



US DOT National
University Transportation Center for Safety

Carnegie Mellon University



Preventing Rear and Side Crashes of Heavy-Duty Tractor Trailer Combinations with Smart Sensors and Vision Systems

Pingbo Tang (PI) <https://orcid.org/0000-0002-4910-1326>

Yorie Nakahira (Co-PI) <https://orcid.org/0000-0003-3324-4602>

Final Report – August 30, 2025

DISCLAIMER

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated in the interest of information exchange. The report is funded, partially or entirely, under [grant number 69A3552344811 / 69A3552348316] from the U.S. Department of Transportation's University Transportation Centers Program. The U.S. Government assumes no liability for the contents or use thereof.

Technical Report Documentation Page

1. Report No. #510	2. Government Accession No.	3. Recipient's Catalog No.
4. Title and Subtitle Preventing Rear and Side Crashes of Heavy-Duty Tractor-Trailer Combinations with Smart Sensors and Vision Systems		5. Report Date August 2025
		6. Performing Organization Code Enter any/all unique numbers assigned to the performing organization, if applicable.
7. Author(s) Enter name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. Form of entry is first name, middle initial (if applicable), last name, and any additional qualifiers. Primary author is listed first. After each author name, enter ORCID (https://orcid.org/) URL, when available. Example: Josiah Carberry, Ph.D. https://orcid.org/0000-0002-1825-0097 Ava Jahanbiglari, Ph.D. student https://orcid.org/0000-0002-9713-0543 Xiyu Deng, Ph.D. student, https://orcid.org/0000-0003-1942-9707 Yorie Nakahira, Assistant Professor https://orcid.org/0000-0003-3324-4602 Pingbo Tang, Associate Professor https://orcid.org/0000-0002-4910-1326		8. Performing Organization Report No. Enter any/all unique alphanumeric report numbers assigned by the performing organization, if applicable.
9. Performing Organization Name and Address Carnegie Mellon University 5000 Forbes Ave, Pittsburgh, PA 15213		10. Work Unit No.
		11. Contract or Grant No. Federal Grant # 69A3552344811 / 69A3552348316

12. Sponsoring Agency Name and Address U.S. Department of Transportation (SPR), 1200 New Jersey Ave SE, Washington, DC 20590	<table border="1"> <tr> <td data-bbox="1045 191 1524 411"> 13. Type of Report and Period Covered Final Report (July 1, 2024 – August 30, 2025) </td> </tr> <tr> <td data-bbox="1045 411 1524 548"> 14. Sponsoring Agency Code USDOT </td> </tr> </table>	13. Type of Report and Period Covered Final Report (July 1, 2024 – August 30, 2025)	14. Sponsoring Agency Code USDOT
13. Type of Report and Period Covered Final Report (July 1, 2024 – August 30, 2025)			
14. Sponsoring Agency Code USDOT			
15. Supplementary Notes Conducted in cooperation with the U.S. Department of Transportation, Federal Highway Administration. Enter information not included elsewhere, such as translation of (or by), report supersedes, old edition number, alternate title (e.g. project name), hypertext links to documents or related information in the form of URLs, PURLs (preferred over URLs - https://archive.org/services/purl/help), DOIs (https://www.doi.org/), insertion of QR codes, copyright or disclaimer statements, etc. Edit boilerplate FHWA statement above if needed.			
16. Abstract Enter a brief factual summary of the most significant information, including the purpose, methods, results, and conclusions of the work. When appropriate, the abstract should include advice on how the results of the research can be used. For guidance, please see ANSI/NISO Z39.14-1997 (R2015) Guidelines for Abstracts (https://www.niso.org/publications/ansiniso-z3914-1997-r2015-guidelines-abstracts). This project, conducted under the Safety21 University Transportation Center, advanced the safety of heavy-duty vehicles (HDVs) through physics-informed modeling, statistical analysis, sensor deployment, and socio-technical analysis. The research focused on two objectives: developing a Physics-Informed Reinforcement Learning (PIRL) framework for near-miss detection in tractor-trailers, and designing customized safety sensor adoption strategies for small motor carriers. Methods included simulation of tractor-trailer dynamics to estimate safety probabilities in rear, side, and lane-keeping scenarios; probabilistic analysis of crash and inspection datasets to identify high-risk vehicle age–region groups; and survey-based studies of inspection practices modeled with an Inverse Contextual Bandit (ICB) framework. Industry interviews highlighted workforce shortages, data integration challenges, and trust in automation as barriers to adoption. Results showed PIRL outperforms standard Deep Q-Networks (DQN) in identifying near-miss risks, sensor adoption should be prioritized for specific high-risk fleets, and inspector decision-making adapts but remains inconsistent. Findings support tailored sensor programs, workforce training, and data interoperability standards to promote proactive, cost-effective HDV safety.			

17. Key Words Heavy-duty vehicles, Near-misses, Crashes, Physics-informed reinforcement learning, Sensors		18. Distribution Statement No restrictions. This document is available through the National Technical Information Service, Springfield, VA 22161. Enter any other agency mandated distribution statements. Remove NTIS statement if it does not apply.	
19. Security Classif. (of this report) Enter security classification of this report (e.g. Unclassified). Reports carrying a security classification will require additional marking giving security and downgrading information as specified by the sponsoring agency.	20. Security Classif. (of this page) Enter the security classification of the form (e.g. Unclassified). When at all possible, this form should remain unclassified. If a classification is required, identify the classified items on the page by an appropriate symbol as per instruction from the sponsoring agency.	21. No. of Pages 52	22. Price Refers to the price of the report. Leave blank unless applicable.

Form DOT F 1700.7 (8-72)

Reproduction of completed page authorized

Chapter 1. Introduction

1.1 Background and Motivation

Heavy-duty vehicles (HDVs), particularly tractor-trailers, are indispensable to the U.S. economy, moving over 70% of freight by weight and underpinning nearly every supply chain (Rezapour Mashhadi et al. 2018; U.S. Department of Transportation Bureau of Transportation Statistics, U.S. Department of Commerce Census Bureau 2012). Yet, their sheer size, weight, and mechanical complexity make them disproportionately involved in fatal and severe crashes. Between 2013 and 2022, the fatal crash rate for large trucks rose from 1.43 to 1.76 per 100 million vehicle miles traveled, a 23% increase in less than a decade (National Center for Statistics and Analysis 2024). In 2022 alone, 5,936 people were killed in crashes involving large trucks, of which 71% were occupants of other vehicles (National Center for Statistics and Analysis 2024). The economic consequences are staggering, with each HDV-involved fatal crash costing an estimated \$3.6 million when accounting for medical, legal, and productivity losses.

Beyond general crash risks, rear and side crashes involving tractor-trailers are particularly devastating. Between 2019 and 2021, fatal rear crashes with large trucks where passenger vehicles underride the trailer increased from 16.8% to 18.0%. Side crashes are also frequent, often caused by vehicles encroaching into truck lanes or remaining in blind zones near the trailer, which accounted for 36% of critical pre-crash events in 2021. Drivers frequently underestimate the safe distance required to follow or overtake heavy vehicles due to their mass, braking dynamics, and road conditions. Moreover, blind spots along the trailer sides exacerbate collision risks, particularly in poor lighting or adverse weather. These dynamics underscore the importance of dynamic safety zones and visibility-aware sensing strategies, forming the foundation of the near-miss PIRL framework developed in this project.

Although driver behavior is a primary contributor, vehicle-related factors account for approximately 10% of HDV crashes, with defective brakes responsible for nearly 29% of such cases (FMCSA (Federal Motor Carrier Safety Administration) 2006). Tire failures and lighting malfunctions are additional leading causes, collectively forming the most common inspection violations recorded during post-crash analyses (NHTSA (National Highway Traffic Safety Administration) 2008). In 2021, brake defects were present in 15.6% of HDVs involved in crashes, while lighting issues accounted for 15.7% and tire problems 7.3%. These findings highlight persistent weaknesses in vehicle maintenance and inspection practices.

Small motor carriers are disproportionately burdened. Although fleets with 30 or fewer vehicles represent only about 8% of registered trucks, they were associated with 76% of HDV-involved fatalities in 2021. Smaller carriers often operate with limited budgets, restricted access to advanced technologies, and fewer dedicated safety personnel compared to large fleets. Regulatory measures,

such as the Electronic Logging Device (ELD) mandate of 2017, improved hours-of-service compliance but failed to reduce crash rates among small carriers; in some cases, accidents even increased, possibly due to unsafe compensatory behaviors like speeding (Scott et al. 2019). These outcomes suggest systemic weaknesses in safety management among small carriers, where financial constraints, operational pressures, and limited workforce expertise amplify risks.

Beyond these organizational disparities, modern HDVs themselves are becoming more complex. The increasing integration of new mechanical elements, electronic subsystems, and multi-modal technologies makes inspection and maintenance more challenging (Barat and Das 2024). Inspectors must contend with vast vehicle state spaces, uncertain fault modes, and variability across manufacturers and configurations. Human inspectors rely on tacit knowledge developed through experience, but this knowledge is unevenly distributed and difficult to transfer, leading to inconsistent inspection outcomes (Johnson et al. 2019). At the same time, AI-driven systems for predictive maintenance and crash prevention are limited by sparse historical datasets, interoperability issues, and workforce adoption barriers (Cantor et al. 2014; Goettee et al. 2010).

This convergence of mechanical failures, organizational vulnerabilities, and socio-technical challenges underscores the urgent need for innovative approaches that combine physics-based modeling, data-driven analysis, and human-centered strategies to improve HDV safety.

1.2 Problem Statement

Traditional road safety strategies rely heavily on reactive measures such as post-crash analysis (Butt and Shafique 2025). While informative, these methods are constrained by under-reporting, exposure bias, and the rarity of crash events (Skaug et al. 2025; Yang et al. 2021). Consequently, they provide limited predictive value for identifying risks before they manifest as accidents.

Near-miss events, close calls that do not result in crashes but carry high collision potential, represent a promising proactive safety indicator (Xu et al. 2024a). These events occur more frequently than crashes, encode rich behavioral and environmental information, and often precede severe incidents. Yet, quantifying near-misses remains a challenge, particularly for HDVs where large blind spots, articulation, and complex dynamics complicate visibility and sensor coverage (Wang et al. 2022a). Current approaches to near-miss analysis often rely on static distance thresholds or handcrafted reward functions in reinforcement learning models, both of which struggle to generalize across diverse environments (Ibrahim et al. 2024).

At the same time, inspection practices remain inconsistent. Studies show that vehicles without valid inspections are up to three times more likely to be involved in serious crashes (Blows et al. 2003). Even when inspections are conducted, variability in inspector expertise and decision-making leads to uneven enforcement (Berthet 2022). Small carriers, in particular, face challenges

in maintaining compliance due to resource limitations, workforce shortages, and weak safety cultures (Cantor et al. 2016).

Finally, while AI and sensor technologies show promise in addressing these challenges, their adoption is constrained by cost, interoperability, durability, and human trust (Bergoffen et al. 2012; Summerskill et al. 2016). A “one-size-fits-all” strategy for sensor deployment fails to account for regional, operational, and fleet-level differences, especially in small carrier contexts where every dollar of investment must be optimized.

Together, these gaps highlight the need for approaches that:

1. Predict risks proactively, using near-miss analysis grounded in physics and data.
2. Customize technology deployment to the unique contexts of small carriers.
3. Address socio-technical barriers in inspection practices, workforce adaptation, and trust in AI systems.

