



US DOT National
University Transportation Center for Safety

Carnegie Mellon University



Smart Trailer Technologies for Data-Driven Safety, Efficiency, and Policy Innovation in Heavy-Duty Fleet Operations

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

Peter Zhang (Co-PI) <https://orcid.org/0000-0002-0422-834X>

Final Report – August 30, 2026

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.

1. Report No. 589	2. Government Accession No.	3. Recipient's Catalog No.
4. Title and Subtitle Smart Trailer Technologies for Data-Driven Safety, Efficiency, and Policy Innovation in Heavy-Duty Fleet Operations		5. Report Date <i>August 30, 2026</i>
		6. Performing Organization Code
7. Author(s) Ava Jahanbiglari, https://orcid.org/0000-0002-9713-0543 Pingbo Tang, Ph.D., https://orcid.org/0000-0002-4910-1326 , Peter Zhang, Ph.D., https://orcid.org/0000-0002-0422-834X		8. Performing Organization Report No.
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		13. Type of Report and Period Covered Final Report (July 1, 2025 – July 31, 2026)
		14. Sponsoring Agency Code USDOT

15. Supplementary Notes

Conducted in cooperation with the U.S. Department of Transportation, Federal Highway Administration.

16. Abstract

This project developed an integrated framework for evaluating smart-trailer safety technologies using naturalistic vehicle data, camera-based ground localization, computer-vision auditing, and benefit-cost analysis. Naturalistic data from one trailer operated with two tractors across 13 trips and 10 collection dates produced 56 hours and 201,600 synchronized one-second observations. Tractor and trailer GNSS/IMU measurements, three exterior camera views, detector activity, and roadway context were used to characterize operating behavior across five environments. Ramp-classified operation showed the clearest trailer-specific departure from arterial operation, with lower trailer-relative yaw activity but higher trailer-relative lateral activity, indicating that tractor-only measurements may incompletely represent trailer response. Camera-specific pixel-to-ground mappings were also developed and evaluated. Strict nested validation showed the strongest performance for the driver-side camera, followed by the curb-side camera, while rear-camera continuous-distance estimates remained provisional. A retrospective audit of 156 videos further showed that the project computer-vision pipeline and recorder alarm boxes provided complementary but incomplete evidence, supporting an evidence-fusion approach rather than reliance on either source alone. A deterministic benefit-cost analysis found national side-guard BCRs of 0.106-0.122 under the modeled narrow NHTSA target population and quantified warning-system and targeted-deployment break-even conditions. Overall, the findings support trailer-aware monitoring, context-specific interpretation, camera-specific uncertainty controls, exposure-normalized alert evaluation, and targeted deployment strategies.

17. Key Words

smart trailers; heavy trucks; tractor-trailers; naturalistic driving; trailer sensing; blind-zone monitoring; computer vision; camera calibration; GNSS; inertial measurement unit; near-miss assessment; underride guards; benefit-cost analysis; roadway context; fleet safety

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)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
		66	

Form DOT F 1700.7 (8-72)

Reproduction of completed page authorized

Chapter 1. Introduction and Project Context

1.1 Transportation Safety Motivation

Large trucks play an essential role in freight movement and economic activity, but crashes involving these vehicles can impose particularly severe consequences because of the differences in mass, geometry, maneuverability, and occupant protection between large trucks and other road users. In 2024, 5,340 people were killed in crashes involving large trucks in the United States. Seventy percent of those killed were occupants of vehicles other than the large truck (United States Department of Transportation and National Highway Traffic Safety Administration 2026). These national statistics encompass multiple types of large trucks and should not be interpreted as applying exclusively to tractor-trailer combinations. Nevertheless, they demonstrate that large-truck safety remains a substantial public concern and that the consequences of these crashes frequently extend beyond truck occupants to passenger-vehicle occupants and other road users.

Articulated tractor-trailer combinations present safety and monitoring challenges that are distinct from those of a single rigid vehicle. Although the tractor and trailer are mechanically connected, they do not necessarily follow the same path or experience the same motion. Their responses depend on speed, articulation geometry, wheelbase, steering input, loading, roadway curvature, lane position, and the conditions under which a maneuver is initiated. Prior heavy-vehicle research has therefore distinguished among low- and high-speed off tracking, swept path, yaw stability, lateral load transfer, rollover resistance, and rearward amplification rather than treating “vehicle

motion” as a single quantity (Fancher and Winkler 2007; Hac et al. 2008; Kamnik et al. 2003; Trigell et al. 2017). These distinctions are important for side- and rear-region safety because a tractor-mounted sensor may describe the motion of the towing unit accurately while incompletely representing the position or response of the trailer.

The trailer’s length, ground clearance, articulation, and visibility envelope can complicate merging, lane changing, turning, backing, positioning, and operation in constrained facilities. These interactions do not all represent imminent crashes, but they create circumstances in which the relative positions of the trailer, adjacent vehicles, workers, pedestrians, or other road users can change rapidly. The side and rear regions are especially important because the driver may have limited direct visibility, the trailer may track differently from the tractor, and the clearance available to nearby road users may not be readily inferred from tractor position alone.

One of the most severe crash mechanisms involving these regions is underride. An underride crash occurs when a passenger vehicle travels beneath the body of a larger truck, such as a tractor-trailer or single-unit truck, because of the height difference between the vehicles. During such a crash, the truck or trailer structure may intrude into the passenger compartment, resulting in severe injuries or fatalities. The U.S. Government Accountability Office reported that, from 2008 through 2017, an average of approximately 219 fatalities from underride crashes involving large trucks were reported annually. GAO concluded that these fatalities were likely underreported because police agencies did not use a consistent definition of underride, state crash-report forms varied, and responding officers received limited information or training on identifying and documenting these crashes. These data limitations affect the ability of transportation agencies to characterize the safety problem and make informed research, regulatory, and countermeasure decisions (U. S. Government Accountability Office 2019). The resulting uncertainty affects not only estimates of the size of the safety problem but also regulatory analysis, technology evaluation, and the economic assessment of proposed countermeasures.

The estimated number of underride fatalities also depends strongly on how the target crash population is defined. In its 2024 *Report to Congress: Side Underride Protection*, NHTSA estimated that an average of 89 fatalities annually fell within the scope of the Bipartisan Infrastructure Law directive: light-passenger-vehicle occupant fatalities in two-vehicle crashes with tractor-trailers in which the light passenger vehicle struck the side of a trailer or semitrailer and underrided it (National Highway Traffic Safety Administration 2024). For its side-guard benefit-cost analysis, however, NHTSA used a narrower target population aligned with the crash configuration and tested performance range of the evaluated side guard. The central analysis estimated 17.7 annual fatalities and 81 annual serious injuries in light-passenger-vehicle–combination-truck side-underride crashes at impact speeds up to 40 mph (United States Department of Transportation and National Highway Traffic Safety Administration 2023). The 2024 Report to Congress presented these values in rounded form as approximately 18 fatalities and 81 injuries. These estimates are not contradictory; rather, they demonstrate how the estimated safety benefit depends on the crash definition, impact-speed range, intervention mechanism,

underreporting adjustment, and evidence of countermeasure effectiveness. Underride represents only one part of the broader trailer-safety problem. Side and rear interactions may also involve lane encroachment, insufficient clearance, backing conflicts, workers near a parked or maneuvering trailer, vulnerable road users beside the vehicle, or vehicles occupying a monitored proximity region without being on a collision path. A safety framework focused only on recorded crashes may therefore overlook important precursors and operating conditions. At the same time, a framework that labels every nearby object or detector activation as a safety-critical event can exaggerate risk and create unnecessary warnings. A defensible smart-trailer approach must operate between these extremes: it must detect and preserve potential safety evidence while using vehicle motion, roadway environment, object trajectory, measurement uncertainty, and exposure to determine what that evidence means.

This layered perspective is consistent with the Safe System principle of redundancy, under which multiple forms of protection are used to prevent crashes and reduce harm when prevention is unsuccessful. Passive structural measures can mitigate the consequences of contact, while active sensing and warning systems may help identify an interaction before contact occurs. Neither layer should be assumed to eliminate the need for the other.

1.2 State of Practice in Passive and Active Trailer-Safety Measures

Trailer-safety interventions can be broadly divided into passive and active measures. Passive measures operate primarily through vehicle structure and crashworthiness. Rear and side underride guards are intended to limit the extent to which a smaller vehicle can travel beneath a trailer during a collision. Active measures use sensing, communication, analytics, or control to detect a potentially hazardous situation and support avoidance before impact. Examples include camera-based visibility systems, blind-zone detection, radar or ultrasonic proximity sensing, telematics, stability monitoring, automatic emergency braking, and data-driven warning or decision-support systems.

The federal regulatory framework is more developed for rear underride protection than for side underride protection. Federal Motor Vehicle Safety Standards Nos. 223 and 224 specify rear-impact-guard performance and installation requirements for applicable trailers and semitrailers. A 2022 final rule strengthened the rear-impact requirements for specified 35-mph full-overlap and 50-percent-overlap passenger-vehicle impacts, with the objective of preventing passenger-compartment intrusion in the covered crash configurations (National Highway Traffic Safety Administration 2022).

Side underride protection presents a more complex regulatory and operational question. NHTSA's 2023 assessment and its 2024 Report to Congress considered potential safety benefits, equipment and fuel costs, trailer compatibility, loading-dock and intermodal operations, freight capacity, and the limited range of independently demonstrated guard performance. NHTSA estimated that costs exceeded monetized safety benefits under the modeled national target population and assumptions. It also identified important effects that could not be fully quantified, including interactions with

infrastructure, loading operations, maintenance, and different trailer configurations. NHTSA consequently deferred a determination on whether to develop side-guard performance requirements while further recommendations and analyses were considered.

These findings demonstrate why an economic conclusion cannot be separated from the scope of the safety problem. A guard may be highly effective in a particular crash configuration and still have a benefit-cost ratio below one when deployed across a broad national fleet with relatively few qualifying target crashes. Conversely, a technology that does not appear cost-beneficial under average national conditions may warrant consideration in a more concentrated operating context, provided that the higher risk concentration can be demonstrated rather than assumed. This project therefore treated economic modeling not simply as a final calculation, but as a means of identifying the effectiveness, cost, exposure, and targeting evidence required for defensible deployment decisions.

Active technologies address a different point in the crash sequence. Trailer-mounted cameras and proximity systems can extend visibility into side and rear regions, while telematics and inertial sensing can characterize vehicle state and motion. Field evaluations of heavy-truck warning technologies have shown the value of combining warning logs with video, vehicle kinematics, and operating context rather than interpreting warning counts alone (Nodine et al. 2011; Schaudt et al. 2014). Naturalistic studies have also shown that proximity-system activations may involve real objects within a monitored zone without representing an imminent collision or a need for evasive action (Ruff 2006; Tomasch and Smit 2023).

This distinction is important for both technical performance and human use. A detector may produce a factually incorrect alert because no relevant object is present, or it may produce an alert that is technically understandable but unnecessary in the current operating context. Human-factors research indicates that false and unnecessary warnings can affect later compliance differently and may reduce trust or contribute to desensitization if the operator cannot understand why the system activated (Naujoks et al. 2016). A high alert count is therefore not automatically evidence of a hazardous route, an effective safety system, or unsafe driving. It may reflect repeated low-speed positioning, objects occupying an intentionally broad detection region, difficult image conditions, detector calibration, or a mixture of these factors.

Active and passive countermeasures should therefore be evaluated as complementary components of a layered safety strategy. A passive guard is intended to reduce injury severity after a collision becomes unavoidable or contact begins. An active system is intended to recognize conditions that may precede a collision and support avoidance or earlier intervention. Their benefits may overlap: a crash avoided by a warning system is no longer available to be mitigated by a guard. Within the project framework, passive and active benefit pathways are evaluated separately to avoid double counting. The guard module represents severity mitigation after contact, while the warning-system module reports the effectiveness and cost conditions required for crash avoidance to break even. Any combined accounting should distinguish crashes avoided by active intervention, residual

crashes mitigated by structural protection, and operational benefits that do not overlap with either pathway.

1.3 Smart Trailers and the Need for Trailer-Aware, Context-Aware Monitoring

The Project proposal identified the absence of a clear and consistently applied definition of a “smart trailer” as a barrier to evaluation, policy development, and deployment planning. The term is used broadly in practice and may refer to technologies as different as asset tracking, cargo monitoring, tire-pressure sensing, brake diagnostics, rear cameras, blind-zone detection, telematics gateways, or advanced analytics. A trailer does not become meaningfully “smart,” however, merely because a connected device has been installed. The safety contribution depends on what is sensed, whether the measurement is valid, how information from different sources is synchronized, how the operating context is represented, and whether the resulting output supports a defensible decision.

For this project, a smart trailer is defined functionally as:

A trailer equipped with an integrated sensing, communication, and analytical capability that can independently observe trailer state or surrounding conditions, relate those observations to tractor and operating context, and convert the combined evidence into information that can support safety, maintenance, operational, or policy decisions.

This definition is operationalized through six connected capabilities:

1. Trailer-specific observation: sensing the state or environment of the trailer rather than assuming that tractor measurements fully represent the articulated combination.
2. Measurement validity: calibrating sensors and defining the conditions under which their outputs can be interpreted physically.
3. Temporal and spatial integration: synchronizing tractor, trailer, video, warning, and roadway information using compatible time and coordinate references.
4. Context-aware interpretation: distinguishing expected operation from unusual response by considering roadway environment, speed, facility conditions, and exposure.
5. Evidence-based event assessment: separating raw system outputs, proximity events, trajectory-based candidates, and validated safety-critical events.
6. Decision support: translating technical evidence into operational guidance, deployment priorities, economic thresholds, or policy recommendations.

This functional definition provides a consistent basis for comparing smart-trailer technologies according to the measurements they produce, the contexts they represent, and the decisions they support.

The need for trailer-specific observation is strongly supported by articulated-vehicle dynamics. Controlled studies and analytical models have long established that trailer response depends on articulation, geometry, loading, maneuver, and speed. Naturalistic evidence is more limited because collecting synchronized measurements from both physical units during ordinary service is difficult. A recent naturalistic study of long combination vehicles showed that the relative performance of the first and last units varied across lane changes, curves, roundabouts, and intersection maneuvers, indicating that no single performance measure is sufficient for all scenarios (Behera et al. 2024). The naturalistic work conducted under this project extends this logic from selected maneuvers to broad operating environments and retains the tractor and trailer as separate measurement units.

