Safety21

INNOVATING FOR SAFETY

US DOT National
University Transportation Center for Safety

Carnegie Mellon University

Penn University of Pennsylvania

The Ohio State University

CCAC Community College of Allegheny County

Community College of Philadelphia

The University of Texas Rio Grande Valley

Morgan State University

Safety21—2025 - 2026 Projects

Research Project Name: Avoiding Collisions in Connected and Autonomous Driving Using Safe Deep Reinforcement Learning Exploration through Control Barrier/Lyapunov Functions

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Bilin Aksun-Guvenc, The Ohio State University; Levent Guvenc, The Ohio State University

Project Partners: City of Marysville and Union County, Ohio; Transportation Research Center

Research Project Funding: Federal Funding—$81,535; Non-Federal Funding—$106,112

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: Connected and autonomous driving with carefully developed crash mitigation methods will help in the reduction of road traffic accidents and the fatalities and injuries they cause, including the ones with vulnerable road users. Based on these facts, this proposal focuses on developing data driven collision avoidance algorithms that learn to avoid collisions with nearby vehicles and vulnerable road users, building upon our approach to pedestrian and bicyclist safety in our Year 1 and Year 2 Safety 21 projects. A deep reinforcement learning safety algorithm with safe exploration will be developed using realistic hardware-in-the-loop (HIL) and vehicle-in-virtual-environment (VVE) platforms.

The current state-of-the-art crash mitigation methods use a catalog of possible collision interaction scenarios using model-in-the-loop simulations for development and evaluation. While this approach will be successful in mitigating collisions for the scenarios in the catalog being used, it is not generic in nature and will not necessarily mitigate crashes in the many other possible dangerous interactions between the road actors. The motivation for the use of the data driven approach in this proposal is this limitation in scope and the difficulty of accurately modeling the interaction dynamics between vehicles and between vehicles and vulnerable road users in dangerous encounters. Our deep reinforcement learning based collision avoidance algorithm will learn from interactions in complex traffic environments, adapt to various traffic scenarios, and satisfy real-time performance requirements, making it a more robust and advanced solution. We will use a hybrid approach in that the trained deep reinforcement learning algorithm will only be activated and work during dangerous and close interactions with other road users. Training will also be based on such interactions using traffic microsimulation to generate realistic training interactions automatically as compared to the fixed nature of the scenario driven approach.

Since reinforcement learning algorithms suffer from the absence of safety guarantees during the learning process, we will integrate our deep reinforcement learning algorithm with control Lyapunov and control barrier functions to limit its training to feasible areas of interaction with other road actors within the road barriers, with realistic vehicle dynamic model constraints as compared to the oversimplified ones used in the literature. Vehicle-to-Vehicle and Vehicle-to-VRU connectivity will also be used as low latency, and high speed near field and over the cloud cellular and other means of communication are available and as they help significantly in non-line-of-sight situations in aiding the perception sensors of autonomous vehicles. Efficacy of the proposed method will be demonstrated using HIL simulations and with a real vehicle using VVE.

US DOT Priorities: Section left blank until USDOT's new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: The research and technology outcome will be the development of the novel deep reinforcement learning algorithm with control Lyapunov/barrier function exploration using more realistic vehicle/VRU interaction models and vehicle dynamics constraints for collision avoidance with other nearby vehicles and vulnerable road users and corresponding conference and journal publications. The algorithm will be developed and demonstrated safely first using model-in-the-loop (MIL) and hardware-in-the loop (HIL) simulations and then with a real vehicle using the VVE method of our research group at the Ohio State University. The scenarios for training and evaluation will be created by injecting other vehicle(s) and vulnerable road users with unexpected motion within realistic traffic generated by a microscopic traffic simulator. The project will also contribute to graduate student theses and undergraduate student senior projects. Project results will be integrated into the relevant parts of the OSU courses ECE 5553 Autonomy in Vehicles (undergraduate and graduate) and ME 8322 Vehicle System Dynamics and Control (graduate). Exemplary project results will be shared using our lab's Youtube site. Anticipated invention disclosure(s) will be followed, if applicable, with potential provisional patent applications and subsequent patent applications by OSU's Technology Commercialization Office.

Outcomes/Impacts: Our deep reinforcement learning with control Lyapunov/barrier function exploration based algorithm will help avoid collisions of connected and autonomous vehicles with other vehicles and vulnerable road users and help automotive OEMs and suppliers in their development of advanced emergency braking algorithms for the FMVSS No. 127 standard by 2029. The developed collision avoidance algorithm will fuse vehicle-to-vehicle and vehicle-to-VRU communication with perception sensor data to improve accuracy including no-line-of-sight encounters. The VVE approach of our Year 1 and 2 projects will be further developed to inject real high collision risk vehicles and vulnerable road users into the development and evaluation environment to safely and realistically test extreme events naturally. This approach has the potential of being used in pre-deployment evaluation of connected and autonomous vehicles and the novel collision avoidance algorithm to be developed has the potential of improving road safety and eventually contributing to the goal of safer roads with less accidents especially those due to driver error. Connected and autonomous vehicle collision avoidance algorithms developed and evaluated in this approach with natural injection of unexpected real vehicle and vulnerable road user motion will also reduce cost and improve efficiency of development and deployment.

Final Research Report: N/A

Research Project Name: Enhancing Transportation Policy and Planning for Safety

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Chris Hendrickson, Carnegie Mellon University; Corey Harper, Carnegie Mellon University; Karen Lightman, Carnegie Mellon University

Project Partners: Regional Industrial Development Corporation (RIDC); Duquesne Light Company

Research Project Funding: Federal Funding—$95,000; Non-Federal Funding—$97,174

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: Transportation policy studies and improved planning are essential for furthering goals of the University Transportation Centers and the US Department of Transportation. This project is intended to build upon long-standing and successful activities in these areas. Three tasks are envisioned.

First, we will continue to produce policy briefs and professional papers on safety and enhancing usage of battery electric vehicles (BEV). In our 2024-2025 Safety21 project, a forthcoming policy brief and papers were developed on light duty BEV safety and policies to enhance adoption of BEV (Naiek 2024, Harper 2024, Hendrickson 2025). Safety concerns addressed included BEV fires, vehicle weight and stopping distance. Enhancing adoption was informed from ownership characteristics compiled from the 2022 National Household Travel Survey and existing purchase subsidies (Wang 2024). In 2025-2026, we propose to extend this work to other vehicle types (notably trucks) and a focus on workforce training, particularly vehicle technicians and first responders to crashes and environmental hazards (such as floods). BEV have particular needs for avoiding electrical shocks and short circuits. We expect to contribute to a series of policy briefs on BEV in the next year.

Second, we will continue activities intended to promote adoption of effective vehicle automation. In our 2024-2025 project, we actively participated in National Academies of Science, Engineering and Medicine (NASEM) studies on road safety research (NASEM 2024), critical issues in transportation (NASEM 2024), a global Council of Academies of Engineering and Technological Societies (CAETS) study on the e-mobility transformation (ICAETS 2024), and an empirical study of partial automation impacts. The weekly Safety21 Smart Transportation Dispatch (produced by project participant Karen Lightman) features safety improvements due to new technologies. In the coming year, we plan to actively comment on proposed federal and state rules regarding automation, respond to requests for information, and track performance of partial automation vehicles in a policy brief. This coming year has several likely rule-making activities, potential vehicle manufacturer announcements on partial automation, examination of the TESLA driving system, and possible implementation of the NASEM 2024 study recommendations. Monitoring and influencing these changes is a challenging task.

Third, project participants will continue to work with Regional Industrial Development Corporation (RIDC) in the planning for Pennsylvania Safety Transportation and Research Track (PennSTART), a safety, training and research facility for autonomous vehicle testing and emergency responders. The results of both tasks 1 and 2 can help inform appropriate test scenarios for PennSTART.

US DOT Priorities: Section left blank until USDOT's new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: We anticipate several professional papers (on BEV truck safety), one or more policy briefs, proposed rule comments, and presentations at professional conferences such as the Transportation Research Board. PennSTART will begin construction and partial operation during the course of this project.

Outcomes/Impacts: This project will have a variety of outcomes and impacts: (1) The literature and policy attention to battery electric vehicle safety is scant. Professional papers and policy briefs from this project will be a major incremental contribution. As the number of battery electric vehicles grows, safety concerns are likely to become important for prospective owners and for policy makers. (2) Most of the media and policy attention to vehicle automation focuses on more extensive and sophisticated automation. However, partial automation is proven to avoid and mitigate crashes. Moreover, partial automation continues to grow in the vehicle fleet. Enhancing this performance and extending to other vehicle types and situations is a goal of this work. (3) Aiding the implementation of PennSTART could help transform the state of use for automated and connected vehicles as well as emergency personnel training.

Final Research Report: N/A

Research Project Name: Safe and Efficient Automated Freeway Traffic Control - Phase 3

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Benjamin Coifman, The Ohio State University

Project Partners: DriveOhio

Research Project Funding: Federal Funding—$73,762; Non-Federal Funding—$64,129

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: Shockwaves are a naturally emerging phenomena in freeway traffic, but they represent one of the largest safety risks on freeways. Freeway drivers do not expect to encounter abrupt drops in speed or stopped traffic, as a result, shockwaves sharply increase the accident rates, particularly in the context of rear end collisions. For example, US interstate highways in 2021 saw the following rear-end collision numbers: Fatality 985, Injury-Only 71,408, Property-Damage-Only 152,011. Rear end collision severity is directly related to the relative speed between the involved vehicles, shockwaves increase these relative speeds, and thus, they also increase accident severity. Shockwaves also reduce freeway capacity and have a detrimental impact on traffic flow.

Connected and autonomous vehicles (CAV) hold the promise to attenuate and eliminate shockwaves (and thus, also reduce the severity and number of accidents), but only if the system is explicitly designed to do so. The very factors that give rise to shockwaves in human driven vehicles (HDV) will also do so in CAV. While CAV offers new ways to manage traffic dynamics, an automated freeway will still be subject to traffic dynamics. The real challenge is designing the CAV system so that it ensures the safest possible operation, and then within those bounds, the greatest operational efficiency (maximizing capacity, minimizing delays, etc.).

This proposal builds from earlier Safety21 research. Phase 1 developed the transition from unstable stop and go traffic to a smoother trajectory for the first CAV. It uses current conditions along the corridor to estimate the optimal trajectory for the CAV encountering congestion. Phase 2 extended the control to a platoon of CAV, suppressing major disturbances, typically with the 15th vehicle exhibiting a nearly constant speed trajectory even while the 1st vehicle in the platoon encountered stop waves.

The proposed research for Phase 3 seeks to address several key aspects. First, extend the methodology to actively manage multi-lane flows, including sequencing and optimizing lane change maneuvers. In all phases of this work, the HDV dynamics ahead of the first CAV come from empirical microscopic data sets. Phase 2 simply accepts the lane changes when and where the HDV happened to change lanes. But these maneuvers disrupt the platoons and degrade performance. Phase 3 will actively handle the lane change maneuvers so that they no longer disrupt platoons and in fact, can be used to robustly support the platoons. Another limitation of Phase 2 is that the empirical data only covers 1/3 of a mile, as a result, it is difficult to match demand of vehicles from upstream with the variable supply of capacity downstream. This research will develop ways to use multiple platoons to buffer against capacity fluctuations and provide finer control, e.g., dynamically advancing vehicles from upstream platoons when downstream capacity exceeds forecasts and slowing platoons when it falls below forecasts. In this way, the proposed research anticipates and develops the management strategies needed to maintain the smoothly flowing queued traffic with CAV.

US DOT Priorities: Section left blank until USDOT's new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: It is anticipated that the quantifiable outputs of this research will include one Master's thesis and several scholarly publications. While harder to quantify, we believe the goal oriented approach, "what do we want CAV to do and how do we get there?" will help the work achieve impactful results. As discussed in the body, we seek to move from simply addressing disturbances as they arise to proscribing the target traffic state over time and space to prevent the disturbances from arising, which will address the primary objective of improving safety (eliminating unexpected speed drops), while also leading to secondary benefits of increasing throughput (stable traffic has a higher capacity than fluctuating traffic states).

Outcomes/Impacts: Since the earliest autonomous vehicle research the focus has primarily been the very challenging task of making the technology work. While the fact that automated vehicles would benefit operations was taken as self-evident. To date, these priorities were well placed, but now that we are at the cusp of finally deploying safe and reliable CAV, it is important to take a step back and critically assess exactly how CAV can and should impact operations. Much of the past motivation was built on the recognition that CAV can maintain tighter headways, which increases vehicle throughput locally. This research seeks to take a step back from such a point perspective, moving to a proactive corridor perspective. Freeway operations are ultimately constrained by a few links that either have the lowest capacity or highest demand. Away from these bottlenecks is excess local capacity that we plan to use to optimize the flow of traffic through the bottlenecks. This research seeks work outward from bottlenecks to both supply demand to meet (potentially dynamic) capacity and manage any excess demand in a stable manner where the waiting vehicles do not give rise to shockwaves. By mitigating shockwaves it will improve safety and increase throughput.

Final Research Report: N/A

Research Project Name: UTC Safety21: Improve Transportation Infrastructure (Geotechnical Asset) Safety by Reducing the Risks of Landslides - Phase 3

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Zhuping Sheng, Morgan State University; Yi Liu, Morgan State University

Project Partners: Maryland DOT/SHA

Research Project Funding: Federal Funding—$292,680; Non-Federal Funding—$292,680

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: Morgan State University team proposes a multi-phase (multi-year) project focusing on safety of transportation infrastructure systems by preventing geohazard, specifically slope failure and landslides and minimizing impacts of geohazard. This project will employ an integrated approach of geotechnical and AI/Machine Learning methods for assessing conditions of geotechnical assets, such as cut slopes and embankment of the DOT SHA and delineating landslides and high-risk areas.

