1. Report No. UB-24-RP-01	2. Government Accession No N/A	3. Recipient's Catalog No. N/A

4. Title and Subtitle Continuous & Low-cost Inspection of Precast Concrete Bridges using Connected Automated Vehicles (CAVs)	5. Report Date December 2025 6. Performing Organization Code N/A
7. Author(s) Chaozhe He, Assistant Professor, 0000-0002-0299-8412	8. Performing Organization Report No. UB-24-RP-01
9. Performing Organization Name and Address [Enter name and address of university conducting project]	10. Work Unit No. N/A 11. Contract or Grant No. 69A3552348333
12. Sponsoring Organization Name and Address Transportation Infrastructure Precast Innovation Center (TRANS-IPIC) University of Illinois Urbana-Champaign. Civil & Environmental Engineering. 205 N Mathews Ave, Urbana, IL 61801 US Department of Transportation Office of Research, Development and Technology (RD&T) 1200 New Jersey Avenue, SE Washington, DC 201590	13. Type of Report and Period Covered Final Report [Enter Month Year – Month Year] 14. Sponsoring Agency Code N/A

15. Supplementary Notes

For additional reports and information visit the TRANS-IPIC website https://trans-ipic.illinois.edu

16. Abstract

This project established a continuous inspection framework for precast concrete bridges. While specialized inspection vehicles can provide vehicle-view inspections, they are limited by the cost and capacity of these vehicles. Leveraging all road vehicles with forward-facing cameras would maximize the inspection potential, yet a feasibility study is missing. This project attempts to realize such a system by establishing three pillars. 1) a communication protocol based on vehicle-to-infrastructure (V2I) communication that enables inspection requests and results transmission 2) cooperative vehicle motion that allows optimal detection motion by the vehicle, and 3) a crack detection algorithm using a vehicle front-facing camera assisted with V2I communication and vehicle motion. We report on our design and performance results, showcase potential while acknowledging gaps. The technologies developed in this work can be adapted to broader transportation infrastructure inspections by involving all road users, creating continuous and low-cost solution that extend longevity of transportation systems.

17. Key Words remote sensing, cooperative sensing, Vehicle to infrastructure communication, connected automated vehicle	18. Distribution Statement No restrictions.

19. Security Classification (of this report) Unclassified.	20. Security Classification (of this page) Unclassified.	21. No. of Pages 21	22. Price N/A

Disclaimer:

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated in the interest of information exchange. The report is funded, partially or entirely, under 69A3552348333 from the U.S. Department of Transportation's University Transportation Centers Program. The U.S. Government assumes no liability for the contents or use thereof.

Transportation Infrastructure Precast Innovation Center (TRANS-IPIC)

University Transportation Center (UTC)

Continuous & Low-cost Inspection of Precast Concrete Bridges using Connected Automated Vehicles (CAVs)
UB-24-RP-01

FINAL REPORT

Submitted by:
PI: Chaozhe He, chaozhe@buffalo.edu
Department of Mechanical and Aerospace Engineering
University at Buffalo

Collaborators / Partners:
None

Submitted to:
TRANS-IPIC UTC
University of Illinois Urbana-Champaign
Urbana, IL

Executive Summary:

This project established a continuous inspection framework for precast concrete bridges. While specialized inspection vehicles can provide vehicle-view inspections, they are limited by the cost and capacity of these vehicles. Leveraging all road vehicles with forward-facing cameras would maximize the inspection potential, yet a feasibility study is missing. This project attempts to realize such a system by establishing three pillars. 1) a communication protocol based on vehicle-to-infrastructure (V2I) communication that enables inspection requests and results transmission 2) cooperative vehicle motion that allows optimal detection motion by the vehicle, and 3) a crack detection algorithm using a vehicle front-facing camera assisted with V2I communication and vehicle motion. We report on our design and performance results, showcase potential while acknowledging gaps. The technologies developed in this work can be adapted to broader transportation infrastructure inspections by involving all road users, creating continuous and low-cost solution that extend longevity of transportation systems.

