CENTER FOR CONNECTED AND AUTOMATED TRANSPORTATION

CCAT-CTS 26-01
January 2026

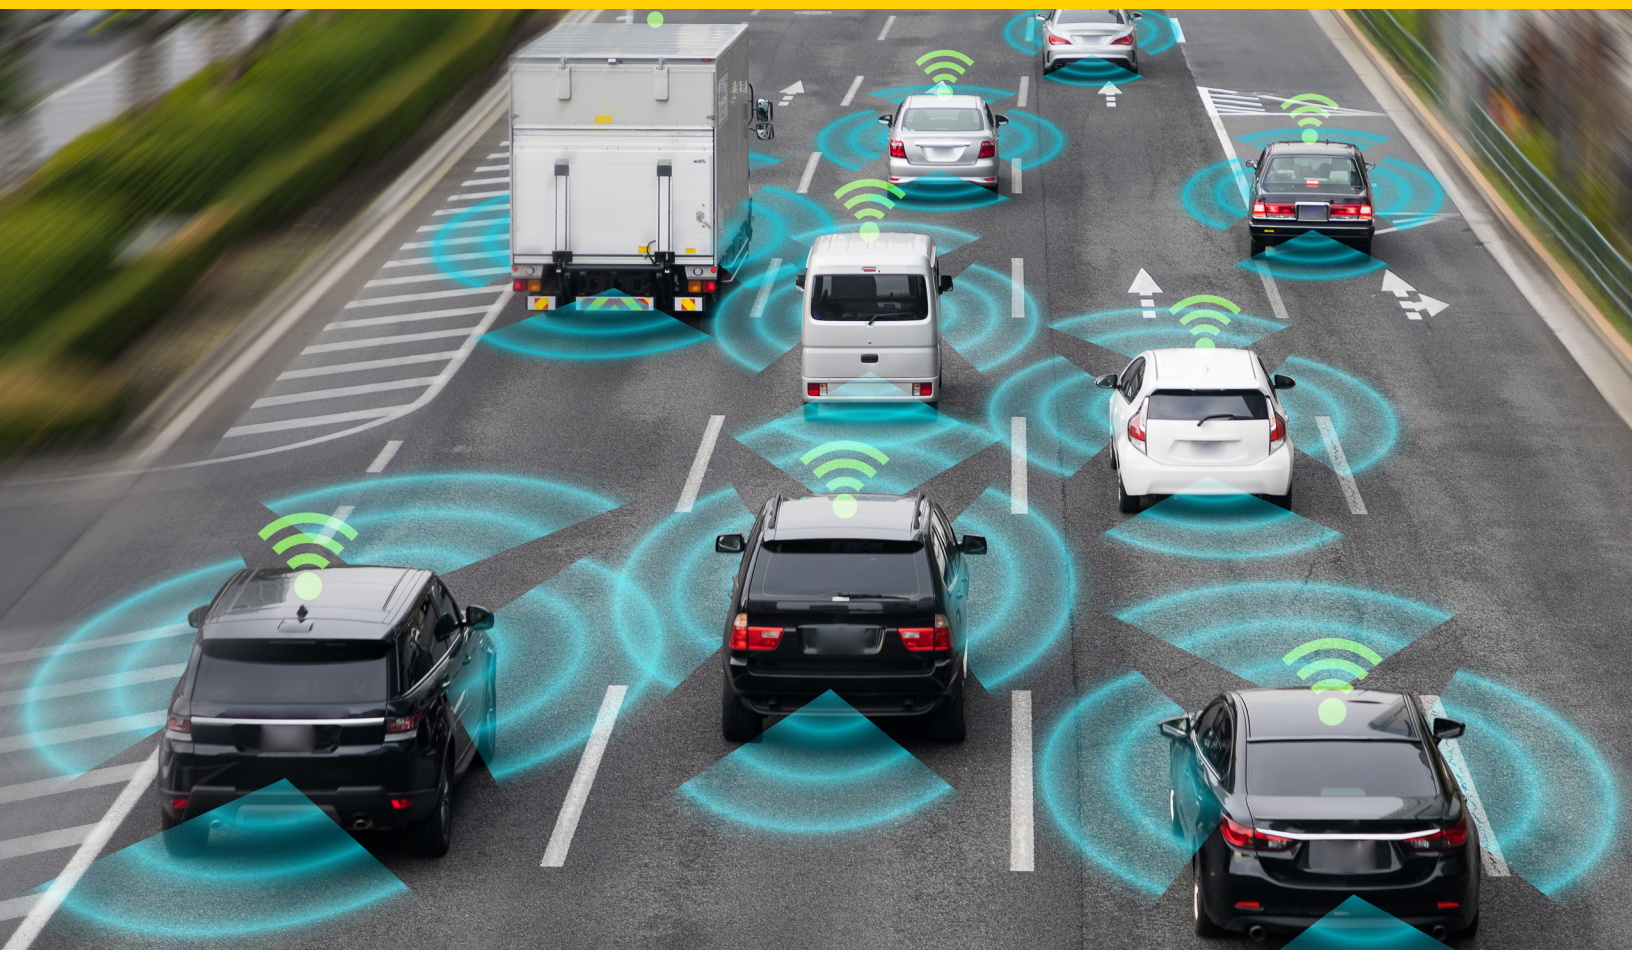

Hardening the CAV Ecosystem to Reduce Cybersecurity Risks — Year One

Zhi-Li Zhang
Z. Morley Mao
Yiheng Feng

Center for Transportation Studies

CENTER FOR CONNECTED AND AUTOMATED TRANSPORTATION

Report No. CTS 26-01 **January 2026**

Project Start Date: 10/01/2023

Project End Date: 09/30/2025

Hardening the CAV Ecosystem to Reduce Cybersecurity Risks – Year One

by

Zhi-Li Zhang, Professor, University of Minnesota

Z. Morley Mao, Professor, University of Michigan

Yiheng Feng, Assistant Professor, Purdue University

CENTER FOR CONNECTED AND AUTOMATED TRANSPORTATION

DISCLAIMER

Funding for this research was provided by the Center for Connected and Automated Transportation under Grant No. 69A3551747105 of the U.S. Department of Transportation, Office of the Assistant Secretary for Research and Technology (OST-R), University Transportation Centers Program. The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated under the sponsorship of the Department of Transportation, University Transportation Centers Program, in the interest of information exchange. The U.S. Government assumes no liability for the contents or use thereof.

Suggested APA Format Citation:

Zhang, Z., Mao, Z.M., & Feng, Y. (2026). Hardening the CAV Ecosystem to Reduce Cybersecurity Risks – Year One. Final Report. USDOT CCAT Project No. CTS 26-01.