Roadway and operational context are also essential. Drivers adapt speed before, during, and after curves and ramps, and those adaptations can affect observed yaw, acceleration, and lane-position behavior (Dhahir and Hassan 2018; Hallmark et al. 2015; Vergara et al. 2024; B. Wang et al. 2017). A motion level that is expected in a parking or staging facility may be interpreted differently on an arterial or freeway. Similarly, a uniform monitoring rule may treat normal ramp-related adaptation as abnormal or overlook trailer-specific response if it relies only on tractor telemetry. An environment-aware operating baseline is therefore needed before sensor data can be used to define anomalies, calibrate alerts, or prioritize intervention.

Naturalistic analysis introduces substantial data-engineering and methodological challenges. Tractor and trailer GNSS/IMU records, recorded video, detector overlays, roadway-network data, and facility geofences may use different sampling rates, timestamps, identifiers, coordinate systems, and definitions of missingness. High-frequency records create a large number of observations but not an equivalent number of independent operating situations. Connected-vehicle and naturalistic-infrastructure research therefore emphasizes timestamp verification, map matching, coordinate harmonization, missing-data screening, trajectory validation, and appropriate treatment of repeated routes, trips, or dates (Howell et al. 2025; Zhou and Bridgelall 2020; X. Wang et al. 2020).

Camera data add another measurement layer. A detector can identify an object in image pixels, but pixel location alone does not provide physical clearance from the trailer. Perspective and fisheye distortion cause the number of pixels per foot to vary substantially across each image. The Project calibration works therefore evaluated camera-specific mappings from an object's visible ground-contact pixel to a local road-plane coordinate. Importantly, the available planar survey did not support trustworthy estimation of physical fisheye intrinsics and full camera pose. The accepted product was instead a set of empirical, fixed-camera pixel-to-ground mappings restricted to each surveyed image region. Strict, nested validation supported controlled use of the driver-side and curb-side mappings under stated conditions, while the rear mapping remained provisional because of substantially greater continuous-distance uncertainty.

Computer vision introduces a related but distinct issue: no single detector output should automatically be treated as ground truth. In the Project retrospective audit, a YOLO-based segmentation and tracking pipeline independently identified and localized some close objects for which no matching recorder alarm box was visible. The reverse also occurred: the recorder displayed boxes around some partially visible or occluded objects for which the general-purpose YOLO model did not produce a matched detection at the sampled frame. These disagreements arose from multiple sources, including image-boundary truncation, occlusion, fisheye domain shift, sampling frequency, detector confidence, colored-overlay extraction, and box-association criteria. The appropriate engineering response is therefore evidence fusion and independent validation, not the assumption that either stream is universally correct.

This observation also requires careful safety terminology. In this report:

- A warning-system output is a displayed or recorded system activation.
- A CV-supported proximity event is an object track with sufficient valid metric support within a specified distance region.
- A trajectory/time-to-contact candidate satisfies a declared temporal and closing-motion rule.
- A validated near miss should be reserved for an event supported by predetermined conflict criteria and independent review.

Proximity alone does not establish that two road users were on a collision path. Similarly, the absence of a recorder box does not prove that the system missed a required warning, and the presence of a box does not prove that an immediate evasive response was necessary. This hierarchy allows the project to preserve potentially important interactions without overstating their safety meaning.

1.4 Research and Deployment Gaps

The existing evidence reveals a connected set of technical, analytical, and deployment gaps. Controlled vehicle-dynamics studies provide valuable measures of articulated behavior but do not fully describe routine fleet operation. Naturalistic-driving studies capture real operating conditions but rarely retain synchronized tractor and trailer states across multiple roadway environments. Warning-system evaluations demonstrate that alerts require context but do not necessarily provide a trailer-specific behavioral baseline or an independently validated object trajectory. Computer-vision systems can identify objects but may not provide defensible physical clearance without camera-specific calibration. Finally, national benefit-cost models provide policy benchmarks but may not determine which fleets, routes, times, or operating environments should be prioritized for deployment. Table 1 summarizes how the project responds to these connected gaps.

Table 1. Project responses to the identified connected gaps

Research or deployment gap	Consequence for safety interpretation	Project response
No consistently applied functional definition of a smart trailer	Technologies are compared by device labels rather than capabilities, evidence, and decisions supported	Develop a working functional definition organized around sensing, validation, integration, context, event interpretation, and decision support
Tractor-only or context-free vehicle monitoring	Trailer-specific response may be obscured, and expected environmental adaptation may be treated as abnormal	Synchronize tractor and trailer GNSS/IMU measurements and establish an environment-aware naturalistic baseline
Warning counts interpreted without exposure or operating context	Alert totals may be mistaken for conflict frequency or safety-system effectiveness	Analyze detector activity as a secondary system state and relate it to speed, roadway environment, monitored zone, and exposure
Uncalibrated wide-angle camera observations	Image detections cannot be converted reliably into physical clearance	Develop and evaluate camera-specific fixed-road-plane pixel-to-ground mappings, with explicit support regions, uncertainty estimates, and camera-specific limitations.
A recorder output or general-purpose CV model treated as ground truth	Misses, nuisance detections, association errors, and unsupported near-miss claims may result	Compare the two visual streams independently, visually audit disagreement cases, and recommend evidence fusion
National-average economic evaluation without deployment targeting	A technology may be rejected or adopted without identifying the risk concentration needed for value	Develop a corrected deterministic BCA, warning-system break-even thresholds, and targeted-deployment risk-concentration conditions

Together, these gaps show that the central challenge is not simply detecting more objects or collecting more telemetry. The challenge is building a traceable evidence chain from measurement to action. Each stage must answer a different question: Was the sensor observation valid? Was the tractor-trailer operating as expected for the environment? Was an external object present and metrically supported? Was it approaching or merely nearby? Did the installed system respond? What intervention would have been relevant? What level of effectiveness and cost would justify deployment? A smart-trailer framework that omits any of these stages risks producing outputs that are technically impressive but operationally or economically uninterpretable.

1.5 Project Purpose, Objectives, and Contributions

The project, Smart Trailer Technologies for Data-Driven Safety, Efficiency, and Policy Innovation in Heavy-Duty Fleet Operations, was initiated to address gaps in trailer-safety definitions, data, modeling, and deployment guidance. During the project period, the research produced several connected technical foundations. This final report organizes those products around five objectives.

Objective 1: Establish a functional and evidence-based smart-trailer framework

The first objective is to define the capabilities that make trailer technology useful for safety decision-making. Rather than defining a smart trailer by the presence of a particular sensor, the project organizes the concept around trailer-specific observation, measurement validity, multimodal data integration, operating context, event interpretation, and deployment decisions.

Objective 2: Characterize tractor-trailer behavior across natural operating environments

The second objective is to establish an environment-aware baseline using synchronized measurements from the tractor and trailer. The analysis distinguishes the likelihood of movement, speed when moving, and paired trailer-minus-tractor yaw-rate and lateral specific-force activity across facility, collector/local, arterial, freeway/interstate, and ramp-classified exposure. Recorded detector activity is incorporated as secondary context rather than as a verified safety outcome.

Objective 3: Develop and evaluate camera-based ground localization for proximity assessment

The third objective is to convert supported ground-contact observations from the curb-side, driver-side, and rear cameras into camera-local road-plane coordinates and approximate clearance distances. The work evaluates multiple empirical mapping methods, uses complete-setup and nested validation to reduce model-selection bias, defines supported image regions, and documents camera-specific limitations. This objective provides the metric foundation needed to move from an image detection to an approximate physical clearance. The resulting evidence supports controlled in-region use of the driver-side and curb-side mappings and provisional interpretation of rear-camera continuous-distance estimates.

Objective 4: Audit computer-vision and installed alarm evidence for proximity and near-miss assessment

The fourth objective is to determine how an independent computer-vision pipeline and the installed recorder's alarm boxes agree and disagree. The project combines object segmentation, temporal tracking, calibrated ground-contact mapping, trajectory analysis, recorder-box extraction, and visual review. The purpose is not to declare one system correct, but to identify complementary detection evidence and define the validation needed before warning effectiveness, near-miss rates, or detector error rates can be estimated.

Objective 5: Evaluate economic feasibility and targeted deployment conditions

The fifth objective is to connect technical performance with economic and policy decisions. The project reconstructs a deterministic side-underride-guard benefit-cost baseline using a clearly defined target population and current USDOT valuation assumptions. It also calculates warning-system break-even thresholds and the degree of risk concentration required for targeted deployment. These quantities are analytical thresholds rather than measured warning effectiveness or proof that a particular fleet will break even.

Taken together, these objectives support the following project-level questions:

1. What capabilities and evidence should define a smart trailer for safety and deployment analysis?
2. Under what operating conditions does trailer-specific sensing provide information beyond tractor-only monitoring?
3. Under what calibration and image-support conditions can camera observations be interpreted as physical proximity?
4. What can be concluded when an independent CV pipeline and an installed warning display agree or disagree?
5. What cost, effectiveness, and risk-concentration conditions would be required for passive or active trailer-safety technologies to be economically justified?

The interpretation of the findings follows the scope of the supporting evidence. The naturalistic results describe the observed trailer, tractors, routes, and collection period and are interpreted as conditional operating associations rather than causal crash-risk estimates. Camera-based distance estimates are limited to the matching physical camera, unchanged image geometry, approximately planar roadway conditions, and surveyed pixel support. The retrospective CV audit is used to compare complementary evidence streams rather than to estimate population-level detector performance. The economic results represent deterministic benefit-cost estimates and break-even thresholds under explicitly defined intervention scopes and assumptions. Together, the five objectives establish an integrated evidence chain from trailer-specific sensing and operating-context analysis to camera-based proximity assessment, visual evidence comparison, economic evaluation, and implementation guidance.

1.6 Organization of the Report

The remainder of this report is organized as follows.

Chapter 2 presents the integrated smart-trailer research framework and connects passive protection, trailer-specific sensing, data integration, context-aware analysis, object and trajectory evidence, economic evaluation, and deployment decisions.

Chapter 3 describes the field deployment, tractor and trailer instrumentation, exterior-camera system, roadway and facility data, project-data inventory, synchronized analytical cohort, calibration survey, and dataset lineage.

Chapter 4 presents the environment-aware naturalistic analysis of tractor-trailer operating behavior, paired dynamic response, and mapped detector activity.

Chapter 5 documents the camera-based ground-localization methodology, strict camera-specific validation, computer-vision detection and tracking workflow, proximity and trajectory definitions, and the audit of disagreements between CV evidence and burned-in recorder boxes.

Chapter 6 presents the benefit-cost analysis of side underride guards, warning-system break-even thresholds, targeted-deployment conditions, and the limitations of combining passive and active intervention benefits.

Chapter 7 synthesizes the implications for fleets, safety-system developers, transportation agencies, and policymakers; documents project products and technology-transfer outcomes; and assesses progress relative to the original proposal.

Chapter 8 summarizes the principal conclusions and presents recommendations for implementation, validation, evidence integration, and targeted smart-trailer deployment.

Chapter 2. Integrated Smart-Trailer Safety and Evaluation Framework

This project continues the previous Safety21 research direction on dynamic safety-zone concepts, smart sensors, and vision-based assessment for heavy-duty tractor-trailer combinations (Jahanbiglari et al. 2025). The present project extends that foundation through naturalistic paired-unit analysis, camera-specific ground localization, retrospective comparison of computer-vision and recorder-alarm evidence, and deterministic economic assessment. The work was organized as an integrated evaluation rather than as an assessment of one device or intervention. Figure 1 summarizes the overall project flow. Naturalistic field data were first synchronized and quality-controlled, then examined through the context-aware vehicle, camera-localization and CV-audit, and economic workstreams. The resulting evidence was integrated to support trailer-aware monitoring priorities, proximity and alarm-evidence interpretation, and implementation and policy decisions. This chapter establishes the common framework that connects those workstreams and defines the limits of the conclusions that can be drawn from each one.




Figure 1. Overall project flow

2.1 Working Functional Definition of a Smart Trailer

The term *smart trailer* is used broadly in transportation practice and can refer to technologies serving substantially different purposes, including trailer location, cargo condition, tire pressure, brake status, lighting, rear visibility, blind-zone monitoring, stability assessment, or predictive maintenance. The Project proposal identified the absence of a clear and standardized definition as a barrier to technology comparison, performance evaluation, policy development, and implementation planning. It therefore called for a framework defining the essential components and capabilities of a smart trailer.

The project materials support a capability-based definition that connects sensing, measurement quality, contextual interpretation, event assessment, and decision support. Accordingly, this report adopts the following project definition:

“A smart trailer is a trailer equipped with an integrated sensing, communication, and analytical capability that can independently observe trailer state or surrounding conditions, relate those observations to tractor and operating context, and convert the combined evidence into information that supports safety, maintenance, operational, or policy decisions.”

This definition is capability-based rather than device-based. Installing an individual connected component may provide an important function, but the safety value of the overall system depends on whether the resulting measurements are valid, synchronized, interpreted in context, and translated into an appropriate decision. For example, a camera can extend visibility around the trailer, but an image detection alone does not establish physical clearance, collision trajectory, or

conflict severity. Likewise, an inertial sensor can measure motion accurately at its mounting location, but tractor measurements do not necessarily represent the response of the trailer during articulated operation. The project's naturalistic analysis found that ramp-classified operation changed the balance between tractor and trailer yaw and lateral activity, demonstrating why trailer-specific observation can add information beyond tractor-only telemetry.

For the purposes of the project, six connected capabilities define the smart-trailer framework.

2.1.1 Trailer-Specific Observation

A smart-trailer system should be able to observe conditions associated with the trailer rather than treating the tractor as a complete proxy for the articulated combination. Relevant observations may include trailer position, speed, acceleration, yaw rate, lateral response, brake or tire condition, surrounding objects, blind-zone occupancy, and rear or side clearance.

The field deployment used separate Global Navigation Satellite System and inertial measurements from the tractor and trailer, together with curb-side, driver-side, and rear video. Retaining these sources separately made it possible to compare trailer response directly with tractor response and to determine whether the information available from the towing unit fully represented the trailing unit. The resulting analysis identified ramp-classified and related merge/diverge contexts as priority settings for trailer-aware monitoring, sensor-performance assessment, and context-specific review.

2.1.2 Measurement Validity and Calibration

Sensor output must be connected to a defensible physical interpretation. This requirement is especially important for wide-angle cameras, where perspective and fisheye distortion cause spatial scale to vary substantially across the image. A fixed pixels-per-foot conversion would therefore be invalid.

The camera work conducted under this project developed separate empirical mappings from image pixels to local road-plane coordinates for the curb-side, driver-side, and rear views. The mappings were evaluated using complete-setup holdouts and strict nested validation. The driver-side mapping provided the strongest current performance, the curb-side mapping supported controlled use with an identified edge limitation, and the rear mapping remained provisional because its continuous-distance error and model-selection instability were substantially larger. The project therefore demonstrates that measurement quality must be assessed separately for each physical camera and that unsupported extrapolation should be rejected rather than converted into an apparently precise distance.