The objectives (tasks) of the proposal include: (1) with AI/Machine Learning approaches assess the risks of landslides based on soil/rock types, weather conditions, mechanical properties of slope materials, stream gage station flow data, pavement material and design, and the status of existing retaining structures along the selected highway sections, using Maryland as case studies and railroad, (2) identify and map the high-risk areas based on controlling factors such as geometry and mechanical properties of soil or rock, and triggering factors, including gravitational and hydraulic forces, using available survey data, remote sensing and LiDAR and InSAR data and other factors like transportation modes, (3) design and test protocols for real time monitoring at selected sites in consultation with agencies and organizations (such as DOT SHA, Federal Railroad Administration), and (4) recommend strategies for reducing the risks of landslides with real-time monitoring for the high-risk areas, and improving the safety of the transportation infrastructure. All the methods and strategies can be transferred to other states or regions with similar geological conditions and engineering configurations. Phase 1 of this project covered task 1 and part of task 2. Phase 2 covers part of task 1 and task 2. Phase 3 will cover task 3 and will also include literature review and summary of landslides along the railroad. This project will primarily complement the ongoing project sponsored by the Maryland DOT SHA (see more information in TRID) led by Zhuping Sheng in collaboration with CMU (Dr. Sean Qian). In this phase we will also expand our collaboration with the CMU team (Dr. Christoph Mertz) by including technology transfer in photographical images processing to build conceptual models and identify slope failures.

US DOT Priorities: Section left blank until USDOT's new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: The following are expected outputs for the Phase 3 of the proposed project: Journal articles on methodology and case studies; Conference presentations at professional conferences and technical meetings; Literature review summary of landslides on the railroads and preliminary findings of distribution patterns; Enhanced GIS coverages (LiDAR and InSAR data) of detail delineation of selected landslides including those on railroads at two or more selected regions based on data availability from USGS by the end of project; Integrated models for landslides risk assessment with physical models and machine learning models (modified or new): Calibrated models will be completed at selected sites (one watershed or selected counties within Maryland based on the preliminary work last two years) by the end of the project; Initial prototype of the monitoring system for early warning and detection of landslides will be deployed at one or two sites in consultation with Maryland DOT/SHA by the end of the project; Training workshop for SHA staff and other professionals as part of workforce development in transportation safety and disaster management; Presentations in webinar series related to Adoption of Cyberinfrastructure and Research Workforce Development for Disaster Management to promote UTC program.

Upon the completion of the whole project, the following deliverables are expected: Maps of landslides risks and make them accessible through DOT SHA or UTC; Curricular materials (methodology, case studies of landslides, monitoring protocols and more) for classroom teaching for slope stability analysis and risk assessment of landslides and improvement of transportation safety for different transportation modes; Protocols for monitoring of slope movement and failure as well as warning systems; Guidelines for mapping risk areas and the real-time monitoring and alarm system for high-risk areas in cooperation with DOT SHA.

Outcomes/Impacts: Besides advance in research and teaching related to transportation infrastructure safety and disaster management, the project is expected to bring broad positive impacts on transportation system in the following aspects: Early detection of slope failure with warning and monitoring systems, especially for those road sections with recurring landslides, could help to design and implement measures to provide safe routing to avoid injury to vehicle occupants and minimize damage to the physical infrastructures. It is expected that early warning system could detect slope failures in advance one week or two weeks, reducing potential failure by 20 to 30%; Reliable assessment of landslide risk with new technology, such as AI/ML approaches and image interpretation, could help with preventive maintenance to avoid slope failures, assuring safe and reliable roadways and saving costs by reducing the impacts of potential failures. Pilot systems can be deployed at 1 or 2 landslides-prone segments; Research findings will help policy makers to improve the roadway design standards and secure resources to assure long-term safety of transportation systems under extreme weather conditions.

Research results, including assessment approaches and design protocols, are transferable to other states and regions with similar geological and hydrological conditions and roadway designs.

Final Research Report: N/A

Research Project Name: ViCARUS: Vision-centric Connected Autonomy for Vulnerable Road Users' Safety

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Vijayakumar Bhagavatula, Carnegie Mellon University

Project Partners: General Motors

Research Project Funding: Federal Funding—$79,247; Non-Federal Funding—$97,174

Project Start and End Date: 7/1/25 to 6/30/26

Project Description:

Motivation

In the United States, pedestrian fatalities have risen dramatically, with a 57% increase from 4,779 deaths in 2013 to 7,522 in 2022 (NHTSA, 2024). Intersections remain particularly dangerous, accounting for approximately 25% of all traffic fatalities and 50% of injuries (FHWA, 2024). Vulnerable Road Users (VRUs) - including pedestrians, cyclists, and motorcyclists - are especially at risk in these scenarios. On the other hand, early road tests show autonomous driving systems (Waymo Safety Report, 2024) have the potential to do better than human drivers safety-wise. However, our previous work under Safety 21 UTC project #440, Connected Vision for Improved Pedestrian Safety (CVIPS) (Bhagavatula et al., 2024) has shown that many state of the art computer vision algorithms, when tested on autonomous driving benchmarks such as nuScenes (nuScenes detection task, 2020) show significant performance degradation when detecting VRUs as shown in Fig 1 of supplementary materials.

Recent studies suggest that connected autonomous vehicles using V2X communication can offer enhanced capability in detecting and responding to potential safety critical scenarios (Gao et al., 2022), particularly in complex urban environments. However, current V2X datasets and methods are limited in their scope to single-vehicle detection. Extending this by benchmarking perception performance for VRUs and other road users in a multi-agent environment is a critical opportunity for improving road safety (Xiang et al., 2024). In addition, As part of CVIPS (Bhagavatula et al., 2024), we simulated the impact of communications limitations (Shenkut & Vijaya Kumar, 2024) such as latency, bandwidth, and communication interruptions. Fig 3 and Fig 4 of supplementary materials show these impacts in different V2X scenarios. As shown in Fig 3, communication interruption from even a single agent causes ~25% performance degradation. Fig 4 illustrates the impact caused by communication latency, where a delay of about 200 ms causes a 10% performance drop in detection. These communication limitations make it hard to consistently detect and track VRUs, especially in scenarios with occluded views and varying weather conditions. Building upon the findings from CVIPS, ViCARUS aims to enhance VRU safety through four primary research thrusts:

Primary Research Thrusts

Investigation of Detection Performance Gap: Extending our CVIPS findings regarding camera-only systems' inferior VRU detection performance by analyzing the disparity in detection performance between camera-only systems and sensor fusion methods and studying specific scenarios where detection gaps are most pronounced and developing mechanisms to minimize these detection gaps through cooperative perception.

Cooperative Perception Framework development: Implementing a connected perception approach where cameras in vehicles and infrastructure units share information via V2X to improve detection performance on VRUs.

Simulation-to-Real Adaptation: Creating comprehensive synthetic datasets using CARLA, developing scenarios not well-represented in existing datasets, using existing real-world datasets such as WTS (Kong et al., 2024), V2X-Real (Xiang et al., 2024), TUMTraf V2X (Zimmer et al., 2024), designing adaptation techniques to bridge the gap between synthetic training data and real-world performance and developing computationally lightweight methods suitable for real-world deployment

V2X Communication Analysis: Using advanced C-V2X simulators (e.g., Fraunhofer IIS C-V2XSim platform (Simulation C-V2X, 2024)) to investigate the impact of critical C-V2X parameters such as Bandwidth, latency, packet losses and error rates, GNSS-related location errors and quantifying tradeoffs between communication constraints and system performance

Implementation Strategy

The project will be implemented through the following tasks: Creation of synthetic image sequences and acquisition of existing real datasets with different VRU scenarios; Development and testing of deep learning-based cooperative perception algorithms on both synthetic and real datasets; Developing methods to adapt algorithms trained on synthetic data to real-world datasets (Sim2Real adaptation); C-V2X parameter impact analysis; Evaluation under challenging imaging conditions and safety improvement demonstration.

US DOT Priorities: Section left blank until USDOT's new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: The anticipated outputs from this project are as follows: Synthetic image and image sequence datasets that include vehicles, pedestrians and other VRUs at intersections equipped with connected cameras, including in challenging imaging conditions; Publications describing the deep learning algorithms for VRU trajectory estimation in images and image sequences collected by multiple cameras at intersections, subject to C-V2X constraints; Software implementations of the developed deep learning algorithms; Technical report quantifying the tradeoffs between multi-agent communication parameters and VRU trajectory estimation accuracy.

Outcomes/Impacts: Outcomes of this work include: Enhanced VRU safety through improved detection and prediction capabilities in real-world scenarios; Robust VRU detection performance in degraded imaging conditions; Practical insights into C-V2X system constraints and their impact; Multi-agent perception dataset with various weather and lighting conditions; Adaptation techniques to bridge the gap between synthetic training data and real-world performance; Computationally lightweight methods suitable for real-world deployment.

These outcomes will bring closer the vision of connected vehicles leading to increased safety, particularly for vulnerable road users.

Final Research Report: N/A

Research Project Name: Humanoid as Human Assessment for the Safety of Vulnerable Road Users

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Ding Zhao, Carnegie Mellon University

Project Partners: PennSTART

Research Project Funding: Federal Funding—$100,000; Non-Federal Funding—$100,000

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: Despite rapid advancements in autonomous driving and connected vehicle technology, significant challenges persist in the safety assessment of the new technologies. Ensuring the functional safety of these complex systems is essential, yet there still needs to be a critical gap in methodologies that can comprehensively cover all necessary testing conditions. This gap is especially pronounced when evaluating scenarios involving vulnerable road users, defined as pedestrians, bicyclists, individuals using personal conveyances, and those in work zones as outlined by the ANSI D16.1-2007 standard and U.S. safety regulations.

Three main challenges highlight the complexity of this task. First, critical scenarios involving vulnerable road users are rare in real-world settings, making it difficult to gather sufficient data for testing. Simulated environments often lack the fidelity to replicate nuanced interactions, complicating the evaluation of autonomous vehicle performance in high-risk situations. Second, despite advances in simulation technologies, a sim-to-real gap remains due to inherent limitations in simulating vehicle behavior and environmental complexity. This requires real-world testing to validate simulation findings, which can otherwise undermine confidence in autonomous vehicle safety. Third, most real-world evaluations rely on simple dummies to represent pedestrians, failing to capture diverse, realistic human behaviors and interactions. This limits the ability to evaluate autonomous systems under realistic conditions thoroughly.

To the best of our knowledge, neither the academic community nor industry leaders have converged on a robust, standardized approach for conducting real-world evaluations that accurately reproduce complex, high-risk scenarios involving vulnerable road users and intricate traffic conditions. Addressing these challenges requires multidisciplinary research and innovative solutions that enhance both simulation fidelity and real-world testing methodologies.

Inspired by advancements in large language models (LLMs), we propose to develop a specialized foundation model for the evaluation of autonomous vehicles, focused on scenarios involving vulnerable road users. We aim to create a platform that integrates humanoid robots capable of simulating these road users' movements to test autonomous vehicle responses. Leveraging our expertise in critical scenario generation, we aim to build a foundation model that produces scenarios posing safety challenges to vulnerable road users. These scenarios will be generated randomly, conditionally based on text prompts, or through retrieval-augmented generation using specific scenarios as input.

We will also develop a platform with humanoid robots programmed to replicate the behaviors of pedestrians, scooter users, and individuals in wheelchairs. Our research will include developing algorithms to simulate the movements of pedestrians with specific physical limitations, ensuring a comprehensive range of real-world testing. Additionally, we will apply our prior work in accelerated evaluation to enhance testing efficiency. This initiative seeks to address current validation challenges and improve the safety and reliability of autonomous systems.

We have built a strong connection to PennSTART who will deploy the vulnerable robots to their test track which will be ready in 2026. Over the next half-decade, while Pittsburgh leverages its strong industrial background and leadership in the autonomy sector, it will face increased national and international competition. By contributing to autonomous driving evaluation and focusing on vulnerable road user safety, this project aims to reinforce Pittsburgh's position as a leader in the industry and support its continued innovation and growth.

US DOT Priorities: Section left blank until USDOT's new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: We plan to transfer the technology to PennSTART through a CMU spin-off company Trustelligence llc led by the PI. Additional outcomes for the project include: A dataset comprising real-world crash data paired with reconstructed scenarios, as well as a dataset of generated critical scenarios that present significant risks to vulnerable road users and include corresponding responsibility analyses based on traffic rules; A foundation model capable of following textual and visual guidance to conduct scenario analysis, identify responsibility according to traffic rules, and create critical scenarios based on multi-modal inputs; An AV evaluation system that leverages the foundation model to generate and execute critical scenarios involving vulnerable road users and process the resulting evaluation data; A uniform robot platform designed to autonomously perform the necessary behaviors of vulnerable road users and execute scenarios for AV evaluation, with built-in capabilities to collect and report evaluation results to the evaluation system.

Outcomes/Impacts: The primary anticipated outcome of this project is the development of a comprehensive foundation model and a uniform VRU platform utilizing a humanoid robot. This integrated system will include a foundation model enabling users to specify scenario requirements; a simulation environment for selecting critical scenarios; a humanoid robot platform capable of executing these scenarios and performing physical evaluations for autonomous vehicles; and an overarching system that connects these components and collects evaluation data.

The foundation model will empower users to create customized scenarios using multi-modal inputs such as images, text descriptions, or structured reference scenarios. The uniform platform will enable the simulation of various types of vulnerable road users, including pedestrians, cyclists, scooter users, and individuals using wheelchairs or strollers. We will collaborate with PennSTART to design and implement the real-world evaluation system, ensuring comprehensive and realistic assessments of autonomous vehicle performance.

