Yes <input type="checkbox"/> No If Yes, Specify RR _____			8. Do Other Railroads Operate Over Your Track at Crossing? <input type="checkbox"/> Yes <input type="checkbox"/> No If Yes, Specify RR _____		
9. Railroad Division or Region <input type="checkbox"/> None		10. Railroad Subdivision or District <input type="checkbox"/> None		11. Branch or Line Name <input type="checkbox"/> None	12. RR Milepost _____ _____ _____ (prefix) (nnnn.nnn) (suffix)
13. Line Segment *		14. Nearest RR Timetable Station *	15. Parent RR (if applicable) <input type="checkbox"/> N/A		16. Crossing Owner (if applicable) <input type="checkbox"/> N/A
17. Crossing Type <input type="checkbox"/> Public <input type="checkbox"/> Private	18. Crossing Purpose <input type="checkbox"/> Highway <input type="checkbox"/> Pathway, Ped. <input type="checkbox"/> Station, Ped.	19. Crossing Position <input type="checkbox"/> At Grade <input type="checkbox"/> RR Under <input type="checkbox"/> RR Over	20. Public Access (if Private Crossing) <input type="checkbox"/> Yes <input type="checkbox"/> No	21. Type of Train <input type="checkbox"/> Freight <input type="checkbox"/> Transit <input type="checkbox"/> Intercity Passenger <input type="checkbox"/> Shared Use Transit <input type="checkbox"/> Commuter <input type="checkbox"/> Tourist/Other	
22. Average Passenger Train Count Per Day <input type="checkbox"/> Less Than One Per Day <input type="checkbox"/> Number Per Day _____					
23. Type of Land Use <input type="checkbox"/> Open Space <input type="checkbox"/> Farm <input type="checkbox"/> Residential <input type="checkbox"/> Commercial <input type="checkbox"/> Industrial <input type="checkbox"/> Institutional <input type="checkbox"/> Recreational <input type="checkbox"/> RR Yard					
24. Is there an Adjacent Crossing with a Separate Number? <input type="checkbox"/> Yes <input type="checkbox"/> No If Yes, Provide Crossing Number _____			25. Quiet Zone (FRA provided) <input type="checkbox"/> No <input type="checkbox"/> 24 Hr <input type="checkbox"/> Partial <input type="checkbox"/> Chicago Excused Date Established _____		
26. HSR Corridor ID <input type="checkbox"/> N/A		27. Latitude in decimal degrees (WGS84 std: nn.nnnnnnn)		28. Longitude in decimal degrees (WGS84 std: -nnn.nnnnnnn)	29. Lat/Long Source <input type="checkbox"/> Actual <input type="checkbox"/> Estimated
30.A. Railroad Use *		31.A. State Use *			
30.B. Railroad Use *		31.B. State Use *			
30.C. Railroad Use *		31.C. State Use *			
30.D. Railroad Use *		31.D. State Use *			
32.A. Narrative (Railroad Use) *			32.B. Narrative (State Use) *		
33. Emergency Notification Telephone No. (posted)		34. Railroad Contact (Telephone No.)		35. State Contact (Telephone No.)	
Part II: Railroad Information					
1. Estimated Number of Daily Train Movements					
1.A. Total Day Thru Trains (6 AM to 6 PM)	1.B. Total Night Thru Trains (6 PM to 6 AM)	1.C. Total Switching Trains	1.D. Total Transit Trains	1.E. Check if Less Than One Movement Per Day <input type="checkbox"/> How many trains per week? _____	
2. Year of Train Count Data (YYYY)		3. Speed of Train at Crossing 3.A. Maximum Timetable Speed (mph) _____ 3.B. Typical Speed Range Over Crossing (mph) From _____ to _____			
4. Type and Count of Tracks Main _____ Siding _____ Yard _____ Transit _____ Industry _____					
5. Train Detection (Main Track only) <input type="checkbox"/> Constant Warning Time <input type="checkbox"/> Motion Detection <input type="checkbox"/> AFD <input type="checkbox"/> PTC <input type="checkbox"/> DC <input type="checkbox"/> Other <input type="checkbox"/> None					
6. Is Track Signaled? <input type="checkbox"/> Yes <input type="checkbox"/> No		7.A. Event Recorder <input type="checkbox"/> Yes <input type="checkbox"/> No		7.B. Remote Health Monitoring <input type="checkbox"/> Yes <input type="checkbox"/> No	
FORM FRA F 6180.71 (Rev. 08/03/2016)		OMB approval expires 11/30/2022		Page 1 OF 2	

U. S. DOT CROSSING INVENTORY FORM