Table of Contents

1. Problem description:

Precast bridge components offer significant advantages in construction speed while maintaining high quality, thanks to rigorous quality control measures in offsite manufacturing facilities. Despite their superior production quality, precast members still require regular inspections, particularly in connection zones, to identify potential joint degradation. Visible signs of damage, such as joint cracking, grout spalling, or reflective cracking, can indicate underlying issues for members such as adjacent beams, decked beams, and precast concrete deck panels. Loss of joint material or joint cracking can lead to durability problems, unintended live load distribution, and joint or member capacity loss. Manual inspections, reliant on visual examination, are labor-intensive and pose safety risks. Although not yet standard practice, frequent inspections are crucial for precast decks, decked girders, and adjacent members prone to reflective cracking. Early damage detection is vital, as deterioration can be visible on the surface of these components. While remote sensing systems using drones and dedicated inspection vehicles are commercially available, their operating costs remain high, making continuous monitoring uneconomical. Currently, there is a lack of cost-effective, automated inspection systems that can detect damage between routine inspections, particularly early signs of degradation. Developing such a system would enable continuous monitoring over time, reducing the risk of undetected damage and enhancing the overall durability of precast bridge components.

This project pioneers a transformative approach to monitoring early damage and degradation in precast concrete components using Vehicle-to-Everything (V2X) technologies. By harnessing the power of crowd sensing through regular road vehicles equipped with forward-facing cameras, our approach offers a cost-effective solution for continuous monitoring. Integrating connected automated vehicles (CAVs) can further enhance inspection and monitoring performance. This V2X-based system has the potential to be extended to other road infrastructure elements, amplifying the safety and durability of existing and future US infrastructure. This project's goal resonates with TRANS-IPIC's mission to develop innovative solutions for precast concrete components condition monitoring and remote sensing, revolutionizing the development and performance of future road infrastructure. It also aligns with US DOT priorities to enhance infrastructure safety and durability. Our project goals complement the recently released national deployment plan for V2X technologies of the US DOT [1], thus can leverage such deployment to boost implementable outcomes. By converging these efforts, we can unlock significant gains in durability, safety, and resource efficiency for repair and replacement.

2. Background (including literature review):

Precast concrete bridge elements offer significantly improved bridge quality and durability, among other benefits. However, routine inspections are still needed, especially at joints between the components. Deterioration at these joints can compromise structural integrity and capacity. Continuous monitoring can identify joint degradation that may be visible through joint cracking, grout spalling, or reflective cracking. Continuous, real-time, and cost-effective inspections could detect early signs of damage in time, allowing early intervention and thus improving bridge durability.

There are 17,607 highway bridges in New York State (NYS) and more than 600,000 in the United States [2]. Federal regulations mandate routine inspections every 24 months to ensure safety; manual inspections remain the predominant means of inspection. Manual inspections typically include road closures, which cause traffic delays. Partial closures can pose safety risks to manual inspectors. Furthermore, manual inspections can be expensive. Studies have shown that manual inspection on a 4-lane divided highway bridge costs $4,600 directly, and this cost could increase to more than $10,000 if the cost due to traffic delays and road closures is considered [3]. The total inspection cost for bridges in NYS alone amounted to 50 million annually, and improving inspection efficiency carries economic benefits. Constrained by cost and procedure, manual inspections cannot be carried out frequently for every bridge.

State-of-the-art intelligent inspection methods, including drone and ground vehicle inspections, leverage a suite of advanced, non-intrusive sensors such as high-resolution cameras, LiDAR, thermal imaging cameras, and ground-penetrating radars (GPRs) [4, 5]. While these approaches have demonstrated impressive inspection performance and significant cost savings, they are better suited for in-depth, targeted inspections - augmenting or replacing manual inspections - rather than continuous, real-time monitoring. Meanwhile, existing continuous monitoring is usually done with sensors overlayed on the bridge infrastructure. The setup complexity and high upfront cost make them suitable for short-term studies rather than long-term inspection purposes. The missing pieces are the interplay between bridge users, such as vehicles, and the infrastructure for bridge inspection.