Contacts

For more information:

Dr. Zhi-Li Zhang
B22 B Walter Library
117 Pleasant St. SE
Minneapolis, MN 55455
Phone: 612-625-8568
Email: zhzhang@cs.umn.edu

CCAT
University of Michigan Transportation Research Institute
2901 Baxter Road
Ann Arbor, MI 48152
uumtri-ccat@umich.edu
(734) 763-2498

UMTRI | The University of Akron | Central State University | Illinois Center for Transportation | University of Minnesota | Northwestern | Purdue University | Washtenaw Community College | Wisconsin – University of Wisconsin-Madison

CENTER FOR CONNECTED AND AUTOMATED TRANSPORTATION

Technical Report Documentation Page

1. Report No. CTS 26-01	**2. Government Accession No.** Leave blank – not used	**3. Recipient's Catalog No.**
4. Title and Subtitle Hardening the CAV Ecosystem to Reduce Cybersecurity Risks – Year One		**5. Report Date** January 2026 **6. Performing Organization Code** Enter any/unique members assigned to the performaing organization, if applicable.
7. Author(s) Zhang, Zhi-Li, Ph.D., https://orcid.org/0000-0001-8584-2319 Mao, Z. Morley, Ph.D., https://orcid.org/0000-0002-9844-2055 Feng, Yiheng, Ph.D., https://orcid.org/0000-0001-5656-3222		**8. Performing Organization Report No.** Enter any/all unique alphanumeric report numbers assigned by the performing organization, if applicable.
9. Performing Organization Name and Address Center for Connected and Automated Transportation University of Michigan Transportation Research Institute 2901 Baxter Road Ann Arbor, MI 48109.		**10. Work Unit No.** **11. Contract or Grant No.** Contract No. 69A3551747105
12. Sponsoring Agency Name and Address U.S. Department of Transportation Office of the Assistant Secretary for Research and Technology 1200 New Jersey Avenue, SE Washington, DC 20590		**13. Type of Report and Period Covered** Final Report (October 2023 – September 2025) **14. Sponsoring Agency Code** OST-R
15. Supplementary Notes Conducted under the U.S. DOT Office of the Assistant Secretary for Research and Technology's (OST-R) University Transportation Centers (UTC) program.		
16. Abstract In this project, we have taken a multi-pronged framework to comprehensively study how to harden and secure the connected and automated vehicle (CAV) ecosystem by simultaneously considering cybersecurity threats posed to CAVs and physical and cyber transportation infrastructure that support CAV operations. We have also explored how to leverage intelligent collaborations among CAVs and between CAVs and physical/cyber transportation infrastructures to develop solutions to defend the CAV ecosystem against both physical and cyber threats in a holistic fashion. More specifically, we have carried out major research activities along four key aspects including CAV ecosystem threat analysis and risk assessment (TARA), shared state approach against cybersecurity attacks on CAV teleoperation, robust collaborative perception under lossy networks, and infrastructure-based anomaly detection.		
17. Key Words Connected and Autonomous Vehicles; Infrastructure; Cybersecurity; Physical Security; Threat Analysis and Risk Assessment; Cooperative Perception; Cellular V2X Communications; 5G Networks; Artificial Intelligence (AI); Data Injection Attack; Adversarial Machine Learning	**18. Distribution Statement** No restrictions.	

19. Security Classif. (of this report) Unclassified	**20. Security Classif. (of this page)** Unclassified	**21. No. of Pages** 21	**22. Price** Leave blank – not used

Form DOT F 1700.7 (8-72) Reproduction of completed page authorized

UMTRI | The University of Akron | Central State University | Illinois Center for Transportation | University of Minnesota | Northwestern | Purdue University | Washtenaw Community College | Wisconsin – University of Wisconsin-Madison

CENTER FOR CONNECTED AND AUTOMATED TRANSPORTATION

Table of Contents

UMTRI | The University of Akron | Central State University | Illinois Center for Transportation | University of Minnesota | Northwestern | Purdue University | Washtenaw Community College | Wisconsin – University of Wisconsin–Madison

List of Figures

1. Introduction

The transportation system is a critical national infrastructure: We use public roads every day to move goods and people. Roadways are vital to our national well-being. Safety is critical to the transportation system, but the United States has long suffered a road safety crisis [1]. According to the US Department of Transportation (USDOT) website [2]: "More than 370,000 people died in transportation incidents over the last decade (2011-2020) in the United States, of which more than 350,000 of them died on our roads." Connected and autonomous transportation is envisaged as the future of transportation that will help achieve zero fatality on roadways.

Ever since the Defense Advanced Research Projects Agency (DARPA) Grand Challenge in 2005, connected and autonomous vehicles (CAVs) have been eagerly anticipated and discussed. Tremendous progress has been made in the development of CAVs, spurred in part by rapid advances in machine learning and artificial intelligence (ML/AI) in recent years. General Motors (GM) SuperCruise [2] and Tesla with "autopilot'' [3] are already roaming on our streets. Waymo robotaxi services are now available in several major US cities including San Francisco, Los Angeles, Phoenix AZ and Austin TX, and are expanding to additional cities including Minneapolis MN. Tesla is piloting its robotaxi service in Austin TX, and Amazon Zoox [4] is planning to roll out robotaxi trials.

New cellular V2X (vehicle-to-everything) -- including V2I (vehicle-to-infrastructure), V2V (vehicle-to-vehicle) -- communication standards as well as emerging 5G networks and next-generation (NextG) networking technologies such as 6G under development will enable exciting new opportunities such as cooperative driving automation (CDA) and tele-operations of CAVs that will bring additional safety benefits. On the other hand, connectivity and dependence of CAVs on other CAVs and/or the physical and cyber (transportation) infrastructures (e.g., a "smart" road sign, traffic light, trajectory tracking running on an edge server or operation commands from a remote human operator) for driving directives or intelligence also introduce new attack surfaces that can be exploited by cyberattacks. For example, an attacker may aim to alter a CAV behavior by attacking the infrastructure sensors or comprising a cyber component instead of directly tampering vehicle onboard sensors.

In this Year 1 CCAT project, we have taken a multi-pronged framework to comprehensively study how to harden and secure the CAV ecosystem by simultaneously considering cybersecurity threats posed to CAVs and physical and cyber transportation infrastructure that support CAV operations. We have also explored how to leverage intelligent collaborations among CAVs and between CAVs and physical/cyber transportation infrastructures to develop solutions to defend the CAV ecosystem against both physical and cyber threats in a holistic fashion. More specifically, we have carried out major research activities along four key aspects.

- **CAV Ecosystem Threat Analysis and Risk Assessment (TARA):** We have conducted extensive literature review and developed a comprehensive TARA framework from multiple perspectives, focusing in particular on new attack surfaces introduced by communications, data sharing and the use of AI models.

- **Shared state approach against cybersecurity attacks:** By leveraging collaboration among trusted entities, for example, a CAV and its remote teleoperator, we advocate a novel shared state approach which emphasizes maintaining (shared) state consistency and integrity and utilizes potential inconsistency in (shared) data to detect anomalies and mitigate cyber threats. As a case study, we have developed a novel lightweight mechanism to secure the CAV teleoperation sensor data delivery pipeline against cyber data injection attacks.