1.3 Research Objectives

This project, conducted under the Safety21 University Transportation Center, set out to address these intertwined challenges through a combination of physics-informed AI methods and socio-technical strategies. The original proposal defined two primary thrusts:

1. Near-miss detection using physics-informed reinforcement learning (PIRL): Develop a dynamic, context-sensitive framework to estimate near-miss probabilities for tractor-trailers under varying sensor configurations and roadway conditions. This thrust advances proactive crash prevention by moving beyond static threshold-based definitions.
2. Customized safety sensor deployment for small carriers: Use crash and inspection datasets to design tailored strategies for deploying brake, tire, and lighting sensors according to vehicle age, operating region, and fleet characteristics, enabling small carriers to maximize safety returns under budget constraints.

Building on these technical objectives, the project expanded in scope to include complementary socio-technical contributions:

1. A systematic review of over 80 studies: Synthesizing challenges and opportunities in HDV safety across domains such as workforce training, human-technology interaction, trust, and data integration.
2. Human-centered studies of inspection practices, including surveys and decision-making experiments, to capture how inspectors prioritize features, refine strategies, and adapt under uncertainty. These studies demonstrate the value of human heuristics in overcoming data limitations and inform collaborative human-AI frameworks for inspection reliability.

1.4 Safety21 Context and Contributions

The Safety21 UTC is dedicated to promoting safety in transportation through cutting-edge technology and human-centered design. This project contributes to that mission by integrating physics-based safety modeling, data-driven sensor strategies, and socio-technical insights into a coherent framework for improving HDV safety. Specifically, it delivers:

- A PIRL-based framework for quantifying near-miss probabilities in tractor-trailers.
- Region- and age-specific recommendations for safety sensor adoption in small fleets.
- A socio-technical analysis of barriers to technology adoption, highlighting the workforce and organizational contexts that shape outcomes.
- Insights into human learning and decision-making in inspections, supporting training and AI collaboration.

The remainder of this report is organized as follows: Chapter 2 reviews the literature and state of practice. Chapter 3 details the research approach and methodologies, including PIRL-based near-miss detection and sensor customization. Chapter 4 integrates industry perspectives. Chapter 5 presents the findings of the two main research thrusts. Chapter 6 highlights an additional study on human learning in inspections. Chapters 7 and 8 summarize the project outputs, outcomes, and recommendations.

Chapter 2. Literature Review

2.1 Socio-Technical Challenges in HDV Safety

Heavy-duty vehicle (HDV) safety is shaped not only by technical reliability but also by the broader socio-technical context of inspections, operations, and workforce practices. A systematic review of recent literature highlights several recurring challenges that constrain the effective adoption of advanced safety measures. The motor carrier industry, a critical component of the global supply chain, has seen significant advancements in safety technologies. However, small motor carriers (defined as those with ≤ 30 vehicles) face unique challenges in adopting these technologies, impacting their safety performance and overall compliance with regulations.

Larger carriers are able to invest substantially in safety technologies and personnel (Miller 2020). In contrast, small carriers often struggle with limited budgets and safety cultures, hindering their ability to adopt advanced safety technologies (Goettee et al. 2010). This gap in technology adoption is not merely a matter of financial constraints but also reflects a lack of tailored safety management strategies suitable for small-scale operations (Bergoffen et al. 2012). In addition, smaller firms often lag to adopt new technologies, primarily due to cost considerations (Cantor et al. 2006). Furthermore, smaller carriers often have higher crash rates (Cantor et al. 2014), which can be attributed to less strict safety practices and lower technology adoption rates, constraints in human and physical capital resources, lack of investment in scientific knowledge, and fewer vehicle maintenance schedules (Cantor et al. 2016).

The adoption of advanced technologies in heavy-duty equipment management is transforming operations, but the integration of human-centric design and usability remains a critical challenge. Effective interaction between humans and technology requires systems that align with workflows, reduce cognitive load, and facilitate decision-making processes. Many studies explore the intersection of usability and stakeholder collaboration, revealing the opportunities and challenges in this domain. recurring issue in heavy equipment operations is the misalignment between user expectations and system design. For example, augmented reality (AR) applications in excavation management have demonstrated potential for enhancing both safety and productivity; however, usability concerns such as interface complexity and cognitive overload hinder widespread adoption (Abdeen et al. 2024). Similarly, systems designed for task automation, such as real-time monitoring applications, often neglect to consider end-user feedback, leading to resistance from operators and managers (Kim et al. 2024). In the construction domain, this disconnect is particularly pronounced due to the diversity of stakeholders, ranging from equipment operators to project managers, each with unique needs and technical proficiencies (Liu et al. 2023).

Data integration and interoperability also pose significant barriers. Safety-relevant data is fragmented across inspection records, crash reports, telematics feeds, and maintenance logs, each collected in different formats and for different purposes. The lack of harmonized data systems

reduces the effectiveness of predictive models and prevents holistic views of fleet risk (Cantor et al. 2014).

Workforce and organizational challenges compound these problems. Inspections depend heavily on tacit knowledge gained through experience, yet this knowledge is unevenly distributed and difficult to transfer to new inspectors (Johnson et al. 2019). Inconsistent inspection practices and workforce shortages lead to variability in outcomes, leaving critical risks undetected.

Finally, predictive maintenance limitations emerge from both technological and human constraints. Sensors are increasingly available for brakes, tires, and lighting, but adoption is uneven, especially among small carriers with limited budgets. Even when deployed, systems can fail without proper maintenance or skilled staff to interpret data. These socio-technical issues underscore why safety solutions must balance technical sophistication with usability, interpretability, and workforce readiness.

At the same time, socio-technical literature warns that even when sensors are deployed, their effectiveness is mediated by workforce practices and organizational capacity. Without adequate training and integration into daily routines, sensors risk being ignored or underutilized. This dual technical and social challenge provides the backdrop for the project's emphasis on tailored, context-sensitive sensor deployment.

2.2 Near-Miss Modeling and Proactive Safety

Visibility Constraints in Tractor-Trailers

Visibility challenges in tractor-trailer systems are a major factor affecting road safety and operational efficiency. These challenges arise from the vehicles' large size, complex articulation, and the limitations of both human and sensor-based perception, especially in dynamic or constrained environments. The multi-unit structure of tractor-trailers creates significant blind spots and makes it difficult for drivers to maintain awareness of the entire vehicle, especially during lane-keeping and turning maneuvers. The size and risk of blind zones increase during right turns, influenced by turning speed, radius, and vehicle configuration. Larger blind zones raise the risk of collisions, especially with vulnerable road users (Wang et al. 2022a). Even with mirrors, some blind spots remain. Enhancing driver awareness through ergonomic design and training is essential for maximizing the effectiveness of visibility systems (Barat and Das 2024). The geometry and articulation points introduce unique constraints that differ from single-unit vehicles, complicating both manual and automated control (Han et al. 2024; Zhao et al. 2023). The large physical size and changing relative positions between tractor and trailer lead to low overlapping fields of view for cameras and sensors, causing perception gaps and misalignment in visual data. This is further complicated by asynchronous vibrations and pose changes between the units (Liang et al. 2025). These limitations increase the likelihood of crashes and near-miss events, especially involving

adjacent vehicles, cyclists, or pedestrians; however, these incidents often go undocumented, leaving a gap in the literature.

Visibility Enhancement Technologies for Tractor-Trailers

Passive solutions, like reflectorized lighting and markings, have shown effectiveness in reducing crash rates; (Burger et al. 1985) observed collision reductions of 16.3% during daylight and 21.2% at night. Active technologies such as Camera/Video Imaging Systems (C/VIS) improve merge performance and nighttime awareness (Camden et al. 2011; Fitch et al. 2010), although challenges like glare and side coverage persist (Summerskill et al. 2016). Blind spot monitoring has also influenced behavior, with (Schaudt et al. 2010) reporting reductions in risky driving and seatbelt violations. Prototype-level integration of deep learning systems (Anwar et al. 2022) shows promise but lacks large-scale fleet deployment data. Despite advancements in these active systems, persistent gaps in sensor coverage and limitations in real-world validation pose significant measurement challenges, particularly in accurately identifying and quantifying visibility-induced near-miss events.

Visibility-Induced Near-Misses

Measuring these near-misses is challenging due to the complexity of visibility fields, dynamic environments, and the limitations of current measurement methods. Visibility failures such as, delayed detection, occlusions, or unexpected pedestrian interactions, are rarely tracked as standalone metrics. (Burger et al. 1980) proposed volumetric visibility maps to guide safety spec development, but direct correlation with near-miss occurrences is lacking. Accurately mapping complex visibility fields is difficult, particularly as the vehicle moves and the environment changes. In addition, operators often need to monitor both forward and rear fields of view, which can be obstructed by the trailer or cargo, making real-time assessment of near-miss situations challenging. Advanced techniques like terrestrial laser scanning can provide detailed, accurate measurements of visibility fields, but require specialized equipment and expertise, and may not capture all real-time operational scenarios (Zvěřina et al. 2022). (Fitch et al. 2010) emphasized that driver behavior, cab geometry, and contextual factors complicate quantifiable evaluations, making surrogate measures (like eye tracking or spatial judgment tests) essential in visibility assessments. Measuring near-misses requires not just static compliance but also dynamic monitoring during actual operation, which is rarely addressed in current standards or measurement protocols (Zvěřina et al. 2022). Therefore, addressing visibility-induced near-misses necessitates novel, integrated measurement approaches that combine advanced sensing technologies, dynamic operational assessments, and human factors analyses beyond current standardized protocols.

Physics-Informed Safety Modeling

To overcome data scarcity around visibility-related near-misses, Physics-Informed Reinforcement Learning (PIRL) offers a viable modeling approach. Safety probabilities can be characterized as solutions to partial differential equations (PDEs) (Chern et al. 2021), which possess well-structured properties such as low-dimensional representations (Wang et al. 2024) and

decomposability (Yasunaga et al. 2024). These PDEs govern how risk propagates across states and time horizons, enabling the extraction of risk information from samples from safe trajectories and near-miss events. To exploit this merit, (Hoshino and Nakahira 2024) introduced a framework that estimates maximal safety probability by enforcing the PDEs information in the objective function in reinforcement learning, allowing agents to learn safe control policies from sparse binary rewards. These physics-based loss functions act as analogs to reward shaping, making safety estimation feasible without dense crash or near-miss data. Their Deep Q-Network (DQN)-based PIRL structure demonstrated accurate safety boundary learning during lane-keeping tasks. (Hoshino et al. 2024) extended PIRL into extreme scenarios such as (Hoshino et al. 2024) autonomous drifting on racing circuits. Using high-fidelity simulations, such as CARLA (Dosovitskiy et al. 2017), agents learned to navigate safely under high-speed, low-traction conditions using only sparse binary rewards. Unlike traditional model-based drifting (Hindiyeh and Christian Gerdes 2014) or reinforcement learning agents requiring complex reward shaping (Cai et al. 2020; Cutler and How 2016), PIRL achieved safe maneuvering without reference trajectories. This result emphasizes PIRL’s ability to learn under high uncertainty, making it an ideal candidate for simulating visibility-induced near-misses where ground truth is hard to define.