The widespread adoption of active safety technology has transformed modern vehicles into capable sensing hubs equipped with cameras and wireless communication devices that enable them to collect, process, and transmit data in real time. Thus, they have evolved from mobile terminals to intelligent data acquisition and processing units. With sensors such as forward-facing cameras continuously monitoring the road, vehicles can inspect infrastructure status based on vision. Existing work along this line focuses only on benefiting the vehicle: estimating road roughness and friction coefficient and detecting pavement cracks and potholes to improve ride comfort [6, 7]. Opportunities to extend such capability to improve infrastructure durability are unexplored. One major obstacle is how to enable direct information sharing between vehicles and infrastructure regarding the inspection results acquired from road vehicles.

V2X technologies allow secure, high-bandwidth, instantaneous communication between vehicles and infrastructure within a perimeter of 1 mile. Vehicles with connectivity can share information with other vehicles and infrastructure to increase the overall situation awareness of all road users. CAVs, furthermore, could navigate safely and actively improve traffic safety and efficiency [8]. Given their proven benefit in improving traffic safety and efficiency, the US DOT has released a national deployment plan for V2X technologies. These prior and ongoing activities enable a new continuous real-time infrastructure inspection system to leverage modern connected vehicles and infrastructure. However, careful design is needed to deliver such a system, as two immediate challenges must be addressed. (1) The state-of-the-art vision-based road surface detection techniques are optimized for generic road pavement detection but not bridge crack inspections. When inspecting bridges with precast components, the area of interest is typically the component connection zones. Effective detections must consider this, which is missing in the current algorithm. (2) The existing V2X application applies to infrastructure and focuses mainly on improving traffic efficiency and safety; the benefits to the infrastructure are ignored. As a result, the established smart sensing system is optimized for vehicle detection rather than road surface detection. A continuous inspection system that leverages all available road vehicles should address all the above challenges while remaining low-cost by maximizing the usage of existing infrastructure and road users' (e.g., vehicles) capabilities. The target continuous inspection system complements the current manual and intelligent inspection techniques surveyed, as its inspection results can be used to schedule more efficient, on-demand, and just-in-time inspections using these more cost-intensive yet detailed inspection methods. The target system and the underlying system can also be extended to infrastructure elements beyond precast bridges. Detecting early damage signs specific to the infrastructure allows early intervention, in-depth inspections, and repairs, thus improving the safety and durability of all transportation infrastructure

To date, there is no standard application for bridge/infrastructure inspections. More specifically, an investigation into the feasibility of vehicular-centric detection using a forward-facing camera is not being conducted.

3. Research scope and objectives:

This project aims to develop a continuous, low-cost data collection system for bridges with precast components prone to reflective cracking. Early damage detection is crucial for durability, but existing inspection methods are costly, primarily when performed continuously. Leveraging existing vehicles on the roads, we propose using vehicle-to-everything (V2X) technologies, where bridges request connected vehicles with standard forward-facing cameras to inspect critical areas. Connected automated vehicles (CAVs) will execute cooperative motion to maximize inspection efficiency. Aggregated data from continuous traffic and inspection results will detect early signs of damage, enabling proactive maintenance.

To deliver such a system, a couple of problems need to be tackled: 1) how to design a robust crack-detection algorithm that can work with standard forward-facing cameras with limited focus on the road surface; 2) how to effectively communicate detection requests and results using CV2X communication; 3) how to coordinate the motion of the CAVs to maximize the detection results.

In the first year of the project, we focus on the viability of the crack detection algorithm approach. Specifically, we aim to achieve the following goals

- Enabling crack detection and estimation in real-time: pipeline establishment
    - Vehicle side crack detection algorithm
    - Vehicle to infrastructure communication protocol
    - Vehicle cooperative control for crack detection.

- Understand the gap in the pipeline, benchmark the performance, and investigate future ways for future improvements.

4. Research description:

To achieve the research objective, we propose to develop techniques that achieve a prototype system summarized in Figure 1. We seek to establish three pillars that allow long-term development to achieve the target continuous inspection of pre-cast concrete bridge. 1) a communication protocol based on vehicle-to-infrastructure (V2I) communication that enables inspection requests and results transmission; 2) cooperative vehicle motion that allows optimal detection motion by the vehicle, and 3) a crack detection algorithm using a vehicle front-facing camera assisted with V2I communication and vehicle motion.