2.1.3 Temporal and Spatial Integration

Smart-trailer evidence commonly originates from systems with different sampling rates, timestamps, identifiers, and coordinate references. A valid analysis therefore requires temporal synchronization, spatial alignment, and explicit treatment of missing or unsupported data.

The naturalistic study synchronized tractor and trailer GNSS/IMU measurements, video-derived detector activity, roadway-network attributes, and facility context at one-second resolution before aggregating the observations into analysis blocks. No missing GNSS or IMU values were interpolated or carried across gaps. The camera analysis similarly required direct verification of camera identities, time-correct selection of survey frames, preservation of the original image geometry, and camera-specific coordinate definitions. These procedures are not merely preprocessing details; they determine whether two observations describe the same vehicle state, image location, or operating interval.

2.1.4 Operating-Context Interpretation

The meaning of a sensor observation depends on where and how the tractor-trailer is operating. Facility positioning, arterial travel, freeway operation, and ramp transitions represent different movement and speed regimes. A uniform threshold applied across all environments may identify expected adaptation as abnormal or overlook trailer-specific behavior that becomes important only in particular contexts.

This project established an arterial-referenced naturalistic baseline across facility, collector/local, arterial, freeway/interstate, and ramp-classified exposure. The analysis distinguished whether the vehicle was moving, how fast it traveled once moving, and how trailer-minus-tractor yaw and lateral activity differed by environment. The environmental estimates were interpreted as operating baselines rather than safety thresholds. A departure from the baseline does not by itself establish an unsafe event; it identifies a context or observation that may warrant closer review with object, geometry, loading, and operator-response information.

2.1.5 Object, Event, and Uncertainty Assessment

A smart-trailer system must distinguish raw detector output from stronger forms of safety evidence. The project materials support a hierarchy in which a displayed alarm box, an independently detected object, a calibrated proximity estimate, a trajectory-based candidate, and an adjudicated near miss represent different evidentiary levels.

The retrospective near-miss audit compared two independent visual streams: a YOLO11n segmentation and BoT-SORT tracking pipeline with calibrated ground-distance estimation, and colored alarm rectangles burned into the recorder video. Across 156 processed videos, the CV pipeline produced 6,637 calibrated trajectory points and 2,270 track summaries. Thirty supported tracks from 21 events entered the project's 0-9-ft proximity screen; eight tracks had a machine-matched recorder box and 22 did not. The reverse disagreement also occurred: clear recorder boxes appeared around some objects that the general-purpose CV model did not match at the sampled frame. These results show that neither visual stream should be treated as complete ground truth by itself.

For this reason, the report distinguishes among a recorder system output, a CV-supported proximity event, a trajectory or time-to-contact candidate, and a validated near miss. The audit

documents recorder-system outputs and CV-supported proximity events and identifies one provisional candidate under the original trajectory/TTC rule. It is therefore interpreted as an evidence-comparison and engineering-audit study rather than as a population-level detector-performance or near-miss-rate analysis.

2.1.6 Decision Support

The final capability of a smart trailer is the translation of technical evidence into an operational, economic, or policy decision. Technical performance alone does not establish that a technology should be deployed nationally, installed on every trailer, or used to generate driver-facing warnings. Deployment decisions also require information about target risk, intervention effectiveness, capital and recurring costs, sensor availability, nuisance-alert burden, maintenance, operating constraints, and the concentration of relevant risk across fleets, routes, and times.

The corrected benefit-cost analysis developed under this project provides a transparent national benchmark for side underride guards, break-even effectiveness thresholds for warning systems, and mathematical conditions for targeted deployment. The analysis intentionally treats warning effectiveness as an unknown threshold rather than importing an effectiveness estimate from a different application. The model quantifies the concentration of relevant risk that a treated fleet subset would need to capture for the modeled benefits to reach break-even under proportional treatment-cost scaling. These values are deployment-screening thresholds rather than evidence that a particular fleet, route, or operating environment already satisfies them. Table 2 summarizes the functional capabilities and the corresponding project evidence.

Table 2. Functional capabilities and the corresponding project evidence

Functional capability	Central question	Project evidence	Interpretation and scope
Trailer-specific observation	Does the system observe the trailer independently of the tractor?	Separate tractor/trailer GNSS and IMU; three trailer-mounted camera views	One trailer and two tractors; broader fleet generalization remains necessary
Measurement validity	Can the output be interpreted physically?	Camera-specific pixel-to-ground mappings with strict nested validation	Driver and curb support controlled use; rear continuous distance remains provisional
Temporal and spatial integration	Do the observations refer to the same time, location, and coordinate system?	One-second synchronization, map matching, facility	Sub second clock drift and complete full-archive exposure remain unresolved

Functional capability	Central question	Project evidence	Interpretation and scope
		geofences, camera-specific frames	
Operating-context interpretation	Is the observation expected for the roadway or facility environment?	Five-environment naturalistic baseline	Conditional associations, not causal crash-risk estimates
Object and event assessment	Was an object present, close, approaching, and safety-relevant?	CV tracking, calibrated distance, trajectory/TTC logic, recorder-box audit	No independent population truth set or validated near-miss rate
Decision support	What cost, effectiveness, and targeting conditions justify deployment?	Side-guard BCA, warning break-even thresholds, targeted-risk conditions	Warning effectiveness and fleet-specific risk concentration are not yet measured

The resulting definition treats “smartness” as a property of the complete evidence chain rather than the number of installed devices. A system becomes more useful for safety decision-making when its measurements are independently meaningful, its limitations are known, and its outputs can be interpreted in relation to trailer state, roadway environment, nearby objects, and deployment consequences.

2.2 Layered Passive and Active Safety Architecture

Within the integrated project flow shown in Figure 1, passive and active trailer-safety measures are treated as complementary layers that address different stages of a hazardous interaction. The original proposal includes the evaluation of side and rear underride guards, proactive warning systems, sensing and telematics, predictive analysis, and policy or deployment alternatives. The project materials do not support the conclusion that one of these approaches can replace all others. Instead, they support a layered architecture in which structural protection, active observation, evidence validation, and deployment decision-making perform different functions.

2.2.1 Passive Structural Protection

Passive trailer-safety measures are intended primarily to mitigate consequences once physical contact begins or can no longer be avoided. Side and rear underride guards are examples of this layer. Their intended function is structural: to reduce the extent to which a smaller vehicle can travel beneath the trailer in specified crash configurations.

The project's benefit-cost work evaluates side guards using the narrowly defined NHTSA light-passenger-vehicle front-to-side target population and the documented performance range of the evaluated guard. It does not claim that guards are ineffective. Rather, it distinguishes conditional safety performance from the economics of equipping a broad national trailer cohort. The corrected analysis finds that the monetized benefits remain below the costs of cohort-wide deployment under the tested target population and assumptions, even though avoided fatalities and serious injuries have substantial social value.

This distinction is important for policy. A structural countermeasure may provide meaningful protection in the crash configurations for which it is designed while still having a benefit-cost ratio below one when its cost is distributed across a much larger population of trailers. The appropriate interpretation is therefore not that structural protection has no value, but that deployment scope, eligible crash population, operating compatibility, lifecycle cost, and risk concentration materially affect the economic conclusion.

2.2.2 Active Observation and Early Recognition

Active measures are intended to identify relevant conditions before contact or before the interaction becomes unavoidable. In this project, this layer includes trailer-mounted cameras, recorded blind-zone activity, tractor and trailer telemetry, computer-vision object detection and tracking, calibrated ground-distance estimation, and trajectory analysis.

The naturalistic analysis showed why active monitoring should retain both trailer state and operating environment. Ramp-classified operation produced the clearest trailer-specific departure from the arterial baseline: the trailer exhibited lower yaw activity but higher lateral activity relative to the tractor. This finding does not establish that ramps have greater crash risk. It identifies ramp-classified and related merge/diverge operation as contexts in which tractor-only measurements may be least complete and where trailer-aware sensing and validation should be prioritized.

The visual audit similarly showed that active sensing is not equivalent to a single detector. The project CV pipeline identified supported close objects without a matching recorder box, while the recorder displayed boxes around some boundary-truncated or occluded objects that YOLO did not match. The correct engineering response is not to select one stream as universally correct. It is to use the streams as complementary proposals, preserve the evidence from either stream, and apply additional temporal, geometric, and manual validation before assigning a safety interpretation.

2.2.3 Context and Evidence Validation

A third layer is required between sensing and intervention. A detector activation should not directly trigger a claim of conflict, near miss, or safety benefit without confirming the conditions under which the output was produced.

This project demonstrates several forms of validation required at this layer:

- Tractor and trailer observations must be synchronized before their responses are compared.

- Roadway environment must be assigned independently of the modeled outcomes.
- Camera measurements must use the correct physical camera, original image geometry, and supported pixel region.
- A detected object must have a defensible ground-contact point before a physical distance is calculated.
- Distance estimates must retain camera-specific uncertainty and must not be extrapolated outside the surveyed hull.
- Proximity must be distinguished from closing motion, trajectory conflict, and an adjudicated near miss.
- Recorder boxes and CV detections must be checked for extraction and association failures before disagreement is labeled as a detector miss.

These requirements explain why the project separates a warning-system output from a safety outcome. In the tractor-trailer behavior analysis, mapped detector activity was associated more clearly with lower-speed operating context than with a detectable change in paired tractor-trailer yaw or lateral response. In the CV audit, close objects and recorder boxes disagreed in both directions. Together, the results show that warning activity should be interpreted with monitored zone, speed, roadway environment, trailer state, object evidence, and exposure rather than as an isolated count.

2.2.4 Intervention and Decision Layer

The final layer determines whether the evidence warrants an operator warning, fleet review, targeted pilot, structural investment, regulatory action, or additional data collection. The completed analyses support post-trip review, sensor-performance assessment, operating-context prioritization, and economic threshold setting. Because the vehicle and CV analyses are observational and retrospective, they are not used as causal estimates of real-time warning effectiveness or crash reduction.

For active systems, the benefit-cost analysis expresses total warning effectiveness as a sequence of necessary components: system availability, object detection, valid-alert probability, driver or vehicle response, and crash avoidance conditional on that response. Weakness in any component reduces the overall causal effect. This decomposition connects the technical work to the economic model and clarifies why detector accuracy alone is insufficient for estimating monetary safety benefit.

For passive systems, the decision layer considers whether the severity reductions expected in the target crash population justify hardware, installation, fuel, maintenance, and operational costs. For combined passive and active deployment, the benefits must be partitioned carefully. A crash avoided through an effective warning is no longer available to be mitigated by a guard.

Accordingly, combined benefit accounting should distinguish crashes avoided by active intervention, residual crashes whose severity is mitigated by structural protection, and operational benefits that do not overlap with either safety mechanism.

The project findings support a staged implementation strategy that prioritizes environments in which trailer-specific measurements add information, uses the most defensible camera mappings, preserves both CV and recorder evidence, and collects comparable exposure and non-alarm observations.

In this framework, passive protection, active sensing, and analytical validation are not competing alternatives. Passive systems provide structural redundancy when avoidance is unsuccessful. Active systems provide earlier evidence about nearby objects and trailer behavior. Contextual and methodological validation determines whether those observations are interpretable. Economic and policy analysis determines where an intervention can reasonably be deployed. The integrated function of these layers, not any single sensor or model, is the central smart-trailer contribution of this project.

Chapter 3. Field Deployment, Data Sources, and Dataset Lineage

3.1 Study Design and Deployment Scope

The project combined a naturalistic field study, a broader video and alarm inventory, a controlled camera-ground survey, a retrospective computer-vision audit, and an economic modeling workstream. These components used different units of observation and different date ranges. The distinctions among them are important because a trip, a synchronized second, a detector episode, a video clip, a computer-vision track, and a unique physical interaction are not interchangeable analytical units.

The naturalistic component used an observational design. Routes, speeds, maneuvers, traffic conditions, and operating environments were observed during routine fleet service and were not assigned by the research team. One trailer was operated with either of two tractors during 13 trips conducted on 10 collection dates from March 30 through April 17, 2026. The synchronized record represented 56 hours of operation and 201,600 unique one-second observations.

The complete synchronized exposure was analyzed rather than selecting observations because an alert or unusually large response occurred. This exposure-wide design distinguishes the naturalistic study from the broader alarm-centered archive. It permits movement, speed, and paired vehicle response to be interpreted relative to the full observed operating record.

A broader project inventory documented 28 calendar dates between March 24 and May 22, 2026. That archive included 1,748 unique alarm records, 5,487 readable video clips, 1,478 alarm records

with coordinates, and 1,282 records with alarm boxes. Because this inventory was organized around alarms and available clips, it was used to characterize the available evidence and support focused review; it was not treated as a complete denominator for estimating warning rates, missed-event rates, or causal safety effectiveness.

The camera-ground survey was designed to relate original image pixels from the rear, curb-side, and driver-side cameras to camera-local coordinates on an approximately planar road surface. The retrospective CV audit then used saved April 17 and April 29 videos and project-generated trajectory outputs. April 17 overlaps the synchronized naturalistic cohort, while April 29 belongs to the broader archive and was not automatically assigned the roadway-environment labels developed for the March 30-April 17 analytical cohort.

3.2 Data Sources and Analytical Roles

The integrated project used physically and administratively separate data systems. Table 3 summarizes the recorded information and the principal role of each source.

Table 3. Project data sources and analytical roles

Data source	Recorded or derived information	Principal analytical role
Tractor GNSS and IMU	Position, ground speed, course, acceleration, and angular rate	Trip identification, movement and speed outcomes, roadway assignment, tractor yaw response, and tractor lateral response
Trailer GNSS and IMU	Position, ground speed, course, acceleration, and angular rate	Temporal pairing and trailer yaw and lateral response
Exterior video	Rear, curb-side, and driver-side views	Visual context, object evidence, ground-contact estimation, and review of system activations
Recorder display states	Burned-in red, yellow, and green alarm rectangles and display activity	Detector-derived system-state episodes and comparison with independent CV evidence
PennDOT RMSSEG	Roadway geometry, route identity, functional attributes, and ramp/link information	Independent assignment of broad roadway environment and recurrent roadway features
Operational-facility geofences	Four predefined parking, staging, and industrial locations	Identification of facility and low-speed operational context

County and urban-rural context	County boundaries and six-level county classification	Exploratory exposure-normalized description of alert occurrence
Camera-ground survey	Surveyed mat corners, checkerboard imagery, camera identity, and setup times	Camera-specific pixel-to-road-plane mapping and strict validation
Regulatory and economic sources	Target crash populations, guard performance, equipment and fuel costs, valuation guidance, and discount rates	Side-guard BCA, warning break-even thresholds, and targeted-deployment screening

3.3 Exterior-Camera Configuration

Three exterior cameras recorded the rear, curb-side, and driver-side regions of the trailer. Camera identity was verified from the burned-in labels rather than inferred from file order. The video data sample is shown in Figure 2.