This research will deliver a solution for real-world evaluation of autonomous vehicles, focusing on the safety of vulnerable road users with greater behavioral fidelity than conventional dummies. Additionally, this work will result in an open-sourced foundation model for autonomous vehicle verification and validation, incorporating specialized capabilities for understanding traffic rules and the behavior of vulnerable road users.
Final Research Report: N/A

Research Project Name: Analysis of Contributing Factors in Crashes Involving Electric Vehicles and Vehicles with Warning Systems and Level 1 Automated Features: A State Level Analysis

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Corey Harper, Carnegie Mellon University; Chris Hendrickson, Carnegie Mellon University

Project Partners: PennDOT; City of Pittsburgh

Research Project Funding: Federal Funding—$100,000; Non-Federal Funding—$100,000

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: Most light-duty vehicle (LDV) crashes occur due to human error. The National Highway Safety Administration (NHTSA) reports that about nine percent of fatal crashes in 2019 were distraction-affected crashes, while close to ninety-four percent of all crashes occur in part due to human error. Crash avoidance features could reduce both the frequency and severity of light and heavy-duty vehicle crashes, primarily caused by distracted driving behaviors and/or human error by assisting in maintaining control or issuing alerts if a potentially dangerous situation is detected. As the automobile industry transitions to partial vehicle automation, newer crash avoidance technologies are beginning to appear more frequently in non-luxury vehicles such as the Honda Accord and Mazda CX-9. Additionally, the market penetration of electric vehicles (EVs) is increasing, in turn increasing the weight and size of vehicles on the road. Although there has been research done to better understand the effects of warning and partial automation systems and EVs on traffic safety at the national level, we are not aware of any analysis that has been at the state or local level. While conducting research at the national level helps to inform federal policy, case studies are needed at the state and local levels to help better inform state and local policy and infrastructure investment decision-making. This project develops a replicable, open, deployable model that can: 1) identify crashes that occurred in Pennsylvania (PA) involving EVs or vehicles with warning or Level 1 systems, 2) assess the distribution of crashes involving EVs or vehicles with warning or Level 1 systems across factors such as weather conditions, vehicle speed, crash severity, pre-crash movement, facility type, and time of day, 3) estimate the relationship between contributing factors and the severity of crashes involving EVs or vehicles with warning or Level 1 features using regression analysis, and 4) make recommendations to the Pennsylvania Department of Transportation (PennDOT) and the City of Pittsburgh on policy and infrastructure improvements to improve traffic safety.

Our hypothesis is that current and past vehicles with warning systems and Level 1 features perform well in certain driving scenarios, facility types, and weather conditions but not others. However, there are not many frameworks and tools to help state and local agencies better understand the causes of crashes involving partially automated features, the locations of crash hotspots, and the infrastructure improvements and policies that could enhance road safety with automated vehicles (AVs). Additionally, this study should help technology providers better understand under which scenarios and conditions further testing and technology development is needed. We also hypothesize that the characteristics and patterns of crashes involving EVs differ from crashes not involving this technology.

Most studies that have assessed the role of warning and partial automation systems on traffic safety have focused on national-level analysis (Harper et al., 2016; Khan et al., 2019). This study builds off of previous UTC projects and adds to the literature by being the first analysis we are aware of to provide a framework for state and local agencies to assess factors contributing to crashes involving vehicles with warning systems and Level 1 features and identify locations and the infrastructure upgrades needed to support today's sensing systems.

In the first part of this project, we will compile crash data to better understand the frequency of crashes with warning systems and Level 1 features across contributing factors (e.g., weather or vehicle speed). Second, we will incorporate this data into a regression model to better understand the relationship between the contributing factors and the frequency and severity of crashes. Third, we will conduct an exploratory analysis on the contributing factors for crashes involving EVs. Finally, we will make recommendations to PennDOT and the City of Pittsburgh on policy and infrastructure improvements to improve traffic safety. To do this analysis, we will use the PennDOT crash dataset, which contains information on all police-reported crashes of all severities for PA.

In this phase of this project, we will mainly focus on contributing factors for crashes involving vehicles with warning systems and Level 1 features and will use similar datasets to conduct a more in-depth analysis on the contributing factors for EV crashes in future iterations. Because of Tesla's dominance in both the EV and AV market space it's important to consider both the role that EVs and AVs play in crash safety. By focusing on both technologies independently, we can better understand the different ways they are impacting crash safety and how to mitigate any negative effects through policy.

US DOT Priorities: Section left blank until USDOT's new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: Anticipated outputs from this proposal include: (1) Novel methods and frameworks for state and local agencies to evaluate the safety implications of EVs and AVs; (2) Data sets, modeling results, and online/offline tools to deploy and validate the proposed methods; (3) Memorandum to the City of Pittsburgh and PennDOT and the development of a policy brief to inform future policy-making; (4) Journal paper to disseminate results to research community.

Outcomes/Impacts: This project will have a variety of outcomes and impacts: (1) Designing the framework to assess the contributing factors for crashes involving AVs and EVs lays the foundation for state and local organizations to better understand the infrastructure improvements and policies that are needed to enhance AV and EV safety; (2) Working closely with the City of Pittsburgh and PennDOT could help inform local and state crash safety policy as the number of EVs and AVs on the road increase. Additionally, this will provide an opportunity for students to network and learn the missions and goals of City of Pittsburgh and PennDOT, providing them with pathways to obtain internships and higher paying jobs after graduation (e.g., director of planning, policy analyst); (3) Journal papers and policy briefs can inform decision makers and researchers about pathways to help achieve safety goals such as zero roadway fatalities.

Final Research Report: N/A

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

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Pingbo Tang, Carnegie Mellon University; Peter Zhang, Carnegie Mellon University

Project Partners: Regional Industrial Development Corporation (RIDC); PITT OHIO; Clarience Technologies, LLC; Safety Emissions Solutions, LLC

Research Project Funding: Federal Funding—$90,000; Non-Federal Funding—$110,000

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: This project addresses challenges by clarifying the definition of smart trailers and establishing a data-driven framework to evaluate their safety, efficiency, and economic benefits. Smart trailers integrate advanced sensors, telematics, and artificial intelligence (AI) to monitor operations, predict risks, and optimize maintenance schedules. A key focus is determining the direct and indirect costs and benefits of equipping trailers with side and rear underride guards and evaluating proactive warning systems or additional technical and policy solutions to reduce trailer-related collisions. Tackling the limited real incident data, the study will integrate data-driven analysis with physics-informed machine learning to guide policymakers in developing practical, safety-enhancing regulations.

To achieve these objectives, the project will synthesize data and physics-based models from multiple sources, including accident reports, vehicle physics models, historical crash data, regulatory documents, and insurance industry records. This information will support the development of simulation-driven models to assess both the safety and economic implications of underride guards or alternative solutions. By leveraging these models, the research will provide actionable insights into the optimal design and implementation of safety measures and how existing smart trailer technologies can support evidence-based regulatory decisions.

Collaboration with key partners will ensure the research's practical relevance and impact. Clarience Technologies will evaluate current smart trailer technologies, identify limitations, and help define a standardized framework. Safety Emissions Solutions will contribute inspection and regulatory data to enhance cost-benefit analyses, while RIDC will assess how smart trailer technologies can address regional transportation challenges and resource allocation challenges. Pitt Ohio's participation in validating the system through real-world testing and operational data collection is tentative and subject to confirmation.

US DOT Priorities: Section left blank until USDOT's new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: Development of a standardized framework defining the essential components and capabilities that constitute a 'smart trailer,' establishing industry benchmarks for technology integration and performance. Creation of physics-informed machine learning models capable of simulating vehicle dynamics and predicting safety outcomes under various scenarios, aiding in the assessment of underride guard effectiveness. Design and implementation of a platform that consolidates data from accident reports, vehicle physics models, highway design specifications, regulatory documents, historical crash data, and insurance reports, facilitating comprehensive analysis and informed decision-making. Formulation of evidence-based guidelines and recommendations for policymakers regarding the implementation and design of side and rear underride guards, supported by rigorous cost-benefit analyses. Additionally, this project will collaborate with national workforce and industry stakeholders to establish a standardized definition of ‘smart trailer,’ ensuring alignment with safety and operational standards, while fostering a skilled workforce capable of advancing smart trailer technology.

Outcomes/Impacts: Comprehensive Evaluation of Smart Trailer Systems. A detailed assessment of existing smart trailer technologies, identifying their capabilities, limitations, and integration challenges within current fleet operations. This evaluation will provide actionable insights to optimize the deployment of smart trailer systems in various operational environments.

Standardized Definition of Smart Trailers. Development of a clear, industry-accepted definition of “smart trailers,” establishing benchmarks for technology integration, performance metrics, and safety standards. This definition will serve as a foundation for advancements in smart trailer technologies and regulatory alignment.

Data and Simulation Model Documentation. Thorough documentation of collected data and simulation models, including methodologies for assessing safety benefits and economic implications of implementing side and rear underride guards. These resources will offer evidence-based insights for stakeholders, enabling data-driven decision-making regarding trailer safety investments.

Predictive Simulation Models for Safety and Cost-Benefit Analysis. Development of simulation models that evaluate vehicle operations in various transportation scenarios, predicting the safety outcomes and cost-benefit implications of deploying underride guards. These models will support policymakers and industry leaders in making informed decisions about trailer safety solutions.

Policy Recommendations and Technical Tools. A policy brief with recommendations for state and local policymakers on integrating underride guards, smart trailer technology, and additional safety solutions to reduce fatalities and improve operational efficiency. Deliverables include actionable policy suggestions, such as “Truck Maps,” to assist government agencies in optimizing truck routes, usage regulations, and operational timeframes, along with work zone guidelines and technical implementation roadmaps.

Strengthened Industry Partnerships. Enhanced collaboration with industry stakeholders. These partnerships will facilitate knowledge exchange, adoption of best practices, and practical implementation.

Training of Future Experts. Contribution to workforce development by training Ph.D., master’s, and undergraduate students in advanced smart trailer systems, data analysis, and safety policy research, preparing the next generation of experts to tackle critical transportation challenges.

Final Research Report: N/A

Research Project Name: Enhanced Crash Risk Estimation in Urban Environments: Integrating Multi-Source Data and Advanced Modeling Approaches for the City of Pittsburgh

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Sean Qian, Carnegie Mellon University

Project Partners: University of Pittsburgh Medical Center (UPMC); Velo.ai

Research Project Funding: Federal Funding—$100,000; Non-Federal Funding—$100,000

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: Crash risk estimation is a critical aspect of transportation safety study, aimed at identifying high-risk locations and mitigating potential hazards/risks for road users. Limited resources can be allocated to those high-risk areas for effectively improving safety for the entire network. As travel demand grows, new mobility options emerge and traffic volumes increase, the complexity of factors influencing crash risk has also grown. Historically, crash risk estimation has been approached using statistical models, such as Safety Performance Functions (SPFs), which predict the expected crash frequency or risk on road segments based on exposure and other factors. However, the dynamic nature of urban environments and the growing availability of diverse data streams, such as video analytics, near-miss reports, real-time traffic data, and social media, present opportunities to enhance traditional methods. In particular, data directly measuring near-misses in addition to observed and rare accidents enable a more holistic and accurate measure of crash risks. These sensing and methodological advancements can improve model accuracy, reduce biases, and provide more granular insights into high-risk locations for both motorized and non-motorized traffic.

By leveraging long-term multi-source data related to crash risks, diverse sources such as police reports, traffic stops, and near-miss incidents from the city of Pittsburgh, and applying both statistical and machine learning approaches, this project aims to enhance the precision of crash risk estimations. It addresses key challenges related to data aggregation, modeling performance, and bias mitigation to provide more reliable safety insights that can inform future interventions.

Our approach is to estimate SPFs for both motorized and non-motorized crash risks in the city of Pittsburgh with multiple data sources, including historical crash data, UPMC urgent care data, police traffic stop reports, 311 near miss report, 911 Dispatch data, Waze reports (social media), Velo.ai video analytics data, and so on. Specifically, we process the video, medical, police and Waze data among all locations and in the past few years to either improve or validate the SPFs (namely current risk maps). We apply both statistical and machine learning models and select the best model based on model performance. We develop a new crash map and compare it with the current risk map network. We explore the potential biases and how to improve the accuracy of crash risk estimation in general. This work will be conducted particularly for the City of Pittsburgh, but can be replicated in any other communities in general.

Step 1: Conduct a comprehensive literature review on risk factors and video analytics algorithm for near-miss detection—We plan to carry out a comprehensive literature review on types of risk factors that make significant contributions to the crash risk (particularly near-misses related to vulnerable road users). This literature review helps collect necessary datasets for building robust SPFs to estimate crash risk for streets in the following task. In addition, we will review the approaches of detecting incidents from the video data, which could help improve our algorithm in detecting near-miss cases from the camera videos.

Step 2: Data collection and video processing for near-miss incidents—In this step, we will collect the related datasets which could benefit our safety performance function estimation from multiple data sources. Table 1 lists potential datasets that could be applied in this research project. Furthermore, we plan to carry out video processing with advanced deep learning techniques to detect the near-miss incidents from the camera videos from selected intersections. The result of the number of near-miss incidents and driving factors could be either used to improve the SPFs or validate the estimated SPFs.

Step 3: Estimate SPFs with both statistical and machine learning modeling approaches—We will estimate two sets of SPFs for motorized crash risk and non-motorized crash risk, respectively, with the data collected and processed in Step 2. We plan to apply both statistical (e.g., linear regression, Poisson regression, and Negative Binomial regression) models and machine learning models (e.g., random forest, gradient boosting decision tree, and neural network). We will compare the model performance and select the best one as our final model,

Step 4: Develop a crash map for all street segments in Pittsburgh, compare it with the original high injury network, and prioritize high-risk locations—We plan to use the estimated SPFs to predict the expected crash risk (both motorized and non-motorized) for all streets in the city of Pittsburgh and develop a crash map dashboard. We will compare the crash risk with those estimated in the high injury network and find the differences. We will explore the sources of bias and study how to improve crash risk estimation. Finally, we will rank the locations and prioritize high-risk ones for further examination.

US DOT Priorities: Section left blank until USDOT's new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: Enhanced knowledge and awareness about the effects of various data sets on estimating crash risks on roads, especially on roads with a history of fatal accidents. Improved understanding of how multi-source data can aid public agencies in achieving greater transparency and making more informed decisions regarding roadway safety operations, especially in the context of evolving technologies and the increasing availability of diverse data sources.

Outcomes/Impacts: The goal of this research is to develop a comprehensive and universally applicable framework of safety risk models and tools for regional transportation networks, leveraging multi-source data and near-misses estimation. The framework will be accessible through open-source codes available online, complemented by a prototype web application. This application will leverage multi-modal data collected over several years from the City of Pittsburgh, offering user interfaces to manage visualizations of risks attributed to each data source. It will also enable visualization of crash risks for selected road segments, categorized by motorized, nonmotorized, by time of day, by weekday and non-weekday, by road level, and by season. A case study in Pittsburgh will assist public agencies in setting guidelines for sensing, estimating and mitigating roadway risks. We plan to actively collaborate with local transportation agencies, including PennDOT, BikePGH, various departments in the City of Pittsburgh and UPMC, to assess their interest in these tools and integrate them into their regular operations.