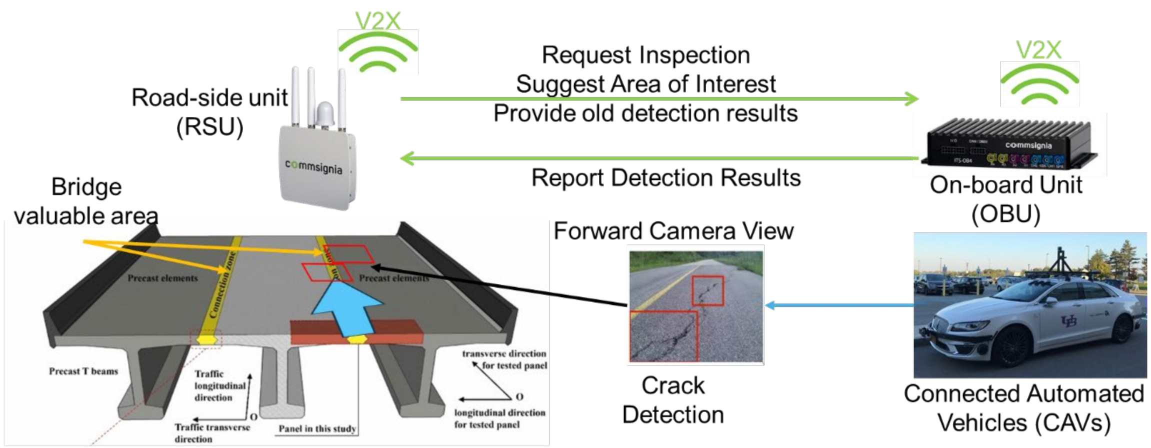

Figure 1 Research Overview

Candidate bridge of Inspection
With the help of the NYS DOT team. We have identified one precast bridge candidate (Bridge Bin# 1045890, [Google Map]) near the UB campus; see Figure 2. We have obtained historical inspection results

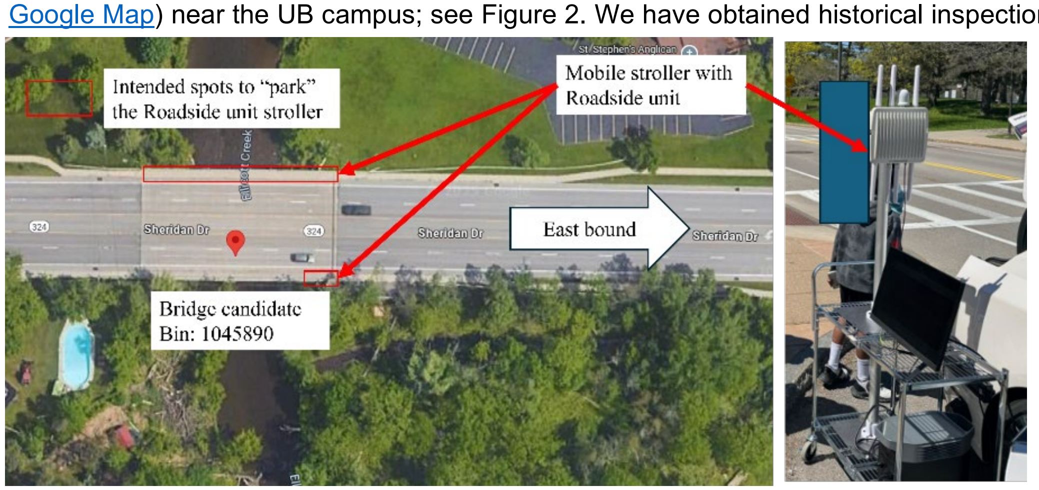

Figure 2 Identified bridge candidate and testing set up

and started planning the demo and evaluation metrics. Throughout the project, the team has made a few visits to the bridge site and collected images of the bridge deck surfaces. to calibrate and test our platform. See Figure 3 for a series of images collected from multiple sources regarding the same reflective crack spot. While the bridge deck had reflective cracks developed over the years, it was sealed in early

September. Due to this maintenance, the planned demonstration was able to detect any cracks. To continue the development and fine-tune effort, the ongoing focus is now shifted to surface crack detection on the main campus site identified from last quarter (Bridge Structure #5511840). The bridge has a reinforced concrete slab, and many visible cracks can be used for both training and testing purposes.