Figure 2. Representative exterior camera views: (a) driver side, (b) rear, and (c) curb side

The video displays also contained detector-generated visual states. These displayed states were extracted and consolidated into monitored-zone episodes. A display activation was treated as a system output. It did not independently establish the presence, identity, trajectory, or criticality of an external object.

3.4 Dataset Lineage and Units of Observation

Table 4 distinguishes the main analytical layers. The counts should not be compared without considering the unit and inclusion rules.

Table 4. Dataset lineage and analytical units

Analytical layer	Unit of observation	Documented size	Primary use
Broader alarm/video inventory	Alarm records and readable clips	1,748 unique alarms; 5,487 readable clips; 1,478 coordinate-linked alarms; 1,282 with alarm boxes	Evidence inventory, data-gap assessment, and case-study source
Synchronized naturalistic cohort	Trips and synchronized seconds	13 trips; 10 dates; 56 h; 201,600 seconds	Exposure-wide operating and paired-response analysis
Naturalistic detector episodes	Consolidated monitored-zone episodes	729 episodes	Secondary operating context within synchronized exposure
Camera-ground survey	Camera setups, mats, and surveyed validation centers	Six setups and four mats per setup for each camera	Pixel-to-ground mapping and camera-specific validation
Retrospective CV audit	Videos, frame-level metric points, tracks, and events	156 videos; 6,637 trajectory points; 2,270 track summaries; 52 alarm events	CV/recorder evidence comparison and proximity screening

The 1,748 alarm records, 729 consolidated episodes, and 30 supported CV proximity tracks represent different processing levels. A raw alarm record can contain multiple clips or display states. Temporally adjacent display activity can be consolidated into one episode. A single physical interaction can generate more than one CV track because of object fragmentation or track reinitialization. For this reason, track counts are reported alongside unique-event counts.

The broader archive is useful for documenting the available project evidence, but alarm-selected clips cannot by themselves estimate event prevalence or warning-system performance. Estimating rates requires a corresponding exposure denominator, such as hours, miles, trips, route segments, or sensor-observable seconds, and a comparison sample that includes non-alarm periods.

3.5 Temporal Synchronization and Trip Construction

Tractor and trailer GNSS and IMU measurements were timestamped in UTC nanoseconds. Observations were assigned to integer UTC seconds by flooring the epoch timestamp. Within each second:

- Valid GNSS position, speed, and course values were summarized by their median.
- Primary IMU acceleration and angular-rate channels were summarized by their mean.

- Source-specific summaries were joined on the union of observed seconds within each trip.
- No GNSS or IMU value was interpolated, forward-filled, or carried through a missing interval.

A new trip was identified after a temporal gap greater than 60 seconds, a change in local collection date, or a change in assigned tractor. Time-difference quantities were calculated only across consecutive seconds in the same trip, preventing data gaps and trip boundaries from being interpreted as physical motion.

Video timestamps were localized to America/New_York and converted to UTC before matching. Frame-level detector states were consolidated into episodes, and every synchronized second intersected by an episode was marked active for the corresponding zone. Because the analytical layer did not retain an independent display-observability flag for every second, an alert value of zero means that no detector-derived episode was mapped; it does not necessarily confirm that the display was visible and inactive.

Before paired tractor-trailer analysis, corresponding IMU channels were transformed into a common vehicle-fixed convention. Temporal consistency was evaluated using GNSS fixes matched within 1.5 seconds. Accepted hourly pairings had a median absolute timestamp difference of 0.0 seconds, and no systematic time shift was applied.

3.6 Analytical Blocks and Eligibility

The synchronized record was partitioned within each trip into complete, nonoverlapping 10-second blocks aligned to absolute UTC boundaries. The 10-second interval balanced preservation of local operating context against short-term variability and serial dependence.

A block was eligible for roadway-context analysis when at least 80% of its seconds had high- or medium-quality roadway assignments or when its modal context was a predefined facility. Low-quality, unmatched, and missing roadway assignments counted against the 80% threshold.

The movement-likelihood and speed-when-moving analyses required at least eight valid tractor-speed observations per block. Paired yaw and lateral analyses required at least eight timestamps with simultaneous valid tractor and trailer measurements for the corresponding channel. Unit-specific paired summaries were calculated from identical timestamps. The overall workflow for naturalistic analysis is shown in *Figure 3*.




Figure 3. Naturalistic-analysis workflow and analytical cohorts

3.7 Operating-Environment Assignment

Facility context was assigned first using four predefined geofences. Non-facility observations were then spatially matched to the archived Pennsylvania Department of Transportation RMSSEG State Roads network, which provided the roadway geometry and functional attributes used in the operating-environment classification (Pennsylvania Department of Transportation 2025). PennDOT polylines were subdivided into pieces no longer than 35 m, and up to 24 nearby candidates were evaluated for each GNSS observation.

For candidate roadway feature (j) at second (t), the map-matching score was

$$S_{tj} = d_{tj} + \lambda \min(\Delta\theta_{tj}, 60^\circ)$$

where:

- d_{tj} is the approximate perpendicular distance in meters;
- $\Delta\theta_{tj}$ is the direction-insensitive difference between vehicle course and roadway-axis orientation; and
- $\lambda = 0.22$ m/degree.

The angular contribution was capped at 60° , limiting its maximum contribution to 13.2 m. Candidate ranking used distance alone below 5 km/h or when course was unavailable. Observations more than 80 m from the selected feature were rejected.

Assignments were classified as high quality when perpendicular distance was no greater than 20 m and the margin from a distinct alternative feature was at least 5 m, or when no finite alternative existed. Remaining assignments within 40 m were medium quality, those within 80 m were low quality, and more distant observations were unmatched.

The final broad categories were:

1. Facility
2. Collector/local
3. Arterial
4. Freeway/interstate
5. Ramp-classified exposure

Broad categories were emphasized because they were more defensible under the available GNSS and network resolution than lane-level or carriageway-specific labels. The term ramp-classified is retained because exact ramp/link and travel-direction interpretations were not manually verified for every observation.

Chapter 4. Environment-Aware Tractor-Trailer Operating Baseline

4.1 Analytical Questions

The naturalistic analysis addressed one overarching question: how do tractor-trailer operating behavior and paired dynamic response vary across distinct operating environments?

The overarching question was separated into three analytical questions provided in Table 5 to avoid conflating normal speed adaptation, trailer-specific response, and detector activity.

Table 5. Naturalistic research questions and interpretation

Question	Primary analysis	Interpretation
RQ1	Movement likelihood and speed when moving across facility, collector/local, arterial, freeway/interstate, and ramp-classified exposure	Environment-specific operating behavior and speed adaptation
RQ2	Same-second trailer-minus-tractor mean absolute yaw-rate and lateral specific-force activity, with arterial exposure as the reference	Speed-adjusted trailer response relative to the tractor; not formal rearward amplification
RQ3	Monitored-zone composition of mapped detector episodes and associations of alert-active fraction with speed and paired response	Secondary system-output context; not a verified conflict, near miss, or calibrated warning-performance measure

Arterial operation was selected as the reference because it was well represented and provided an intermediate public-road regime between lower-speed facility and collector/local operation and higher-speed freeway/interstate travel.

4.2 Movement and Speed Outcomes

For block b , the moving-state indicator was defined as

$$M_b = \mathbb{I}(\bar{v}_b \geq 5 \text{ km/h})$$

where $\bar{v}_b$ is the mean tractor ground speed among valid observations in the block. The 5 km/h threshold separated stationary or crawling operation from sustained movement.

The analysis first modeled whether a block was moving and then modeled mean speed conditional on movement. This two-component structure was especially important for facility exposure, where extended stationary and low-speed intervals occurred (Farewell et al. 2017).

4.3 Paired Tractor-Trailer Response

For each 10-second block, yaw-rate and lateral specific-force responses were calculated separately for the tractor and trailer. Only timestamps for which both units had valid measurements of the corresponding motion component were included. The block-level response for each unit was then calculated as the mean absolute value of its measurements over those common timestamps.

$$R_{bk}^c = \frac{1}{n_{bk}^\cap} \sum_{t \in T_{bk}^\cap} |x_{tk}^c|$$

In the equation above, R_{bk}^c represents the mean absolute response of vehicle unit, either the tractor or the trailer, for motion component k, either yaw rate or lateral specific force, during block b. The term $T_{bk}^\cap$ represents the common timestamps at which both the tractor and trailer had valid measurements, and $n_{bk}^\cap$ is the number of those timestamps. The variable x_{tk}^c is the recorded motion measurement for unit c at timestamp t. Absolute values were used so that the response represented the magnitude of motion without distinguishing between leftward and rightward directions.

The trailer-relative response was calculated by subtracting the tractor response from the trailer response:

$$\Delta R_{bk} = R_{bk}^{\text{trailer}} - R_{bk}^{\text{tractor}}$$

Positive values indicate greater trailer-relative activity, negative values indicate greater tractor-relative activity, and values near zero indicate similar activity. Yaw responses were expressed in degrees per second and lateral specific force in (g).

Eligible paired observations received the coverage weight

$$w_{bk} = \frac{n_{bk}^\cap}{10}$$

The paired differential is not formal rearward amplification. Rearward amplification generally uses maneuver-specific trailer-to-tractor peak ratios, while ΔR_{bk} represents a difference in average dynamic activity during short naturalistic exposure blocks.

4.4 Detector-Activity Measure

Overall detector activity for block (b) was defined as the fraction of synchronized seconds intersecting at least one mapped curb-side, driver-side, or rear episode:

$$A_b = \frac{1}{10} \sum_{t \in b} I_t$$

where $I_t=1$ when any monitored zone was active during second (t). Simultaneous activation of multiple zones contributed one active second to A_b . The variable was interpreted as a detector-derived system-state measure rather than evidence of conflict severity.

4.5 Statistical Models

The movement model was

$$\text{logit}\{P(M_b = 1)\} = \alpha_M + \mathbf{E}_b^\top \boldsymbol{\beta}_M + \mathbf{C}_b^\top \boldsymbol{\gamma}_M + \mathbf{Z}_b^\top \mathbf{u}_M$$

where:

- $\mathbf{E}_b$ contains operating-environment indicators;

- $\mathbf{C}_b$ contains unpenalized adjustment terms; and
- $\mathbf{Z}_b$ contains trip and recurrent-roadway-feature indicators.

Environment coefficients were exponentiated and reported as odds ratios. Adjusted differences in predicted movement probability were also calculated relative to arterial exposure.

Speed when moving and the paired yaw and lateral outcomes were estimated using coverage-weighted Huber regression, which limits the influence of unusually large residuals while retaining the observations in the analytical sample (Sun et al. 2020).

$$Y_b = \alpha + \mathbf{E}_b^\top \boldsymbol{\beta} + \eta_1 \tilde{v}_b + \eta_2 \tilde{v}_b^2 + \theta \tilde{A}_b + \mathbf{C}_b^\top \boldsymbol{\gamma} + \mathbf{Z}_b^\top \mathbf{u} + \varepsilon_b$$

For the speed model, ($Y_b = \bar{v}_b$) and the speed-adjustment terms were omitted. For paired motion, ($Y_b = \Delta R_{bk}$), and standardized mean tractor speed and its square were included. The primary environment models omitted alert-active fraction; secondary models included it to describe concurrent association rather than causal mediation.

The models adjusted for daylight condition, local time period, tractor identity, collection date, and standardized time since trip start. Trip and recurrent roadway-feature indicators were ridge penalized to stabilize unequal observational support.

Huber loss used the conventional tuning constant of 1.345, which provides high efficiency under approximately normal residuals while reducing the influence of large residuals (Sun et al. 2020). Primary uncertainty was evaluated at the collection-date level because observations recorded on the same date were not treated as statistically independent. Delete-one-date jackknife standard errors were used with term-specific Student-t confidence intervals. A complementary date-cluster bootstrap requested 999 resamples and retained converged model fits with finite coefficients (MacKinnon et al. 2025, 2023).

4.6 Analytical Cohorts and Environmental Support

The 201,600 synchronized seconds formed 20,160 complete 10-second blocks. Table 6 reports the progression to each model-specific cohort.

Table 6. Naturalistic analytical cohorts

Analytical stage	Blocks
Complete nonoverlapping 10-s blocks	20,160
Environment-supported blocks	15,818
Moving-probability model	14,151

Speed-when-moving model	5,676
Same-second paired yaw model	11,469
Same-second paired lateral model	11,469

The 15,818 environment-supported blocks included:

- 9,998 facility blocks (63.2%)
- 4,135 arterial blocks (26.1%)
- 1,400 freeway/interstate blocks (8.9%)
- 181 collector/local blocks (1.1%)
- 104 ramp-classified blocks (0.7%)

The paired-response sample contained 5,969 facility, 3,892 arterial, 1,342 freeway/interstate, 173 collector/local, and 93 ramp-classified blocks. The lower ramp and collector/local support is reflected in wider uncertainty and more cautious interpretation.


4.7 Environment-Dependent Movement and Speed

Facility exposure was the only environment with clearly lower movement likelihood than arterial exposure: odds ratio [OR] = 0.00885; 95% confidence interval [CI]: 0.00323-0.02425.


Point estimates for collector/local, freeway/interstate, and ramp-classified movement likelihood differed from arterial operation, but their confidence intervals crossed the no-difference reference.

All four nonarterial environments differed clearly from arterial operation in speed among moving blocks. Figure 4 summarizes the operational behavior and the vehicle dynamic responses.


A. Likelihood of Moving



B. Speed When Moving



C. Trailer-Relative Yaw Activity



D. Trailer-Relative Lateral Activity




Figure 4. Adjusted differences from the arterial operating baseline

The speed contrasts establish that the five environments represent meaningfully different operating regimes. They should not be interpreted as rankings of crash risk. Relative to arterial operation, freeway/interstate speed when moving was 19.64 km/h higher (95% CI: 17.39–21.89 km/h), whereas ramp-classified speed was 25.06 km/h lower (95% CI: –32.33 to –17.79 km/h).