- **Robust collaborative perception under lossy networks:** Collaborative Perception enables multiple agents, such as autonomous vehicles and infrastructure, to share sensor data via vehicular networks so that each agent gains an extended sensing range and better perception quality. Despite its promising benefits, realizing the full potential of such systems faces significant challenges due to inherent imperfections in underlying system layers, consisting of network layer imperfections and hardware-level noises. Such imperfections and noises include packet loss in vehicular networks, localization errors from GPS measurements, and synchronization errors caused by clock deviation and network latency. To address these challenges, we propose a novel end-to-end collaborative perception framework, SCORPION, that harnesses the AI co-design of the application layer and system layer to tackle the aforementioned imperfections. SCORPION consists of three main components: lost bird's eye view feature reconstruction (L-BEV-R) recovers lost spatial features during lossy V2X communication, while deformable spatial cross attention (DSCA) and temporal alignment (TA) compensate for localization and synchronization errors in feature fusion. Experimental results on both synthetic and real-world collaborative 3D object detection datasets demonstrate that SCORPION advances the state-of-the-art collaborative perception methods by 5.9 - 13.2 absolute AP on both standard and noisy scenarios.

- **Infrastructure-based anomaly detection:** Current research predominantly focuses on detecting attacks from the vehicle's perspective [5], while the role of transportation infrastructure in cybersecurity remains underexplored. As transportation infrastructure supports real-time traffic monitoring and data collection, it offers a unique angle for detecting and mitigating cyber-attacks targeting CAVs. We develop a novel infrastructure-based anomaly detection framework to identify cyber-attacks on CAVs under time interference attacks and V2X communication attacks. The optimal attack strategies are generated using a Pareto optimization that jointly maximizes safety risk and stealthiness. To detect the attacks, a Transformer-based trajectory prediction model is developed to predict normal driving behaviors. An XGBoost classifier is then developed to detect anomalies by comparing predicted and observed vehicle trajectories. We validate the proposed framework using real-world trajectory data. Results show that the proposed anomaly detection model achieves high accuracy with low false positive and

false negative rates in both offline and online settings.

2. Findings

2.1 CAV Ecosystem Threat Analysis and Risk Assessment (TARA)

The development of secure and safe automated vehicles requires i) performing threat analysis and risk assessment (TARA) determines the likelihood and impacts of various cyberattacks in terms of financial damage, safety, privacy, and operational cost, and ii) devising proper security mitigation solutions. In doing so, we need to fully understand the attacks that can be mounted from organized crimes to technical novices. To this end, we have conducted a comprehensive literature review of cybersecurity to analyze existing vulnerabilities and identify potential new attacks against the CAV ecosystem, with a particular focus on sensor attacks, V2X communication attacks and their interplays. We have also quantified attack feasibility, investigated system robustness under attacks, and studied their impacts on the transportation system. Some example case studies are presented in Sections 2.2, 2.3 & 2.4.

Threat Models. Following the standard TARA methodology, we have considered several potential threat agents against next-generation connected and automated transportation systems (NextG-CATS): they can be individual hackers or criminals vs. organized crime syndicates vs. extremist organizations vs. nation state adversaries. Further, these threat agents can be from outside vs. insider (e.g., a disgruntled employee). Depending on the threat agents, their motivations will be very divergent, and thus will likely seek different kinds/levels of impact. As such, the capabilities of the threat agents will also be vastly different, and therefore the opportunity and likelihood of successfully launching attacks will also differ. When developing defense mechanisms against cyberattacks, it is hence important to explicitly define the threat model, lay out the assumptions regarding the motivations and capabilities of threat so as to clearly identify the opportunity and likelihood of the attacks under consideration and quantify the potential impact.

Securing and Hardening Complex NextG-CATS Ecosystem: Two Perspectives. In order to secure and harden next-generation connected and automated transportation systems (NextG-CATS) in a holistic manner, we adopt two key differing yet complementary perspectives: **Constituent System Perspective** vs. **Key Building Block/Asset Perspective,** as shown in Figure 1 below.

Taking the constituent system perspective, the key research questions can be boiled down to the following: 1) Can we – and how do we – secure & harden each constituent system? Perhaps more importantly, can we – and how do we – prevent insecurity of or attack on one component from threatening the security of the entire system? Taking the key building block/assess perspective, the key research questions can be boiled down to the following: 1) How can we leverage many existing defense mechanisms and techniques to harden various software modules & hardware components? 2) What new security countermeasures are needed to secure (V2X) communications, data and AI models that are critical to NextG-CATS?

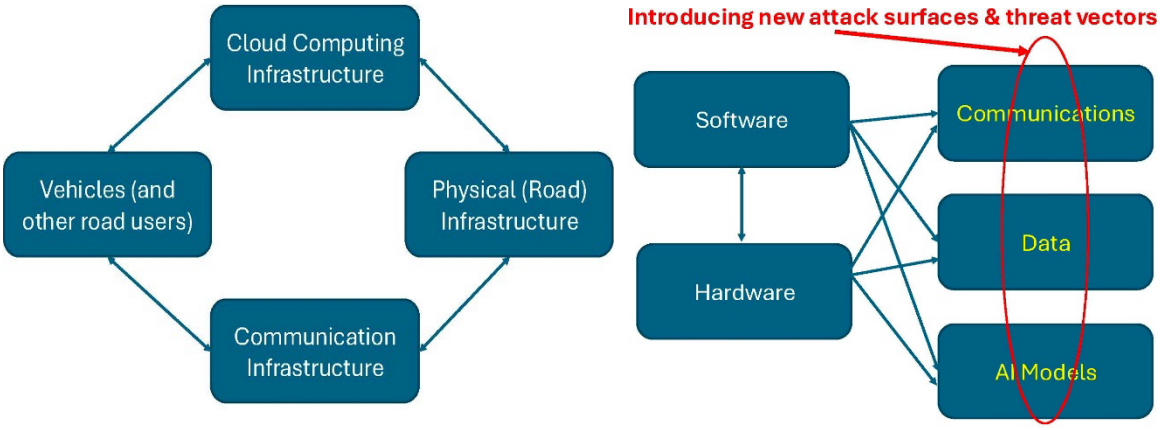

Figure 1 Securing and Hardening Complex NextG-CATS Ecosystem: Two Perspectives.

New Attack Surfaces/Vectors and Defenses. Based on our systematic analysis and threat assessment, we have identified new attack surfaces and vectors against NextG-CATS that require in-depth research. Because of their reliance on communications such as 5G/NextG networks and V2X communication technologies as well as artificial intelligence (AI) models for shared precipitation and autonomous decision making, safeguarding the NextG-CATS against new (physical and cyber) threats such as attacks against the communication infrastructure, cloud services and AI models, e.g., communication channel hijacking, man-in-the-middle attacks, data injection attacks, fake shared sensor data, physical environment manipulation and other adversarial machine learning attacks, is essential. In the following subsections, we present several case studies.

2.2 Shared state approach against cybersecurity attacks on CAV teleoperation

NextG CATS are complex, diverse system-of-systems that operate in highly dynamic environments. Securing them requires a holistic approach: as pointed out earlier in Section 2.1, we not only need to harden and secure the individual constituent components and systems, but also ensure that vulnerabilities from one constituent component/system do not propagate to other components/systems. Furthermore, protecting them against new attack surfaces and threat vectors, such as attacks against the communication and cloud infrastructure, data sharing and AI model, is especially important.