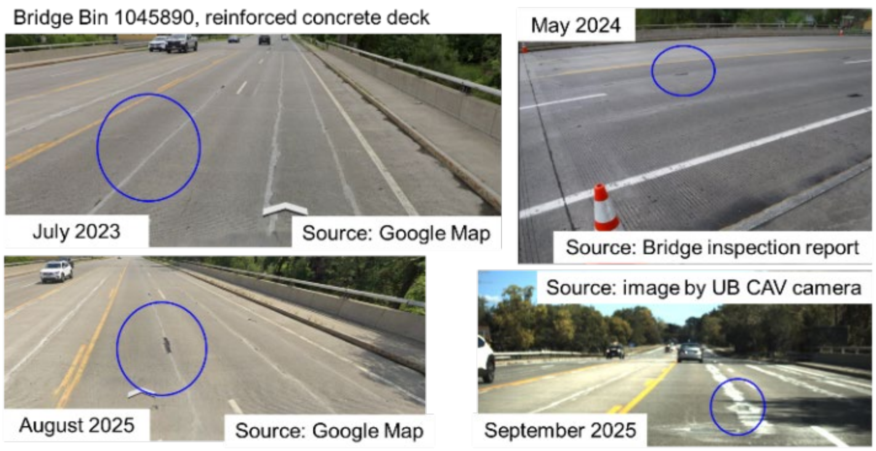

Figure 3 Pilot Bridge Deck Reflective Crack Evolution

5. Project results:

Over this project period, the team first established the hardware platform, which includes both the vehicle and the infrastructure, enabling the development of a connectivity-based inspection system. Building on the system, the team investigated thoroughly crack detection models (5.2), designed, evaluated, and implemented an inspection communication protocol (5.3) and a cooperative control strategy (5.4). Finally, the team establishes the complete pipeline and evaluates performance in real-world bridge and road crack conditions (5.5).

5.1. Hardware Preparation

Over the project period, the following list of hardware has been integrated to establish the prototyping cooperative inspection systems. The connected automated vehicle (CAV) in Figure (c) Equipped with a

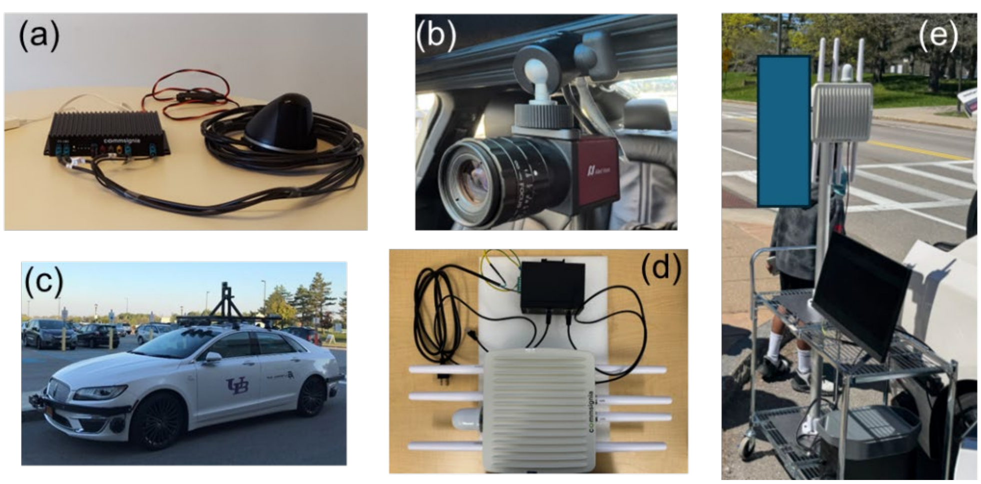

Figure 4 Hardware Preparation

drive-by-wire system allows speed control and path following, as well as RTK GPS, which allows centimeter accuracy of localization. It serves as the base development vehicle for data collection, detection pipeline development, and testing. It is integrated with a C-V2X onboard unit (OBU), shown in Figure (a), is manufactured by Commsignia, and broadcasts standard messages according to SAE J2735 [9]. The CAV is also retrofitted with a Mako G-319 CMOS camera with a resolution of 2064(H) x 1544(V), which provides raw images for surface crack detection. On the infrastructure side, a roadside unit (RSU), shown in Figure (d), which is also manufactured by Commsigina, serves as the inspection coordination unit that resides in the infrastructure (e.g., precast concrete bridges). To facilitate development and testing at different sites, a portable stroller, shown in Figure (e), was built to house the RSU as well as a portable charging station and monitor.