Comparative Safety Assessment Frameworks

Across multiple studies, physics-informed models consistently outperform data-driven baselines, especially when historical data is limited or noisy. (Geng et al. 2023) reported up to 70% reduction in trajectory prediction error using physics-informed machine learning in driver modeling, while (Huang and Agarwal 2022) achieved 7.9% improvement in traffic state estimation. (Jurj et al. 2021) validated physics-guided reinforcement learning in adaptive cruise control, observing improved time headway and reduced conflicts. These hybrid approaches integrate domain knowledge with neural learning, offering more robust and interpretable safety estimates than purely empirical models. The convergence of visibility technology and physics-informed modeling holds substantial potential. By simulating how visibility systems interact with environmental uncertainty, researchers can evaluate near-miss likelihood without waiting for incidents to occur. PIRL methods enable this evaluation by learning control boundaries through sparse event feedback, while visibility systems supply real-time sensory data. Future research should focus on combining volumetric blind spot data with PIRL-based safety estimators. This integration will turn passive visibility tools into active decision-making aids that can predict and prevent near-misses.

2.3 Human Learning and Inspection Practices

Studies reveal that vehicles without valid inspections face significantly higher crash risks. In New Zealand, such vehicles are three times more likely to cause serious incidents (Blows et al. 2003). Similarly, Australian research links higher inspection failure rates to increased mechanical-failure-related crashes (Assemi and Hickman 2018). These patterns highlight the importance of regular and reliable inspections in reducing road risks (Assemi et al. 2021).

Despite the critical role of inspections, ensuring their consistency and effectiveness remains a challenge, as inspection outcomes often depend on human decision-making strategies shaped by heuristics, experience, and contextual factors. Understanding these strategies is crucial for improving inspection processes, especially in high-uncertainty environments where traditional analytical methods may fall short. Research on smart heuristics in civil engineering demonstrates how these simple yet effective decision-making strategies can help individuals navigate complex tasks by focusing on the most relevant information and adapting to the context (Love 2025).

Human decision-making strategies, particularly in safety-critical contexts often depend on fast-and-frugal heuristics, simple, adaptive rules that prioritize efficiency over exhaustive information processing (Wang et al. 2022b). These heuristics enable inspectors to focus on important features, and disregard less relevant details, allowing for timely and effective decisions. As (Gigerenzer et al. 2022). emphasize, such smart heuristics are grounded in ecological rationality, aligning decision-making strategies with the structure of the task environment to optimize outcomes under constraints like time, uncertainty, or incomplete data. These insights are particularly valuable for vehicle inspections, where inspectors must navigate sparse or inconsistent datasets and identify the most impactful features to ensure safety and reliability.

Human heuristics provide critical insights into decision-making processes, yet identifying and quantifying the features prioritized by inspectors remains a challenge. Inverse Contextual Bandits (ICB) has been applied in domains such as healthcare and application areas like recommender systems to model non-stationary behaviors and evolving decision-making patterns (Hüyük et al. 2022; Xu et al. 2024b). While reinforcement learning (RL) and cognitive modeling approaches have been explored for decision-making analysis, these methods often require extensive labeled training data and predefined reward structures, making them less adaptable to dynamic, heuristic-based tasks like vehicle inspections. The BICB framework is particularly suited for capturing evolving human heuristics in environments with limited supervision, aligning well with the complexities of HDV inspections. By focusing on feature prioritization in dynamic environments, ICB provides a structured framework that allows researchers to integrate human heuristics with AI tools through systematic analysis, thereby refining inspection strategies and enhancing reliability.

2.4 Synthesis and Research Needs

The literature points to three interconnected needs that frame this project:

1. Proactive safety modeling: Traditional crash-based approaches lack predictive power. Near-miss detection, supported by PIRL and grounded in physical dynamics, provides a promising alternative for HDVs.
2. Customized sensor adoption: Small carriers require risk-based, region- and age-specific strategies for adopting brake, tire, and lighting sensors, balancing cost constraints with safety imperatives.

3. Socio-technical integration: Human learning, inspection practices, and organizational readiness must be considered alongside technical innovation to ensure effective adoption and consistent safety outcomes.

Together, these needs establish the foundation for this project. By integrating PIRL-based near-miss detection, customized sensor deployment, and socio-technical insights, the project advances both the technical and organizational dimensions of HDV safety.

Chapter 3. Research Objectives and Approach

3.1 Research Objectives

This project aims to address persistent safety challenges in heavy-duty vehicles (HDVs) by integrating physics-informed artificial intelligence with data-driven safety sensor strategies. Based on the proposal and gaps identified in the literature, two primary objectives were defined:

1. Develop a proactive framework for near-miss detection in tractor-trailers. This objective aimed to move beyond static crash-based safety analysis by applying Physics-Informed Reinforcement Learning (PIRL) to estimate safety probabilities in rear and side crash scenarios, capturing how speed, clearance, and trailer geometry interact to create risk.
2. Design customized safety sensor deployment strategies for small motor carriers. This objective focused on developing data-driven recommendations for allocating brake, tire, and lighting sensors based on vehicle age and operating region, enabling small carriers to maximize safety returns under budget constraints.

Together, these objectives were designed to advance the Safety21 mission of promoting proactive and equitable safety improvements by combining new AI methods with practical deployment strategies.

3.2 Research Approach

Near-Miss Detection with PIRL:

A simulation framework was developed to model tractor-trailer dynamics and interactions with obstacles in rear-end, side, and lane-keeping scenarios. PIRL was used to estimate safety probabilities under sparse event data, embedding vehicle dynamics into the learning process to generate interpretable reported safety maps. Comparative experiments with standard Deep Q-Networks (DQN) highlighted PIRL's superior performance in capturing non-linear safety boundaries, particularly when trailer-side sensors were included.

Customized Sensor Deployment:

Crash and inspection datasets were analyzed to identify high-risk vehicle populations based on age and operating region. Using probabilistic methods, the study revealed that vehicles aged 6–23 in rural regions face the highest risks across all three sensor categories, while vehicles aged 24–29 in urban regions are especially vulnerable to brake and lighting violations. These results demonstrate that targeted sensor deployment can achieve meaningful safety improvements without requiring comprehensive fleet-wide adoption, which is often infeasible for small carriers.

3.3 Integration and Extensions

While these two branches address the core objectives of the project, technical innovations must be situated within human and organizational contexts. To extend the socio-technical perspective, an additional line of research was pursued on human learning and decision-making in inspections. This work, presented in a separate chapter, analyzes how inspectors adapt feature prioritization strategies and how human heuristics can be integrated into AI frameworks to improve inspection consistency. By linking near-miss modeling, sensor deployment, and inspection practices, the project forms an integrated research agenda that combines technical, data-driven, and socio-technical insights into HDV safety as illustrated in Figure 1.




Figure 1. Integrated framework connecting sensor customization, PIRL-based near-miss detection, and human-centered inspection analysis under proactive and human-centered HDV safety

3.4 Anticipated Impacts

The approach was designed to produce impacts across three dimensions:

- **Technical:** New AI models for proactive near-miss detection.
- **Operational:** Practical, cost-sensitive recommendations for sensor adoption in small fleets.
- **Socio-Technical:** Insights into workforce learning and adoption that support more reliable and trusted safety practices.

The following chapters detail the methodologies, industry perspectives, and findings for each research branch.

Chapter 4. Methodology

This project employs a dual methodological approach to advance heavy-duty vehicle safety through both physics-informed modeling and statistical analysis. First, we develop a Physics-Informed Reinforcement Learning (PIRL) framework to estimate safety probabilities in tractor-trailer systems under rear-end, side, and lane-keeping scenarios, embedding vehicle dynamics and physical constraints into the learning process to quantify near-miss risks. Second, we analyze crash and inspection datasets to identify region and age-specific patterns of brake, tire, and lighting violations in crashes, enabling the design of customized safety sensor deployment strategies for small motor carriers. Together, these complementary approaches link proactive risk modeling with practical sensor allocation, forming an integrated framework for enhancing HDV safety.

Near-miss detection and safe zone with PIRL

Problem formulation

We consider a stochastic dynamical system with w -dimensional Brownian motion $\{W_t\}_{t \in \mathbb{R}_+}$ starting from $W_0 = 0$. The stochastic differential equation (SDE) governing the system is given by

$$dX_t = f(X_t, U_t) dt + \sigma(X_t, U_t) dW_t.$$

Here, $X_t \in \mathbb{X} \subset \mathbb{R}^n$ is the system state, and $U_t \in \mathbb{U} \subset \mathbb{R}^m$ is the control input. We assume that the functions f and σ satisfy the necessary regularity conditions to ensure that the SDE admits a unique strong solution. The magnitude of $\sigma(X_t, U_t)$ captures uncertainties arising from disturbances, unmodeled dynamics, and prediction errors of environmental variables.

To approximate numerical solutions of the SDE and address optimal control problems, we consider the state X_t at discrete time steps $t \in \mathbb{T}$ with step size Δt , where $\mathbb{T} = \{0, 1, \dots, \tau\}$ and τ is the time horizon. The discretized system is given by

$$X_{t+1} = F^\pi(X_t, \Delta W_t),$$

where $\Delta W_t := W_{(t+1)\Delta t} - W_{t\Delta t}$, and F^π is the state transition function derived from the control policy $\pi: [0, \infty) \times \mathbb{X} \rightarrow \mathbb{U}$. From an optimal control perspective, using a Markov policy is not restrictive when the value function has sufficient smoothness under appropriate technical conditions.

Safety of the system is defined using a safe set $\mathbb{C} \subset \mathbb{X}$. For the discretized system and a given control policy π , the safety probability Φ^π of the initial state $X_0 = x$ over the time horizon $\mathbb{T}$ is characterized as the probability that the state X_t remains within the safe set $\mathbb{C}$:

$$\Phi^\pi(\tau, x) := \mathbb{P}[X_t \in \mathbb{C}, \forall t \in \mathbb{T} \mid X_0 = x, \pi].$$

Consider the system starting from an initial state $x \in \mathcal{C}$. The maximal safety probability is defined as

$$\Phi^*(\tau, x) := \sup_{\pi \in \Pi} \Phi^\pi(\tau, x),$$

where Π is the class of bounded and Borel measurable Markov control policies, and the corresponding optimal safe control policy is $\pi^* := \operatorname{argsup}_{\pi \in \Pi} \Phi^\pi$

Here the authors build upon the Physics-Informed Reinforcement Learning (PIRL) framework for safety probability estimation from (Hoshino and Nakahira 2024) and (Hoshino et al. 2024). They first summarize the essential formulations here and then describe their specific application for HDV.

With the backbone of standard Deep Q-Network (DQN), the authors employed a Physics-Informed Neural Network (PINN) as the function approximator that penalizes the discrepancy from a PDE condition that the safety probability should satisfy. Here, they considered the augmented state space $\mathbb{S} := \mathbb{R} \times \mathbb{X} \subset \mathbb{R}^{n+1}$ and the augmented state $S_k \in \mathbb{S}$ where they denote the first element of S_k by H_k and the other elements by X_k , i.e.,

$$S_k = [H_k, X_k^\top]^\top,$$

where H_k represents the remaining time before the outlook horizon τ is reached and X_k denotes the vehicle (and environment) states. Then, consider the stochastic dynamics starting from the initial state $s := [\tau, x^\top]^\top \in \mathbb{S}$ for all $k \in \mathbb{Z}_+$,

$$S_{k+1} = \begin{cases} \tilde{F}^\pi(S_k, \Delta W_k), & S_k \notin \mathcal{S}_{\text{abs}}, \\ S_k, & S_k \in \mathcal{S}_{\text{abs}}, \end{cases}$$

with the function $\tilde{F}^\pi$ given by

$$\tilde{F}^\pi(S_k, \Delta W_k) := \begin{bmatrix} H_k - 1 \\ F^\pi(X_k, \Delta W_k) \end{bmatrix},$$

and the set of absorbing states $\mathcal{S}_{\text{abs}}$ given by

$$\mathcal{S}_{\text{abs}} := \{[\tau, x^\top]^\top \in \mathbb{S} \mid \tau < 0 \vee x \notin \mathcal{C}\}.$$

The absorbing states represent scenarios where the HDV enters unsafe conditions, such as departing from the roadway, experiencing collisions, or exceeding safe operational limits.