Final Research Report: N/A

Research Project Name: Cost-Effective Radar-Based Depth Perception and Scene Interpretation

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): John Dolan, Carnegie Mellon University

Project Partners: Blue Fusion; Robotics Institute Summer Scholars (RISS)

Research Project Funding: Federal Funding—$36,622; Non-Federal Funding—$36,622

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: One of the current challenges for autonomous vehicles (AV) is providing provable safety in urban driving and highway driving. LIDAR provides accurate depth and shape information, but remains quite expensive. In this research we explore the viability of fusing 4D radar images with camera images as an alternative to using LIDAR. Our research aims to establish the data fusion of radar and camera images in autonomous driving as a replacement for data fusion using traditional LIDAR, radar, and camera images. Our current approach demonstrates that a trained radar detector can achieve approximately 3.5 meters of Unidirectional Chamfer Distance (UCD) against ground truth LiDAR data on a dataset recently released by Delft University of Technology (DUT) in the Netherlands, which is 28% better than the state of the art (SOTA).

Our approach integrates the images obtained from deep neural network (DNN) 4D radar detectors with conventional camera RGB images. It introduces a novel pixel positional encoding algorithm inspired by Bartlett's spectrum estimation technique. This algorithm transforms radar depth maps and RGB images into a unified pixel image subspace facilitating the learning of their potential correspondence. Our method effectively leverages high-resolution camera images to train radar depth map generative models, addressing the limitations of conventional radar detectors in complex vehicular environments while sharpening the radar output. We develop spectrum estimation algorithms tailored for radar depth maps and RGB images, a comprehensive training framework for data-driven generative models, and a camera-radar deployment scheme for AV operation.

Our research plan is as follows:

(1) Develop a deeper analysis of the training framework to develop more advanced deep learning-based radar detectors

(2) Scale current results on about 100 seconds of representative AV driving from a DUT dataset to a significantly larger number of frames, including a wider range of training and testing scenarios

(3) Quantify additional performance benefit from real-time (instead of current offline) fusion schemes and deployment frameworks for 4D radars and cameras

(4) Use the existing data to measure the underlying latency in the proposed data fusion of 4D radar and camera images. This will help quantify the capabilities and limitations of our scheme

(5) Conduct a detailed error analysis to understand the distribution of error between the point cloud we generate with our data fusion approach against the ground truth. This will help establish any existing bias in the error we obtain in the UCD

(6) DUT’s 4D radar training data set was generated in the Netherlands. We will investigate the augmentation of these data by data collected in collaboration with our deployment partner Blue Fusion, who will allow us the use of one or more of their high-performance radars. This effort will address one of the important pitfalls of machine learning systems, distribution shift or domain adaptation.

The results will be instrumental in providing a cost-effective and weather-resilient solution to the perception part of the provable safety challenge in L4 and L5 Autonomous Vehicles.

US DOT Priorities: Section left blank until USDOT’s new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: Outputs of this project include: Creation of a novel Bartlett's spectrum-based algorithm for fusing 4D radar and camera data; Creation of a radar-only system with depth estimation performance good enough to replace LIDAR; Creation of new datasets for evaluation of camera-radar fusion-based perception.

Outcomes/Impacts: Outcomes of this project include: Creation of a camera-radar fusion method that reduces perception cost for autonomous vehicles while maintaining high performance, thereby extending the availability of safe autonomous driving technology; Validation of the system in simulation and a broad range of datasets; Documentation of the results in a form appropriate for hand-off to a partner capable of broader deployment.

Final Research Report: N/A

Research Project Name: Safe Driving Across Domains

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Peter Zhang, Carnegie Mellon University; Pingbo Tang, Carnegie Mellon University

Project Partners: Tier IV and Autoware Foundation; Regional Industrial Development Corporation (RIDC); Cheikh Anta Diop University

Research Project Funding: Federal Funding—$99,828; Non-Federal Funding—$128,621

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: This project tackles the critical research question: *How can pre-trained autonomous driving models be seamlessly and dynamically adapted across diverse operational design domains (ODDs) to achieve safe and robust deployment?* ODDs define the specific conditions under which an autonomous driving system can operate safely—spanning factors like road types, traffic density, weather patterns, and geographic features. For instance, a model trained to navigate Pittsburgh's urban intersections must also adapt to Pittsburgh's rural routes with sparse signage and rolling hills, or poorly paved, mixed-traffic roads. What is the right way for a robot to, in real time, smoothly switch between different planning trajectories to ensure safety and consistency amid changing driving conditions?

By focusing on adapting pre-trained models rather than retraining them from scratch, this project offers a scalable approach to ensure safe deployment. This is particularly critical in data-limited settings—collecting training data for every possible scenario can cost millions and take years. For example, equipping autonomous systems to operate in rural areas avoids costly new data collection while preserving high safety standards. A single deployment failure in an untested environment could cost lives and erode public trust, emphasizing the importance of rigorous adaptation. With an estimated 1.35 million road traffic deaths annually worldwide in drastically different driving conditions, this project aims to ensure that autonomous driving technologies can be deployed safely. In other words, this research algorithmically and proactively catalyzes translation of new technologies to data-limited areas, instead of passively relying on commercial technology spillover – the latter would take years or even decades.

We leverage diverse datasets to ensure robust testing. These include driving data from TIER IV in Japan, and over 10 million miles of publicly available driving data from Waymo, collected across 25 U.S. cities and capturing diverse urban environments. In Pittsburgh, we have access to truck-mounted camera data that covers mixed urban and rural routes, while stationary cameras funded through a PITA partnership with RIDC capture real-time trajectory data in high-traffic industrial areas. Simulation-based fine-tuning complements these datasets, addressing data gaps in under-documented scenarios, such as rural roads with limited signage. This rich dataset ecosystem provides the foundation for testing and validating model adaptations in diverse operational design domains.

The project employs open source, cutting-edge algorithms, including inverse reinforcement learning, guided diffusion, blueprints, and fine-tuning. Inverse RL provides an overall framework to complement rule-based modules to directly and flexibly extract algorithm parameters (sometimes implicitly) from expert (human or simulated) driving. Guided diffusion allows for precise and smooth generation and adjustments of such parameters in constrained domains, such as adapting autonomous driving algorithms from Pittsburgh's dense urban streets to its hilly rural roads. Blueprinting ensures scalable and structured design of template models that are transferable across diverse contexts, while fine-tuning efficiently adapts pre-trained models to novel environments, particularly in low-data settings. Open source models already implement these algorithms. We adapt these powerful emerging methods to address challenges in rural areas and infrastructure-limited regions.

US DOT Priorities: Section left blank until USDOT's new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: This project will produce several outputs to advance autonomous driving technology. Such as adaptable autonomous trajectory generation and selection modules for the Autoware open source auto-driving stack, optimized for diverse operational design domains (ODDs), including urban, rural, and infrastructure-limited environments. These models will be developed and refined using cutting-edge methods such as inverse reinforcement learning, guided diffusion, and fine-tuning to ensure robust performance and transferability. The project will release these open-source software modules into the Autoware driving stack on Github, featuring advanced motion planning and risk-aware optimization algorithms. A comprehensive dataset of simulation and field-testing results will be shared with the research community, offering a valuable resource for studying the adaptation of autonomous systems across ODDs.

The research findings will also be prepared for publication in high-impact journals, such as Transportation Research Part A: Policy and Practice, INFORMS Journal on Applied Analytics, and IEEE journals on robotics and vehicular technologies. These publications will disseminate the project's results to both academic and practitioner audiences, accelerating adoption of the methodologies and insights. Additionally, the project may generate invention disclosures or patent filings related to risk-aware optimization and adaptation processes. Together, these outputs will drive safety and scalability in autonomous vehicle deployment.

Outcomes/Impacts: This project will produce adaptable autonomous driving models, open-source software modules integrated into the Autoware stack, and simulation-tested deployment frameworks for diverse operational design domains (ODDs). These outputs will improve safety by reducing collision rates in urban intersections, enhancing navigation accuracy on rural roads, and ensuring reliable operation in infrastructure-limited environments, such as Senegal's mixed-traffic conditions. By leveraging pre-trained models and fine-tuning instead of retraining, the project also reduces costs.

The research will directly influence policy decisions by promoting CAV deployment practices, ensuring certain areas are prioritized for rigorous testing before rollout. It will provide key stakeholders, including USDOT, state transportation agencies, and international partners, with actionable guidance to prevent unsafe deployment of autonomous vehicles, fostering safety in transportation systems.

The project's findings will benefit the transportation system by enhancing reliability through robust real-time adaptability, durability by addressing edge cases and diverse road conditions, and cost efficiency through scalable, low-resource deployment strategies. These advancements support a safer transportation future while setting a foundation for global regulatory frameworks and commercialization opportunities.

Final Research Report: N/A

Research Project Name: Deep Reinforcement Learning Based Driving Strategies Bootstrapped by Suboptimal Policies

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Keith Redmill, The Ohio State University; Ekim Yurtsever, The Ohio State University

Project Partners: DriveOHIO

Research Project Funding: Federal Funding—$83,927; Non-Federal Funding—$71,862

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: Designing autonomous driving systems requires a modular approach. This includes vehicle sensing, perception, localization, scene representation, path planning, decision-making, and vehicle control. Our study focuses on decision-making in complex driving scenarios. Traditional methods often use model-based, heuristic, or rule-based controllers. Heuristic and rule-based controllers rely on human expertise to handle different driving situations. They include safety measures and rules to avoid deadlocks. However, these methods face limitations as designing rules to cover every possible driving scenario is highly challenging and resource intensive. Reinforcement learning (RL), on the other hand, offers a self-supervised approach, learning control policies through interaction with the environment. With advancements in neural networks, deep reinforcement learning (DRL) has shown the ability to efficiently map state-action transitions in high-dimensional environments.

Driving environments encompass a range of challenging scenarios, such as roundabouts, overtaking, merging, intersections, and highway exits. These scenarios involve delayed rewards, which discourage exploration and complicate the RL learning process. Examples of delayed rewards include the cumulative benefit of using a free lane to achieve optimal cruising speed, waiting for the right moment to enter a roundabout safely, or merging into traffic while avoiding collisions—even if it requires temporarily stopping in the lane. Failure to explore and identify these states during training can lead to sub-optimal policies or even unsafe behaviors, such as remaining stuck behind a slow vehicle, rushing into a roundabout prematurely, or misjudging merges in traffic.

In this study, we propose to address these limitations by combining the stability of rule-based sub-optimal controllers with the adaptive learning capabilities of DRL. We aim to utilize sub-optimal policies to bootstrap efficient training during the RL process and to overcome exploration barriers. Specifically, we leverage a rule-based controller to prevent deadlocks and guide RL agents toward promising states, improving their ability to learn optimal driving strategies. This approach not only improves collision avoidance but also enhances training efficiency by guiding agents through challenging scenarios where rewards are delayed. By integrating suboptimal strategies, our method strikes a balance between simplicity and adaptability, enabling agents to learn more effectively.

Our experiments are conducted in a complex driving scenario where the driving policy must optimize long-term rewards. The key contributions of this study are summarized as follows: (1) The sub-optimal policies: We design sub-optimal controllers capable of making driving decisions under reasonable motivations; (2) Guidance for RL agents: The sub-optimal controllers guide RL agent toward exploring states associated with delayed rewards, improving exploration and learning efficiency; (3) Validation in complex driving scenarios: We validate the effectiveness of our method in complex driving scenarios, demonstrating its ability to optimize long-term rewards and handle challenging conditions effectively.

US DOT Priorities: Section left blank until USDOT's new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: This research project will produce novel methods, open-source software, publications about the proposed hybrid framework that combines rule-based driving strategies with reinforcement learning for robust automated driving: (1) The findings, methodologies, and analysis of this research will be thoroughly documented in a research paper. The paper will describe the hybrid framework that combines rule-based strategies with reinforcement learning. It will also explain the methods used to address exploration challenges and provide experimental results. This will serve as a useful resource for studying autonomous driving systems and reinforcement learning; (2) To promote further research and enable replication, the project will include the release of source code for the developed models and methods. Comprehensive documentation and user guides will accompany the code, making it accessible to the research community and supporting future advancements in driving decision-making systems; (3) To demonstrate the practical relevance of this research, a case study and customized simulation environment will be developed. These will illustrate the application of the hybrid framework in scenarios such as lane changing, efficient navigation, and collision avoidance, providing a foundation for testing and refining autonomous driving strategies in realistic scenarios.

Outcomes/Impacts: This research aims to produce significant advancements in the field of autonomous driving by combining rule-based strategies with reinforcement learning to improve decision-making in driving scenarios. The key outcome is a hybrid approach that enhances exploration and long-term reward optimization, addressing challenges faced by conventional reinforcement learning methods. By incorporating sub-optimal rule-based controllers, the framework facilitates safer and more efficient driving strategies, particularly in tasks such as overtaking and collision avoidance.

The results of this study are expected to have practical applications in autonomous vehicle systems, providing a robust solution for navigation and decision-making. The proposed approach aims to improve vehicle safety, reliability, and efficiency, offering a scalable strategy that can adapt to various traffic scenarios. Simulations and case studies will demonstrate the effectiveness of this method in real-world-inspired environments, highlighting its ability to deliver practical improvements to autonomous driving performance.

Beyond its application in complex driving scenarios, this research contributes to the broader understanding of how reinforcement learning can be enhanced through the integration of human driving logic. By focusing on the intersection of rule-based strategies and machine learning, the work has the potential to set new standards in designing safer and more reliable transportation systems.