5.2. Crack Detection Model Investigation

Model Selection

In this project period, the main focus of model selection was to leverage available state-of-the-art concrete crack detection models and adapt them to vehicle applications. As a result, the main effort has been spent on evaluating existing models and focusing on real-world data collection and performance evaluation. While there are many crack detection models reported in the literature, based on a recent survey of crack detection models for vehicular applications [10], the team has identified a popular model with pre-trained weights, DeepCrack [11], as the baseline. The LECSFormer [12] is selected for its good balance of performance and real-time potential [10]. As the pretrained weights are not available for this mode, the team focused on dataset preparation and on developing advanced image augmentation techniques to train and improve the model's performance for our vehicle-view crack inspection.

Dataset Preparation

Facing the data scarcity, especially from the vehicle forward-facing camera view for precast concrete bridges, the team has collected and prepared a dataset of road pavement with cracks from the vehicle forward-facing camera view. We carefully selected road crack data of different types of surfaces and different crack locations, road geometry, and extended the standard dataset. This is to enhance the vehicle's

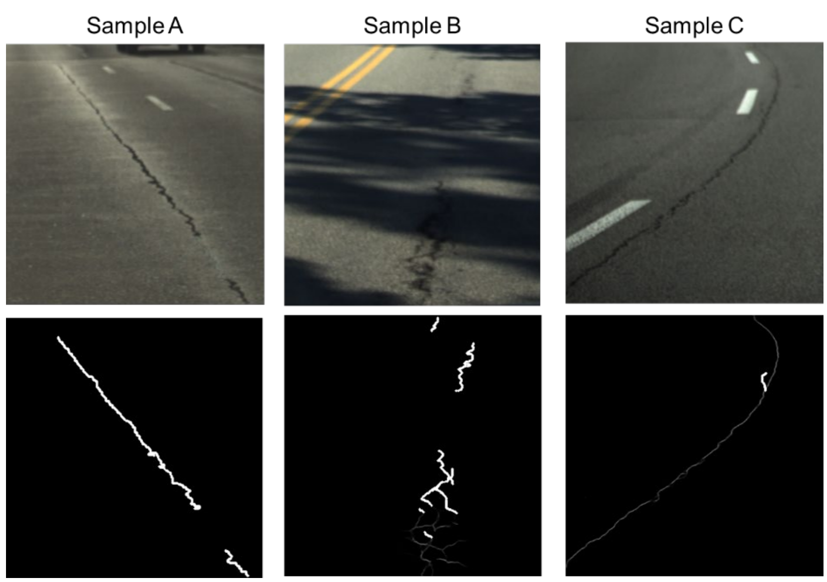

Figure 5 Sample of established road crack dataset: raw (top row) vs masks (bottom row)

forward-facing camera view. All images are collected with vehicle's forward-facing camera

In Figure , we present 3 samples in the collected dataset. The data collection for this project used the public CrackTree260 Dataset as a baseline and two phases of a Real-World Dashcam Dataset, which was collected using a vehicle's dashcam and presented real-world challenges such as motion blur and glare. While the initial 111 real-world images were fully hand-annotated, the larger dataset (332 images) was annotated using a more efficient hybrid semi-automatic workflow to combat the time-consuming and inconsistent nature of manual labeling (see the hybrid annotation shown in the sample masks). This process began with Blackhat filtering to automatically generate a preliminary crack mask, which was then manually refined by annotators in a custom interface to add missing cracks and remove false detections, ultimately yielding higher-quality, more consistent ground-truth masks. This dataset of 443 images (332+111) images we acquired consist of the final customized dataset we prepared for model training and testing.