To tackle these challenges, we propose a (shared) state-centric approach to harden and secure NextG-CATS. This approach is developed based on the recognition that CATS is a dynamic, distributed and highly interactive system. For example, autonomous vehicles (AVs), such as cooperative and teleoperated AVs, perceive the environment, interact with objects, including other vehicles, pedestrians, traffic signals, etc., in the environment, share sensor and other data

when needed, and make dynamic decisions (using AI or via teleoperators) based on the perceived environment, share data and remotely receive command and control data (e.g., from a teleoperator or an AI agent running at an edge cloud). The same also applies to intelligent traffic control systems. Hence maintaining (shared) state consistency and integrity (among various entities such as AVs, teleoperators, AI agents, intelligent traffic control systems) is essential. This will help secure against "state manipulation" and detect anomalies, thereby throttling new attacks such as cyber-based data injection, sensor or physical object tampering, model parameter poisoning. The major challenge lies in how to maintain (shared) state consistency and integrity in a timely manner, especially under compute and network constraints. To address this challenge, we use AV teleoperation as a specific case study.

A Case Study: Securing the AV Teleoperation Streaming Pipeline against Data Injection Attacks. As an illustration of our proposed shared state approach, we consider securing the AV teleoperation streaming pipeline against cyber-based data injection attacks as a case study. Figure 2 below depicts the overall pipeline for the sensor data streaming (5G uplink) from an AV to a remote teleoperation station and the command-and-control data delivery (5G downlink) from the teleoperation station to the AV, based on a teleoperation system platform we have developed at the University of Minnesota. In the same figure, we also illustrate how a malicious hacker can potentially inject falsified data through the 5G networks or cloud services.

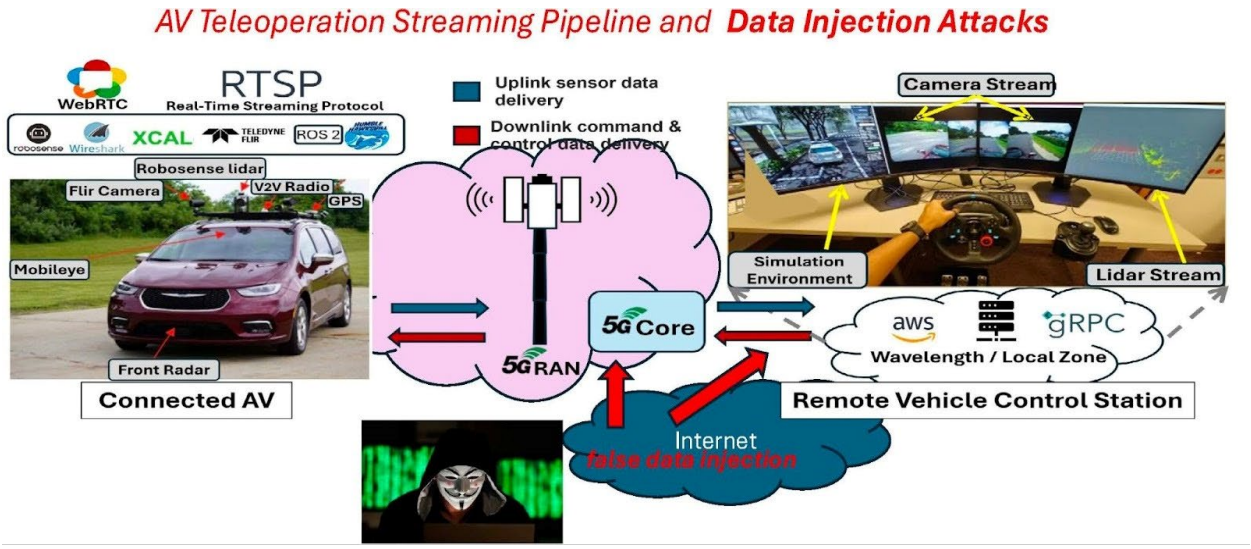

Figure 2 AV Teleoperation Streaming Pipeline and Potential Cyber-based Data Injection Attacks

In order to secure the AV teleoperation data streaming pipeline, one straightforward and naive solution is to simply encrypt all the data delivered through the pipeline. This makes sense for, e.g., the downlink command-and-control data delivery, since the amount of data is relatively small. Thus, encrypting them incurs little overhead. On the other hand, encrypting all sensor data such as video streams can potentially

incur significant computational costs and furthermore, increase the network latency. We therefore need a lightweight approach that minimizes compute overhead and network latency. Applying the shared state principle, our approach is to encode and encrypt a small amount of "water marking" data directly onto video frames – these watermarking data include initial (shared) state synchronization data, e.g., a random secret, sequence number as well as frame identifier, timing, GPS location, and other state information. The "piggy-backed" state information provides the added benefits that it allows us to embed, calculate and feedback (frame-level) network quality-of-service (QoS) metrics as well as application-level quality-of-experience (QoE) metrics. Figure 3 provides a snapshot of the AV teleoperation streaming pipeline operation in a real-world driving experiment.

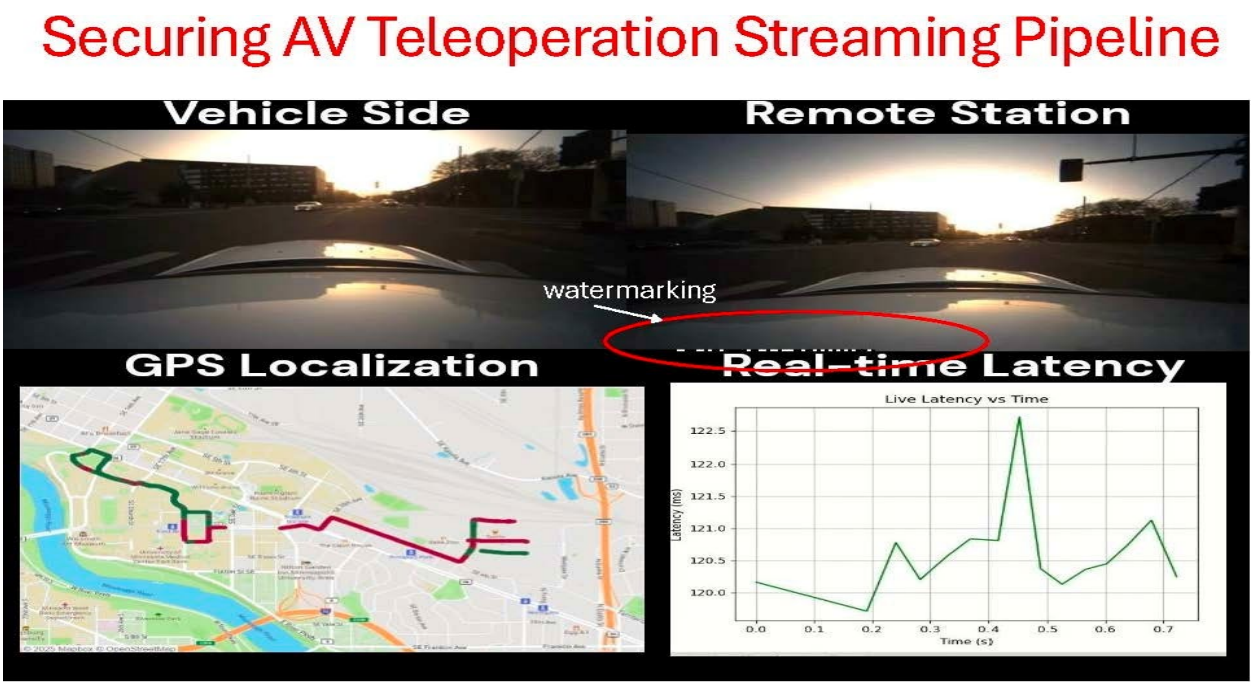

Figure 3 A Snapshot of the AV Teleoperation Streaming Operation in a Real-World Driving Experiment: The figure on top left shows a video frame captured by the onboard front view camera; the figure on top right shows the video frame displayed at the teleoperation state, where the bottom of the video contains the embedded watermarking data. The bottom left figure shows the route of the driving experiment. The borrow right figure shows the one-way frame-level latency calculated using the embedded timing information contained in the watermarking data.