A. Revision Date (MM/DD/YYYY)		PAGE 2		D. Crossing Inventory Number (7 char.)	
Part III: Highway or Pathway Traffic Control Device Information					
1. Are there Signs or Signals? <input type="checkbox"/> Yes <input type="checkbox"/> No		2. Types of Passive Traffic Control Devices associated with the Crossing			
2.A. Crossbuck Assemblies (count)		2.B. STOP Signs (R1-1) (count)	2.C. YIELD Signs (R1-2) (count)	2.D. Advance Warning Signs (Check all that apply; include count) <input type="checkbox"/> None	
				<input type="checkbox"/> W10-1 _____ <input type="checkbox"/> W10-3 _____ <input type="checkbox"/> W10-11 _____ <input type="checkbox"/> W10-2 _____ <input type="checkbox"/> W10-4 _____ <input type="checkbox"/> W10-12 _____	
2.E. Low Ground Clearance Sign (W10-5) <input type="checkbox"/> Yes (count _____) <input type="checkbox"/> No		2.F. Pavement Markings <input type="checkbox"/> Stop Lines <input type="checkbox"/> Dynamic Envelope <input type="checkbox"/> RR Xing Symbols <input type="checkbox"/> None		2.G. Channelization Devices/Medians <input type="checkbox"/> All Approaches <input type="checkbox"/> Median <input type="checkbox"/> One Approach <input type="checkbox"/> None	2.H. EXEMPT Sign (R15-3) <input type="checkbox"/> Yes <input type="checkbox"/> No
2.I. ENS Sign (I-13) Displayed <input type="checkbox"/> Yes <input type="checkbox"/> No					
2.J. Other MUTCD Signs <input type="checkbox"/> Yes <input type="checkbox"/> No Specify Type _____ Count _____ Specify Type _____ Count _____ Specify Type _____ Count _____		2.K. Private Crossing Signs (if private) <input type="checkbox"/> Yes <input type="checkbox"/> No	2.L. LED Enhanced Signs (List types)		
3. Types of Train Activated Warning Devices at the Grade Crossing (specify count of each device for all that apply)					
3.A. Gate Arms (count)	3.B. Gate Configuration <input type="checkbox"/> 2 Quad <input type="checkbox"/> Full (Barrier) Resistance <input type="checkbox"/> 3 Quad <input type="checkbox"/> 4 Quad <input type="checkbox"/> Median Gates	3.C. Cantilevered (or Bridged) Flashing Light Structures (count) Over Traffic Lane _____ <input type="checkbox"/> Incandescent Not Over Traffic Lane _____ <input type="checkbox"/> LED		3.D. Mast Mounted Flashing Lights (count of masts) _____ <input type="checkbox"/> Incandescent <input type="checkbox"/> LED <input type="checkbox"/> Back Lights Included <input type="checkbox"/> Side Lights Included	3.E. Total Count of Flashing Light Pairs
3.F. Installation Date of Current Active Warning Devices: (MM/YYYY) ____/____/____ <input type="checkbox"/> Not Required		3.G. Wayside Horn <input type="checkbox"/> Yes Installed on (MM/YYYY) ____/____/____ <input type="checkbox"/> No		3.H. Highway Traffic Signals Controlling Crossing <input type="checkbox"/> Yes <input type="checkbox"/> No	3.I. Bells (count)
3.J. Non-Train Active Warning <input type="checkbox"/> Flagger/Flagman <input type="checkbox"/> Manually Operated Signals <input type="checkbox"/> Watchman <input type="checkbox"/> Floodlighting <input type="checkbox"/> None				3.K. Other Flashing Lights or Warning Devices Count _____ Specify type _____	
4.A. Does nearby Hwy Intersection have Traffic Signals? <input type="checkbox"/> Yes <input type="checkbox"/> No	4.B. Hwy Traffic Signal Interconnection <input type="checkbox"/> Not Interconnected <input type="checkbox"/> For Traffic Signals <input type="checkbox"/> For Warning Signs	4.C. Hwy Traffic Signal Preemption <input type="checkbox"/> Simultaneous <input type="checkbox"/> Advance	5. Highway Traffic Pre-Signals <input type="checkbox"/> Yes <input type="checkbox"/> No Storage Distance * _____ Stop Line Distance * _____	6. Highway Monitoring Devices (Check all that apply) <input type="checkbox"/> Yes - Photo/Video Recording <input type="checkbox"/> Yes - Vehicle Presence Detection <input type="checkbox"/> None	