Augmentation

We established a robust starting point for our work by adopting the **baseline augmentation** method, representing a standard technique often seen in the literature [10] [11] [12]. This approach is highly modular and repeatable, employing a fixed structure that generates 135 augmented variations per image for a total factor of 136. While the baseline method successfully introduces enhanced diversity by sequentially applying extended augmentations like brightness, contrast, blur, and noise, its core limitation remains its fixed-count nature. It provides rich variety in effects but cannot strategically adjust data supply to address the severity of class imbalance. To overcome this, we developed and proposed our method, which moves beyond a fixed count to an intelligent, adaptive approach. Our proposed **augmentation method** implements smart balancing by first categorizing images based on their crack density. Crucially, it replaces the fixed multiplier with a variable augmentation multiplier, targeting rare crack categories with a 2x boost. This ensures that the most challenging, underrepresented images generate 20 to 40 augmented samples, while more common types generate 10 to 30. This strategy, combined with adaptive cropping and elastic deformations, directly solves the imbalance issue, making our proposed method a significantly superior and more effective training methodology compared to the fixed-count baseline.

Performance metrics

Following the standard practice in the literature [10], precision (Pre) and recall (Rec) are used as metrics for the binary segmentation task, calculated from the model prediction results and manually labeled ground truth. F-measure is an overall metric for evaluating model performance, defined as follows:

$$\mathrm{Pre} = \frac{\mathrm{TP}}{\mathrm{TP}+\mathrm{FP}}, \quad \mathrm{Rec} = \frac{\mathrm{TP}}{\mathrm{TP}+\mathrm{FN}}, \quad \mathrm{F1} = \frac{2\times\mathrm{Pre}\times\mathrm{Rec}}{\mathrm{Pre}+\mathrm{Rec}}$$

We use F-measure-based metrics for model performance evaluation for crack detection. The crack detection model predicts crack probability maps and needs a threshold to produce results. We use two strategies to set the threshold, one is the Optimal Dataset Scale (ODS), which uses a fixed threshold to calculate the best F-measure for all test set. ODS can be defined as follows:

$$\mathrm{ODS} = \max\left\{\left\{\frac{1}{N}\sum_{k}^{N}\mathrm{F1}_T^k\right\} : \forall T \in \{0.01, \ldots, 0.99\}\right\}$$

Where N is the number of test set images and T denotes the confidence threshold. In the study, we primarily use ODS to evaluate the performance of the crack detection model. We remark that while other metrics, such as Optimal Image Scale (OIS) (which selects the optimal threshold for each image in the test set to calculate the F-measure) and Average Precision (AP) (which is the area under the precision-recall curve), are also used in the literature, we noticed similar trend in model evaluation results thus skipped them in this report.

Trained model

Here we present four different models for comparison. The DeepSeek model with pretrained weights is referred to as baseline one. This one is pretrained using a baseline augmentation method based on the CrackTree260 dataset, as reported in [11]. This is referred to as **baseline 1**. We trained the LECSFormer using the CrackTree260 dataset with baseline augmentation to get results similar to those reported in [12]. This is referred to as **baseline 2**. To show the benefit of proposed augmentation methods, we trained the LECSFormer using the same CrackTree260 dataset with the proposed augmentation methods, and this is referred to as **our model v1**. Lastly, we trained a LECSFormer model with proposed augmentation methods and further on a combination of the CrackTree260 dataset and our dataset. This is referred to as **our model v2.** All models' information are summarized in Table 1.

Results

The achieved performance, and we compared our models against the baselines over the public dataset and our separately collected dataset of 100 images from vehicle forward-facing camera view (named 100 Real-time dataset). These images are not included in the training set of any models.
The results are highlighted in Figure 6 and Table 1.
Specifically, on the widely used public dataset CRKWH100, our models achieve superior performance with augmentation, as evidenced by the highest ODS score, outperforming the baseline reported in the original paper. We demonstrate that the augmentation indeed significantly improves LECSFormer's performance by boosting the F1 score and reducing its sensitivity to threshold selection, as seen from the left panel in the Figure. Our models are able to perform even better and flatter F1 score curve.