2.3 Robust collaborative perception under lossy networks

Our proposed system SCORPION: Robust Spatial-Temporal Collaborative Perception Model on Lossy Wireless Network addresses the critical challenges of collaborative perception in connected and autonomous vehicles (CAVs). While sharing sensor data between agents (vehicles and

infrastructure) extends sensing ranges and improves safety, realizing this potential is hindered by inherent system imperfections. These imperfections include packet loss in Vehicle-to-Everything (V2X) networks, localization errors from GPS, and synchronization errors caused by clock deviations and network latency, as shown in Figure 4. Existing research often assumes ideal conditions or addresses these issues in isolation. To bridge this gap, we propose SCORPION, a novel end-to-end framework that co-designs the application and system layers. It introduces a Bird's-Eye View (BEV) based fusion model featuring Lost BEV Feature Reconstruction (L-BEV-R), Deformable Spatial Cross-Attention (DSCA), and Temporal Alignment (TA) to robustly handle noisy, lossy, and asynchronous real-world environments. The overall architecture of SCORPION is shown in Figure 5.

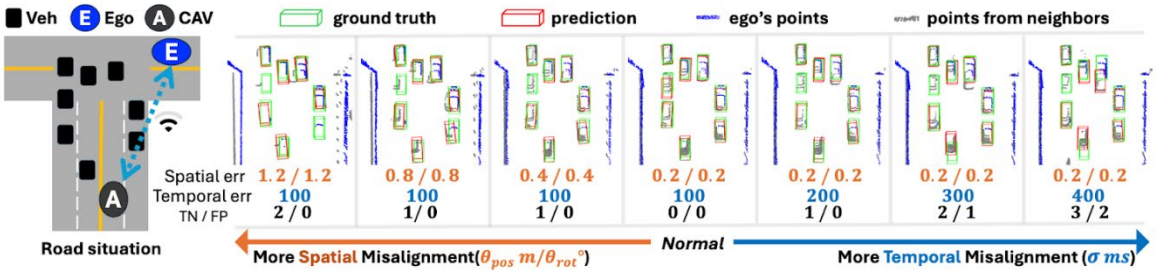

Figure 4 Motivating examples of various degrees of spatial and temporal misalignments. Spatial misalignments occur due to sensing errors (⋆) or dropped network packets, while temporal misalignments arise from sensor asynchronization (⋆) and network delays. For simplicity and clarity, we focus on the cases marked with (⋆) in this figure. This scenario is generated using the CoBEVT [6] model on the V2XSet Dataset [7]. TN and FP represent the True Negative and False Positive detections made by the model.

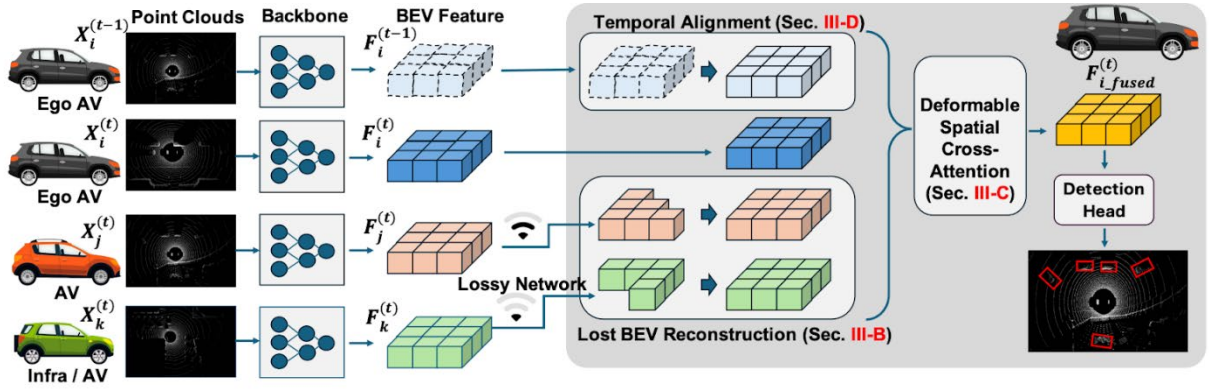

Figure 5 The overall architecture of SCORPION. The framework consists of three major components: a lost BEV feature reconstruction (L-BEV-R) module, a deformable spatial cross-attention (DSCA) module, and a temporal alignment (TA) module.

Superior Robustness to Network Loss: SCORPION consistently outperforms state-of-the-art (SOTA) methods under all levels of network loss. At an 80% loss rate, it achieves 5.1%, 7.5%, and 3.1% higher accuracy than the second-best approach on V2XSet, OPV2V, and DAIR-V2X datasets, respectively.
Resilience to Localization Errors: The model surpasses existing pose-error aware methods, improving average perception performance by 6.8%, 1.3%, and 1.0% across three datasets under varying localization noise. Compared to the robust baseline SCOPE, SCORPION demonstrates 5.5% to 16.9% better performance under high noise conditions.
Effective Handling of Synchronization Errors: Unlike competitors, SCORPION improves perception accuracy by 1.1% to 11.6% across datasets in the presence of synchronization delays.
Performance in Realistic "Combined Noise" Scenarios: In realistic tests where network loss, localization errors, and synchronization errors coexist, SCORPION outperforms baselines by over 4%, 2%, and 5% in absolute Average Precision (AP) for V2XSet, OPV2V, and DAIR-V2X datasets.
Ablation Validation: Systematic removal of components (L-BEV-R, DSCA, TA) consistently resulted in accuracy drops, validating that each module significantly contributes to the model's robustness.

2.4 Infrastructure-based anomaly detection

An overview of the proposed detection framework is shown in Figure 6. When the victim CAV (blue vehicle) is under attack (in the red dashed box), due to compromised perception system and/or communication data, the vehicle may misjudge its surrounding vehicles' status (e.g., the yellow vehicle). As a result, its trajectory planning model may generate trajectories that are not consistent with normal driving behavior (e.g., interactions with surrounding vehicles). On the other hand, the infrastructure is able to detect all vehicles and predict their future trajectories in normal scenarios without attacks (in the blue dashed box). In our study, a Transformer-based prediction model is developed. The predicted trajectory can be used as a baseline to compare with the observed trajectories of the two vehicles. Then an anomaly classifier can be applied to determine whether the observed trajectories are under attack based on a list of high-level features. The proposed method is generic and can be applied to various attack types, as long as abnormal driving behaviors appear in the victim vehicle. In the case study, we test the proposed detection framework under two types of attacks. The first one targets onboard camera sensors and the other one targets V2X communication systems. We employ Pareto optimization to determine "optimal" attack scenarios under the two types of attacks, considering both high safety risk and stealthiness. All attack scenarios are derived from real world vehicle trajectories collected from the Mcity 2.0 dataset [8] using a roundabout merging driving scenario. The "optimal" attack scenarios are then used to train and test the anomaly classifier. The main components of the detection framework and the findings are introduced below. More details can be found in our recent publication [9].