4.8 Paired Tractor-Trailer Dynamic Response

Ramp-classified exposure produced the clearest trailer-specific departure from arterial operation. The trailer-minus-tractor mean absolute yaw-rate differential was lower, while the corresponding lateral specific-force differential was higher. Relative to arterial operation, the ramp-classified trailer-minus-tractor yaw-rate differential was –0.1444 deg/s (95% CI: –0.2518 to –0.0370 deg/s), while the lateral specific-force differential was 0.0172 g higher (95% CI: 0.0027–0.0317 g).

The opposing directions are important. The trailer did not become uniformly more or less dynamic. Lower trailer-relative rotational activity coexisted with higher trailer-relative side-to-side activity. Tractor yaw alone would therefore provide an incomplete description of the trailer’s lateral response under the observed ramp-classified conditions.

The contrasts remained after adjustment for tractor speed and its squared value. The result is not summarized solely by the lower speed chosen during ramp-classified operation. It remains

observational, however, because articulation angle, steering input, loading, lane position, detailed curvature, traffic interactions, and prior maneuvers were not experimentally controlled.

4.9 Mapped Detector Activity

Within the synchronized cohort, 729 consolidated detector-derived episodes were classified by monitored zone indicated in Figure 5.




Figure 5. Mapped detector-episode composition within the synchronized analytical cohort

A 0.10 increase in alert-active fraction, equivalent to one additional mapped alert-active second within a complete 10-second block, was associated with a 2.736 km/h lower speed among moving blocks (95% CI: -4.321 to -1.151 km/h). In contrast, the estimated associations between alert-active fraction and paired tractor-trailer motion were small and uncertain. The association with the trailer-minus-tractor yaw-rate differential was -0.00156 deg/s (95% CI: -0.00421 to 0.00108 deg/s), while the association with the trailer-minus-tractor lateral specific-force differential was -0.000694 g (95% CI: -0.001812 to 0.000423 g). Because both paired-response confidence intervals included zero, the available data did not provide clear evidence that mapped alert activity was associated with a change in the relative yaw or lateral response of the tractor and trailer. These estimates describe concurrent associations and should not be interpreted as evidence that an alert caused a speed reduction or represented a conflict or near miss.

4.10 Robustness and Sensitivity

The paired-response findings were evaluated using mean absolute response, root-mean-square response, and the absolute value of the signed block mean. The ramp-classified yaw contrast remained negative and the lateral contrast remained positive across all definitions. Same-second

pairing retained the full paired sample and produced coefficients numerically indistinguishable from the separate-count approach.

The principal ramp and conditional-speed interpretations also remained substantively consistent across complete 5-, 10-, and 20-second blocks as shown in Figure 6. Broad environment assignments were stable across map-matching heading weights from 0 to 1.00 m/degree, with broad-category agreement ranging from 99.76% to 99.93%.




Figure 6. Robustness of adjusted ramp-classified paired-response contrasts relative to arterial operation

4.11 Transportation Interpretation

The naturalistic results provide operating baselines, not safety thresholds. A departure from the fitted baseline does not independently indicate unsafe operation. Safety interpretation requires additional information regarding road geometry, articulation, loading, surrounding objects, lane position, and operator response.

The transportation value of the ramp finding is that it identifies a context in which tractor-only monitoring may be least complete. Ramps, merges, and diverges combine curvature, speed adaptation, lane-position changes, articulation, and interactions with adjacent traffic. Trailer-mounted inertial sensing can therefore complement tractor telemetry when fleets, developers, or agencies evaluate these roadway transitions.

Raw alert totals should likewise be stratified by monitored zone, movement state, speed, roadway environment, facility context, and exposure duration. A large number of activations may represent recurrent normal proximity, a broad detection zone, an image or association problem, a demanding interaction pattern, or a mixture of those factors.

Chapter 5. Camera-Based Ground Localization and Computer-Vision Evidence Audit

5.1 Camera-Ground Localization Objective

The calibration objective was to transform an observed ground-contact pixel from each physical camera into camera-local road-plane coordinates and a camera-appropriate clearance distance. For camera (c), the mapping is

$$\hat{\mathbf{g}} = F_c(\mathbf{p}), \quad \mathbf{p} = \begin{bmatrix} u \\ v \end{bmatrix}, \quad \hat{\mathbf{g}} = \begin{bmatrix} \hat{X} \\ \hat{Y} \end{bmatrix}$$

where (u,v) is the original 1280×720 image pixel and $(\hat{X}, \hat{Y})$ is the estimated location on the camera-local road plane.

Separate models, supported regions, validation predictions, and error statements were retained for the rear, curb-side, and driver-side cameras. Observations were not combined across cameras.

5.2 Survey Design and Coordinate Systems

Each camera had six survey setups and four checker mats per setup. Each mat had a 5×7 square pattern, corresponding to 24 internal checker intersections after orientation. Two opposite outer corners of each mat were surveyed.

The common trailer frame placed its origin at the rear curb-side trailer corner on the ground. Global (X) pointed forward, global (Y) pointed from curb side toward driver side, and global (Z) pointed upward. Trailer width was not required to estimate each camera's local clearance independently. It would be required for common global lateral coordinates, cross-camera fusion, or complete trailer and object footprints.

For each mat, the two surveyed opposite corners and the known 160:120 aspect ratio defined the road-plane rectangle. Derived internal ground-control points were calculated as

$$\mathbf{G}_{ij} = \mathbf{B} + \frac{s_i}{160} W \mathbf{e}_x + \frac{t_j}{120} H \mathbf{e}_y$$

where $\mathbf{B}$ is a rectangle reference corner, W and H are the rectangle dimensions scaled to the measured diagonal, and s_i and t_j describe the internal checker coordinates.

The derived internal corners increased spatial constraint but did not represent 24 independent physical surveys. Validation therefore used independently surveyed mat centers with complete setups held out.

5.3 Frame Selection and Quality Control

Frame selection was based on decoded absolute time near the documented survey setup time. This prevented sharp but physically incorrect frames, such as frames captured while a mat was being moved, from being paired with a survey measurement describing a different placement.

A checker observation was accepted only when:

- All 24 internal corners were detected.
- Ordering and grid-geometry checks passed.
- Region-of-interest coordinates were restored to the original image system.
- The board was not materially occluded.
- The frame and survey described the same physical setup.
- The original 1280×720 image geometry was preserved.

Curb setup 5, mat 2 was retained as a held-out validation target but excluded from final training because it occupied only 382 pixels squared in a difficult upper-left region. Rear setups 3 and 4 remained resolution limited, with too few complete boards at the documented survey time.

5.4 Mapping Methods

Multiple empirical mapping families were evaluated independently for each camera, including affine mapping, robust planar homography, radial-warp homography, quadratic and cubic polynomial surfaces, ridge-regularized polynomial surfaces, rational quadratic mapping, thin-plate splines, piecewise affine interpolation, local robust inverse-distance weighting, and fixed ensembles. Ridge-regularized polynomial models were included to stabilize coefficient estimation when the expanded polynomial features were correlated or weakly supported (Hoerl and Kennard 1970). Thin-plate splines were also evaluated as smooth scattered-data mappings capable of representing nonlinear spatial deformation between image and road-plane control points (Bookstein 1989).

For standardized image feature matrix (Φ), ground-coordinate matrix (G), penalty matrix (P), and ridge coefficient (λ), regularized coefficients were estimated as

$$\hat{\mathbf{B}} = (\Phi^T \Phi + \lambda \mathbf{P})^{-1} \Phi^T \mathbf{G}$$

The curb-side frozen development model used a cubic ridge surface with ($\lambda = 1$). The driver-side model used a mean ensemble of smooth models. The rear development model used a component wise median ensemble.

5.5 Physical Fisheye Calibration Assessment

Planar-target camera calibration requires observations that provide sufficient variation in target orientation relative to the camera, while fisheye cameras require a projection model appropriate for wide-angle image formation (Kannala and Brandt 2006; Zhang 2000). Zhang’s method uses repeated observations of a planar calibration pattern at different orientations, and the Kannala-Brandt model provides a generic formulation for conventional, wide-angle, and fisheye lenses.

A conventional OpenCV fisheye calibration and a constrained fisheye-plus-pose fit were attempted separately for each camera. Neither approach passed the project’s acceptance checks.

All checkerboards were positioned on essentially the same physical road plane. Translation and in-plane rotation did not provide sufficient independent three-dimensional orientation diversity to separate focal length, distortion, and camera pose reliably. A low reprojection error alone was therefore not accepted as proof of physically meaningful intrinsics.

The accepted engineering product is an empirical fixed-camera pixel-to-road-plane mapping restricted to the surveyed image hull. It is not a general intrinsic calibration transferable to a moved camera, changed image-processing pipeline, nonplanar terrain, or arbitrary three-dimensional point.

5.6 Validation Design

Randomly splitting individual checker corners would place highly correlated observations from one physical mat into both training and testing sets. The workflow instead held out every corner and center from one complete setup:

$$T_s = \{i: \text{setup}(i) \neq s\}, \quad V_s = \{i: \text{setup}(i) = s\}$$

Approximately 45 model and quality-policy combinations were considered during development. Reporting performance from the same folds used to select the model would produce an optimistically biased estimate of the complete selection procedure (Varma and Simon 2006). The primary accuracy result therefore used nested leave-one-complete-setup-out validation: the outer setup was excluded from both model selection and fitting, and model selection occurred only within the remaining setups.

For measured coordinate (X, Y) and prediction ($\hat{X}, \hat{Y}$),

$$\Delta X = \hat{X} - X, \quad \Delta Y = \hat{Y} - Y, \quad e = \sqrt{(\Delta X)^2 + (\Delta Y)^2}$$

The principal error summaries were

$$\text{MAE}_{2D} = \frac{1}{n} \sum_{i=1}^n e_i, \quad \text{RMSE}_{2D} = \sqrt{\frac{1}{n} \sum_{i=1}^n e_i^2}$$

Directional bias retained sign. A negative alarm-axis bias means predicted clearance was smaller than measured clearance. This direction is conservative relative to a distance boundary, although excessive conservatism may also increase unnecessary warnings.

5.7 Safety-Zone Definition

The project used three road-plane distance categories:

$$\text{zone}(d) = \begin{cases} \text{red}, & 0 \leq d < 5 \text{ ft}, \\ \text{yellow}, & 5 \leq d < 9 \text{ ft}, \\ \text{outside}, & d \geq 9 \text{ ft}. \end{cases}$$

For side cameras, (d) is perpendicular outward clearance. For the rear camera, (d) is perpendicular rearward distance. Figure 7 indicates the distances mentioned above.




Figure 7. Camera-specific projected ground grids and working 5-ft and 9-ft boundaries

5.8 Calibration Results

Table 7 reports the primary strict nested-validation results.

Table 7. Strict nested camera-ground-localization performance

Camera	Validation centers	2-D MAE (ft)	2-D RMSE (ft)	2-D P95 (ft)	Maximum (ft)	Alarm-axis RMSE (ft)
Curb side	18	0.965	1.373	2.329	4.386	0.835 outward
Driver side	19	0.610	0.822	1.798	2.334	0.714 outward
Rear	14	2.412	4.485	11.474	11.544	3.905 rearward

The driver-side camera is the strongest current mapping. The curb-side camera supports controlled in-hull use with caution near the far-forward and far-outward edge. The rear mapping remains provisional for continuous distance because of its larger RMSE, 11.474-ft P95 error, and two approximately 10-ft far-rear errors.

Zone accuracy was 94.4% for curb side, 94.7% for driver side, and 92.9% for rear. Each camera had one yellow-to-red classification error. These errors were conservative relative to the 5-ft boundary. Zone accuracy should nevertheless remain secondary because validation classes were imbalanced and the driver-side and rear validation sets contained no red-zone centers.

5.9 Runtime Use Conditions

A camera-based distance should be retained only when:

1. The correct physical camera calibration is used.
2. The image remains 1280×720 with no unaccounted crop, resize, or rotation.
3. The estimated ground-contact point lies inside the surveyed pixel hull.
4. The object is not truncated at the image boundary.
5. The road surface and trailer attitude are sufficiently close to the surveyed planar condition.
6. Camera mounting and image processing remain unchanged.
7. Camera-specific uncertainty is retained in interpretation.

Pixels outside the surveyed hull should be flagged for review rather than converted into an apparently precise distance. Rear values should not be used as high-confidence hard thresholds.

5.10 Computer-Vision Detection and Tracking

The retrospective audit used YOLO11n instance segmentation to identify cars, trucks, buses, motorcycles, bicycles, and people. BoT-SORT supplied temporal track identities across sampled frames (Aharon et al. 2022).

When a segmentation mask was available, candidate ground-contact points were obtained from a bottom band of the mask. When no mask was available, nine points along the lower edge of the bounding box were evaluated. The accepted contact point was transformed using the matching camera calibration.

A metric sample was accepted only when:

- The ground intersection was valid.
- The contact point was inside the surveyed pixel hull.
- The object was not truncated by the frame boundary.
- The camera identity was known.
- The original image geometry was preserved.

At least three valid metric points were required to create a track-level trajectory summary.

5.11 Proximity and Trajectory Measures

For retained distance sequence (d_t), the project used

$$d_{\min} = \min_t \{\text{median}_3(d_t)\}$$

$$\text{closing}_t = \max\left(0, -\frac{dd_t}{dt}\right)$$

and, when closing speed was positive,

$$\text{TTC}_{\text{plane},t} = \frac{d_t}{\text{closing}_t}$$

The original trajectory/TTC rule combined uncertainty-adjusted distance, closing speed, plane TTC, persistence, and approach-then-separation behavior.

The source audit also applied a later 0-9 ft reporting rule. In this final report, that rule is described as a distance-based proximity screen rather than a validated near-miss definition. An object can remain within 9 ft while stationary or separating; closeness alone does not prove a collision path.

5.12 Recorder-Box Extraction and Association

Red, yellow, and green rectangles were extracted from source frames using HSV color masks and rectangle-geometry checks. Each candidate recorder rectangle was compared with YOLO boxes using intersection-over-union:

$$\text{IoU}(B_{\text{CV}}, B_{\text{alarm}}) = \frac{\text{area}(B_{\text{CV}} \cap B_{\text{alarm}})}{\text{area}(B_{\text{CV}} \cup B_{\text{alarm}})}$$

Association was accepted when the maximum IoU was at least 0.20. Otherwise, the recorder rectangle was retained as an alarm-overlay-only observation. A failed IoU association can reflect different box geometry rather than a pure object-detection failure.

5.13 Audit Inventory and Results

The audit processed 156 videos from April 17 and April 29 shown in Table 8.