Final Research Report: N/A

Research Project Name: Monitoring of Urban Roadway Safety Hazards from Existing Bus-based Video Imagery: Phase 3

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Rabi Mishalani, The Ohio State University; Mark McCord, The Ohio State University; Keith Redmill, The Ohio State University

Project Partners: The Ohio State University (OSU) Transportation and Traffic Management

Research Project Funding: Federal Funding—$102,237; Non-Federal Funding—$99,358

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: Traffic safety is diminished by drivers' changing lanes in queued traffic and mixes of pedestrians and vulnerable vehicles (scooters, bikes, and motorcycles) with large vehicles in the traffic stream. Assessing locations with these recurring but dynamic hazardous conditions requires extensive and ongoing monitoring. Traditional monitoring methods rely on sensors at permanent or temporary fixed locations, which is costly, labor intensive, and provides limited spatial coverage. Moreover, the location of these sensors may be influenced by factors other than optimal sampling, such as requests from well-organized constituencies. Therefore, traditional monitoring of hazards could miss high-risk conditions resulting in decreased safety.

Transit buses operate regularly over wide networks, and most bus fleets are already equipped with cameras that record the environment inside and outside buses for liability, security, and safety purposes. Consequently, the imagery is available for other uses at low marginal cost. Coupling the available imagery with the extensive spatial and repeated temporal coverage of transit fleets has been shown by the principal investigators (PIs) to be effective in monitoring traffic volumes across space and time. In the phase 1 project, the PIs investigated the use of available imagery recorded by cameras mounted on transit buses in regular operation to identify hazardous locations associated with recurrent vehicle queuing on urban streets. The phase 2 project is building on the phase 1 investigations by identifying locations with recurring noticeable volumes of vulnerable vehicles to serve as indications of hazardous conditions.

The disparity in size and weight between large vehicles and pedestrians and vulnerable vehicles contributes to the increased risk of accident occurrences that lead to serious injuries and fatalities. A deployment partner commented that near misses of such interactions "happen all the time." Similarly, and as was recently highlighted in The Economist Sept 7, 2024 issue titled "America's killer cars", the increasing size and weight of passenger vehicles is leading to an exponentially increasing risk of injuries to and fatalities of pedestrians and passengers of vulnerable vehicles. Therefore, identifying locations on the roadways where such risks are relatively high is critical to informing mitigation efforts. The proposed phase 3 of the study focuses on this identification. The representative quoted above also noted: "We don't record near misses, so we don't have any documentation but [they are] frequent", thus clearly highlighting the need. Therefore, phase 3 builds on phase 2 by considering the presence of near misses between pedestrians or vulnerable vehicles and large vehicles to investigate the potential of using bus imagery to identify locations of recurring high-risk combinations of these elements.

The OSU campus will again be used as a living lab. The size and diversity of land uses make the campus representative of urban areas, and the close and ongoing collaboration with OSU's Transportation and Traffic Management, the deployment partner and contributor to the success of the OSU campus transportation living lab, offers insights on safety concerns and the practical validations of the developing safety monitoring scheme.

US DOT Priorities: Section left blank until USDOT's new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: General research: New methods and technologies to identify safety hazards along urban roadways. Demonstration of quantifying vehicle type mix safety factors from video imagery.

Research based outreach: Discussion of identified safety vulnerabilities and higher risk locations on the OSU campus during regular meetings with the deployment partner, OSU's Transportation and Traffic Management (who will provide bus imagery). These regular discussions would increase the practical impact of the project's results and allow user grounded investigations. More specifically, in the medium term, mitigation measures could be taken by the deployment partner at higher risk locations identified, and the impacts could be assessed by the research team.

Outcomes/Impacts: The ultimate impact of this project will be the ability to determine the location and times of recurring roadway traffic conditions that could compromise safety across urban roadways using an existing and widespread source of available data, namely, video imagery obtained from cameras mounted on buses in regular transit service. More specifically, in the proposed phase 3 of this study, the focus is on identifying higher risk locations stemming from disparities in the traffic mix between large vehicles and vulnerable vehicles or pedestrians. The ability to systematically determine and monitor hazardous roadway conditions would provide input to policymakers and designers developing measures aimed at reducing the likelihood and severity of traffic and traffic-related accidents.

Final Research Report: N/A

Research Project Name: Generating Unsafe Road Activity from LLMs and Road Imagery to Learn Better Safety

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Robert Tamburo, Carnegie Mellon University; Srinivasa Narasimhan, Carnegie Mellon University

Project Partners: dashcam.bike

Research Project Funding: Federal Funding—$90,000; Non-Federal Funding—$99,630

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: The FHWA zero death vision is a strategic goal to eliminate all traffic fatalities and severe injuries. One of the pillars of this strategy is to advance life-saving technology in vehicles and infrastructure. Many such in-vehicle technological advances have been developed and deployed in the past decade such as backup cameras, lane departure warning systems, adaptive cruise control, emergency brake systems, driver attention monitoring, etc. However, deployments of technology in the infrastructure have lagged. Most infrastructure advances have centered around redesigning roads to control traffic and improve visibility of vulnerable road users (VRUs), i.e., anyone in the road environment unprotected by a vehicle's frame. Ideally, the infrastructure should also monitor road activity for unsafe conditions and communicate those conditions to drivers and VRUs. To accomplish this goal, the infrastructure needs to be instrumented with sensors to capture data for real-time analysis and devices for relaying information to all road users. With prior UTC support, we have addressed some of the building blocks needed for a 'smart' infrastructure such developing and deploying methods for reliably detecting and tracking vehicles, classifying anomalous activity, and estimating 3D vehicle parameters (speed and direction). These methods were developed using images collected from cameras deployed at intersections. The main challenges that we faced were difficulty identifying a variety of 'unsafe' events because they rarely occurred at the deployment locations and difficulty tracking pedestrians and bicyclists because they were rarely captured.

In the proposed work, we will address these challenges by developing a system to simulate unsafe vehicle, pedestrian, bicyclist, and road worker activity from previously acquired images and large language models (LLMs). The proposed work also enables safety technology development without exhausting resources for camera installation. Focus will be on generating visual data of normal and anomalous activities such as near misses, erratic driving, blocked bicycle lanes, etc. Methods will be developed to model realistic driver behavior based on observations from datasets acquired from prior UTC projects in collaboration with deployment partners City of Pittsburgh Department of Mobility & Infrastructure and Shaler Township. Realistic bicyclist behavior will be modeled using bicyclist-centric data provided by partner dashcam.bike. Methods will be developed to leverage these models with LLMs to generate unsafe activity involving vehicles and VRUs based on text entered by users. The system will be deployed to a website so that anyone can create their own dataset. The benefit of such a road activity system will enable safety research and permit data driven evaluation of road or intersection safety without installing cameras and storing/analyzing hours upon hours of images per camera.

US DOT Priorities: Section left blank until USDOT's new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: We anticipate the following outputs: Methods to simulate unsafe road activity based on real visual data of vehicles and bicycles; Easy to use web-based interface to simulate data from text prompts.

Outcomes/Impacts: It is anticipated that the road activity simulator will be a useful tool for infrastructure safety planning and developing road safety methods without acquiring or using images. Safety is generally evaluated by reviewing crash reports, but the standard for reporting is fairly high. Thus, some crashes are not reported and unsafe interactions are never reported. A text-based simulation environment provides an easy-to-use way of evaluating safety for any road environment.

Final Research Report: N/A

Research Project Name: Evaluating Autonomous Vehicles' Pedestrian Safety Benefits in Road Networks

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Carlee Joe-Wong, Carnegie Mellon University; Osman Yagan, Carnegie Mellon University

Project Partners: City of Pittsburgh; Transportation Research Center; Southwestern Pennsylvania Commission

Research Project Funding: Federal Funding—$90,000; Non-Federal Funding—$100,000

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: Pennsylvania experienced more than 3000 pedestrian-related injuries in 2022, according to PennDOT, and the Governor's Highway Safety Association estimates that Pennsylvania ranked 10th among all 50 states in 2022 for the most pedestrian fatalities. This project will explore and evaluate the potential for connected autonomous vehicles (CAVs) to reduce these pedestrian-involved incidents. While many past works have explored using CAV capabilities to detect and avoid pedestrians, we will specifically focus on ensuring that these benefits are felt where and when they are most needed.

Recent studies show that pedestrian fatalities in the United States are especially common at night and on roads that do not have infrastructure designed to support pedestrians. By increasing the concentration of CAVs, which may be safer for pedestrians, on such dangerous roadways, we may be able to increase overall pedestrian safety. Many pedestrian-related incidents, for example, involve human driver distraction, and CAVs can make use of non-visual (e.g., radar) environmental data to detect pedestrians, especially at night. However, CAVs may also be more dangerous for pedestrians if their detection systems fail. Moreover, intentionally diverting CAVs to areas that are dangerous for pedestrians may have implications for safety and traffic flow through the road network as a whole. For example, proactively routing CAVs to these areas may create more traffic congestion and potentially increase accidents involving vehicles, leading human-driven vehicles to avoid this area. Indeed, CAVs will likely be deployed initially in mixed-autonomy settings with human-driven vehicles, which adds an additional layer of complexity in analyzing how pedestrians will affect safety or traffic flow.

We plan to build on our 2023-4 Safety21 project investigating CAVs' safety effects in road networks. In particular, we will make use of our mixed autonomy simulator from the previous project, which includes models of vehicle-involved accidents but does not explicitly seek to improve pedestrian safety. Our work will comprise four parts: (i) integrating pedestrian fatality and injury statistics into our mixed autonomy simulator, (ii) developing analytical models for traffic flow and pedestrian injuries/fatalities given CAV routing, (iii) designing algorithms for CAVs to minimize pedestrian fatalities and injuries, and (iv) using the simulator and our analytical models to evaluate whether CAVs can effectively reduce pedestrian injuries and fatalities, as well as any tradeoffs with increasing traffic congestion.

US DOT Priorities: Section left blank until USDOT's new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: Our deliverables will be (i) new algorithms for CAVs to improve pedestrian safety by reducing injuries in fatalities at times and locations where they are more likely to occur; (ii) new models of CAVs' effects on pedestrian injury and fatality rates in mixed-autonomy settings; (iii) quantitative estimates of how much CAVs can reduce pedestrian injuries and fatalities, given a realistic road network and CAV prevalence; (iv) quantitative estimates of any performance tradeoffs between improving pedestrian safety and worsening traffic congestion; (v) regulatory implications of these quantitative findings; and (vi) extensions of our existing mixed autonomy simulator to include better models of pedestrian behavior, injuries, and fatalities.

We will work with our deployment partner, DOMI, to ensure the usefulness of our project outputs, through regular meetings for feedback. We will also seek feedback from deployment partner, the Transportation Research Center, as to whether our CAV models are realistic. We will consider leveraging DOMI's existing public participation processes to receive community feedback on our work. We expect that this feedback will allow us to account for metrics that are important for all communities affected by CAV deployments, helping to ensure that we address concerns associated with CAV deployments.

Outcomes/Impacts: Regulatory agencies, including our partner the City of Pittsburgh's Department of Mobility and Infrastructure, may be able to use our results to forecast the effects of gradually increasing CAV deployment, particularly with regards to pedestrian safety. Similar to our prior "Big Idea" project with Mobility21 and current project with Safety21, on which this project builds, our work could inform the policies that agencies may wish to take in regulating CAVs (e.g., developing guidelines for how CAVs should route themselves around a city). It will further provide quantitative validation for (or evidence against) claims that CAVs can benefit traffic and traffic safety as a whole. Finally, our simulator may aid such agencies and other researchers in studying CAV and human behavior in mixed-autonomy systems. We further plan to explore collaborations with the Southwestern Pennsylvania Commission, which was a partner for our prior work, as they may also find our results to be of interest.

CAV manufacturers and service providers, like our partner the Transportation Research Center, may also be able to benefit from this work, as it will help them quantify the importance of developing CAVs that interact "well" with humans and design learning algorithms that help CAVs to do so. These new learning techniques may enhance the benefits of CAVs and increase their potential for deployment. They may also help CAV service providers earn public trust of CAVs, by demonstrating that CAVs may have benefits for pedestrian safety.

Final Research Report: N/A

Research Project Name: Cyber Resilience of Connected and Autonomous Transportation Systems (Phase II): Game-theoretic Security Assurance of Network-level Traffic Signal Control

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Mohamadhossein Noruzoliaee, University of Texas - Rio Grande Valley

Project Partners: University of Texas - Rio Grande Valley

Research Project Funding: Federal Funding—$32,431; Non-Federal Funding—$54,577

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: Two trends are transforming transportation systems. On the bright side, the increasing complexity of cyber-physical technological advances enable connectivity and autonomy through seamless integration of computation, communication, sensing, and control, holding promise for societal and economic impacts. On the dark side, however, these tightly coupled cyber-physical interdependencies can be self-defeating as they can pose new vulnerabilities to accelerating disruptions in cyber space (i.e., cyber-attacks), raising concerns about the safety, security, and privacy of the transportation system users. In view of these two overlapping trends, it is crucial to bolster the transportation cyber-physical resilience. Building on the findings from a comprehensive literature review of transportation cyber resilience in Phase I of this project, Phase II focuses on security assurance for network-level (connected) and AI-powered (autonomous) traffic signal control systems.

Learning-based control methods, such as deep reinforcement learning (DRL), have gained traction in traffic signal control due to their computational capability to efficiently deal with real-time changes in traffic flow dynamics. Despite growing popularity, DRL has been recently found to be vulnerable to adversarial attacks, which precludes its use in safety-critical traffic signal control applications unless its vulnerabilities are mitigated. From a methodological standpoint, mitigating the security issues facing DRL is still a daunting if not impossible task. The well-established methods in the area of (un)supervised machine learning security do not methodologically suit DRL security, since DRL deals with sequential decision-making, i.e., non-fixed distribution of training data, based on interactions with an environment. The interdependence of the current decision on the ones taken in previous time steps increases the degrees-of-freedom of adversarial attacks on DRL, e.g., an adversary can perturb multiple state observations of the environment over time.

Given the above research gap in DRL security assurance, most relevant studies resort to empirical (aka heuristic) methods, in which attack models sample only some of the possible disturbances (e.g., self-crafted perturbed data) and the devised defense strategies do not generalize out of the experienced attacks. Leaping from empirical security toward certified security, this project proposes a game-theoretic approach to robustifying DRL-based traffic signal control at the network level. The game-theoretic approach provides computational security assurance through explicitly modeling an adversary's destabilizing perturbations rather than sampling only some of the possible perturbations. The network-level scope of the proposed project prevents higher-order cyber risks compared to isolated traffic signal control, through the explicit consideration of the underlying network effects, such as inter-agent (intersection) interactions over networks.