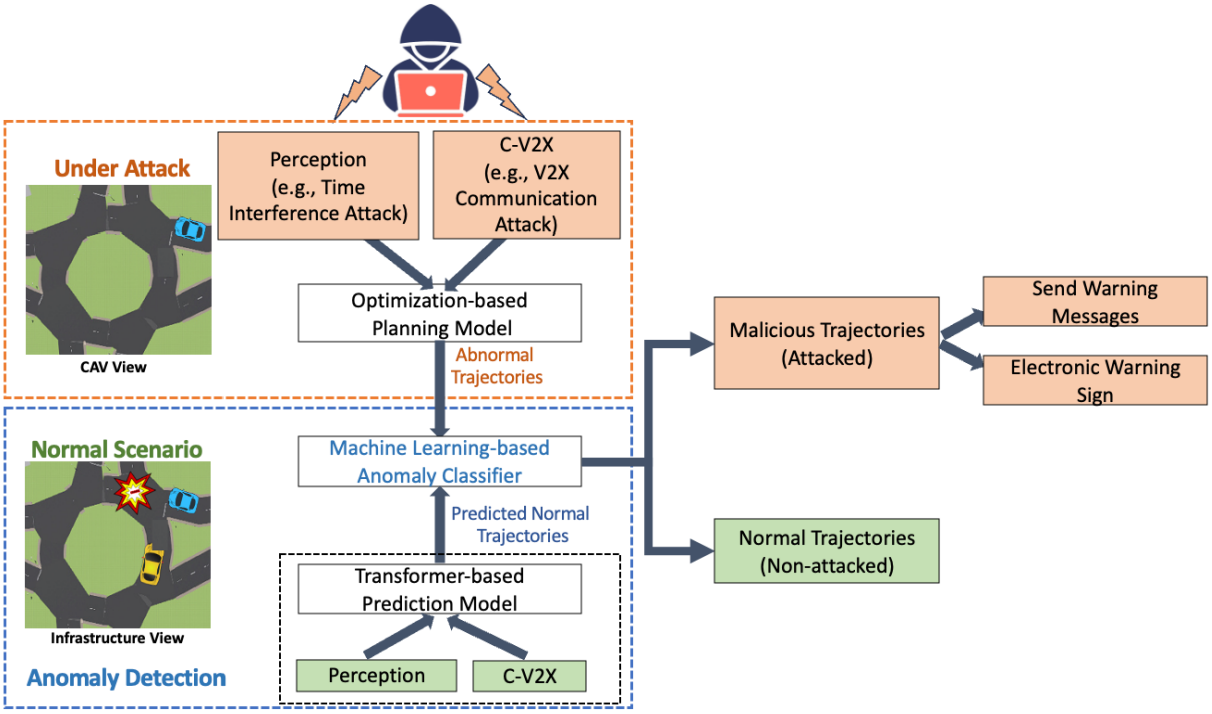

Figure 6 Infrastructure-based anomaly detection framework

CAV Trajectory Planning Model: Trajectory planning for CAVs is typically aimed at ensuring traffic safety [10] and enhancing operational efficiency [11][12]. We apply a simple optimization-based trajectory planning model in this study. The objective function comprises three components representing safety, efficiency, and smoothness. The first term minimizes the vehicle's acceleration, promoting smooth driving. The second term enhances efficiency by encouraging the vehicle to maintain a speed close to the design velocity within the roundabout. The third term introduces an exponential penalty for desired time headway deviation with other vehicles, thereby prioritizing both safety and efficiency during the merging process. Constraints include boundaries of speed, acceleration, jerk, and initial conditions.

Time Interference Attack Threat Model: The SlowPerception attack is a latency-based attack on camera systems within autonomous driving (AD) systems, proposed by Ma et al. [13]. This attack introduces adversarial perturbations (e.g., using a projector) to generate phantom objects on real-world surfaces. It significantly increases computational latency in the object detection pipeline—reportedly reaching an average of 2.5 seconds across various scenarios. When a CAV is under SlowPerception attack, the vehicle fails to detect surrounding vehicles. Therefore, the objective function of the trajectory planning model should be modified by removing the interaction term with the surrounding vehicle.

V2X Communication Attack Threat Model: Reliable V2X communication underpins essential safety functions such as cooperative perception, hazard detection, and coordinated maneuvers (e.g., intersection crossing, lane merging). However, V2X systems are inherently susceptible to communication-layer threats, particularly DoS attacks, which can severely impair their availability and reliability. In normal operations, CAVs rely on V2X messages for decision-making, particularly in environments where onboard sensors' coverage or detection capability is limited (e.g., occlusion). However, under a DoS attack, the Packet Delivery Ratio (PDR) can drop to as low as 10%, while End-to-End (E2E) communication latency may rise to as high as 1000 ms, according to [14]. Under such attack scenarios, the objective function of the trajectory planning model is revised to include a delay term to represent the communication latency.

"Optimal" Attack Strategy Identification: The timing and duration of an attack significantly influence driving behaviors and consequences. We develop a systematic approach to identify "optimal" attack scenarios under the two threat models. We consider two factors to quantify the "optimality" of attacks including safety risks and stealthiness. Time-to-Collision (TTC) has been widely implemented as a safety indicator, with lower TTC values indicating higher risk of imminent collisions. Conversely, shorter attack durations represent higher stealthiness, as mild disruptions are more difficult to detect and respond. Therefore, identifying optimal attack cases can be considered as a multi-objective optimization problem. After the Pareto fronts are identified, we further quantify minimum TTC and attack duration into two cost terms to select the unique optimal solution.

Transformer-based Trajectory Prediction: We apply a Transformer-based trajectory prediction framework to predict CAV trajectories under normal roundabout merging scenarios without attack. We select nine features in the Transformer-based prediction model. The features can be categorized into three groups: road characteristics including the difference between vehicle heading and road direction; vehicle motion including x and y coordinates and headings of the CAV and the in-roundabout vehicle, respectively; and their interactions including TTC and distances to the collision point of the two vehicles, respectively.

XGBoost-based Anomaly Classifier: An XGBoost model is trained as the anomaly classifier, utilizing statistical features derived from key interaction variables of the CAV and the in-roundabout vehicle. Specifically, the classifier incorporates the maximum, minimum, and standard deviation of six interaction variables, including 1) the speed difference between the predicted and observed velocity of the CAV; 2) the distance differences (between predicted and observed) traveled by the CAV from its starting position up to time t; 3) the observed TTCs; 4) the observed net distances; 5) the observed accelerations; and 6) the observed acceleration deviation differences. The anomaly classifier is implemented in both offline detection and online detection settings.