Table 8. CV and recorder-alarm audit results

Measure	April 17	April 29	Total
Processed videos	87	69	156
Calibrated trajectory points	3,363	3,274	6,637
Track summaries	1,315	955	2,270
Alarm events	29	23	52
Supported tracks below 9 ft	20	10	30
Unique events containing supported close tracks	12	9	21
Supported tracks with matched recorder box	3	5	8
Supported tracks without matched recorder box	17	5	22

Of the 30 supported tracks that entered the project’s 0–9-ft distance screen, 22 did not have a machine-matched recorder box. These 22 observations are track-level records rather than 22 independent physical interactions or validated near misses. A single event may contain multiple objects, and one physical object may generate multiple track identities if tracking is interrupted or reinitialized. Track counts should therefore be reported together with unique-event counts.

Under the original trajectory/time-to-contact logic, the April 17 data contained one provisional rear-view candidate without a visible recorder box. The minimum distance retained within the surveyed calibration hull was 9.70 ft. Later estimated distances of approximately 4–7 ft occurred outside the surveyed hull and were excluded from the metric evidence. Because the rear-camera mapping had substantially greater continuous-distance uncertainty than the side-camera mappings, this event was interpreted as a provisional trajectory candidate rather than a precisely measured near miss.

The reverse mismatch was also observed. Clear recorder boxes were visually verified around a pickup and partially visible workers for which the sampled YOLO output did not provide a matching object. Plausible mechanisms identified during the audit included image-boundary truncation, object occlusion, extreme fisheye perspective and model-domain differences, 5-fps

sampling, limitations in detector recall, and the IoU-based box-association rule. An unmatched recorder box therefore does not necessarily indicate a pure YOLO detection failure; it may also reflect sampling or association differences.

The automated processing produced 125 alarm-overlay-only track fragments. These records were treated as machine-generated review flags and were not reported as 125 confirmed YOLO failures. Six high-ranked cases balanced across the two audit dates were manually reviewed. Four contained a clear recorder rectangle around a real object, one was a false colored-shape detection caused by red trailer lettering, and one was ambiguous near an image boundary.

Representative examples of both mismatch directions are shown in Figure 8. Panels (a) and (b) show a person and a pickup detected by the project CV pipeline without a visible recorder box around the target object. Panels (c) and (d) show recorder boxes around a pickup and a worker, respectively, for which the sampled CV output did not provide a matching object.




Figure 8. Complementary visual evidence across person and vehicle detected by CV or recorder box where they disagree

5.14 Engineering Interpretation

The CV pipeline and recorder boxes should be treated as parallel evidence proposals. Activation by either stream should retain the event for contextual review. Corroboration across streams can strengthen confidence, but requiring both to activate would discard informative cases.

A robust implementation should:

- Preserve full-rate source video for second-stage review.
- Use recorder boxes as proposals for targeted higher-recall CV inference.
- Replace pure IoU matching with a score combining overlap, containment, center distance, temporal continuity, class plausibility, and contact-point consistency.
- Fine-tune object models on fisheye truck-mounted imagery, especially truncated people, adjacent tractor sides, glare, darkness, and boundary regions.
- Keep object-presence, recorder-box, CV-detection, and supported-distance labels as separate fields.
- Reject out-of-hull metric predictions.
- Report unique events as well as tracks.
- Preserve camera-specific uncertainty and observability.

The audit demonstrates a functioning computer-vision pipeline and a completed evidence-comparison workflow. It does not provide population-level precision, recall, false-negative rate, near-miss prevalence, or crash-reduction effectiveness.

Chapter 6. Benefit-Cost Analysis and Targeted Deployment

6.1 Scope

The passive-safety module evaluated side underride guards for NHTSA's light-passenger-vehicle front-to-side target crashes at impact speeds up to 40 mph. It compared the present value of monetized reductions in fatalities and serious injuries with hardware, installation, and fuel-related costs for an annual cohort of eligible new trailers.

The primary side-guard module was intentionally limited to the crash population and performance range supported by the NHTSA analysis. It did not monetize rear-underride crashes, vulnerable-road-user events, crashes above the evaluated speed boundary, property damage, productivity effects, maintenance, cargo protection, loading-dock interference, axle-position changes, payload

effects, or broader warning-system benefits. These exclusions define the scope of the reported BCR and should not be interpreted as evidence that the excluded effects have no value.

The analysis is deterministic. It reports named scenarios, sensitivity analyses, and break-even thresholds rather than assigning probability distributions unsupported by the available literature.

6.2 Final Modeling Assumptions and Methodological Refinements

Table 9. Corrected BCA assumptions and interpretation

Modeling element	Final treatment	Rationale
Cohort exposure	260,000 / 6,600,000 = 3.939%	Aligns benefits with the modeled annual new-trailer cohort
Dollar year	2024 dollars	Places costs and benefits on a common basis
Discounting	7% primary; 3% sensitivity	Makes the present-value basis explicit
Serious injury value	Normalized severity-weighted MAIS 3-5 value	Aligns injury valuation with the target injury-severity distribution
Target population	Paired central and high fatality/injury scenarios	Preserves coherent fatality and serious-injury combinations
Warning effectiveness	Break-even threshold only	Avoids transferring an unsupported effectiveness estimate from another application
Uncertainty	Deterministic sensitivity	Reports the effect of alternative assumptions without assigning unsupported probabilities

The 2024 dollar base year, the 7% primary real discount rate, and the treatment of present-value benefits and costs followed USDOT's *Benefit-Cost Analysis Guidance for Discretionary Grant Programs: 2026 Update* (U.S. Department of Transportation 2025).

6.3 Cohort Exposure and Avoided Outcomes

The annual new-trailer cohort share, denoted by s , was calculated by dividing the annual cohort of 260,000 eligible new trailers by the reference fleet of 6.6 million trailers:

$$s = \frac{260,000}{6,600,000} = 0.0393939$$

For national annual target fatalities (F), serious injuries (I), fatality effectiveness ($e_F = 0.97$), and serious-injury effectiveness ($e_I = 0.85$), avoided outcomes for the modeled cohort were

$$F_{\text{avoided}} = F s e_F, \quad I_{\text{avoided}} = I s e_I$$

Using the central national target population of 17.7 fatalities and 81 serious injuries per year, the model estimated approximately 0.676 avoided fatalities and 2.712 avoided serious injuries per cohort-year.

F = annual national target fatalities

I = annual national target serious injuries

s = annual new-trailer cohort share

e_F = fatality-reduction effectiveness, 0.97

e_I = serious-injury-reduction effectiveness, 0.85

F_{avoided} = annual fatalities avoided for the modeled cohort

I_{avoided} = annual serious injuries avoided for the modeled cohort

6.4 Injury Valuation

The serious-injury value was constructed from the NHTSA MAIS 3-5 severity mix and USDOT fractions of the value of a statistical life:

$$V_I = \text{VSL} \frac{\sum_{m=3}^5 p_m q_m}{\sum_{m=3}^5 p_m}$$

where (p_m) is the target-population proportion at severity level (m), (q_m) is the corresponding VSL fraction, and (VSL=\$13.7) million in 2024 dollars. The MAIS 3-5 fractions were obtained from USDOT's departmental injury-valuation guidance, while the \$13.7 million VSL and 2024-dollar basis followed the current BCA parameter guidance (Departmental Guidance 2021; U.S. Department of Transportation 2025)

Using the severity proportions 0.764, 0.172, and 0.063; the corresponding VSL fractions 0.105, 0.266, and 0.593; and a VSL of \$13.7 million, the normalized serious-injury value was approximately \$2,239,875 per injury in 2024 dollars.

$$B_{\text{annual}} = F_{\text{avoided}} \text{VSL} + I_{\text{avoided}} V_I$$

6.5 Discounting and Economic Criteria

For real discount rate (r) and service life (L), the finite-life annuity factor was

$$AF(r, L) = \frac{1 - (1 + r)^{-L}}{r}$$

At a 7% real discount rate r and a 15.25-year service life L , the annuity factor was 9.1948. The 7% rate was used as the primary discount-rate assumption in accordance with the applicable USDOT BCA guidance (U.S. Department of Transportation 2025).

Present-value benefit was

$$PV_B = B_{\text{annual}} AF(r, L)$$

The primary cost included the 2020 hardware and installation estimate of \$2,990 per trailer, converted to 2024 dollars using a factor of 1.19, plus NHTSA’s discounted 450-lb or 800-lb fuel-cost case. The resulting present value of cohort costs is denoted by PV_C .

Benefit-cost ratio and net present value were


$$BCR = \frac{PV_B}{PV_C}, \quad NPV = PV_B - PV_C$$

6.6 Primary Side-Guard Results

Table 10. Primary side-guard BCA results

Guard case	Avoided fatalities/year	Avoided serious injuries/year	PV benefit	PV cost	BCR	NPV
450 lb	0.68	2.71	\$141.1M	\$1.16B	0.122	−\$1.02B
800 lb	0.68	2.71	\$141.1M	\$1.33B	0.106	−\$1.19B

Table 10 shows that both primary side-guard cases produced benefit–cost ratios substantially below 1.0. Figure 9 compares these primary results with the most favorable deterministic



sensitivity scenario evaluated in the study. The horizontal line at a BCR of 1.0 represents economic break-even, at which monetized benefits equal costs.

Figure 9. Primary and favorable side-guard benefit-cost ratios relative to the 1.0 break-even line

As shown in Figure 9, the primary 450-lb and 800-lb guard cases produced BCRs of 0.122 and 0.106, respectively. Under the favorable sensitivity assumptions, a higher target crash population, a 3% real discount rate, and a 20% reduction in hardware cost, the BCRs increased to 0.259 for the 450-lb case and 0.214 for the 800-lb case. Nevertheless, all four values remained well below the 1.0 break-even threshold. The tested changes to target population, discounting, and hardware cost therefore did not reverse the national annual-cohort conclusion within the modeled scope.

These results should not be interpreted as evidence that side guards provide no safety value. Rather, they indicate that the monetized safety benefits did not exceed the costs of equipping the modeled national new-trailer cohort when the analysis was restricted to NHTSA’s narrow front-to-side target crash population and the tested cost assumptions. The result also does not preclude different economics for alternative guard designs or for targeted fleets with demonstrably concentrated relevant risk.

6.7 Sensitivity Analysis

Table 11. Selected deterministic side-guard sensitivities

Scenario	450-lb BCR	800-lb BCR	Interpretation
Central target, 7%, base cost	0.122	0.106	Primary result
High target, 7%, base cost	0.175	0.153	Higher target population increases benefit
High target, 3%, 20% hardware reduction	0.259	0.214	Most favorable tested combination

Reasonable changes to target population, discounting, and hardware cost did not reverse the national-cohort conclusion within the modeled scope.

6.8 Warning-System Break-Even Logic

Warning effectiveness was not estimated from the alarm-selected clips. Instead, the model calculated the total causal effectiveness required for monetized benefits to equal lifecycle cost.

Overall effectiveness can be decomposed as

$$\eta_{\text{total}} = a \times d \times v \times h \times c$$

where:

- (a) = system availability;
- (d) = relevant-object detection probability;
- (v) = probability that an issued alert is valid for the target event;
- (h) = probability of an appropriate human or automated response; and
- (c) = probability that the response avoids the target crash.

The break-even effectiveness threshold is

$$\eta^* = \frac{PV_C}{PV_{B,100\%}}$$

where ($PV_{B,100\%}$) represents the monetized target benefit under 100% total causal effectiveness.

Table 12. Warning-system break-even thresholds

Capital cost/trailer	Annual O&M/trailer	PV cohort cost	Required total effectiveness	Feasible at $\leq 100\%$?
\$0	\$25	\$59.8M	38.9%	Yes
\$250	\$25	\$124.8M	81.3%	Yes
\$500	\$0	\$130.0M	84.7%	Yes
\$1,000	\$0	\$260.0M	169.3%	No
\$1,000	\$75	\$439.3M	286.1%	No

The table shows that a relatively costly warning package cannot break even using only the narrow side-override benefit set, even under perfect effectiveness. Lower technology cost or a broader validated benefit scope would be required.

6.9 Targeted-Deployment Conditions

For a national BCR below one, the treated subset must carry greater-than-average relevant risk if costs scale proportionally with the number of treated trailers.

The required risk-concentration multiplier is

$$m^* = \frac{1}{BCR}$$

For treated fleet share (f), the required share of national target risk is

$$q^* = f m^* = \frac{f}{\text{BCR}}$$

Table 13. Required target-risk concentration under proportional treatment-cost scaling

Treated share	fleet Required national risk share: 450-lb case	Required national risk share: 800-lb case	Mathematically feasible
5%	41.0%	47.1%	Yes
10%	82.1%	94.2%	Yes
25%	205.1%	235.6%	No

These values are deployment-screening thresholds rather than empirical findings about a specific fleet. Demonstrating the required concentration would require geocoded exposure, validated event definitions, and comparable operating data across treated and untreated fleet segments.

6.10 Passive and Active Benefit Accounting

Passive and active intervention benefits were considered separately to avoid double counting overlapping safety outcomes.

A side override guard primarily mitigates crash severity after physical contact begins. A warning system is intended to prevent or alter the interaction before impact. Consequently, a crash avoided through an active warning system is no longer available to be mitigated by a guard. Combined benefit accounting should therefore distinguish:

1. Target crashes avoided by an active intervention.
2. Residual crashes whose severity is reduced by structural protection.
3. Non-overlapping operational benefits.
4. Capital and recurring costs shared across interventions.
5. False or unnecessary warning burden.
6. Installation, maintenance, fuel, downtime, and compatibility effects.

The BCA therefore provides a transparent national baseline and a set of break-even thresholds against which observed project evidence, independently validated technology performance, and documented industry costs can be compared.

Chapter 7. Integrated Deployment, Policy, and Project Outcomes

7.1 Integrated Interpretation of the Workstreams

The project's four technical workstreams answer different parts of the same deployment question stated in Table 14.

Table 14. Integrated contributions of the project workstreams

Workstream	Primary question	Main project output	Decision use
Naturalistic paired-unit analysis	Where and under what operating conditions does trailer-specific information differ from tractor information?	Environment-aware movement, speed, yaw, lateral-response, and detector-context baseline	Prioritize contexts for trailer-aware monitoring
Camera-ground localization	When can an image observation be interpreted as approximate physical clearance?	Camera-specific mappings, supported image hulls, strict error estimates, and zone interpretation	Establish physical meaning and uncertainty for visual observations
CV/recorder audit	What can be concluded when independent CV and installed detector evidence agree or disagree?	Proximity tracks, trajectory/TTC candidates, evidence mismatch categories, and visually reviewed examples	Support evidence fusion and algorithm refinement
Benefit-cost analysis	What cost, effectiveness, and risk-concentration conditions are required for economic justification?	Guard BCRs, warning break-even thresholds, and targeted-risk thresholds	Distinguish blanket deployment from targeted implementation

The workstreams should not be collapsed into a single universal risk score because each uses a different unit of observation, validation basis, and interpretation. Instead, their contributions form a traceable evidence chain. Raw measurements are first subjected to synchronization, calibration, and data-quality checks and are then interpreted within the corresponding roadway and operating context. Camera and computer-vision outputs are subsequently evaluated as object, proximity, and trajectory evidence. Finally, the resulting technical evidence is connected to benefit-cost thresholds and implementation decisions, as summarized in Figure 1.