US DOT Priorities: Section left blank until USDOT's new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: The anticipated outputs of this research project include a game-theoretic data-driven model for adversarially robust network-level traffic signal control as well as model implementation results.

Outcomes/Impacts: The expected outcome of this research project is the capability enhancement in enhancing transportation safety through bolstering the cyber resilience of connected and autonomous transportation systems, exemplified through the context of network-level (connected) and AI-powered (autonomous) traffic signal control systems.

The anticipated impacts of this research project are 1) improved understanding of security assurance for safety-critical traffic signal control at the network scale; and 2) enlargement of the pool of trained transportation professionals at the nexus of transportation safety and cyber resilience.

Final Research Report: N/A

Research Project Name: Planning and Policy for Safer Roads with Autonomous Vehicles: Decision Making Behavior in Dilemma-inducing Situations - Phase II

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Fatemeh Nazari, University of Texas - Rio Grande Valley

Project Partners: University of Texas - Rio Grande Valley

Research Project Funding: Federal Funding—$50,495; Non-Federal Funding—$28,352

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: Despite the advances made in the technical availability of autonomous vehicles (AVs) for road tests and even limited commercial mobility services, a critical challenge in the integration of AVs into the real-world transportation systems and our lives is establishing safety and ethical regulations. In particular, principles are required for an AV to make safe and ethical decisions in high-impact low-frequency situations that induce dilemmas, such as whether to hit pedestrians to spare passengers onboard? The absence of such policies can concern the potential users about safety aspects of AVs ,which is identified to negatively affect the public acceptance of the AV technology, as highlighted by the findings of the PI's previous research.

This project is focused on designing these policies derived from the research on the potential autonomous system users' expectations and perception of a safe AV through exploring how the users themselves would make decisions that entail a dilemma-inducing situation. This project will be the second Phase of the ongoing project from the previous year. In particular, this proposed work will be built on and complement the first Phase through accomplishing three tasks explained below.

Task 1 - Launching Stated Preferences (S.P.) Survey to Collect Sample Datasets on Two Groups of Individuals. The S.P. survey is designed in Phase I using experimental design methods which is demonstrated to be a rigorous and efficient tool for collecting datasets on individual persons' opinions and preferences. In Phase II, the results of the pilot dataset collected in Phase 1 will be utilized for updating the survey. Then, the survey will be fully launched on two population groups with distinctive heterogenous attitudes. The PI has the experience of designing and dissemination S.P. surveys in a different context.

Task 2 - Estimating Modeling Frameworks Capturing Heterogeneity of Behavior on the Datasets Collected in Task 1. This task focuses on estimating methods developed in Phase I which are built on behavioral economics, mathematical psychology, and cognitive science. These methods are capable of capturing the morality dimension of decisions as well as the heterogeneous behavior of individual persons. One of such methods is the choice theory-based model of latent class choice, which is capable of capturing "reason-based" morality and at the same time endogenously classifies persons into groups to be homogenous with respect to attitudes within the groups and heterogeneous across the groups. The PI has the experience of developing and estimating such models in different contexts. The models are then empirically estimated on the datasets collected in Task 1.

Task 3 - Integrating Research into Teaching. The research findings from Tasks 1 and 2 will be presented to the students of the course on Transportation Engineering. The purpose is familiarizing undergraduate students with advanced research topics through presenting the findings of this study. This will then be followed by a discussion session by groups of students who will be sharing their ideas and discussion results to all their peers.

US DOT Priorities: Section left blank until USDOT's new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: The expected outcomes can provide: (i) policy implications which can inform stakeholders and industries as well as policy makers, as emphasized by recent studies on transportation policy and autonomous vehicle (AV) safety; (ii) a critical component of travel behavior analysis investigating users' trust in and acceptance of AV; and (iii) an essential input for the research efforts beyond transportation area searching for a decision-making compass to equip artificial intelligence employed in AV.

Outcomes/Impacts: By eliciting the public preferences for a moral autonomous vehicle (AV), the expected outcomes of this research project will be beneficial to society through providing policy implications which can be beneficial to stakeholders, industries, and policy makers. The research findings will be disseminated to the broader community via publications in journals and presentations at conferences. Finally, this project will integrate the research on understanding human decision making in the era of automation into education in two ways including: (i) recruiting and mentoring the next generation engineering workforce at UTRGV; and (ii) integrating the research study into teaching through presenting the study findings to the students of the course on Transportation Engineering and moderating a follow-up discussion session.

Final Research Report: N/A

Research Project Name: Validating Safe Intersection Crossing Technology in the Field

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Stephen Smith, Carnegie Mellon University

Project Partners: pathVu; City of Pittsburgh

Research Project Funding: Federal Funding—$50,000; Non-Federal Funding—$50,000

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: This project proposes to configure a local signalized intersection to demonstrate and validate a low-cost approach to deployment of PedPal, a smartphone app designed to support safe intersection crossing app for persons with disabilities. Recent work with PedPal has designed and prototyped solutions to overcome practical limitations and obstacles to widespread deployment, including (1) the use of prepositioned Ultra-wide band (UWB) beacons at the intersection to overcome the inherent localization accuracy of smartphones in urban canyons and allow corner detection and pedestrian tracking at a few centimeters accuracy, (2) a broadening of capabilities to include pedestrian-friendly route planning and wayfinding and support for the 'complete trip', and (3) specification of a PedPal system variant that eliminates its dependence on the surtrac adaptive traffic signal system and instead is capable of operating with the conventional, fixed-timing plan, traffic signal control systems that typify most urban environments], providing a much more economical basis for PedPal deployment. These collective advances have led to a recent project award to Sacramento County within DOT's Advanced Transportation Technologies and Innovative Mobility Deployment (ATTAIN) program to produce a hardened version of the complete trip app prototyped in that operates with Sacramento County's Centracs Edaptive traffic control system, and carry out a large-scale two-year pilot test over a 31-intersection region of the county (nominal start date: first quarter, 2025). Our goals in this proposed project are threefold: (1) to consolidate the above mentioned advances into a hardware/software system that is deployable with a conventional traffic signal system running a fixed timing plan and demonstrate its operation, (2) to provide a local site for testing and validating components of the hardware/software system that are applicable to the Sacramento deployment, and (3) to establish a basis for subsequently deploying PedPal in Pittsburgh in the future.

In more detail, the City of Pittsburgh's Department of Mobility and Infrastructure (DOMI) has agreed to give us access to an intersection in the vicinity of the CMU campus that currently operates with a higher-end 2070 hardware controller. Use of a 2070 controller equipped intersection will enable us to utilize the standard NTCIP communication protocol and provide a version of PedPal that will operate with any NEMA-compliant controller, ensuring a measure of generality in the solution that is developed across the controllers of different manufacturers. We will augment the equipment in the cabinet at the intersection to include a separate processor to act as a surrogate for the surtrac system in the original PedPal system architecture. This surrogate process will communicate with the PedPal cloud server on one end and invoke US DOT's V2I-Hub software on the other end to interact with the 2070 controller. Realization of this surrogate process, which we refer to as the PedPal roadside unit (P-RSU), will constitute the proposed effort's principal software development task. To enable WIFI connectivity between the PedPal cloud server and the P-RSU, we will also install a router in the cabinet. Finally, an Estimote UWB tag (or equivalent) will be installed at a known location at each corner of the intersection. Once the intersection is fully configured and the P-RSU is operational, we will run a series of tests to verify the performance of all app capabilities. We will provide DOMI with access to the final PedPal app for subsequent adoption. Currently, there are a number of 2070s scattered around the city and this number is likely to grow over time as the Smart Spines project progresses.

US DOT Priorities: Section left blank until USDOT's new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: This project will produce a field tested and validated version of the PedPal safe intersection crossing app that is capable of operating with conventional traffic signal systems that run pre-programmed signal timing plans, which are by far the most prevalent type of traffic signal control system currently in use in urban settings. The project will also enable local refinement and testing of PedPal architectural components that will be required for subsequent remote deployment of an integrated PedPal/wayfinding app in Sacramento CA (e.g., enhanced smartphone app localization via use of UWB beacons) which is being funded through a 2024 award to Sacramento County from the US DOT ATTAIN program. Finally, the variant of PedPal that is produced will be made available to the City of Pittsburgh to provide a basis for future deployment here.

We anticipate that the technology results of the proposed project will result in an invention disclosure and subsequent patent application.

Outcomes/Impacts: We anticipate that the proposed research will significantly lower the cost barrier to municipalities interested in providing their local disability communities with advanced technology support for safe intersection crossing and wayfinding, by eliminating the pre-requisite need to have a sophisticated (and expensive) adaptive traffic signal control infrastructure in place and instead enabling operation of such mobility aids with conventional traffic signal systems that run fixed signal timing plans. Further, by providing a safe intersection crossing app architecture that is capable of operating with any NEMA-compliant hardware controller, the proposed work will enable deployment in a broad range of urban environments. Ultimately, the proposed work will significantly enhance the mobility of persons with disabilities in urban environments.

Final Research Report: N/A

Research Project Name: Operating Autonomous Vehicles Safely in Common "Edge" Cases

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Raj Rajkumar, Carnegie Mellon University

Project Partners: Regional Industrial Development Corporation (RIDC)

Research Project Funding: Federal Funding—$300,000; Non-Federal Funding—$300,000

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: Autonomous vehicles (AVs) continue to make good progress but widespread deployment still seems years away. A primary reason behind this lack of large-scale adoption of AVs despite significant effort and investments is that there are many common "edge" scenarios that arise during day-to-day occurrences but cannot be navigated safely yet by AVs. These scenarios relate to the large number of operating conditions that can be encountered: bad weather conditions, poor lighting conditions (with streetlights being present or not at light), road conditions (e.g. wet/icy/snow-covered driving surfaces, potholes, speed bumps, work zones) and traffic conditions (e.g. dense merging traffic, misbehaving/disabled vehicles, flagmen and jaywalking pedestrians). Some of these conditions like heavy rain or heavy snow may not be encountered in arid and warm regions of the planet but are very common across many temperate and colder regions of the world. Our work at CMU has been systematically trying to tackle and handle each of these dimensions (various weather phenomena as well as different lighting, traffic and road conditions). In this proposed effort, we will continue to advance safe autonomous navigation of work zones and other unsafe scenarios. The key elements of our proposed work include having contextual awareness of (a) the operating environmental driving conditions and (b) sensory and behavioral limitations of the AV under these operating conditions. We therefore propose to add to AVs the ability to recognize and react safely to the operating environment. We believe that this is a critical step to full safe autonomy, which represents our focused pathway to the global adoption of AVs in the coming years.

US DOT Priorities: Section left blank until USDOT's new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: The holy grail of "Vision Zero" with zero fatalities and zero injuries, given human nature to be distracted and not always be focused on the roadways, can be (nearly) accomplished only when vehicles can drive themselves. In turn, members of society who are legally blind or otherwise unable to drive themselves can experience a much better standard of living by not being dependent on others for their transportation needs. From an economic perspective, when vehicles drive themselves, commute times for those driving to and from work can enhance their work productivity and experience the benefits of having a personal chauffeur. However, the journey to full autonomy has been longer than many even in the industry anticipate - meaning that there is a strong need to rethink how the all-too-common "edge" cases are handled. Our proposed effort aims to lay strong foundations and demonstrable capabilities that will enable safe operations and adoption of AVs. Most importantly, our proposed innovations must be realized on a functional but highly complex platform that is capable of driving on real-world road networks. The PI's group has testing permits from both state and city government agencies to test its AVs on public roads. As such, we are uniquely positioned to demonstrate our capabilities to Safety21 sponsors, the research and user communities, as well as visitors to Safety21 and CMU. We will work closely with the Autoware Foundation, a non-profit focused on open-source AV software, to develop and promote the transformational capabilities we proposed to research, develop, test and deploy.

Outcomes/Impacts: AVs being inherently safety-critical systems; self-driving vehicular technologies must not only be technologically reliable but also garner societal acceptance by being trustworthy. The roadway transportation infrastructure that we drive on was originally designed and deployed by humans for humans - meaning that it is heavily dependent on human vision. Correspondingly, cameras are used by AVs to recognize cues, signs and artifacts in the road environment and, machine-learning (using AI) is how AVs interpret camera data. Unfortunately, today's state-of-the-art AI-based computer vision techniques can yield false positives (leading to perilous phantom braking incidents) and false negatives (causing dangerous crashes and fatalities) today as well as in the foreseeable future, One of the key objectives of the proposed effort is to understand and mitigate the causes of false positives and false negatives leading to more trustworthy AI techniques. A complementary objective of our effort is to use other non-camera-based sensors to make an AV safer. The use of such redundancy is fundamental to safety-critical systems like nuclear power plants, spacecraft and aviation. Last but not the least, an AV must know that it has entered a zone of operating conditions where it cannot guarantee safe driving behaviors and takes actions to mitigate the situation as soon as viable (e.g. by slowing down, pulling over to the side of the road, and stopping with flashers on). We will also supplement the system using vehicular communications technologies (like CV2X) promoted by the US DOT to warn other surrounding vehicles and infrastructure to make the system safer in case of safety concerns that can arise (such as disabled vehicles, malfunctioning traffic signals and work zones to clear debris from an automotive crash).

Final Research Report: N/A

Research Project Name: Advancing Accessible Bus Transit: Human-Machine Interfaces for Inclusive Communication and Safe Boarding of Passengers with Disabilities

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Sarah Fox, Carnegie Mellon University; Nikolas Martelaro, Carnegie Mellon University

Project Partners: AFL-CIO; ATU; TWU; Pittsburghers for Public Transit

Research Project Funding: Federal Funding—$90,000; Non-Federal Funding—$100,000

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: This research focuses on co-designing and testing innovative rider communication tools to enable passengers with disabilities to effectively signal their intent to board public bus transit. People with disabilities often face challenges in visually or physically signaling their needs at bus stops, leading to missed rides, inadequate assistance, and reduced confidence in transit systems. By developing human-machine interfaces (HMIs) that facilitate clear communication between riders and drivers, this project aims to improve access to public transit.