Results: The Mcity 2.0 dataset is applied to validate the proposed models, which collected from the real-world roundabout at Ellsworth and State St, Ann Arbor, MI. We select 278 vehicle trajectory pairs, under which adversarial perturbations are more likely to induce safety-critical events, to identify the "optimal" attack strategies. In the 278 pairs of ground truth trajectories (without attack), there are 157 conflict cases (TTC < 2.0 seconds) with no near-crash (TTC < 0.5 seconds) nor crash (TTC < 0.1 seconds) cases. In contrast, under the time interference attack, there are 251 near-crash events, of which 49 progress to crashes. Similarly, under the V2X communication attack, we observe 188 near-crash events, including 10 crashes. To evaluate the effectiveness of the proposed detection framework, only cases with a minimum TTC below 0.5 seconds under the two threat models are considered (251 under time interference attack and 188 under V2X communication attack), as our focus is on scenarios involving significant safety risks. Seventy percent of the cases are used for training (176 cases under time interference attack and 132 under V2X communication attack), and the remaining 30% are used for testing (76 and 57 cases, respectively). Under time interference attacks, the proposed detection framework achieved strong performance in both offline and online settings. In offline detection, the model attained 100% accuracy, precision, recall, and F1 score. To mitigate potential bias from class imbalance, the numbers of true positives (TP) and true negatives (TN) in the dataset are randomly sampled to achieve approximate balance. Both the false positive rate (FPR) and the false negative rate (FNR) are 0.00% (0 out of 38 and 0 out of 37, respectively), indicating perfect classification performance. The online detection achieved 96% accuracy, with a FPR of 2.70% (1 out of 37) and a FNR of 5.26% (2 out of 38). The mean attack success time - defined as the time at which TTC first drops below 0.5 seconds from the start of the anomaly, is 2.7 seconds. The mean detection time is 1.3 seconds, yielding an average lead time of 1.4 seconds for mitigation. Under V2X communication attacks, the offline detection performance remains consistently high, achieving 96.43% accuracy, 96.66% precision, 96.43% recall, and 96.43% F1 score. The FPR is 7.14% (2 out of 28), while the FNR is 0.00% (0 out of 28), indicating reliable detection with minimal misclassifications. In the online setting, the model achieves 92.86% accuracy, with both the FPR and FNR at 7.14% (2 out of 28). The mean attack success time is 3.0 seconds, and the mean detection time is 2.6 seconds, resulting in an average response margin of 0.4 seconds before the attack leads to a critical event.

3. Recommendations

As next-generation connected and automated transportation systems, are complex, diverse system-of-systems that operate in highly dynamic environments, it is imperative that we take a holistic approach to not only harden and secure the individual constituent components and systems, but also ensure that vulnerabilities from one constituent component/system do not propagate to other components/systems. New security mechanisms to protect such systems against new attack surfaces and threat vectors, such as attacks against the communication and cloud infrastructure, data sharing and AI models are particularly needed.

In this project, we comprehensively study how to harden and secure the CAV ecosystem with four key aspects. Although the developed models are tested under specific applications (e.g., teleoperation, collaborative perception), the same methodologies can be applied in a wide range of CAV and CATS applications to enhance cybersecurity and safety.

Specifically, we have advocated a novel *shared state-centric* framework to protect next-generation connected and automated transportation systems against new attacks, and used the problem of securing the sensor data streaming pipeline for teleoperated autonomous vehicles as a case study.

Furthermore, we propose *SCORPION*, which expands AI's role in co-designing the hardware-network-AI software stack by addressing persistent issues like lossy communication and sensor noise in networking and hardware. SCORPION demonstrates improved accuracy and robustness across extensive experiments, including normal and challenging scenarios that expose typical system issues. Future collaborative perception systems should move beyond domain-specific solutions and adopt AI co-design across the hardware, network, and software layers to handle persistent real-world issues. Systems should utilize reconstruction networks (like L-BEV-R) that specifically account for network packet structures and Maximum Segment Size (MSS) to recover spatial semantics lost during V2X transmission. To mitigate the inevitable localization and synchronization errors in dynamic environments, perception models should employ deformable attention mechanisms to allow for flexible, global feature interaction across agents. The framework should be extended to support additional sensor modalities, such as monocular and stereo cameras, to further enhance robustness.

We propose a novel *infrastructure-based* anomaly detection framework for identifying cyber-attacks targeting CAVs. By leveraging infrastructure-side sensing, monitoring and prediction capabilities, the framework enabled independent anomaly detection without relying on compromised vehicle data. This study provides valuable insights in transportation infrastructure's role in transportation system cybersecurity. There are several promising directions for future research. First, future work can explore how detected anomalies can be integrated into the decision-making system to enable more responsive and active traffic management. Second, future work can move beyond binary anomaly detection to classify specific types of cyber-attacks—such as spoofing, jamming, or replay to support more targeted mitigation strategies. Third, incorporating human-in-the-loop mechanisms for real-time verification and decision-making. For example, the system could issue warnings to nearby vehicles with human drivers or prompting CAV systems to request human intervention upon detection of anomalous behavior, enabling proactive risk mitigation.

4. Outputs, Outcomes, and Impacts

The following outputs are generated during the performance of this project:

Presentation at CCAT 2024 Global Symposium, May 21 2024

The three PIs of this project delivered a joint presentation at the CCAT 2024 Global Symposium. They presented a proposed security analysis framework to study various cybersecurity risks in the CAV ecosystems. In particular, they discussed new attacks on AI (artificial intelligence) algorithms and systems used for cooperative driving, and propose potential new directions in mitigating such threats.

Poster Presentation and Demonstration at MAASTO 2025, June 16 2025

Several members of the research team, led by PI Zhang, participated in the 2025 MAASTO (Mid America Association of State Transportation Officials) Connected and Automated Vehicle (CAV) Summit, held June 16–18 at the Mall of America Event Center in Minneapolis, MN (see https://housmanassociates.swoogo.com/maastocav2025/8261354). The summit brought together transportation professionals, innovators, researchers, and policymakers to exchange knowledge and address emerging opportunities and challenges in CAV deployment. The team actively engaged with industry experts, government representatives, and fellow researchers during the Technology Showcase Reception, where they explored innovative solutions and shared perspectives on the future of CAV technologies. As part of the showcase, researchers from PI Zhang's group demonstrated teleoperation use-case scenarios utilizing both an experimental vehicle and an autonomous robot. These demonstrations highlighted the team's work on bridging autonomy and need for human intervention through teleoperation, and sparked meaningful discussions with public- and private-sector representatives on potential applications and collaborative opportunities. Throughout the summit, team members attended technical sessions and breakout discussions covering topics such as connected infrastructure, safety assurance, spectrum allocation, and mobility equity. Their participation enabled knowledge sharing on challenges such as bandwidth constraints for real-time communication, cybersecurity risks, and interoperability across different CAV systems.

MN CAV Summit Technology Showcase, September 24, 2025.

The research team, led by PI Zhang, participated in the 2025 MNCAV Tech Showcase, an event bringing together transportation agencies, private industries, AV startups, and technical experts to highlight emerging innovations in connected and automated vehicles (CAVs). As part of the showcase, the team presented demonstrations featuring the MNCAV experimental vehicle to highlight teleoperation use cases, along with autonomous robots designed to illustrate advanced capabilities in automation and remote operation. These demonstrations provided attendees with a practical view of how teleoperation and autonomous systems can complement one another in real-world transportation scenarios. In addition to the live demonstrations, team members joined discussions with industry leaders, policymakers, and innovators on the challenges and opportunities in deploying CAV technologies. The showcase served as a platform for the team to

share its research, foster collaboration, and contribute to shaping the future of safe and connected mobility.