7.2 Implications for Fleet Operators

Fleet operators can use the naturalistic baseline for post-trip review, sensor assessment, and operating-context prioritization. The paired ramp result indicates that trailer-mounted inertial sensing adds information in roadway transitions where tractor and trailer responses differ.

Recommended fleet practices include:

- Preserve tractor and trailer unit identities separately.
- Retain raw timestamps and documented synchronization procedures.
- Stratify detector activity by monitored zone, speed, operating environment, and exposure.
- Review ramp, merge, and diverge segments as priority trailer-aware contexts.
- Retain non-alarm samples and sensor-uptime records.
- Track camera mount, resolution, crop, and maintenance changes.
- Use side-camera metric estimates only within supported regions and with engineering margins.
- Treat rear-camera continuous distance provisionally.
- Collect installed hardware, communications, subscription, repair, calibration, and downtime costs.
- Use targeted pilots only after documenting the treated group's share of relevant exposure and risk.

7.3 Implications for Technology Developers

The project results show that warning-system development requires more than maximizing object detections.

Recommended development practices include:

- Maintain separate tractor and trailer state estimates.
- Use roadway and facility context when evaluating alarm logic.
- Preserve both CV and recorder evidence rather than requiring agreement.
- Process alarm-proposed regions at full source frame rate when practical.
- Improve association using temporal and geometric gating rather than IoU alone.
- Fine-tune vision models on truck-mounted fisheye imagery.
- Evaluate ground-contact estimation separately from object classification.

- Reject unsupported out-of-hull metric predictions.
- Propagate camera-specific uncertainty into distance-zone decisions.
- Report object presence, system box, CV detection, distance support, and event interpretation as separate labels.
- Evaluate unnecessary-alert burden and observability in addition to detector recall.
- Connect technical performance to the multiplicative effectiveness components used in economic assessment.

7.4 Implications for Transportation Agencies and Policymakers

Transportation agencies commonly use speed and location data from tractor-based probes. The paired-unit findings indicate that tractor-only data may not fully represent the trailing unit during selected transitions. Paired trailer observations can therefore complement roadway geometry, traffic data, and conventional probes when agencies investigate ramps, interchanges, truck routes, or constrained operating areas.

Policy interpretation should distinguish:

- National-average economics from targeted deployment.
- Raw alert counts from exposure-normalized rates.
- Proximity from conflict.
- Camera zone classification from continuous-distance accuracy.
- Internal robustness from fleet-level generalizability.
- A system activation from a demonstrated safety benefit.

Public-facing geographic products should use validated, exposure-normalized layers. The project provides several map-ready components:

- Geocoded synchronized exposure.
- Broad roadway-environment classification.
- Operational-facility context.
- Monitored-zone activity.
- Roadway-assignment quality.
- Targeted-deployment risk thresholds.
- Requirements for separating alarm frequency from opportunity exposure.

These components define the technical requirements for route- and time-based prioritization without presenting unvalidated alarm density as a public safety-risk map.

7.5 Relationship to the Project Deployment Plan

Table 15. Relationship between the deployment plan and report evidence

Deployment-plan component	Evidence represented in this report	Resulting project contribution
Multi-source safety data and gap analysis	Naturalistic cohort, broader alarm/video inventory, roadway context, calibration survey, and economic parameter ledger	Dataset lineage, evidence boundaries, and traceable analytical inputs
Exploratory underride and near-miss framework	Passive/active architecture, paired vehicle baseline, proximity zones, trajectory/TTC logic, and BCA	Integrated measurement-to-decision framework
Sensor-integration strategies	Separate tractor and trailer GNSS/IMU, three cameras, detector states, roadway context, and camera-local mappings	Trailer-aware multimodal integration approach
Initial field testing and model results	56-hour naturalistic study, calibration survey, CV audit, and statistical modeling	Empirical operating baseline and camera/CV performance evidence
Policy and fleet roadmap	BCR thresholds, targeted-risk conditions, report recommendations, and implementation sequence	Evidence-based guidance for targeted monitoring and technology assessment
Technical tools and algorithms	Synchronization, map matching, calibration, CV tracking, mismatch audit, and BCA workflows	Reproducible analytical package
Public-facing planning support	Geocoded environment, exposure, system-state, and targeting requirements	Technical specification for exposure-normalized route/time prioritization

7.6 Reproducible Project Outputs

Table 16. Principal analytical and reproducibility outputs

Output	Function
--------	----------

Naturalistic synchronization and map-matching workflow	Aligns tractor, trailer, detector, and roadway records and creates analytical blocks
Paired tractor-trailer models	Estimates movement, speed, trailer-relative yaw, and trailer-relative lateral response
Camera calibration pipeline	Inventories data, aligns survey frames, detects checkers, and fits candidate mappings
Camera-specific calibration files	JSON Store runtime mapping and supported pixel hull for each camera
Nested-validation predictions and summaries	Preserve held-out camera-ground predictions and strict error statistics
CV/BoT-SORT near-miss workflow	Detects and tracks objects, estimates ground contact, maps distances, and evaluates trajectories
CV/alarm-box mismatch audit	Compares independent CV tracks with recorder-system boxes
Corrected deterministic workflow	BCA Calculates guard scenarios, warning break-even thresholds, and targeted-deployment thresholds
Parameter ledgers and validation outputs	Preserve source, units, dollar year, scope, and automated consistency checks

7.7 Recommended Implementation Sequence

The project findings support a staged implementation strategy.

Stage 1: Preserve Auditable Data

Retain raw timestamps, unit identifiers, camera channels, image geometry, sensor uptime, missingness, roadway-assignment quality, and operating exposure. A dashboard summary should not replace the underlying data needed to verify a later safety claim.

Stage 2: Establish Context-Specific Baselines

Characterize movement, speed, and paired tractor-trailer response before defining anomalies. Use broad roadway and facility categories that are defensible under the available location accuracy.

Stage 3: Validate Object and Proximity Evidence

Use camera-specific mappings, supported pixel regions, object tracking, and uncertainty. Keep recorder-system outputs and independent CV observations separate until they are compared and reviewed.

Stage 4: Evaluate Event Relevance

Distinguish stationary proximity, ordinary passing, closing interactions, trajectory/TTC candidates, and adjudicated near misses. Evaluate system observability and non-alarm intervals as well as alarm-selected events.

Stage 5: Connect Performance to Economics

Measure installed and recurring costs, availability, relevant-object detection, valid-alert probability, operator or automated response, and crash avoidance conditional on response. Compare those measurements with break-even thresholds.

Stage 6: Target Deployment

Prioritize trailers, routes, times, and operating environments according to validated exposure and risk concentration. Report sensitivity to cost, effectiveness, coverage, and operational constraints.

7.8 Interpretation Controls

Table 17. Principal interpretation controls across project workstreams

Evidence issue	Why it matters	Required interpretation
One trailer and two tractors	Limits direct fleet-level generalization	Findings characterize the observed configuration, routes, and dates
Ninety-three paired ramp-classified blocks	Ramp estimate has less support than facility, arterial, or freeway estimates	Emphasize robust direction and uncertainty rather than a universal magnitude
Observational design	Roadway context was not experimentally assigned	Interpret coefficients as conditional associations
Broad pointwise map matching	Exact lane and link assignments can be ambiguous	Retain the term ramp-classified
One-second synchronization	Does not estimate subsecond clock drift or propagation delay	Use same-second average response rather than transient propagation claims
Camera survey on one date	Does not test remounting, vibration, or processing changes	Apply mappings only under unchanged geometry

Approximately planar road assumption	Slopes, curbs, trailer pitch, and roll can alter pixel-ground geometry	Treat distances as local planar estimates
Rear-camera RMSE and outliers	Continuous rear estimates are less reliable	Interpret rear distance provisionally
CV audit on two selected dates	Does not represent the full operating population	Use as an engineering case study
No independent frame-by-frame truth set	Prevents absolute precision and recall estimates	Report verified examples and evidence categories
General CV run at 5 fps	Brief boxes or detections may fall between samples	Retain source video for full-rate review
Alarm-centered broader archive	Does not provide an unbiased exposure denominator	Do not estimate warning rates or effectiveness from alarm clips alone
Narrow BCA target population	Excludes other possible safety and operational benefits	Interpret BCR within the modeled scope
Break-even warning thresholds	Are not empirical effectiveness estimates	Compare only with independently measured performance

Chapter 8. Conclusions and Recommendations

8.1 Conclusions

This Safety21 project established an integrated framework for evaluating smart-trailer safety evidence from field sensing through economic interpretation. The project operationalized the smart-trailer concept as a combination of trailer-specific observation, measurement validity, temporal and spatial integration, operating-context interpretation, object and event assessment, and decision support.

The naturalistic deployment demonstrated the value of preserving tractor and trailer states separately. One trailer operated with two tractors during 13 trips on 10 collection dates, yielding 56 hours and 201,600 synchronized one-second observations. The analysis constructed complete exposure before selecting detector or dynamic events, assigned broad roadway environments independently of the outcomes, and accounted for repeated dates and recurrent operating contexts.

Operating environments differed broadly in speed. Freeway/interstate operation was faster than arterial operation, while facility, collector/local, and ramp-classified operation were slower. These patterns established distinct operating baselines rather than safety rankings.

Ramp-classified operation produced the clearest trailer-specific dynamic departure. Relative to arterial operation, the trailer had lower yaw activity but higher lateral specific-force activity relative to the tractor. The opposing directions were robust across response definitions, temporal aggregations, pairing rules, and model specifications. The result indicates that tractor-only telemetry can incompletely represent trailer response in ramp-classified and related merge/diverge contexts.

Detector-derived activity provided useful but secondary evidence. Rear-zone episodes represented 57.6% of the 729 mapped episodes, but those shares do not establish that rear interactions were more common or more dangerous because camera-observable exposure was not separately measured by zone. Detector activity was associated with lower speed but not with a detectable general change in paired tractor-trailer response.

The calibration work established a defensible method for camera-specific local ground localization. The driver-side and curb-side mappings support controlled in-hull use under unchanged camera, image, and roadway-plane conditions. Strict nested validation showed that the rear mapping had substantially greater continuous-distance uncertainty and should be interpreted provisionally.

The CV audit confirmed that the project executed a functioning independent object-detection, tracking, ground-contact, distance, and trajectory workflow. CV-supported close tracks and recorder-system boxes disagreed in both directions. Neither stream was complete by itself. The appropriate technical response is evidence fusion, improved association, full-rate review of selected regions, camera-support checks, and independent labeling.

The corrected benefit-cost analysis found that national annual-cohort side-guard deployment did not break even under the narrow modeled NHTSA target population and tested costs. Primary BCRs were 0.122 for the 450-lb case and 0.106 for the 800-lb case. The result identifies the economic challenge of blanket deployment and provides a transparent baseline for comparing alternative costs, benefits, and targeting strategies.

Warning-system thresholds demonstrated that technical performance must be connected to availability, detection, alert validity, response, and conditional crash avoidance. Targeted-deployment analysis showed that a small treated fleet share must capture a disproportionately large share of relevant risk to reach break-even.

8.2 Recommendations

Recommendation 1: Preserve Trailer-Specific State

Future fleet data systems and safety evaluations should retain trailer measurements independently of tractor telemetry. At minimum, unit identity, timestamps, position, speed, and relevant trailer motion channels should remain traceable.

Recommendation 2: Prioritize Roadway Transitions

Ramp-classified, merge, and diverge contexts should receive priority in trailer-aware sensor review, data collection, and performance assessment because the observed trailer–tractor response balance differed most clearly in those environments.

Recommendation 3: Interpret Alerts Through Exposure and Context

Warning and detector outputs should be normalized by observable operating exposure and stratified by monitored zone, speed, roadway environment, facility context, and sensor uptime. Raw counts should not be used as safety-performance measures.

Recommendation 4: Retain Camera-Specific Calibration and Uncertainty

Each camera should use only its corresponding mapping and surveyed support region. Resolution, crop, rotation, mount, and processing changes should trigger coordinate transformation or recalibration. Rear continuous-distance estimates should remain provisional unless stronger independent ground control is available.

Recommendation 5: Separate Proximity From Near Miss

Objects within a distance zone should be reported as proximity events unless closing behavior, trajectory conflict, and independent adjudication support a stronger label. Track counts and unique-event counts should both be reported.

Recommendation 6: Fuse Independent Evidence Streams

The CV pipeline and recorder-system boxes should be retained as parallel proposals. Disagreement should trigger targeted review rather than automatic dismissal of one stream. Association should use temporal and geometric evidence beyond IoU alone.

Recommendation 7: Use Deterministic and Traceable Economic Assumptions

Economic analyses should preserve a parameter ledger identifying source, value, unit, dollar year, scope, and whether an input is observed, literature-based, or a scenario assumption. Unsupported probabilities of economic success should not be reported.

Recommendation 8: Distinguish Blanket and Targeted Deployment

A negative national-average BCR should not be generalized automatically to every operating context. Targeted deployment should be considered only when the treated fleet, route, or time period can be shown to contain sufficient relevant exposure or risk concentration.

Recommendation 9: Connect Warning Performance to Causal Components

Warning effectiveness should be decomposed into system availability, relevant-object detection, valid-alert probability, driver or vehicle response, and crash avoidance conditional on response. Detector accuracy alone is not a safety-effectiveness estimate.

Recommendation 10: Publish Only Exposure-Normalized Geographic Evidence

Route or Truck Maps should display opportunity-normalized and validated measures rather than alarm density alone. Map layers should preserve roadway-assignment quality, exposure, monitored zone, operating environment, and event-definition confidence.

8.3 Final Statement

The project demonstrates that smart-trailer safety is not determined by the presence of one camera, sensor, detector, or guard. Safety value arises from a complete evidence chain: trailer-specific measurement, data-quality control, synchronization, roadway and operational context, calibrated object and proximity interpretation, and economically meaningful deployment criteria.

The naturalistic baseline identifies when tractor-only monitoring may be incomplete. The camera analysis establishes when image observations can be interpreted metrically. The CV audit shows why independent detector streams must be fused and reviewed. The BCA shows what costs, effectiveness, and risk concentration are required before a deployment claim becomes economically defensible. Together, these contributions provide a practical foundation for trailer-aware monitoring, targeted implementation, and data-driven transportation policy.