Inspired by advancements in V2X communication, the U.S. Department of Transportation has demonstrated the potential of rider-to-vehicle communication through initiatives like the Integrated Dynamic Transit Operation System (IDTO) (Meng et al. 2018). In this system, riders could use an app to request connection protection, relayed to drivers via dispatchers, reducing wait times at transfer points by an average of seven minutes. Similar systems, such as flex-routing, allow riders to request pickups near fixed routes, addressing access in rural areas (e.g., Skagit Transit). While these approaches demonstrate the value of pre-trip communication, they primarily target specific scenarios and do not fully address the needs of passengers with disabilities, particularly those who cannot visually discern approaching buses or signal drivers effectively (Greer et al. 2014, Steinfeld et al. 2019).

To address these gaps, we will conduct participatory design workshops with transit riders, drivers, and advocacy groups to guide the creation of rider communication systems. These workshops will provide collaborative spaces for sharing experiences, identifying unmet needs, and developing solutions through activities like brainstorming, low-fidelity prototyping, and role-playing real-world scenarios. Insights from these sessions will inform functional prototypes, such as mobile apps or physical interfaces at bus stops. Simulated experiments will evaluate the impact of these tools on driver workload, operational safety, and rider satisfaction. These experiments will also assess how HMIs help drivers assist passengers while minimizing distractions and maintaining safety standards. For example, we will analyze how notifications via mobile apps or visual indicators at bus stops can aid drivers in identifying and assisting riders while avoiding information overload that could compromise situational awareness. Additionally, by accounting for various scenarios—such as varying weather conditions, crowded bus stops, and different disabilities—the tools will aim to address real-world challenges. We will also evaluate the timing and content of communications to optimize their effectiveness for drivers and riders alike.

Our participatory approach seeks to create communication systems that reduce barriers to mobility for passengers with disabilities, ensuring access to public transit. This research contributes to the broader goal of building safer transit systems.

US DOT Priorities: Section left blank until USDOT's new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: Outputs from this project include: Two academic papers based on our design activities and simulation based experiments, prepared during the grant period. These papers will be published in relevant human-computer interaction and transportation research venues; A white paper translating our academic findings into recommendations for transit agencies, advocacy organizations and government officials seeking to support public transportation; Open-source the simulation environments and test procedures. We will circulate research materials, simulation tools, and findings as publicly available materials in an open access toolkit at codesigningtransit-dot-com.

Outcomes/Impacts: Centering safety, this project aims to develop driver- and rider-driven human-machine interaction that could further enhance operations. This work will produce novel system designs aimed at introducing communication between passengers with disabilities and drivers regarding assistance needs in advance of boarding. We will empirically evaluate these designs with particular focus on the timing of information exchange, to avoid overloading the driver. Finally, this research will produce methodological contributions in the form of participatory, multi-stakeholder hazard analysis, designed to broaden who is involved in the process of defining risks and potential harms of emergent systems. We expect this will lead to more robust and complete hazard analysis to more effectively avoid risks.

Final Research Report: N/A

Research Project Name: The Effects of Curb Extensions and Stormwater Infrastructure on Traffic Safety in Philadelphia

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Erick Guerra, University of Pennsylvania

Project Partners: City of Philadelphia Water Department; City of Philadelphia Office of Transportation and Infrastructure Systems; PA Safe Roads

Research Project Funding: Federal Funding—$90,000; Non-Federal Funding—$90,000

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: In order to meet the goals, set out in the City's landmark stormwater management plan, and to satisfy state and federal regulations, the Philadelphia Water Department (PWD) is planning, designing, and constructing a portfolio of stormwater infrastructure systems. These systems capture, filter, and store stormwater run-off utilizing sub-surface trenches, and vegetated features, from rain gardens to street trees, mimicking natural processes where possible. As part of this program PWD has developed 156 stormwater stormwater management bumpouts, and anticipates bringing 342 more online within the next 5 years. These bumpouts have vegetated curb extensions that protrude into the street at mid-block or intersection locations. When located at crosswalks, curb extensions are thought to improve traffic safety by reducing the crossing time of pedestrians, improving visibility at intersections, and slowing drivers by reducing the turning radii at intersections.

This proposed project will examine the effects of these curb extensions over time on traffic collisions, injuries and fatalities relative to similar roadways where curb extensions will but have not yet been installed. Although pre-existing crash data featured in the selection of treated locations, the focus on stormwater management and long list of treated and untreated intersections provide unique, quasi-experimental conditions under which to examine the effects of curb extensions. Although curb extensions are widely accepted roadway design interventions in the fields of traffic calming and traffic safety, little is known about their effects on crashes or how these effects may vary by roadway conditions, location, or other features.

While data processing is in progress, the City and Water Department have geocoded the locations of the existing and proposed curb extensions. Crashes from PennDOT's crash reporting system will be to the nearest intersection nodes of the eligible bump-outs. From here, the number of crashes, fatal and serious injury crashes, and injury crashes within each buffer can be counted and aggregated by year, with results associated with their corresponding intersection nodes and curb extensions.

US DOT Priorities: Section left blank until USDOT's new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: The project will result in peer-reviewed journal publication as well as a report to the City of Philadelphia.

Outcomes/Impacts: This project will help City officials determine whether and to what extent its stormwater infrastructure system has created traffic safety benefits beyond the benefits of better managing storm water and preventing system overflow. The results will also provide findings about the effects of curb extensions on traffic safety using a causally valid quasi-experimental research design. These results can be used to help determine whether, where, and by how much curb extensions can be expected to improve traffic safety outcomes. The research may also provide insights into potential synergies between storm water management and traffic safety engineering.

Final Research Report: N/A

Research Project Name: StreetSavvi: Complete Street AI Companion for Digital Twin Driving

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Erick Guerra, University of Pennsylvania; Xiaoxia Dong, University of Pennsylvania; Helen Loeb, University of Pennsylvania

Project Partners: Jitsik; The Bicycle Coalition of Greater Philadelphia; City of Philadelphia; Delaware Valley Regional Planning Commission; SEPTA

Research Project Funding: Federal Funding—$35,000; Non-Federal Funding—$35,000

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: Our project StreetSavvi is an effort to bring the power of AI, Digital Twin technology, advanced simulation and visualization to local Stakeholders such as the City of Philadelphia, PennDOT and SEPTA for a partnership during the Route for Change Project which will redevelop the Philadelphia Roosevelt Boulevard.

In an ever-increasing Digital Age, the US DOT has put great effort into developing a large number of databases that encompass infrastructure, travel behavior and safety, and context as a way to answer the formidable challenge of moving people and goods in a safe and efficient way. Interestingly, commercial satellite and street imagery such as those produced by Google and other geospatial data providers emerge as an opportunity for an up-to-date inventory of transportation infrastructure. In addition, the acceleration of Artificial Intelligence through Machine Learning and its inference processes for data processing or the Large Language Models generative processes offer phenomenal opportunities to take a global approach to the field of transportation.

In this context, the team proposes the development of StreetSavvi, an advanced Digital Twin development platform, which applies AI computer vision processes to geospatial data, namely satellite imagery, street imagery, and camera data, to develop photorealistic 3D models of neighborhoods and develop a scalable, efficient, and user-centric digital twin platform for urban transportation planning.

This project expands on a previous Safety21 University Transportation Center feasibility project. It also expands on the installation of speed cameras on Roosevelt Blvd. Our proposed Streetsavvi project will closely follow the design alternatives developed by PennDOT, the City of Philadelphia, and SEPTA as they include components such as Neighborhood boulevard, Partially-capped expressway, Bus Rapid Transit (BRT), Light Rail Transit (LRT), Subway into the new design. Our simulation and visualization tools will allow the research team to partner with the stakeholders to assess the impact of various designs on multiple metrics such as traffic congestion, safety, cost, etc. Our virtual models will allow an efficient assessment of the proposed infrastructural changes, such as lane removals, enhanced bicycle/pedestrian facilities, and exclusive bus only lanes. We propose with this project to automate the ideation and simulation process by creating a friendly AI based user interface that is easy to use by state, local, and tribal transportation agencies who are stakeholders in the design of our streets.

US DOT Priorities: Section left blank until USDOT's new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: The proposed work's main anticipated outputs include:

- Research Impact: The proposed work will develop novel semi automatic approaches for the development of driveable Digital Twins. This project will integrate various tools such as OpenStreetMap, RoadRunner, Cesium, Unity and Unreal Engine API. The perspective of automating Digital Twin cities roadways generation is a paradigm shift for road and even urban planning with solid impact in transportation.
- Publications: Research developed through the proposed work will be disseminated through publications at major international venues such as IEEE, TRB Annual Conference, TRR Transportation Journal, SAE, etc.
- Data: Our generated data will be integrated to OpenStreetMap database so it can be used as a public database.
- Software: Source code and associated software systems will be developed.
- Student Training: Master and undergraduate students will be trained during the course of the research. The students will actively collaborate with the transportation stakeholders.
- Town Halls demonstration: our software will be made available through demonstrations at the various town halls organized to discuss the Route for Change initiative.

Outcomes/Impacts: Digital twins have a transformative role in transportation by enabling real-time monitoring, simulation, and optimization of transport systems. For transportation infrastructure, digital twins create dynamic models of roads, bridges, and tunnels, allowing cities to simulate traffic flow, identify bottlenecks, and plan improvements based on real-time data. This can lead to smarter traffic management, better route planning, and safer environments for drivers and pedestrians.

Our use case of the Philadelphia Roosevelt Boulevard with a direct participation of stakeholders such as the City of Philadelphia, SEPTA and Penndot will ensure our results are well understood and used by the decision makers who have direct authority over the Route for Change project.

Final Research Report: N/A

Research Project Name: Guardian Angel - Building Trust in Shared Autonomy Vehicles with Safe Human-Machine Interaction

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Erick Guerra, University of Pennsylvania; Rahul Mangharam, University of Pennsylvania; Helen Loeb, University of Pennsylvania

Project Partners: Jitsik LLC; The Autoware Foundation

Research Project Funding: Federal Funding—$35,000; Non-Federal Funding—$35,000

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: Shared autonomy in self-driving vehicles refers to a collaborative control model where both humans and autonomous systems work together to make decisions in real-time. This concept is crucial for the safe deployment of autonomous vehicles (AVs), as it blends the strengths of both human intuition and machine precision.

One of the primary reasons shared autonomy is vital for self-driving cars is safety. Although autonomous systems are designed to reduce human error, they are not infallible. Situations may arise where the vehicle's sensors or algorithms struggle to interpret complex conditions, such as unpredictable weather, erratic behavior from other road users, or road construction zones. In such cases, human intervention becomes necessary. This transition between human and machine control should be seamless and intuitive to avoid confusion or delay during critical moments.

Moreover, shared autonomy fosters trust between humans and autonomous systems. For passengers who are new to self-driving technology, knowing they can take control if needed provides a sense of security. This trust is essential for the adoption of AVs, as people may be hesitant to relinquish control entirely to a machine. As autonomous systems continue to improve, the level of human intervention required may decrease, but in the interim, shared autonomy allows for a balanced approach that builds confidence in the technology while ensuring safety. It is a key factor in encouraging both regulatory bodies and the public to embrace autonomous vehicles on a large scale.

Additionally, shared autonomy is crucial for the interaction between self-driving cars and non-autonomous vehicles. Self-driving cars must be able to communicate with human drivers and adapt to their behavior. For example, the vehicle’s system may recognize that a human driver is acting unpredictably or entering a dangerous situation, allowing the car to adjust its driving behavior to mitigate potential risks. Shared autonomy helps establish this kind of adaptive interaction, ensuring a smoother transition as the roadways evolve to accommodate both human drivers and autonomous technology.

For this project , we propose to leverage an immersive driving simulator, developed with the support of previous Safety21 projects, to expose a cohort of 20 drivers to shared autonomy. We will recruit a cohort of 10 young drivers, ages 25 to 30, who are more likely to embrace self-driving, and a second cohort of 10 more mature drivers, aged 40 to 50 who might be more skeptical of self-driving. We will expose them to a 15 minute challenging scenario so they experience manual, full self-driving and a hybrid model where a Guardian lets them drive manually while being supervised by the automated agent. We will collect survey data as well as simulator data. Shared autonomy is a cornerstone for the development of self-driving cars, combining the reliability of autonomous systems with the experience and judgment of human drivers. By ensuring safety, fostering trust, and enabling effective interactions between autonomous and non-autonomous vehicles, shared autonomy paves the way for a more integrated, efficient, and secure future of transportation.

US DOT Priorities: Section left blank until USDOT's new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: The proposed work's main anticipated outputs include:

- Research Impact: The proposed work will develop Guardian Angel techniques that will allow drivers to drive "manually" while benefiting from the expertise of an Automated Driving Agent. The perspective of blending manual and automated driving is key to the wide adoption of self-driving, which is seen as one of the best ways to improve safety on our roads.
- Publications: Research developed through the proposed work will be disseminated through publications at major international venues such as IEEE, TRB Annual Conference, TRR Transportation Journal, SAE, and ITS.
- Software: Source code and associated software systems will be developed.
- Student Training: Master and undergraduate students will be trained during the course of the research.

Outcomes/Impacts: Shared control in autonomous driving refers to a system where both the vehicle's autonomous system and the human driver collaborate in decision-making and control. This approach offers significant benefits, enhancing both safety and user experience. One key benefit is increased safety. In situations where the autonomous system encounters uncertainty—such as complex traffic or adverse weather conditions—the human driver can intervene if necessary, ensuring that decisions align with human judgment and intuition. This helps reduce the risk of accidents that could arise from an over-reliance on automation, particularly in unpredictable or ambiguous scenarios. Shared control also improves driver comfort and trust. By maintaining an active role in the driving process, even if limited, the human driver feels more in control, which can alleviate anxiety or frustration with fully autonomous systems. This sense of partnership between the vehicle and the driver fosters trust in the technology, encouraging more widespread acceptance of self-driving cars. We anticipate the use of Shared Control will make it easier for stakeholders to develop legislation to accommodate self-driving in all future vehicles.