The major advantages of PIRL are its sparse reward learning capability and physics-informed generalization, eliminating the need for complex hand-crafted reward design. Consider the system starting from an initial state s and the reward function $r: \mathbb{S} \rightarrow \mathbb{R}$ given by

$$r(S_k) = \begin{cases} 1, & H_k = 0 \wedge S_k \notin \mathcal{S}_{\text{abs}} \\ 0, & \text{otherwise} \end{cases}$$

Then, the value function v^π for a given control policy π is defined as:

$$v^\pi(s) := \mathbb{E} \left[\sum_{k=0}^{N_f} r(S_k) \mid S_0 = s, \pi \right],$$

where $N_f := \inf\{j \in \mathbb{Z}_+ \mid S_j \in S_{\text{abs}}\}$, which takes a value in $[0,1]$ and is equivalent to the safety probability $\Phi^\pi(\tau, x)$:

$$v^\pi(s) = \Phi^\pi(\tau, x).$$

We formulate an episodic RL problem where the action-value function under a policy π is:

$$q^\pi(s, a) = \mathbb{E} \left[\sum_{k=0}^{N_f} r(S_k) \mid S_0 = s, U_0 = a, \pi \right],$$

with rewards defined to reflect whether the HDV remained within safety margins.

As an extension of DQN with the PINN, the loss function is given by three components:

$$L = L_D + \lambda L_P + \mu L_B,$$

where L_D is the standard DQN data loss, L_P is the physics loss that enforces the PDE constraint, L_B is the boundary loss that enforces the boundary conditions, and λ and μ are weighting coefficients for the physics-informed regularization terms.

The data loss follows standard DQN methodology. Each episode initializes with an augmented state $s_0 = [h_0, x_0^\top]^\top$ sampled from distribution P_D :

$$P_D(s_0) = \begin{cases} \frac{1}{|\Omega_D|}, & h_0 = \tau_D \wedge x_0 \in \Omega_D \\ 0, & \text{otherwise} \end{cases}$$

where $\tau_D \in \mathbb{R}_+$ represents the data acquisition time interval, $\Omega_D \subset \mathbb{X}$ defines allowable initial vehicle states, and $|\Omega_D|$ denotes the volume of this set. Experience transitions (s_k, a_k, r_k, s'_k) are stored in replay memory D . Target values for a minibatch S_D are calculated using target Q-function $\hat{Q}$:

$$y_j = \begin{cases} r_j, & \text{if } s_{j'} \in S_{\text{abs}} \\ r_j + \max_a \hat{Q}(s_{j'}, a; \hat{\Theta}), & \text{otherwise} \end{cases}$$

where S_{abs} represents absorbing states corresponding to episode termination or unsafe conditions. The data loss is computed as:

$$L_D(\Theta) = \frac{1}{|S_D|} \sum_j (y_j - Q(s_j, a_j; \Theta))^2$$

The physics loss enforces PDE constraints derived from Hamilton-Jacobi-Bellman theory. States $s_l = [h_l, x_l^\top]^\top$ are sampled from distribution P_P :

$$P_P(s_l) = \begin{cases} \frac{1}{\tau|\Omega_P|}, & h_l \in [0, \tau] \wedge x_l \in \Omega_P \\ 0, & \text{otherwise} \end{cases}$$

where $\Omega_P \subset \mathcal{C}$ specifies the PDE enforcement domain within the safe set. For each sample, the greedy action is determined as:

$$a_l^* = \operatorname{argmax}_a Q(s_l, a; \Theta)$$

The physics loss quantifies PDE residual violations:

$$L_P(\Theta) = \frac{1}{|S_P|} \sum_l (W_P(s_l, a_l^*; \Theta))^2$$

where the residual function is:

$$W_P(s_l, a_l^*; \Theta) = \partial_s Q(s_l, a_l^*; \Theta) \tilde{f}(s_l, a_l^*) + \frac{1}{2} \operatorname{tr}[\tilde{\sigma}(s_l, a_l^*) \tilde{\sigma}(s_l, a_l^*)^\top \partial_s^2 Q(s_l, a_l^*; \Theta)]$$

Here, $\tilde{f}$ and $\tilde{\sigma}$ represent augmented system dynamics incorporating the time dimension, and $\partial_s Q$ and $\partial_s^2 Q$ denote the gradient and Hessian of the Q-function with respect to the augmented state.

The boundary loss ensures proper Q-function behavior at domain boundaries. Boundary states $s_m = [h_m, x_m^\top]^\top$ are sampled from:

$$P_B(s_m) = \begin{cases} \frac{1}{2|\Omega_P|}, & h_m = 0 \wedge x_m \in \Omega_P \\ \frac{1}{2\tau|\Omega_B|}, & h_m \in [0, \tau] \wedge x_m \in \Omega_B \\ 0, & \text{otherwise} \end{cases}$$

where $\Omega_B \subset \partial\mathcal{C}$ represents the spatial boundary of the safe set. The boundary loss is formulated as:

$$L_B(\Theta) = \frac{1}{|S_B|} \sum_m (W_B(s_m, a_m^*; \Theta))^2$$

with residual function:

$$W_B(s_m, a_m^*; \Theta) = Q(s_m, a_m^*; \Theta) - l_\epsilon(x_m)$$

The target boundary value $l_\epsilon(x)$ provides a smoothed approximation:

$$l_\epsilon(x) = \max\left(1 - \frac{\operatorname{dist}(x, \mathcal{C}_\epsilon)}{\epsilon}, 0\right)$$

where $C_\epsilon \subset C$ represents a contracted safe set with margin $\epsilon > 0$, and $\text{dist}(x, C_\epsilon)$ denotes the Euclidean distance from state x to the interior of C_ϵ . This formulation ensures continuous boundary conditions while maintaining numerical stability during training. Detailed descriptions about the loss are shown in (Hoshino and Nakahira 2024).

Experiment

We consider a tractor system modeled using a bicycle model approach, where the tractor and trailer are represented as interconnected rigid bodies with single-track dynamics shown in Figure 2.




Figure 2. Tractor and trailer nonlinear dynamic system.

The complete tractor-trailer state is characterized by five key variables that capture the positioning, orientation, and motion of the system:

$$x_v = [x_{tl}, y_{tl}, \theta, \beta, v]^T \in \mathbb{R}^5,$$

where x_{tl} and y_{tl} denote the global coordinates of the center of the trailer's rear axle, θ is the orientation angle of the trailer in the global frame (with $\theta = 0$ corresponding to the east direction), β represents the articulation angle between the tractor and trailer ($\beta = 0$ indicating perfect alignment), and v is the longitudinal velocity of the vehicle. The geometric configuration of the

vehicle is characterized by several parameters: the hitch length M (the distance from the tractor's rear axle to the trailer's front axle), the length and width of the tractor (L_1 and W_1), the length and width of the trailer (L_2 and W_2), and the diameter and width of the wheels (L_{wheel} and W_{wheel}). The dynamics is described as:

$$\begin{aligned}\dot{x} &= v \cos \beta \left(1 + \frac{M_1}{L_1} \tan \beta \tan \delta \right) \cos \theta \\ \dot{y} &= v \cos \beta \left(1 + \frac{M_1}{L_1} \tan \beta \tan \delta \right) \sin \theta \\ \dot{\theta} &= v \left(\frac{\sin \beta}{L_2} - \frac{M_1}{L_1 L_2} \cos \beta \tan \delta \right) \\ \dot{\beta} &= v \left(\frac{\tan \alpha}{L_1} - \frac{\sin \beta}{L_2} + \frac{M_1}{L_1 L_2} \cos \beta \tan \delta \right) \\ \dot{v} &= \alpha\end{aligned}$$

For other subjects in the environment such as other vehicles, trees, traffic lights, and lane boundaries, we model them as geometric entities with predefined safety zones. Here, we consider n surrounding objects, and for each specific object $X_o^{(j)}$, we define the states as:

$$x_o^{(j)} = [x_o^{(j)}, y_o^{(j)}, \omega^{(j)}, v_o^{(j)}, d_{\text{th}}^{(j)}, d_{\text{tk}}^{(j)}, d_{\text{tl}}^{(j)}]^\top \in \mathbb{R}^7,$$

where $x_o^{(j)}$ and $y_o^{(j)}$ are the world coordinates of the j -th object's center, $\omega^{(j)}$ is the object's heading angle, $v_o^{(j)}$ is the object's velocity, and $d_{\text{th}}^{(j)}$, $d_{\text{tk}}^{(j)}$, $d_{\text{tl}}^{(j)}$ represent the threshold distance, the distance from the object to the tractor's center, and the distance from the object to the trailer's center, respectively. We consider that the vehicle enters an unsafe region if $d_{\text{tk}}^{(j)} < d_{\text{th}}^{(j)}$ or $d_{\text{tl}}^{(j)} < d_{\text{th}}^{(j)}$.

In the PIRL setting, the states s should consist of the vehicle state x and the outlook horizon τ as:

$$s = [\tau, x^\top]^\top,$$

and x can be further decomposed into the above x_v and x_o except we should get rid of the world coordinates as the learning algorithm focuses on relative positioning and local interactions rather than absolute global positioning. Therefore, the vehicle state x becomes:

$$x = \begin{bmatrix} \theta, \beta, v, \{\omega^{(j)}, v_o^{(j)}, d_{\text{th}}^{(j)}, d_{\text{tk}}^{(j)}, d_{\text{tl}}^{(j)}\}_{j \in [1, 2, \dots, n]} \\ \check{x}_v' & \check{x}_o^{(j)'} \end{bmatrix}^\top,$$

where $x'_v, x_o^{(j)'}$ denote the modified vehicle and object states respectively, with global coordinates removed to focus on relative spatial relationships. As the HDV is the only one we can control, the control action a is given by

$$a = [\delta, \alpha]$$

where δ is the steering angle and α is the longitudinal acceleration.

To investigate the impact of sensor configurations on safety probability estimation, we consider two different observability scenarios representing current industry practice and enhanced sensing capabilities:

- **Agent 1: Standard industry configuration (tractor-only sensors):** This represents typical commercial vehicle setups with sensors primarily on the tractor, providing limited trailer position awareness, which is a common constraint in current industry practice. In this case, we set d_{tl} for each surrounding object to a large number (e.g., infinity) to represent the lack of trailer position sensing. By setting $d_{tl} = \infty$, the safety condition $d_{tl}^{(j)} < d_{th}^{(j)}$ can never be triggered, effectively removing trailer-based collision detection from the safety assessment and forcing the agent to rely solely on tractor-based sensing for safety decisions.
- **Agent 2: Proposed enhanced configuration (tractor + trailer sensors):** This incorporates trailer-mounted sensors providing full observability of all state variables, including trailer-to-object distances, representing potential future HDV sensing capabilities.