Acknowledgments

The authors thank Paul Sniegocki and Michael Speranza of Clarience Technologies for collecting the real-world heavy-duty vehicle data used in the naturalistic study and for their technical discussions regarding vehicle telematics and sensing technologies. The authors also acknowledge the Safety21 University Transportation Center and the U.S. Department of Transportation University Transportation Centers Program for supporting this research.

References

- Aharon, Nir, Roy Orfaig, and Ben-Zion Bobrovsky. 2022. “BoT-SORT: Robust Associations Multi-Pedestrian Tracking.” arXiv:2206.14651. Preprint, arXiv, July 7. <https://doi.org/10.48550/arXiv.2206.14651>.
- Behera, Abhijeet, Sogol Kharrazi, and Erik Frisk. 2024. “How Do Long Combination Vehicles Perform in Real Traffic? A Study Using Naturalistic Driving Data.” *Accident Analysis & Prevention* 207: 107763. <https://doi.org/10.1016/j.aap.2024.107763>.

- Bookstein, F. L. 1989. "Principal Warps: Thin-Plate Splines and the Decomposition of Deformations." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 11 (6): 567–85. <https://doi.org/10.1109/34.24792>.
- Departmental Guidance. 2021. *Treatment of the Value of Preventing Fatalities and Injuries in Preparing Economic Analyses*. <https://doi.org/https://www.transportation.gov/sites/dot.gov/files/2021-03/DOT%20VSL%20Guidance%20-%202021%20Update.pdf>.
- Dhahir, Bashar, and Yasser Hassan. 2018. "Studying Driving Behavior on Horizontal Curves Using Naturalistic Driving Study Data." *Transportation Research Record* 2672 (17): 83–95. <https://doi.org/10.1177/0361198118784384>.
- Fancher, P., and C. Winkler. 2007. "Directional Performance Issues in Evaluation and Design of Articulated Heavy Vehicles." *Vehicle System Dynamics* 45 (7): 607–47. <https://doi.org/10.1080/00423110701422434>.
- Farewell, V. T., D. L. Long, B. D. M. Tom, S. Yiu, and L. Su. 2017. "Two-Part and Related Regression Models for Longitudinal Data." *Annual Review of Statistics and Its Application* 4 (Volume 4, 2017): 283–315. <https://doi.org/10.1146/annurev-statistics-060116-054131>.
- Hac, Aleksander, Daniel Fulk, and Hsien Chen. 2008. *Stability and Control Considerations of Vehicle-Trailer Combination*. April 14. <https://doi.org/10.4271/2008-01-1228>.
- Hallmark, Shauna L., Samantha Tyner, Nicole Oneyear, Cher Carney, and Dan McGehee. 2015. "Evaluation of Driving Behavior on Rural 2-Lane Curves Using the SHRP 2 Naturalistic Driving Study Data." *J Safety Res* 54: 17–27. <https://doi.org/10.1016/j.jsr.2015.06.017>.
- Hoerl, Arthur E., and Robert W. Kennard. 1970. "Ridge Regression: Biased Estimation for Nonorthogonal Problems." *Technometrics* 12 (1): 55–67. <https://doi.org/10.1080/00401706.1970.10488634>.
- Howell, Fletcher J., Azhaginiyal Arularasu, David B. Logan, and Sjaan Koppel. 2025. "Naturalistic Driving Study Data Applied to Road Infrastructure: A Systematic Review." *Journal of Safety Research* 92: 346–74. <https://doi.org/10.1016/j.jsr.2024.11.022>.
- Jahanbiglari, Ava, Xiyu Deng, Yorie Nakahira, and Pingbo Tang. 2025. *Preventing Rear and Side Crashes of Heavy-Duty Tractor-Trailer Combinations with Smart Sensors and Vision Systems*. <https://rosap.ntl.bts.gov/view/dot/86407>.
- Kamnik, R., F. Boettiger, and K. Hunt. 2003. "Roll Dynamics and Lateral Load Transfer Estimation in Articulated Heavy Freight Vehicles." *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering* 217 (11): 985–97. <https://doi.org/10.1243/095440703770383884>.
- Kannala, J., and S. S. Brandt. 2006. "A Generic Camera Model and Calibration Method for Conventional, Wide-Angle, and Fish-Eye Lenses." *IEEE Transactions on Pattern*

- Analysis and Machine Intelligence* 28 (8): 1335–40.
<https://doi.org/10.1109/TPAMI.2006.153>.
- MacKinnon, James G., Morten Ørregaard Nielsen, and Matthew D. Webb. 2023. “Fast and Reliable Jackknife and Bootstrap Methods for Cluster-Robust Inference.” *Journal of Applied Econometrics* 38 (5): 671–94. <https://doi.org/10.1002/jae.2969>.
- MacKinnon, James G., Morten Ørregaard Nielsen, and Matthew D. Webb. 2025. “Cluster-Robust Jackknife and Bootstrap Inference for Logistic Regression Models.” *Econometric Reviews* 0 (0): 1–29. <https://doi.org/10.1080/07474938.2025.2515161>.
- National Highway Traffic Safety Administration. 2022. *Federal Motor Vehicle Safety Standards; Rear Impact Guards, Rear Impact Protection; Final Rule. U.S. Department of Transportation*. <https://www.nhtsa.gov/sites/nhtsa.gov/files/2022-06/Final-Rule-FMVSS-223-224-Rear-impact-protection-web.pdf>.
- National Highway Traffic Safety Administration. 2024. *Report to Congress-Side Underride Protection*. Text. Reports to Congress-NHTSA.
<https://www.nhtsa.gov/sites/nhtsa.gov/files/2024-07/report-to-congress-june-2024-side-underride-protection.pdf>.
- Naujoks, Frederik, Andrea Kiesel, and Alexandra Neukum. 2016. “Cooperative Warning Systems: The Impact of False and Unnecessary Alarms on Drivers’ Compliance.” *Accident Analysis & Prevention* 97: 162–75. <https://doi.org/10.1016/j.aap.2016.09.009>.
- Nodine, Emily, Andy Lam, Wassim G. Najm, Bruce Wilson, and John Brewer. 2011. “Integrated Vehicle-Based Safety Systems Heavy-Truck Field Operational Test Independent Evaluation.” <https://rosap.nhtl.bts.gov>.
- Pennsylvania Department of Transportation. 2025. “RMSSEG (State Roads).” https://data-pennshare.opendata.arcgis.com/datasets/d9a2a5df74cf4726980e5e276d51fe8d_0/about.
- Ruff, Todd. 2006. “Evaluation of a Radar-Based Proximity Warning System for off-Highway Dump Trucks.” *Accident Analysis & Prevention* 38 (1): 92–98.
<https://doi.org/10.1016/j.aap.2005.07.006>.
- Schaudt, William A., Darrell S. Bowman, Rebecca L. Olson, et al. 2014. “Federal Motor Carrier Safety Administration’s Advanced System Testing Utilizing a Data Acquisition System on the Highways (FAST DASH): Safety Technology Evaluation Project #1 Blindspot Warning: Final Report.” <https://doi.org/10.21949/1502710>.
- Sun, Qiang, Wen-Xin Zhou, and Jianqing Fan. 2020. “Adaptive Huber Regression.” *Journal of the American Statistical Association* 115 (529): 254–65.
<https://doi.org/10.1080/01621459.2018.1543124>.
- Tomasch, Ernst, and Stefan Smit. 2023. “Naturalistic Driving Study on the Impact of an Aftermarket Blind Spot Monitoring System on the Driver’s Behaviour of Heavy Goods

- Vehicles and Buses on Reducing Conflicts with Pedestrians and Cyclists.” *Accident Analysis & Prevention* 192: 107242. <https://doi.org/10.1016/j.aap.2023.107242>.
- Trigell, Annika Stensson, Malte Rothhämel, Joop Pauwelussen, and Karel Kural. 2017. “Advanced Vehicle Dynamics of Heavy Trucks with the Perspective of Road Safety.” *Vehicle System Dynamics* 55 (10): 1572–617. <https://doi.org/10.1080/00423114.2017.1319964>.
- U. S. Government Accountability Office. 2019. “Truck Underride Guards: Improved Data Collection, Inspections, and Research Needed | U.S. GAO.” April 15. <https://www.gao.gov/products/gao-19-264>.
- United States. Department of Transportation and National Highway Traffic Safety Administration. 2023. *Side Impact Guards for Combination Truck-Trailers: Cost-Benefit Analysis*. DOT HS 813 404. <https://rosap.nhtl.bts.gov/view/dot/79828>.
- United States. Department of Transportation and National Highway Traffic Safety Administration. 2026. *Traffic Safety Facts 2024 Data: Large Trucks*. DOT HS 813 816. <https://doi.org/https://doi.org/10.21949/ey01-bq04>.
- U.S. Department of Transportation. 2025. *Benefit-Cost Analysis Guidance for Discretionary Grant Programs*. <https://doi.org/https://www.transportation.gov/sites/dot.gov/files/2025-12/Benefit%20Cost%20Analysis%20Guidance%202026%20Update%20%28Final%29.pdf>.
- Varma, Sudhir, and Richard Simon. 2006. “Bias in Error Estimation When Using Cross-Validation for Model Selection.” *BMC Bioinformatics* 7 (1): 91. <https://doi.org/10.1186/1471-2105-7-91>.
- Vergara, Eduardo, Juan Aviles-Ordóñez, Yuanchang Xie, and Mohammadali Shirazi. 2024. “Understanding Speeding Behavior on Interstate Horizontal Curves and Ramps Using Networkwide Probe Data.” *Journal of Safety Research* 90: 371–80. <https://doi.org/10.1016/j.jsr.2024.05.003>.
- Wang, Bo, Shauna Hallmark, Peter Savolainen, and Jing Dong. 2017. “Crashes and Near-Crashes on Horizontal Curves along Rural Two-Lane Highways: Analysis of Naturalistic Driving Data.” *Journal of Safety Research* 63: 163–69. <https://doi.org/10.1016/j.jsr.2017.10.001>.
- Wang, Xingmin, Shengyin Shen, Debra Bezzina, James R. Sayer, Henry X. Liu, and Yiheng Feng. 2020. “Data Infrastructure for Connected Vehicle Applications.” *Transportation Research Record* 2674 (5): 85–96. <https://doi.org/10.1177/0361198120912424>.
- Zhang, Z. 2000. “A Flexible New Technique for Camera Calibration.” *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22 (11): 1330–34. <https://doi.org/10.1109/34.888718>.

Zhou, Yun, and Raj Bridgelall. 2020. “Review of Usage of Real-World Connected Vehicle Data.” *Transportation Research Record* 2674 (10): 939–50.
<https://doi.org/10.1177/0361198120940996>.

Appendix A

A.1. Journal Publications:

Ava Jahanbiglari, Seongeun Park, Erin Chiou, Pingbo Tang, “Human-AI trust for safer heavy-duty equipment management: Insights from socio-technical literature review and interviews”, *Developments in the Built Environment*, Vol. 27, Article 100983, Oct. 2026, <https://doi.org/10.1016/j.dibe.2026.100983>

Ava Jahanbiglari, Pingbo Tang, “Characterizing Tractor-Trailer Behavior and Paired Response Across Operating Environments: A Naturalistic Baseline for Proactive Operational Safety”, Submitted to *Transportation Research Board (TRB) and Transportation Research Record (TRR)*, 2026

A.2. Conference Publications:

Xiyu Deng, Yuantao Tang, Ava Jahanbiglari, “CAST-Guided Large Language Model Framework for Causal Analysis and Safety Improvement in (Autonomous) Transportation Systems”, Accepted, *Proceedings of the ASCE International Conference on Computing in Civil Engineering*, 2026

Ava Jahanbiglari, Xiyu Deng, Yorie Nakahira, Pingbo Tang, “Proactive Near-Miss Prediction for Tractor-Trailers through Physics-Informed Reinforcement Learning”, Presented at the *Transportation Research Board (TRB)*, 2026

Ava Jahanbiglari, Ying Shi, Ruoxin Xiong, and Pingbo Tang, “Customized Implementation of Advanced Safety Sensors for Small Motor Carriers through Vehicle Inspection and Crash Data Analytics” *Proceedings of the ASCE International Conference on Computing in Civil Engineering*, 2024, <https://doi.org/10.1061/9780784486139.087>

A.3. Presentations:

Poster Presentation at National Adaptation Forum: Toward Safer, More Equitable, and Climate-Resilient Freight Transportation with Physics-Informed AI, May 2026

Presentation at the Graduate Seminar at Carnegie Mellon University, November 2025 and February 2026

Poster Presentation at Transportation Research Board: “Proactive Near-Miss Prediction for Tractor-Trailers through Physics-Informed Reinforcement Learning”, January 2026

Poster Presentation at Safety21 University Transportation Center Deployment Partner Consortium Symposium, November 2025

A.4. Undergraduate and Graduate Student Research Projects and Software Products

Aryana Singh. *Regulatory Framework and Liability Allocation for AI-Driven Commercial Motor Vehicle (CMV) Safety Triage Systems*. Undergraduate research project, Carnegie Mellon University, Summer 2026. Examined regulatory and liability considerations for using AI-assisted systems to support heavy-duty vehicle inspection decisions.

Beibei Sun. *Inferring HDV Risk Hotspots via SUMO Physics Simulation: Pittsburgh's West End Bridge Corridor Case Study*. Research project, Carnegie Mellon University, Summer and Fall 2026. Used traffic simulation and roadway/crash data to identify potential heavy-duty vehicle risk hotspots in Pittsburgh's West End Bridge corridor.

Jessica Cai. *Historical Pittsburgh Crash Heat Map*. Summer Undergraduate Research Apprenticeship (SURA), Carnegie Mellon University, 2025. Developed geospatial analyses of historical crash data to identify heavy-duty vehicle crash-risk patterns. [GitHub repository: Pittsburgh-Crash-Heat-Map-2025](#).

Jane Fleischman. *Automating the Creation of 2D Roadway Representations and Training Data for Safety Simulations Using Public Geospatial Data*. Presented at Meeting of the Minds 2025, Carnegie Mellon University. Developed simulation-ready roadway scenarios using publicly available crash and roadway data.

Niki Lin. *Heavy-Duty Vehicle Safety Simulation*. Undergraduate research project, Carnegie Mellon University, 2025. Explored heavy-duty vehicle safety risks using truck simulations.

A.5. Datasets

Multimodal Heavy-Duty Vehicle Operational Dataset (in preparation). Multimodal field data from heavy-duty tractor-trailer operations, including synchronized vehicle motion, roadway, environmental, and video-derived information. The dataset is currently being prepared for research dissemination.