Final Research Report: N/A

Research Project Name: RoboRacer.AI - Hands-on Autonomous Vehicle Design, Testing and Deployment

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Rahul Mangharam, University of Pennsylvania; Insup Lee, University of Pennsylvania

Project Partners: The Autoware Foundation; Robotics Automation Society

Research Project Funding: Federal Funding—$120,000; Non-Federal Funding—$120,000

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: This project focuses on developing a training community for engineering and ethical skills for developing future autonomous vehicles. This project includes three components - (1) autonomous driving course development with multiple scale vehicles from 1/18th-scale, 1/10th-scale, 1/5th-scale to full scale vehicles where students learn advanced algorithms and software development for perception, planning and control of autonomous driving; (2) Community Activities spanning 89 universities which have one or more RoboRacer platforms and participate in the international autonomous racing competitions. We will host a minimum of 4 competitions in the top robotics, transportation and cyber-physical systems conferences; (3) Development of an ethical framework for using machine learning in life-critical systems.

Contribution: An autonomous vehicle hardware and software platform, called RoboRacer, is developed for teaching autonomous systems hands-on. This project will design and develop the education modules and software stack for teaching at various educational levels with the theme of "racing" and competitions that replace exams. We have expanded the previous F1Tenth effort to now include multiple scale vehicles that use the same autonomous driving software. The important aspect here is that it allows students to mature through more complex driving platforms, sensor suites, algorithm pipelines and safety scenarios - all the way to full scale autonomous vehicles.

Background: College-level robotics courses often focus on theory, while most hardware platforms for robotics teaching are low-level toys aimed at younger students at middle-school levels. The RoboRacer platforms fill the gap between research platforms and low-end toy cars and offers hands-on experience in learning the topics in autonomous systems.

Intended Outcomes: The RoboRacer vehicles offer a modular hardware platform and its related software for teaching the fundamentals of autonomous driving algorithms. From basic reactive methods to advanced planning algorithms, the teaching modules enhance students' computational thinking through autonomous driving with the RoboRacer vehicle.

Application Design: Over 89 universities have adopted the teaching modules for their semester-long undergraduate and graduate courses for multiple years. Student feedback is used to analyze the effectiveness of the RoboRacer platform. This project's focus is to maintain and grow this community through education, outreach and K-12 training events.

Findings: More than 80% of the students strongly agree that the hardware platform and modules greatly motivate their learning, and more than 70% of the students strongly agree that the hardware enhanced their understanding of the subjects. The survey results show that more than 80% of the students strongly agree that the competitions motivate them for the course.

US DOT Priorities: Section left blank until USDOT's new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: Theme I: Safe Autonomy - This thrust will enable AV controllers that combine the performance and generalization abilities of machine learning with the safety guarantees afforded by formal and semi-formal verification. Researchers in this theme develop fast verification methods that scale to run in real-time on-board the vehicle through a combination of formal methods and testing. A cloud-based simulator will enable scalable verification which combines robust testing and falsification with reachability analysis for real systems.

Theme II: Efficient Autonomy - Researchers in this theme develop the hardware and software architectures for power-efficient and timing-guaranteed execution of autonomy algorithms. These include computer vision, motion planning, and neural network inference engines.

Theme III: Coordinated Autonomy - This thrust will enable a fleet of AVs to coordinate on-the-fly to achieve fleet-wide safety, higher transportation network efficiencies and enable exploration of new mobility and ridesharing services.

Theme IV: Secure Autonomy - Researchers in this thrust develop models of cyber-physical attacks, and resilient estimation and control schemes to guard against them and mitigate their effects in the field.

Outcomes/Impacts: The course instructors will develop 10 reference platform vehicles and demonstrate them working at high speeds of 8-35 mph.

We will host 5 international competitions with over 80 academic partners worldwide. These will be hosted at top conference venues for transportation, control systems, robotics and cyber-physical systems.

It will provide training in autonomous vehicle development to over 2000 persons worldwide and extend the state of the art for machine perception, motion planning and safe adaptive control research.

This effort will raise the awareness of ethical issues with using machine learning in future mobility systems.

Final Research Report: N/A

Research Project Name: Reliable Vehicle Sensing in All Weather Conditions

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Rahul Mangharam, University of Pennsylvania; Mingmin Zhao, University of Pennsylvania

Project Partners: The Autoware Foundation; SEPTA

Research Project Funding: Federal Funding—$110,000; Non-Federal Funding—$110,000

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: Safety remains a critical challenge for transportation systems, particularly in environmental conditions where traditional sensors like cameras and LiDAR fail due to fog, rain, snow, or poor lighting. This proposal addresses DOT's safety grand challenge by developing a robust RF-based perception system that enables reliable vehicle sensing in all weather conditions. Our key insight is that mmWave radar and vision sensors have complementary strengths - while cameras offer superior spatial resolution for detailed scene understanding, radar provides unmatched capabilities in adverse weather conditions with precise depth measurements. Through our next-generation RF imaging system design, our goal is to achieve the best of both worlds: the high spatial resolution traditionally only available with optical sensors and the all-weather robustness of RF sensing. Our RF imaging system achieves this by rotating a compact mmWave radar to create a synthetic aperture, precisely tracking the platform's motion to maintain coherent imaging, and leveraging machine learning to enhance resolution and recover fine scene details. This approach achieves LiDAR-comparable 3D sensing accuracy while being inherently robust to environmental conditions. Unlike existing automotive radar systems that provide only crude point clouds, our high-resolution RF imaging delivers detailed 3D scene reconstruction with high-fidelity object detection. Through this work, we will develop and evaluate an integrated perception system that enhances vehicle safety through: 1) High-resolution 3D RF imaging with quantified uncertainty, 2) Robust object detection and semantic understanding, and 3) Real-time operation for automotive deployment. Our preliminary results demonstrate successful sensing and object detection through fog where traditional sensors fail. Field testing will be conducted at selected urban intersections and highway segments where adverse weather conditions frequently impact traffic safety, allowing us to validate the system's performance in real-world scenarios. By focusing on challenging weather conditions where traditional sensors struggle, we aim to demonstrate a robust perception solution that can significantly enhance vehicle safety year-round.

US DOT Priorities: Section left blank until USDOT’s new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: Research Publications: The proposed work will advance fundamental knowledge in high-resolution RF imaging and automotive safety systems. Research findings will be published in top transportation and sensing venues.

Open-Source Software: Core algorithms including motion estimation, uncertainty quantification, and 3D RF imaging will be released to accelerate adoption of robust perception technology. We will integrate the sensors with Autoware to demonstrate it on an autonomous vehicle driving on roads.

Systematic Evaluation: Systematic evaluation data will document system performance in various weather conditions, providing quantitative evidence of reliability compared to existing sensors.

Student Training: Graduate students will be trained in transportation safety systems and robust perception technology, contributing to the future workforce in autonomous vehicle development.

Outcomes/Impacts: The proposal directly addresses automotive safety by developing robust perception capabilities for adverse weather conditions. Our high-resolution RF imaging technology will establish a new paradigm for reliable sensing in fog, rain, and snow where traditional sensors fail. The proposed work will create a foundation for all-weather vehicular safety systems using low-cost, commercially available radar hardware.

The research will advance the state-of-the-art in RF imaging by demonstrating LiDAR-comparable resolution through mechanical rotation, precise motion tracking, and learning-based enhancement. This represents a significant leap beyond current automotive radar capabilities, bringing RF sensing resolution close to that of LiDAR while maintaining operation in all weather conditions.

The proposed work will shape future automotive safety requirements and influence sensor specifications for advanced driver assistance systems.It will demonstrate reliable perception in challenging conditions using commodity hardware and establish practical paths for widespread adoption of all-weather radar technology.

Final Research Report: N/A

Research Project Name: How and to what extent do Vision Zero plans address safety for different road user groups?

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Xiaoxia Dong, University of Pennsylvania; Erick Guerra, University of Pennsylvania; Megan Ryerson, University of Pennsylvania

Project Partners: Office of Transportation and Infrastructure Systems in the City of Philadelphia

Research Project Funding: Federal Funding—$36,580; Non-Federal Funding—$36,580

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: Since the early 2010s, many U.S. cities have adopted vision zero plans to enhance traffic safety. Vision zero plans aim to eliminate traffic fatalities and severe injuries. A critical question remains, for whom are vision zero plans improving safety? Answering this question is a crucial first step toward enhancing traffic safety for different road user groups. So far, however, there has been little research on how and to what extent vision zero plans address safety for different road user groups.

In this study, we propose a systematic review of vision zero plans of the 100 largest U.S. cities to evaluate (1). How the plan-making process included different road user groups, including vulnerable road users. (2). How the plans accounted for differences in traffic risks through its analysis. (3). How the plans' recommendations accounted for different road user groups within the Safe System framework. And (4). How the plans accounted for different road user groups in their outcome evaluations. The goal of the project is to (1). Identify the best practices of including different road user groups in vision zero plans. And (2). Identify the challenges of project deployment strategies to ensure safety for different road user groups. To complement our plan review, we will also interview planners and experts who were involved in vision zero plan-making for additional insight, including planners in the Office of Transportation and Infrastructure Systems in the City of Philadelphia.

Our findings will help cities reevaluate their approaches to traffic safety planning and design better strategies to ensure the deployment of projects that are aimed to enhance safety for different road user groups.

US DOT Priorities: Section left blank until USDOT's new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: We plan to publish our findings in a transportation journal and present our findings to the Office of Transportation and Infrastructure Systems in the City of Philadelphia.

Outcomes/Impacts: We expect the finding to show that there is great variation in the extent vision zero plans incorporated different road user groups. We also expect to find common challenges facing cities when deploying safety-improvement strategies for different road user groups.

Final Research Report: N/A

Research Project Name: Autonomous Driving in Hazardous Conditions: Pushing the Limits of Vehicle Control

Recipient/Grant (Contract) Number: Carnegie Mellon University (Grant #: 69A3552344811)

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Rahul Mangharam, University of Pennsylvania; Antonio Loquercio, University of Pennsylvania

Project Partners: Delaware Valley Regional Planning Commission; The Autoware Foundation

Research Project Funding: Federal Funding—$110,000; Non-Federal Funding—$110,000

Project Start and End Date: 7/1/25 to 6/30/26

Project Description: Extreme weather events often render even the most carefully planned routes dangerous to navigate. For example, the heavy rainfall from Hurricane Milton triggered the formation of large sinkholes, columns of debris from collapsed buildings, and partially flooded areas along a large portion of Florida's western coast, severely compromising roadways and creating dangerous conditions for emergency responders and motorists alike. This raises an important question: can we design systems to enhance driver safety in such hazardous conditions? In this proposal, we aim to study the feasibility of such systems and take the first steps toward deploying them in real-world, hazardous driving scenarios.

To achieve this goal, we will develop a control framework for vehicles operating in hazardous conditions. In such cases, the vehicle's behavior is highly nonlinear and not constrained to 2D motions, assumptions that usually lead traditional control methods to fail. This is because the vehicle's model at the core of these approaches is insufficient to capture the complexity of the dynamics encountered during extreme driving. Our key technical insight is that, while such models may be insufficient for precise planning and control, they can serve as abstract, imperfect simulators for training learning-based controllers. In our prior work on simulation-to-reality transfer for high-speed drone racing, published in *Nature* (2023), we addressed this challenge by augmenting imperfect models with stochastic, data-driven components. This hybrid model enabled us to train control policies using reinforcement learning, achieving robust performance despite the limitations of the underlying simulator.

In this proposal, we aim to extend these findings to ground vehicles, where frequent interaction with the terrain introduces dynamics more unpredictable than for drones. We plan to address this difficulty by developing a novel model-based reinforcement learning approach based on sampling multiple hypotheses and selecting, at test time, which hypothesis best fits the current state. Similarly to other model-based approaches, this control framework will be more data efficient than the model-free approaches used for drones, at the cost of a limited drop in performance. However, differently from other model-based approaches, it will be conditioned on the uncertainty and predicted accuracy of the model, which will make it less susceptible to unpredictable dynamics.

US DOT Priorities: Section left blank until USDOT's new priorities and RD&T strategic goals are available in Spring 2026.

Outputs: The proposed work's key anticipated outcomes include:

Research Impact: This work will extend the frontier of what is possible with only an abstract and imperfect model of a real-world system. We will develop a fundamentally new approach to combine first-principle methods with data-driven techniques for grounded vehicle navigation in extreme conditions. Instead of designing our approach to be completely autonomous, we will focus on systems that support drivers in such extreme conditions. While completely autonomous driving is still a far-fetched goal, assisted driving systems are already a reality today. This has important short-term implications for transportation and driver's safety.

Publications: Our research will be disseminated through publications at major venues including conferences.

Data: Prior art has focused on collecting large datasets in normal driving conditions. While these datasets are large, they contain only very limited long-tail events. We will collect and open-source a new dataset that especially focuses on hazardous driving conditions, which could be a great complement to existing datasets.

Software: We will release the source code and associated software of our novel control framework to handle hazardous driving conditions.

Patent Filings: The proposed work will lead to patent filings to protect the intellectual property on the developed assisted driving technology.

Student Training: One PhD student will be trained during the course of the research. The student will collaborate with our stakeholders.

Pilots developments: The proposed work will lead to pilot developments coordinated with DVRPC's emergency response and evacuation project planning. We will locate off-road driving test sites in the Greater Philadelphia region (e.g. Philly Pump Track) to demonstrate the capabilities of the system in adverse environment and road conditions.

Outcomes/Impacts: The proposed work will be implemented and piloted through strategic collaborations with disaster response agencies (DVRPC) and industry partners (The Autoware Foundation), who will provide critical support, including access to simulation platforms, testing facilities, and operational expertise. The proposed work will demonstrate adaptive and superhuman driving capability on uneven off-road terrain with sand, gravel rocks, slush and other hazardous environment conditions. We will demonstrate safe planning and robust control of the vehicle in adverse driving conditions. The proposed work will engage with DVRPC to identify hazard situations and the traversability requirements. We will work with the Autoware Foundation to integrate the work in real vehicle platforms.

Final Research Report: N/A