In our experiments, we focus on zoomed micro-level scenarios rather than global map-based simulations. This approach is motivated by the structure of PIRL, which is designed to operate on compact. In particular, PIRL emphasizes learning from and responding to risky situations that are present at the onset of each scenario, as opposed to relying on the occurrence of rare or hazardous events during long rollouts on a full map. Therefore, our experimental setup does not include a full driving map; instead, we formulate each trial as a single, critical decision-making episode starting from a high-risk initial condition. This design allows us to directly evaluate how effectively a policy can navigate through specifically crafted risky situations, aligning with the strengths of PIRL and enabling more targeted safety assessment.

To evaluate the effectiveness of the proposed Physics-Informed Reinforcement Learning framework in modeling near-miss safety probabilities for HDVs, the authors conducted experiments across three distinct as shown in Figure 3 and compared PIRL with DQN to see how the reported safety probabilities look like. The first scenario involves a static and dynamic obstacle on the side of the tractor-trailer. The second setting involves a dynamic obstacle (e.g., passenger car) at the rear of the tractor-trailer. In the third setting, they considered lane keeping problem. In all three scenarios, the authors estimate the safety probability as a function of the trailer-to-object distance (d_{tl}) and vehicle velocity. These parameters are chosen to reflect the critical spatial and dynamic interactions that influence safety in real-world HDV operations, particularly under limited visibility.




Figure 3. Three distinct settings (left to right): 1) Near-miss involving dynamic and static obstacles on the side of the tractor-trailer 2) Near-miss involving a dynamic obstacle at the rear of the tractor-trailer 3) Lane-keeping

Sensor Customization for Small Carriers

In this section, the authors collected and analyzed vehicle inspection and crash datasets to capture how the inspection outcomes can serve as indicators of safety-critical sensors for vehicles widely used by small motor carriers, and what regions and environments of vehicle operations need customized safety sensor use. The authors collected a national motor carrier performance dataset with crash and vehicle information, and a local vehicle inspection dataset with vehicle mileage and detailed inspection results compared with the national dataset. A summary of the data content used in this research is provided in

The authors then employed an analytical approach, focusing on the relationship between vehicle age, regional factors, and crash incidences in Pennsylvania. The methodology encompassed data preprocessing, probabilistic analysis, and data visualization, using a combination of automated decoding, data grouping, and statistical calculations. **Error! Reference source not found.** illustrates the overall research process.

Table 1.

The authors then employed an analytical approach, focusing on the relationship between vehicle age, regional factors, and crash incidences in Pennsylvania. The methodology encompassed data preprocessing, probabilistic analysis, and data visualization, using a combination of automated decoding, data grouping, and statistical calculations. **Error! Reference source not found.** illustrates the overall research process.

Table 1. Data source summary

Source	Local Vehicle Inspection Dataset	National Crash and Inspection Violation Datasets
Description	Annual vehicle inspection records from Pennsylvania	Annual crash report and roadside inspection violation records of the United States
Coverage Period	2005-2021	2017-2021
Types of Data Collected	VIN, make, model, model year, carrier location and ZIP code, component inspection outcomes, odometer readings	VIN, make, model, model year, inspection violation category, DOT of the motor carriers, carrier location

Decoding the Department of Transportation (DOT) numbers was the initial step to ascertain the power units of each motor carrier. This step was crucial to support the claims regarding the safety issues made by small fleets.

Figure 5 shows the process of comparing the safety performance of small carriers with large and mid-size carriers.

The authors used a region classification scheme provided by the Center for Disease Control's National Center for Health Statistics (NCHS 2017) to classify the carrier regions based on

population density, ranging from large central metros (coded as 1) to non-core areas (coded as 6). This classification was derived from state and county codes in the datasets. Figure 6 shows the distribution of different region types in Pennsylvania.

Then, the authors calculated the crash and inspection violations (for brake, tire, and lighting) probabilities using the equation below. The local vehicle inspection source for Pennsylvania is considered as the total population, and only Pennsylvania-specific records are used from the national datasets, representing the observed events.

$$P(E) = \frac{\text{Number of observed events (crashes or violations within each group)}}{\text{Total number of opportunities (total vehicles in each group)}}$$

In this analysis, the authors applied the Fixed Quantile Threshold method to define high-risk groups based on crash probabilities and inspection violation probabilities, identifying significant deviations from the norm and typical patterns. Thresholds set at the 80th and 90th percentiles distinguished groups with high probabilities - exceeding 5% for crashes, 37% for brakes, 23% for tires, and 19% for lighting violations - highlighting areas in need of urgent intervention.




Figure 4. Research process designation




Figure 5. Separating motor carriers based on their fleet size and comparing their safety performance




Figure 6. Region classification of Pennsylvania based on the population density

Chapter 5. Industry Perspectives

To complement the technical methodologies, this project incorporated direct feedback from industry stakeholders, including fleet managers, inspection experts, and maintenance professionals. Their insights highlight that the barriers to adopting advanced safety and predictive maintenance technologies are not purely technical, but deeply socio-technical.

From the interviews on human factors, several recurring themes emerged. Fleet operators emphasized that operational efficiency, cargo security, compliance, and uptime are the most pressing challenges. Driver retention and workforce shortages were cited as major social issues, compounded by an aging workforce with limited appeal to younger generations. Technically, while IoT and telematics devices can provide real-time insight into vehicle health, fleets often lack the resources and expertise to consolidate and act on this data. Interoperability remains a persistent problem: multiple communication protocols, incompatible data streams, and short hardware lifecycles make it difficult to build unified fleet management systems.

Interviews reinforced these points with concrete examples. Inspection experts noted inconsistencies in manual inspections; technicians under time pressure often miss critical components, leading to unsafe vehicles remaining in service. Automated inspection pads and sensor-based systems have demonstrated significant improvements, yet adoption is slowed by cost concerns and resistance among technicians who fear job loss or increased monitoring. Fleet managers reported that while sensors for brakes, tires, and steering can transform safety outcomes, many companies still rely heavily on mechanic intuition rather than automated alerts. This creates a gap where failures occur between scheduled inspections, leaving vehicles vulnerable to undetected risks.

Privacy and trust also surfaced as recurring concerns. The deployment of in-cab cameras or telematics often meets resistance from drivers who perceive them as surveillance rather than safety tools. Industry leaders explained that acceptance improves when the purpose is clearly communicated and when the technology demonstrates value; for example, exonerating drivers in crash investigations or preventing claims. However, distrust lingers when technologies cause false alerts, such as automated braking triggered by roadside objects, undermining both driver confidence and organizational buy-in.

Interviewees agreed that predictive maintenance and AI-based decision support hold promise, but that cultural and organizational barriers remain. Training, transparency, and incremental integration are critical for building trust. Without these, even well-designed systems risk rejection. Importantly, industry experts stressed that solutions must be co-designed with practitioners, not imposed from academia or manufacturers alone.

These findings align closely with the project's emphasis on socio-technical integration. By embedding industry feedback into the research framework, the project ensures that advances in

PIRL-based near-miss detection and customized sensor deployment are not only technically sound but also realistic, scalable, and acceptable in practice. In doing so, the research bridges the gap between state-of-the-art technology and the workforce realities that ultimately determine its success.

Chapter 6. Results and Findings

Results from Safety probability heatmaps

Figure 7 illustrates a comparative analysis of safety probability heatmaps learned by the standard Deep Q-Network (DQN) and the proposed Physics-Informed Reinforcement Learning (PIRL) models across three critical driving scenarios: rear-end, side obstacle, and lane-keeping. Each figure visualizes the estimated probability of remaining within the safe set as a function of trailer distance (d_{tl}) and vehicle velocity (v), with yellow indicating high safety and blue indicating low safety. These probability maps define the safe region as the subset of states with safety probabilities close to 1, and enable a continuous representation of dynamic near-misses, characterized by steep gradients at the boundary between safe and unsafe regions. Rather than using fixed thresholds, this approach captures how proximity and speed interact to form context-specific safety risks.

In the rear-end scenario, the DQN agent outputs a nearly uniform heatmap, lacking strong gradients and showing minimal responsiveness to changing d_{tl} and v . In contrast, PIRL captures meaningful interactions between distance and speed: safety probability increases with greater distance and decreases with higher velocity. The sharper and physically consistent gradients indicate PIRL’s ability to reflect longitudinal dynamics and temporal risk accumulation.

The side scenario further underscores PIRL’s advantage. The DQN result remains flat and uninformative, while PIRL reveals a broader and more nuanced safety field. Safety improves as lateral clearance increases and is sensitive to velocity, effectively modeling interactions with roadside or adjacent-lane obstacles. This is especially important in limited visibility conditions, addressed by our visibility-aware framework.

In the lane-keeping scenario, DQN again fails to differentiate risk across spatial deviations and speeds. PIRL, however, displays a clear decline in safety at high speeds and large lateral deviations. This aligns with well-known rollover and off-tracking risks in articulated vehicles, and demonstrates PIRL’s ability to encode complex vehicle geometry and stability margins.

Overall, PIRL significantly outperforms DQN in all three scenarios. It learns sharper, more interpretable safety maps by integrating physical knowledge, sparse rewards, and geometric awareness. These findings validate PIRL’s scalability and generalization capability. PIRL not only enables proactive safety estimation in data-sparse environments but also offers a principled alternative to static threshold-based definitions of near-misses.

In addition, to study the impact of trailer-side sensors on safety estimation, we compare the PIRL agent’s safety probability maps in the side-obstacle scenario under two observability settings: with and without access to the trailer-to-object distance d_{tl} .


As shown in Figure 9, the presence of trailer-sensing significantly enhances the agent’s ability to detect unsafe configurations. When d_{tl} is available as input, the PIRL agent produces a safety

heatmap that reflects intuitive risk boundaries: safety decreases rapidly as the trailer approaches the obstacle, especially at high speeds. The safety probability drops below 0.2 for $d_{tl} < 1.5$ m at $v > 20$ m/s, aligning with real-world risk expectations in blind-spot invasion scenarios.


To illustrate the comparative performance of the proposed PIRL framework against a baseline DQN, Figure 8 shows the average episode reward during training for the side-crash near-miss scenario, computed as a running average with a window of 100 episodes. The results indicate that while both agents are capable of achieving high rewards, their stability and learning efficiency differ significantly. PIRL converges more quickly to near-optimal performance and sustains consistently high rewards, reflecting the benefits of embedding vehicle dynamics into the learning process. In contrast, the DQN baseline exhibits substantial variability, with frequent drops in performance and slower recovery, suggesting greater sensitivity to stochasticity and less reliable policy behavior. This comparison highlights PIRL’s advantage in generating stable and interpretable safety-aware policies in side-crash contexts, where blind spots and lateral clearance are critical for heavy-duty vehicle safety.

By contrast, when d_{tl} is not shown to the agent during training, the resulting safety field becomes diffuse and insensitive to lateral proximity. The agent fails to recognize spatial risk patterns, and safety appears to depend primarily on velocity. Without trailer-side sensors, the agent fails to recognize safe conditions even when the trailer is located at a sufficient distance from the obstacle. This lack of situational awareness can result in inappropriate actions, ultimately leading to unsafe outcomes.


This comparison clearly illustrates that PIRL relies on sensor observability to accurately model nonlinear spatial risk, and that trailer-mounted sensors provide essential information for learning interpretable and context-aware safety boundaries. Without sensors, the agent underestimates side risks, compromising predictive power in near-miss conditions as shown in Figure 9.