Part IV: Physical Characteristics					
1. Traffic Lanes Crossing Railroad Number of Lanes _____		<input type="checkbox"/> One-way Traffic <input type="checkbox"/> Two-way Traffic <input type="checkbox"/> Divided Traffic	2. Is Roadway/Pathway Paved? <input type="checkbox"/> Yes <input type="checkbox"/> No	3. Does Track Run Down a Street? <input type="checkbox"/> Yes <input type="checkbox"/> No	4. Is Crossing Illuminated? (Street lights within approx. 50 feet from nearest rail) <input type="checkbox"/> Yes <input type="checkbox"/> No
5. Crossing Surface (on Main Track, multiple types allowed) Installation Date * (MM/YYYY) ____/____/____ Width * _____ Length * _____ <input type="checkbox"/> 1 Timber <input type="checkbox"/> 2 Asphalt <input type="checkbox"/> 3 Asphalt and Timber <input type="checkbox"/> 4 Concrete <input type="checkbox"/> 5 Concrete and Rubber <input type="checkbox"/> 6 Rubber <input type="checkbox"/> 7 Metal <input type="checkbox"/> 8 Unconsolidated <input type="checkbox"/> 9 Composite <input type="checkbox"/> 10 Other (specify) _____					
6. Intersecting Roadway within 500 feet? <input type="checkbox"/> Yes <input type="checkbox"/> No If Yes, Approximate Distance (feet) _____			7. Smallest Crossing Angle <input type="checkbox"/> 0° - 29° <input type="checkbox"/> 30° - 59° <input type="checkbox"/> 60° - 90°		8. Is Commercial Power Available? * <input type="checkbox"/> Yes <input type="checkbox"/> No
Part V: Public Highway Information					
1. Highway System <input type="checkbox"/> (01) Interstate Highway System <input type="checkbox"/> (02) Other Nat Hwy System (NHS) <input type="checkbox"/> (03) Federal AID, Not NHS <input type="checkbox"/> (08) Non-Federal Aid		2. Functional Classification of Road at Crossing <input type="checkbox"/> (0) Rural <input type="checkbox"/> (1) Urban <input type="checkbox"/> (1) Interstate <input type="checkbox"/> (5) Major Collector <input type="checkbox"/> (2) Other Freeways and Expressways <input type="checkbox"/> (3) Other Principal Arterial <input type="checkbox"/> (6) Minor Collector <input type="checkbox"/> (4) Minor Arterial <input type="checkbox"/> (7) Local		3. Is Crossing on State Highway System? <input type="checkbox"/> Yes <input type="checkbox"/> No	4. Highway Speed Limit _____ MPH <input type="checkbox"/> Posted <input type="checkbox"/> Statutory
7. Annual Average Daily Traffic (AADT) Year _____ AADT _____		8. Estimated Percent Trucks _____ %	9. Regularly Used by School Buses? <input type="checkbox"/> Yes <input type="checkbox"/> No Average Number per Day _____		10. Emergency Services Route <input type="checkbox"/> Yes <input type="checkbox"/> No
Submission Information - This information is used for administrative purposes and is not available on the public website.					
Submitted by _____ Organization _____ Phone _____ Date _____					
Public reporting burden for this information collection is estimated to average 30 minutes per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed and completing and reviewing the collection of information. According to the Paperwork Reduction Act of 1995, a federal agency may not conduct or sponsor, and a person is not required to, nor shall a person be subject to a penalty for failure to comply with, a collection of information unless it displays a currently valid OMB control number. The valid OMB control number for information collection is 2130-0017. Send comments regarding this burden estimate or any other aspect of this collection, including for reducing this burden to: Information Collection Officer, Federal Railroad Administration, 1200 New Jersey Ave. SE, MS-25 Washington, DC 20590.					