Journal and Conference Papers

Rostand A. K. Fezeu, Jason Carpenter, Rushikesh Zende, Sree Ganesh Lalitaditya Divakarla, Nitin Varyani, Faaiq Bilal, Steven Sleder, Nanditha Naik, Duncan Joly, Eman Ramadan, Ajay Kumar Gurumadaiah, Zhi-Li Zhang, " Teleoperating Autonomous Vehicles over Commercial 5G Networks: Are We There Yet?" Submitted to IEEE/ACM Transactions on Networking.

Ruiyang Zhu, Minkyoung Cho, Shuqing Zeng, Fan Bai, Z. Morley Mao, "SCORPION: Robust Spatial-Temporal Collaborative Perception Model on Lossy Wireless Network." IROS 2025.

C. Zhang, J. Ying, and Y. Feng, "Smart Infrastructure-based Anomaly Detection Under Cyber Attacks", IEEE Intelligent Transportation Systems Magazine, 2026, (accepted)

New Courses

Led by PI Zhang, the CCAT UMN team has developed a new interdisciplinary seminar on connected and automated transportation systems which were co-listed and offered to graduate students in the Departments of Computer Science, Electrical and Computer Engineering, Mechanical Engineering, and Civil Engineering as well as Humphrey School of Public Affairs in Spring 2025.

Websites

AV teleoperation: https://avteleoperation.umn.edu/
Video demonstration of SCORPION: https://www.youtube.com/watch?v=6zCAlhH7pGw

The following outcomes are generated during the performance of this project:

- **Outcome 1: Improves the operation and safety of the transportation system.** The research has led to the development of an AV teleoperation streaming pipeline with light-weight security mechanisms to safeguard the operation and safety of AV teleoperations.
- **Outcome 2: Increases in the body of knowledge:** The research establishes that robust collaborative perception requires a holistic view of the system stack. It proves that interpreting network packet loss as "masked images" in a Vision Transformer allows for effective data recovery, advancing knowledge in both networking and computer vision.
- **Outcome 3: Improved processes and technologies:** The development of the DSCA and TA modules provides a proven methodology for mitigating GPS and clock errors, which can

be adopted by future autonomous driving systems to improve reliability.

- **Outcome 4: Improved awareness of transportation infrastructure in cybersecurity**: The implementation of infrastructure-based anomaly detection system under two CAV cyberattack scenarios demonstrates the role and importance of transportation infrastructure in defensing the CAV ecosystems.

The following impacts are generated during the performance of this project:

- **Impact 1: Improves the operation and safety of the transportation system:** By improving object detection accuracy by up to 13.2% in realistic, lossy network conditions, the SCORPION technology directly contributes to safer autonomous navigation and reduces the risk of accidents caused by communication failures or sensor noise. With a highly accurate detection rate, the infrastructure-based anomaly detection framework can be effectively used to protect CAVs under various cyberattacks and improves transportation system safety.

- **Impact 2: Increases the body of knowledge and technologies:** The research sets a new baseline for "robust" collaborative perception, encouraging the research community to move away from ideal simulations and focus on practical, real-world deployment challenges involving hardware and network constraints.

- **Impact 3: Improved processes and technologies.** The research has produced a teleoperation system platform that has been deployed on the MnCAV vehicle – a level-3 research AV at the UMN. The instrumented MnCAV vehicle has been used for various real-world research experimentation to improve AV teleoperation processes and technologies.

References

[1] US Department of Transportation, “Our nation’s roadway safety crisis,” 2023. https://storymaps.arcgis.com/stories/9e0e6b7397734c1387172bbc0001f29b, Last accessed: December 23, 2025.
[2] G. Pressroom, “Gm expands super cruise network to 750,000 hands-free miles, largest in north America.” https://media.gm.com/media/me/en/gmc/news/news_archive.detail.html/content/Pages/news/us/en/2024/feb/0215-supercruise.html, Last accessed: December 23, 2025.
[3] Tesla Autopilot and Full Self-Driving Capability. https://www.tesla.com/support/autopilot. Last Accessed 12-23-2025
[4] Michael Liedtke, "Amazon’s Zoox launches its robotaxi service in Las Vegas", AP News, Updated 12:59 PM CST, September 10, 2025. https://apnews.com/article/amazon-zoox-robotaxis-las-vegas-bd5cb24602fb16243efcba05c7fe518f, Last Accessed 12-23-2025
[5] S. Dasgupta, M. Rahman, M. Islam, and M. Chowdhury, “Prediction-based GNSS spoofing attack detection for autonomous vehicles,” arXiv preprint arXiv:2010.11722, 2020.
[6] R. Xu, Z. Tu, H. Xiang, W. Shao, B. Zhou, and J. Ma, “Cobevt: Cooperative bird’s eye view semantic segmentation with sparse transformers,” arXiv preprint arXiv:2207.02202, 2022.
[7] R. Xu, H. Xiang, Z. Tu, X. Xia, M.-H. Yang, and J. Ma, “V2x-vit: Vehicle-to-everything cooperative perception with vision transformer,” in European conference on computer vision. Springer, 2022, pp. 107– 124.
[8] R. Zhang, Z. Zou, S. Shen, and H. X. Liu, “Design, implementation, and evaluation of a roadside cooperative perception system,” Transportation Research Record, vol. 2676, no. 11, pp. 273–284, 2022.
[9] C. Zhang, J. Ying, and Y. Feng, "Smart Infrastructure-based Anomaly Detection Under Cyber Attacks", IEEE Intelligent Transportation Systems Magazine, 2026, (accepted)
[10] R. Ren, H. Li, T. Han, C. Tian, C. Zhang, J. Zhang, R. W. Proctor, Y. Chen, and Y. Feng, “Vehicle crash simulations for safety: Introduction of connected and automated vehicles on the roadways,” Accident Analysis & Prevention, vol. 186, p. 107021, 2023.
[11] X. Hu and J. Sun, “Trajectory optimization of connected and autonomous vehicles at a multilane freeway merging area,” Transportation Research Part C: Emerging Technologies, vol. 101, pp. 111–125, 2019.
[12] C. Cui, Z. Yang, Y. Zhou, J. Peng, S.-Y. Park, C. Zhang, Y. Ma, X. Cao, W. Ye, Y. Feng et al., “On-board vision-language models for personalized autonomous vehicle motion control: System design and real-world validation,” arXiv preprint arXiv:2411.11913, 2024.
[13] C. Ma, N. Wang, Z. Zhao, Q. A. Chen, and C. Shen, “Slowperception: Physical-world latency attack against visual perception in autonomous driving,” arXiv preprint arXiv:2406.05800, 2024.
[14] T. Li, M. Shang, S. Wang, and R. Stern, “Detecting subtle cyberattacks on adaptive cruise control vehicles: A machine learning approach,” IEEE Open Journal of Intelligent Transportation Systems, 2024.