DQN Side Scenario




PIRL Side Scenario




DQN Rear-end Scenario



PIRL Rear-end Scenario



DQN Lane-Keeping



PIRL Lane-Keeping

Figure 7. Safety Probability Heatmaps across scenarios with sensors installed: Rear-end, Side, and Lane-Keeping for DQN and PIRL Models







Figure 8. Learning curves comparing baseline DQN and Physics-Informed Reinforcement Learning (PIRL). PIRL achieved more stable and consistently high average episode rewards, while DQN exhibited greater variability and performance drops.




With Sensor Side Scenario




Without Sensor Side Scenario




With Sensor Rear-end Scenario



Without Sensor Rear-end Scenario



With Sensor Lane-keeping Scenario



Without Sensor Lane-keeping Scenario

Figure 9. Comparison of PIRL safety probability heatmaps with and without sensor-enabled trailer-side visibility

Results from statistical analysis for sensor customization

Over the five years from 2017 to 2021, the heatmaps in Figure 10 illustrate a varying landscape of crash probabilities across different age groups and regions in Pennsylvania, shedding light on critical trends that inform safety sensor deployment strategies.

The heatmaps depicting crash probabilities from 2017 to 2021 in Figure 10 suggest that Region 6 has consistently been a focal point for crashes across multiple age groups, particularly for those aged 18 to 29. However, in 2020 and 2021, this trend does not hold as strongly, indicating a potential shift in crash probabilities towards younger vehicles. The data supports the need for region-specific and age-targeted safety sensor deployments.

In addition, based on motor-carrier level analysis for crashes in 2021, the authors found that the rate of crashes per power unit for small motor carriers in Region 6 is 11%, while for larger carriers this rate is 1%. The result suggests that although small carriers own much fewer vehicles in comparison to larger carriers, they contribute to more crashes in non-core areas of Pennsylvania.

The analysis of brake violation probabilities in crashes across various regions and age groups from 2017 to 2021, as demonstrated in Figure 11, shows that vehicles in Region 6, aged 6 to 23 exhibit higher probabilities of brake violations. This outcome is closely followed by vehicles within the 24 to 29 age range in Region 1 and the 18 to 23 age group in Region 3, with brake violation probabilities of 0.38 and 0.37, respectively. Similarly, Figure 12, which indicates the probabilities of tire violations in crashes during the same period, illustrates that vehicles aged 12 to 17 in Region 6 have the highest risk with a probability of 0.36, followed by the oldest vehicles, those above 30 years old, in Region 1 with a probability of 0.3. Additionally, vehicles aged 6 to 11 and 18 to 23 in Region 6 are other critical groups in terms of tire sensor adoption. Moreover, lighting violation probabilities in crashes, as outlined in Figure 13, suggest that vehicles aged 6 to 23 in Region 6 with average lighting violation probabilities of 0.24, 0.23, and 0.20, and vehicles aged 24 to 29 in Region 2 with violation probability of 0.21, should be prioritized for lighting sensor implementation.

Overall, the patterns reveal that Region 6, particularly for vehicles aged 6 to 23, stands out as a critical risk area for all three types of inspection violations. This consistent finding emphasizes the need for targeted adoption of brake, tire, and lighting sensors for these specific age groups in non-core areas to enhance vehicular safety and compliance.




Figure 10. Crash probabilities in different age groups and regions from 2017 to 2021




Figure 11. Probabilities of brake violations in different age groups and regions in crashes from 2017 to 2021




Figure 12. Probabilities of tire violations in different age groups and regions in crashes from 2017 to 2021




Figure 13. Probabilities of lighting violations in different age groups and regions in crashes from 2017 to 2021

Chapter 7. Human-Centered Analysis of Inspection Practices

In addition to the proposal’s two primary research thrusts (physics-informed near-miss detection and customized sensor adoption) this project pursued an extended line of inquiry into inspection practices from a human-centered perspective. The rationale is straightforward: even the most advanced safety sensors or predictive algorithms ultimately depend on inspectors and fleet personnel for effective deployment. Understanding how humans learn, prioritize, and make decisions in inspection settings therefore provides critical insights into the adoption and reliability of technical solutions. This chapter presents the methodology and findings of that study, highlighting how human heuristics and adaptive learning can be leveraged to design more consistent and trustworthy inspection frameworks.

Methodology

This study examines how human inspectors identify critical features and adapt strategies under uncertainty. Extracting this knowledge is challenging, as it requires capturing both semantic reasoning and behavioral insights. To address this challenge, the authors integrated a structured survey with the Bayesian Inverse Contextual Bandit (BICB) framework to analyze and model dynamic decision-making processes.

Survey Design and Data Collection

The survey required 20 participants (11 male and 9 female) from engineering backgrounds, none of whom held an inspection license, to complete 40 trials across five rounds, evaluating eight vehicles in each trial. Nine key safety-related features, illustrated in Table 2. Features used in vehicle inspections described each vehicle. These features were chosen based on data availability constraints and their established relevance in prior studies on HDV inspections and align with key indicators identified by industry experts and regulatory agencies as critical for assessing vehicle safety and maintenance needs. Participants made binary decisions-to inspect or not inspect a vehicle-based on these features, while the ground truth determined whether an inspection was actually required. The iterative structure of the survey encouraged participants to refine their feature prioritization and decision-making strategies through feedback and repeated attempts. After each round, participants selected the features they believed most influenced their decisions, allowing the capture of explicit reasoning.

Figure 14 shows an example of a vehicle from the survey for which participants decided whether to inspect it or not. Participants needed to achieve a score of at least 75 out of 100 in each round to progress to the next round. Those who failed had to revise their decisions and repeat the round until they met the threshold.

Before beginning the trials, participants completed two training sessions. The first session introduced them to the vehicle features and provided an overview of current regulations for vehicle inspections. In the second session, participants practiced decision-making by evaluating four

sample vehicles, deciding whether to inspect each one, and receiving immediate feedback on their decisions. After each round, participants selected the features that influenced their decisions. This task captured how their reasoning evolved over time. To create diverse inspection conditions and reduce bias, the survey randomized vehicle assignments across trials, prompting participants to develop flexible and adaptive decision-making strategies.

Table 2. Features used in vehicle inspections

Features	Description
Vehicle Make	The brand of the vehicle
Vehicle Body	The type of vehicle, such as Truck or Truck-Tractor
Gross Vehicle Weight	Classification based on weight (Class 7 or 8)
Region	The region where the vehicle operates (e.g., Large fringe metro or non-core area)
Vehicle Age	The age of the vehicle in years
Odometer in the Last Inspection	Distance traveled since the last inspection, measured in miles
Mileage Driven in the Last Year	Distance covered by the vehicle in the past year, measured in miles
Brake Location	Specifies the location of the brake, such as Left Front, Rear Left
Brake Pad Thickness	Measurement of the brake pad's thickness during the last inspection

Property	Value
Vehicle Make	MACK
Vehicle Body	Truck
Gross Vehicle Weight	Class 8: 33,001 lb and above (14,969 kg and above)
Region	Large fringe metro
Vehicle Age	19 years
Odometer in the last inspection	125459 miles
Mileage driven in the last year	3431 miles
Brake Location	Left Front
Brake Pad Thickness in the last inspection	10 / 32 inch

Inspect the **Brake**?

Inspect **other components**? (click the items if applicable)

Figure 14. Example of a vehicle used for inspection decisions in the survey

Framework for Decision-Making Analysis

This study employs the BICB framework to model decision-making dynamics during vehicle inspections. The BICB framework is particularly well-suited to this task because it captures how decision-making strategies evolve over time and adapt to the sequential nature of the task (Hüyük et al. 2022). In this study, participants made binary decisions-whether to inspect or not inspect vehicles-over five rounds of trials. The framework infers participants' feature prioritization and belief updates based on observed actions, enabling a deeper understanding of their decision-making processes. The vehicle inspection task was formulated as a contextual bandit problem, where the key components were defined as follows:

- **State (x_t):** A set of nine key features representing each vehicle, as detailed in **Error! Reference source not found..**
- **Action (a_t):** Binary decisions made by participants for each vehicle: inspect ($a_t = 1$) or not inspect ($a_t = 0$).
- **Reward (r_t):** The reward captures the alignment of participants' actions with inferred feature importance, providing feedback on their prioritization of vehicle features rather than decision correctness.

The BICB framework models participants' beliefs about the importance of vehicle features as a multivariate Gaussian distribution. These beliefs are iteratively updated as new actions and rewards are observed. The belief parameters—mean (μ_t) and covariance (Σ_t)—are updated using equations ((1) and (2), respectively. Here, μ_t represents the inferred feature importance vector at time t , and Σ_t quantifies uncertainty in these beliefs. These updates enable the framework to capture how participants refine their understanding of vehicle features across rounds.

$$\mu_{t+1} := \Sigma_{t+1} \left(\Sigma_t^{-1} \mu_t + \frac{1}{\sigma^2} r_t x_t[a_t] \right) \quad (1)$$

$$\Sigma_{t+1} := \left(\Sigma_t^{-1} + \frac{1}{\sigma^2} x_t[a_t] x_t[a_t]^T \right)^{-1} \quad (2)$$

The decision policy in this framework balances exploration and exploitation, allowing participants to both experiment with different strategies and utilize their learned feature prioritization, where α is the softmax temperature parameter that controls the stochasticity of decisions. A higher α encourages exploration, while a lower α favors exploitation of the inferred feature importance. The reward function reflects participants' perceived prioritization of vehicle features, providing feedback that informs their learning process. It is formulated as:

$$\pi^*(x_t)[a_t] = \frac{\exp(\alpha R(x_t, a_t))}{\sum_{a' \in A} \exp(\alpha R(x_t, a'))}, \quad (3)$$

$$r_t = \sum_{i=1}^K (x_t[a_t])_i \cdot \rho_{\text{env}}[i] + \sigma \cdot \eta, \quad (4)$$

where $x_t[a_t]$ is the feature vector corresponding to the chosen action a_t , ρ_{env} represents the inferred feature importance vector, σ is the noise level, and $\eta \sim \mathcal{N}(0, 1)$ is a standard normal random variable. The agent dynamically updates its beliefs about feature importance by sampling from a posterior distribution, incorporating prior knowledge and observed actions and rewards. The logic of the agent is rooted in Bayesian updating, where the feature importance vector is treated as a multivariate Gaussian distribution. The agent refines its posterior beliefs iteratively as it observes new actions and rewards. Table 3. Hyperparameters used for training the agent lists the hyperparameters used for training the BICB agent.

Table 3. Hyperparameters used for training the agent

Hyperparameter	Value
α (Softmax temperature parameter)	20
σ (Noise level for reward simulation)	0.10
$\mathcal{E}$ (Regularization term for matrix operations)	1×10^{-6}
Iterations	100
Learning rate	0.001

Results

The BICB framework’s inferred feature prioritization aligns with self-reported survey data, with brake pad thickness (0.77), vehicle age, and mileage consistently receiving the highest importance scores. Figure 15. Feature importance reported by the participants Figure 15 and Figure 16 confirm this alignment, emphasizing their role in inspection decisions. However, discrepancies in feature rankings reveal complexities in human decision-making. The model assigns higher importance to region and brake location in some rounds, potentially detecting implicit learning or exploratory behaviors that participants did not consciously report. These variations may also stem from the small sample size, limiting the model’s ability to capture broader decision-making trends. On the other hand, the variance in self-reported data could result from individual biases, such as memory gaps or misunderstandings of the survey questions (Jensen et al. 2019). Experts emphasize the importance of key features such as brake pad thickness, mileage, vehicle age, and odometer readings, with additional features like region

and brake location considered only in challenging cases. Participants, based on self-reported data and their decision trends, appear to align with experts in identifying these critical features, reflecting an understanding of their importance in decision-making.

As participants improved their feature prioritization decisions in the final rounds—aligning with expert opinions on critical features such as brake pad thickness and mileage—the learning curve in Figure 17 illustrates their progress. The average number of attempts required to reach the passing score of 75% steadily decreases from approximately six attempts in Round 1 to just over three in Round 5, highlighting the effectiveness of the training process. Figure 18 provides the individual learning curves across the five rounds. Several participants, such as Participant 1 and Participant 12, demonstrated consistent improvement with fewer attempts per round, reflecting steady progress in learning.

In contrast, others, like Participant 3 and Participant 4, exhibited greater variability in their attempts, potentially due to different strategies or fluctuations in task understanding. Notably, some participants required a significantly higher number of attempts initially, such as Participant 10, who made 25 attempts in Round 1, but subsequently improved in later rounds. These observations underline the diversity in participants' learning trajectories and decision makings and highlight the need for better understanding and addressing inconsistencies in inspection strategies.

Conclusion

This study highlights the challenges of achieving consistent and reliable heavy-duty vehicle inspections due to variability in human decision-making and the increasing complexity of vehicle systems. By conducting a survey and analyzing feature prioritization through the BICB framework, the study offers insights into how human heuristics can be captured to enhance inspection strategies. Participants' alignment with expert-identified critical features, such as brake pad thickness and mileage, underscores the importance of integrating tacit knowledge into data-driven approaches. The observed 33% reduction in inspection attempts demonstrates the utility of structured analysis in improving decision-making processes. Given the diverse backgrounds of participants, variability in decision-making strategies was expected. To mitigate potential biases, vehicle assignments were randomized across trials to ensure that no participant consistently encountered similar inspection conditions. Additionally, participants underwent structured training before the trials to standardize baseline knowledge. While some participants exhibited faster learning curves than others, the overall trend across rounds indicated convergence toward expert-identified feature prioritization. However, this study did not include expert inspectors, limiting direct validation of participants' decision strategies against real-world inspection practices. Future research will address this by conducting interviews and structured evaluations with expert inspectors to compare their decision-making patterns with those of study participants. These interviews will help identify key differences in feature selection, learning speed, and prioritization strategies. These findings support the development of a collaborative human-AI framework that leverages human insights to address

data limitations and refine inspection strategies. While the findings remain valuable for understanding inspector decision-making, future research will expand this survey to include a larger and more diverse participant pool and employ advanced analytical tools, such as knowledge graphs, to capture and explore learning behaviors in greater detail.




Figure 15. Feature importance reported by the participants




Figure 16. Feature importance across 5 rounds captured by BICB




Figure 17. Average learning curve across five rounds




Figure 18. Learning curves of each 20 participants across five rounds

Chapter 8. Conclusions and Recommendations

8.1 Conclusions

This project advanced the Safety21 mission by integrating physics-informed modeling, data-driven sensor strategies, and socio-technical insights to improve heavy-duty vehicle (HDV) safety. The research demonstrated that:

1. **Physics-Informed Reinforcement Learning (PIRL) enables proactive safety modeling.**

By embedding vehicle dynamics and physical constraints into reinforcement learning, PIRL produced interpretable safety probability maps for tractor-trailers under rear, side, and lane-keeping scenarios. Results showed that PIRL outperformed standard DQN by capturing realistic risk gradients based on clearance, speed, and articulation. Importantly, trailer-side sensor inputs proved critical for accurate estimation, underscoring the value of enhanced observability in reducing blind-spot-related near-misses.

2. **Customized sensor deployment strategies improve small-carrier safety outcomes.**

Analysis of crash and inspection datasets revealed significant variations in brake, tire, and lighting violations across vehicle age groups and operating regions. For example, Region 6 vehicles aged 6–23 exhibited high violation rates across all three components, while Region 1 and 2 vehicles aged 24–29 were most vulnerable to brake and lighting issues. These findings highlight the inadequacy of one-size-fits-all mandates and support tailored, risk-based sensor adoption policies that prioritize high-risk carriers while respecting financial constraints.

3. **Human-centered analysis of inspections reveals adaptive but inconsistent decision-making.**

Survey-based experiments demonstrated that inspectors improve feature prioritization with experience, aligning more closely with expert risk factors such as brake pad thickness, mileage, and vehicle age. The use of Inverse Contextual Bandit (ICB) models captured these learning trajectories, showing convergence toward expert strategies while also highlighting individual variability. These insights suggest that human heuristics can be formalized and integrated with AI to enhance inspection reliability.

4. **Industry perspectives confirm socio-technical barriers.**

Interviews with fleet managers and inspectors emphasized barriers such as workforce shortages, data integration challenges, privacy concerns, and resistance to change. While predictive maintenance and AI-based decision support hold promise, successful adoption requires transparency, training, and incremental integration into existing practices.

Together, these findings show that HDV safety cannot be addressed by technology alone. Solutions must integrate technical advances with organizational readiness, workforce practices, and trust-building measures.

8.2 Recommendations

Based on the results, several recommendations are proposed for research, policy, and practice:

1. **Integrate telematics and real-world fleet data into PIRL models.**
Extending PIRL beyond simulation requires combining telematics, event-based video, and inspection data to validate safety probability estimations under real-world operating conditions. Partnerships with fleet operators and DOTs will be critical for scaling.
2. **Pilot customized sensor adoption programs with small carriers.**
Federal and state agencies should incentivize targeted deployments of brake, tire, and lighting sensors in high-risk regions and age groups. Pilot programs should test cost-sharing models, ensuring that small carriers can participate without disproportionate financial burdens.
3. **Develop workforce training programs aligned with predictive maintenance.**
Inspectors and fleet personnel require training that combines tacit knowledge with data-driven tools. Simulation-based training and AR-enhanced modules could improve workforce adoption and reduce variability in inspections.
4. **Strengthen data interoperability standards.**
Regulators and technology providers should prioritize common data formats and communication protocols to address fragmentation across telematics, inspection, and maintenance records. Improved interoperability will enable more robust predictive models and cross-fleet benchmarking.
5. **Build trust and transparency into AI-based safety systems.**
Driver and inspector acceptance depends on clear communication of system benefits, safeguards against misuse, and demonstrable value (e.g., exonerating drivers in crashes, preventing costly breakdowns). Explainable AI models and transparent feedback mechanisms are essential to overcoming resistance.

8.3 Future Research

Future work should focus on:

- Expanding PIRL validation with real-world sensor data from fleet partners.
- Exploring multi-agent PIRL for interactions between HDVs and passenger vehicles.
- Investigating knowledge graph methods to integrate human heuristics, inspection data, and predictive maintenance models into unified decision-support frameworks.

References:

- Abdeen, F. N., R. N. Gunatilaka, S. M. E. Sepasgozar, and D. J. Edwards. 2024. "The usability of a novel mobile augmented reality application for excavation process considering safety and productivity in construction." *Constr. Innov.*, 24 (4): 892–911. Emerald Publishing. <https://doi.org/10.1108/CI-07-2022-0168>.
- Anwar, Q., M. Hanif, D. Shimotoku, and H. H. Kobayashi. 2022. "Driver awareness collision/proximity detection system for heavy vehicles based on deep neural network." *J. Phys. Conf. Ser.*, 2330 (1): 012001. IOP Publishing. <https://doi.org/10.1088/1742-6596/2330/1/012001>.
- Assemi, B., and M. Hickman. 2018. "Relationship between heavy vehicle periodic inspections, crash contributing factors and crash severity." *Transp. Res. Part Policy Pract.*, 113: 441–459. <https://doi.org/10.1016/j.tra.2018.04.018>.
- Assemi, B., M. Hickman, and A. Paz. 2021. "Relationship between Programmed Heavy Vehicle Inspections and Traffic Safety." *Transp. Res. Rec. J. Transp. Res. Board*, 2675 (10): 1420–1430. <https://doi.org/10.1177/03611981211016458>.
- Barat, A., and S. Das. 2024. "A systematic literature review of driver behaviour in blind spot: Assessing risk and influencing factors for enhanced ergonomics in vehicle design." *Ergonomics*, 67 (12): 1765–1781. Taylor & Francis. <https://doi.org/10.1080/00140139.2024.2350667>.
- Bergoffen, G., R. R. Knipling, S. V. Burks, and K. C. Nelson. 2012. *Safety management in small motor carriers*. Washington, D.C.: National Academies Press.
- Berthet, V. 2022. "The Impact of Cognitive Biases on Professionals' Decision-Making: A Review of Four Occupational Areas." *Front. Psychol.*, 12: 802439. <https://doi.org/10.3389/fpsyg.2021.802439>.
- Blows, S., R. Q. Ivers, J. Connor, S. Ameratunga, and R. Norton. 2003. "Does periodic vehicle inspection reduce car crash injury? Evidence from the Auckland Car Crash Injury Study." *Aust. N. Z. J. Public Health*, 27 (3): 323–327. <https://doi.org/10.1111/j.1467-842X.2003.tb00401.x>.
- Burger, W. J., M. U. Mulholland, and R. L. Smith. 1985. "Improved commercial vehicle conspicuity and signalling systems -- Task III field test evaluation of vehicle reflectorization effectiveness. Final report."
- Burger, W. J., M. U. Mulholland, R. L. Smith, and M. C. Bardales. 1980. "MEASUREMENT OF BLIND AREAS ON THE RIGHT FRONT OF HEAVY TRUCKS."
- Butt, M. S., and M. A. Shafique. 2025. "A literature review: AI models for road safety for prediction of crash frequency and severity." *Discov. Civ. Eng.*, 2 (1): 99. <https://doi.org/10.1007/s44290-025-00255-3>.
- Cai, P., X. Mei, L. Tai, Y. Sun, and M. Liu. 2020. "High-Speed Autonomous Drifting With Deep Reinforcement Learning." *IEEE Robot. Autom. Lett.*, 5 (2): 1247–1254. <https://doi.org/10.1109/LRA.2020.2967299>.
- Camden, M. C., G. M. Fitch, M. Blanco, and R. J. Hanowski. 2011. "Commercial Driver Acceptance of Heavy Vehicle Camera/Video Imaging Systems." *Proc. Hum. Factors Ergon. Soc. Annu. Meet.*, 55 (1): 1899–1903. SAGE Publications Inc. <https://doi.org/10.1177/1071181311551395>.
- Cantor, D. E., T. M. Corsi, and C. M. Grimm. 2006. "Safety technology adoption patterns in the U.S. motor carrier industry." *Transp J*, 45 (3): 20–45. <https://doi.org/10.2307/20713642>.

- Cantor, D. E., T. M. Corsi, C. M. Grimm, and P. Singh. 2016. "Technology, firm size, and safety: Theory and empirical evidence from the US motor-carrier industry." *Transp J*, 55 (2): 149–167. <https://doi.org/10.5325/transportationj.55.2.0149>.
- Cantor, D. E., E. Osborn, and P. Singh. 2014. "A firm size and safety performance profile of the US motor carrier industry." *Natl. Transp. Libr.* Accessed November 20, 2023. <https://rosap.nhtl.bts.gov/view/dot/28560>.
- Chern, A., X. Wang, A. Iyer, and Y. Nakahira. 2021. "Safe Control in the Presence of Stochastic Uncertainties." *2021 60th IEEE Conf. Decis. Control CDC*, 6640–6645.
- Cutler, M., and J. P. How. 2016. "Autonomous drifting using simulation-aided reinforcement learning." *2016 IEEE Int. Conf. Robot. Autom. ICRA*, 5442–5448.
- Dosovitskiy, A., G. Ros, F. Codevilla, A. Lopez, and V. Koltun. 2017. "CARLA: An Open Urban Driving Simulator." arXiv.
- Fitch, G., M. Blanco, R. Hanowski, P. Rau, and C. Flanigan. 2010. "Field Demonstration of a Camera/Video Imaging System for Heavy Vehicles - Driver Lane Change Performance Preliminary Results." *SAE Tech. Pap. Ser.* 400 Commonwealth Drive, Warrendale, PA, United States: SAE International.
- FMCSA (Federal Motor Carrier Safety Administration). 2006. "Report to Congress on the large truck crash causation study." *U. S. Dep. Transp.* Accessed January 13, 2024. <https://www.fmcsa.dot.gov/sites/fmcsa.dot.gov/files/docs/ltrcs-2006.pdf>.
- Geng, M., J. Li, Y. Xia, and X. (Michael) Chen. 2023. "A physics-informed Transformer model for vehicle trajectory prediction on highways." *Transp. Res. Part C Emerg. Technol.*, 154: 104272. <https://doi.org/10.1016/j.trc.2023.104272>.
- Gigerenzer, G., J. Reb, and S. Luan. 2022. "Smart Heuristics for Individuals, Teams, and Organizations." *Annu. Rev. Organ. Psychol. Organ. Behav.*, 9 (Volume 9, 2022): 171–198. Annual Reviews. <https://doi.org/10.1146/annurev-orgpsych-012420-090506>.
- Goettee, D., C. Campanian, and W. Spiegel. 2010. "Fostering a Safety Culture in Small Motor Carriers." *51st Annu. TFR*. Arlington, VA.
- Han, S., K. Yoon, G. Park, and K. Huh. 2024. "Robust Lane Keeping Control for Tractor With Multi-Unit Trailer Under Parametric Uncertainty." *IEEE Trans. Intell. Veh.*, 9 (1): 2333–2347. <https://doi.org/10.1109/TIV.2023.3272655>.
- Hindiye, R. Y., and J. Christian Gerdes. 2014. "A Controller Framework for Autonomous Drifting: Design, Stability, and Experimental Validation." *J. Dyn. Syst. Meas. Control*, 136 (051015). <https://doi.org/10.1115/1.4027471>.
- Hoshino, H., J. Li, A. Menon, J. M. Dolan, and Y. Nakahira. 2024. "Autonomous Drifting Based on Maximal Safety Probability Learning." arXiv.
- Hoshino, H., and Y. Nakahira. 2024. "Physics-informed RL for Maximal Safety Probability Estimation." *2024 Am. Control Conf. ACC*, 3576–3583.
- Huang, A. J., and S. Agarwal. 2022. "Physics-Informed Deep Learning for Traffic State Estimation: Illustrations With LWR and CTM Models." *IEEE Open J. Intell. Transp. Syst.*, 3: 503–518. <https://doi.org/10.1109/OJITS.2022.3182925>.
- Hüyük, A., D. Jarrett, and M. van der Schaar. 2022. "Inverse contextual bandits: Learning how behavior evolves over time." arXiv.
- Ibrahim, S., M. Mostafa, A. Jnadi, H. Salloum, and P. Osinenko. 2024. "Comprehensive Overview of Reward Engineering and Shaping in Advancing Reinforcement Learning Applications." *IEEE Access*, 12: 175473–175500. <https://doi.org/10.1109/ACCESS.2024.3504735>.

- Jensen, H. A. R., M. Davidsen, A. I. Christensen, and O. Ekholm. 2019. "Inconsistencies in self-reported health conditions: results of a nationwide panel study." *Int. J. Public Health*, 64 (8): 1243–1246. <https://doi.org/10.1007/s00038-019-01287-0>.
- Johnson, T. L., S. R. Fletcher, W. Baker, and R. L. Charles. 2019. "How and why we need to capture tacit knowledge in manufacturing: Case studies of visual inspection." *Appl. Ergon.*, 74: 1–9. <https://doi.org/10.1016/j.apergo.2018.07.016>.
- Jurj, S. L., D. Grundt, T. Werner, P. Borchers, K. Rothemann, and E. Möhlmann. 2021. "Increasing the Safety of Adaptive Cruise Control Using Physics-Guided Reinforcement Learning." *Energies*, 14 (22): 7572. Multidisciplinary Digital Publishing Institute. <https://doi.org/10.3390/en14227572>.
- Kim, B., E.-J. An, S. Kim, K. R. Sri Preethaa, D.-E. Lee, and R. R. Lukacs. 2024. "SRGAN-enhanced unsafe operation detection and classification of heavy construction machinery using cascade learning." *Artif. Intell. Rev.*, 57 (8). Springer Nature. <https://doi.org/10.1007/s10462-024-10839-7>.
- Liang, H., Z. Dong, H. Li, Y. Yue, M. Fu, and Y. Yang. 2025. "Unified Vertex Motion Estimation for integrated video stabilization and stitching in tractor–trailer wheeled robots." *Robot. Auton. Syst.*, 191: 105004. <https://doi.org/10.1016/j.robot.2025.105004>.
- Liu, D., C. Sun, J. Chen, and L. Liu. 2023. "Multisensory and BIM-Integrated Digital Twin to Improve Urban Excavation Safety." *J. Comput. Civ. Eng.*, 37 (5). American Society of Civil Engineers (ASCE). <https://doi.org/10.1061/JCCEE5.CPENG-5354>.
- Love, P. E. D. 2025. "Smart Heuristics in Civil Engineering: Unearthing the Unexploited." *ASCE-ASME J. Risk Uncertain. Eng. Syst. Part Civ. Eng.*, 11 (1): 06024006. American Society of Civil Engineers. <https://doi.org/10.1061/AJRUA6.RUENG-1440>.
- Miller, J. 2020. "Why are larger motor carriers more compliant with safety regulations?" *Transp J*, 59 (1): 28–72. <https://doi.org/10.5325/transportationj.59.1.0028>.
- National Center for Statistics and Analysis. 2024. *Large trucks: 2022 data*. Washington, DC.
- NCHS (National Center for Health Statistics). 2017. "Data access - Urban-rural classification scheme for counties." *Cent. Dis. Control Prev*. Accessed January 13, 2024. https://www.cdc.gov/nchs/data_access/urban_rural.htm.
- NHTSA (National Highway Traffic Safety Administration). 2008. "National motor vehicle crash causation survey: Report to Congress." *Natl. Tech. Rep. Libr*. Accessed January 19, 2024. <https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/811059>.
- Rezapour Mashhadi, M. M., S. Wulff, and K. Ksaibati. 2018. "A Comprehensive Study of Single and Multiple Truck Crashes Using Violation and Crash Data." *Open Transp. J.*, 12: 43–56. <https://doi.org/10.2174/1874447801812010043>.
- Schaudt, W. A., D. S. Bowman, J. Bocanegra, R. J. Hanowski, and C. Flanigan. 2010. "Enhanced Rear Signaling (ERS) for Heavy Trucks: Mitigating Rear-end Truck Crashes Using Visual Warning Signals." *J. Wash. Acad. Sci.*, 96 (3).
- Scott, A., A. Balthrop, and J. Miller. 2019. "Did the electronic logging device mandate reduce accidents?" *Available SSRN 3314308*.
- Skaug, L., M. Nojournian, N. Dang, and A. Yap. 2025. "Road Crash Analysis and Modeling: A Systematic Review of Methods, Data, and Emerging Technologies." *Appl. Sci.*, 15 (13): 7115. Multidisciplinary Digital Publishing Institute. <https://doi.org/10.3390/app15137115>.
- Summerskill, S., R. Marshall, A. Paterson, and S. Reed. 2016. "Understanding direct and indirect driver vision in heavy goods vehicles."

- U.S. Department of Transportation Bureau of Transportation Statistics, U.S. Department of Commerce Census Bureau. 2012. “Transportation: 2012 Commodity Flow Survey. Technical Report Table 1 Shipment Characteristics by Mode of Transportation for the United States.” Accessed July 27, 2025. <https://www.bts.gov/>.
- Wang, Q., J. Sun, N. Wang, Y. Wang, Y. Song, and X. Li. 2022a. “Exploring the Influencing Factors and Formation of the Blind Zone of a Semitrailer Truck in a Right-Turn Collision.” *Sustainability*, 14 (16): 9805. Multidisciplinary Digital Publishing Institute. <https://doi.org/10.3390/su14169805>.
- Wang, Y., S. Luan, and G. Gigerenzer. 2022b. “Modeling fast-and-frugal heuristics.” *PsyCh J.*, 11 (4): 600–611. <https://doi.org/10.1002/pchj.576>.
- Wang, Z., R. Keller, X. Deng, K. Hoshino, T. Tanaka, and Y. Nakahira. 2024. “Physics-Informed Representation and Learning: Control and Risk Quantification.” *Proc. AAAI Conf. Artif. Intell.*, 38 (19): 21699–21707. <https://doi.org/10.1609/aaai.v38i19.30169>.
- Xu, C., J. Gao, F. Zuo, and K. Ozbay. 2024a. “Estimating Urban Traffic Safety and Analyzing Spatial Patterns through the Integration of City-Wide Near-Miss Data: A New York City Case Study.” *Appl. Sci.*, 14 (14): 6378. Multidisciplinary Digital Publishing Institute. <https://doi.org/10.3390/app14146378>.
- Xu, Y., W. Shen, X. Zhang, and J. Xu. 2024b. “IBCB: Efficient Inverse Batched Contextual Bandit for Behavioral Evolution History.” arXiv.
- Yang, D., K. Ozbay, K. Xie, H. Yang, F. Zuo, and D. Sha. 2021. “Proactive safety monitoring: A functional approach to detect safety-related anomalies using unmanned aerial vehicle video data.” *Transp. Res. Part C Emerg. Technol.*, 127: 103130. <https://doi.org/10.1016/j.trc.2021.103130>.
- Yasunaga, R., Y. Nakahira, and Y. Hori. 2024. “Orthogonal Modal Representation in Long-Term Risk Quantification for Dynamic Multi-Agent Systems.” *IEEE Control Syst. Lett.*, 8: 3177–3182. <https://doi.org/10.1109/LCSYS.2024.3522949>.
- Zhao, H., W. Chen, S. Zhou, F. Zheng, and Y.-H. Liu. 2023. “Localization and Motion Planning of Industrial Tractor–Trailers Vehicles.” *IEEE Trans. Control Syst. Technol.*, 31 (6): 2928–2936. <https://doi.org/10.1109/TCST.2023.3275497>.
- Zvěřina, L., M. Cibulka, L. Staněk, and R. Ulrich. 2022. “Analysis of Field View from Tractor and Trailer Units Using the Method of Terrestrial Laser Scanning.” *Forests*, 13 (9): 1383. Multidisciplinary Digital Publishing Institute. <https://doi.org/10.3390/f13091383>.