Table 1: Model Comparison Summary

Model	Description	Trained Dataset	ODS (100 Real-time)	ODS (CRKWH100)
Baseline 1	Deep Crack Pretrained Model	N/A	0.3374	0.9089
Baseline 2	LECSFormer Trained with baseline data augmentation	CrackTree260	0.3411	0.9045
Our Model v1	LECSFormer Trained with proposed data augmentation	CrackTree260	0.3744	**0.9184**
Our Model v2	LECSFormer Trained with proposed data augmentation	CrackTree260 + Extend 343 dataset	**0.5267**	0.9019

On the other hand, all the baselines performed poorly on the 100 Real-time dataset, which consist of challenging surface crack images. With the introduction of proper dataset, our model v2 is able to increase the ODS score and the F1-curve performance significantly.

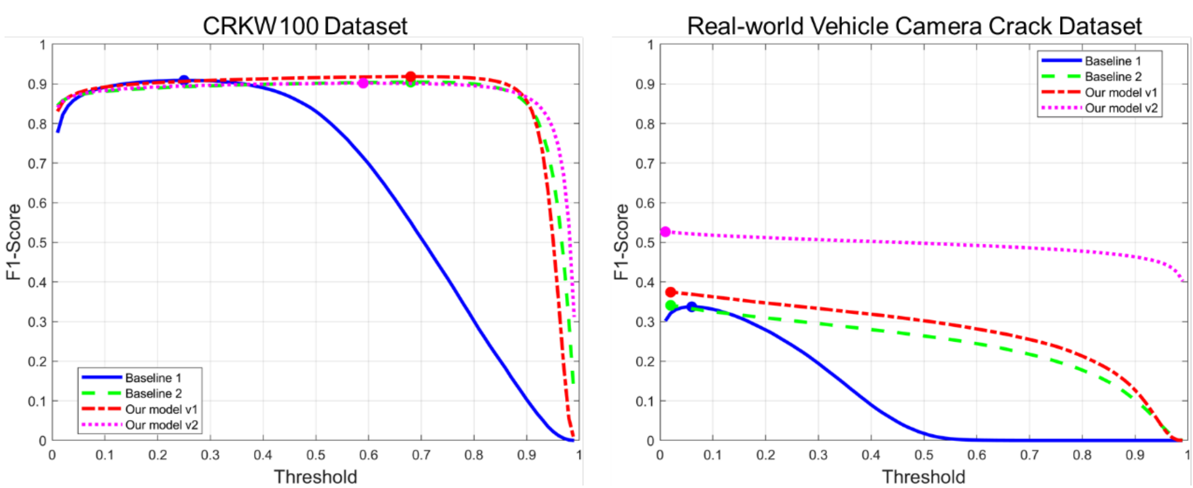

Figure 6 Model Comparison: F1-score

While showing encouraging results, we acknowledge that it is still very challenging to detect surface cracks using a vehicle's front-facing camera. Besides improving model structure as one possible future improvement direction, in this project, we also seek to use CV2X communication to guide the crack detection zoom-in windows, which will be detailed in the upcoming sections.

5.3. C-V2X based Inspections Protocol Design

In this section, we detail the communication protocol designed in the project to facilitate inspections of concrete surface cracks by CAVs. We will first outline the protocol and then present the performance evaluation.

Protocol design

The CV2X communication protocol centers on a Roadside Unit (RSU) - On-Board Unit (OBU) parallel scheme overlayed to the existing SAE J2735 standard [9] and is illustrated in Figure . Here, the OBU is integrated with the CAV, while the RSU resides within the infrastructure (i.e., precast concrete bridges). To enable inspection cooperation between the bridge and the vehicle, a parallel scheme based on Robotics Operation System (ROS2) is used, enabling safe and efficient concurrent execution of tasks. As a result, independent of the designed protocol, both RSU and OBU will maintain broadcasting of basic safety messages (BSM), while logging all received BSM [9].

The RSU is responsible for initiating the inspection process by sending requests to all the vehicles within communication ranges. Once the RSU receives a message from an OBU acknowledging that it's "ready" to detect a crack, it sends a message containing metadata about the area of interest containing the crack.

The OBU transmits a 'detection in progress' status to ensure the RSU remains in range, while, in a parallel stream, it collects the detection results and packages and releases them once the whole process is complete. This scheme allows more robust communication between the RSU and OBU in the event of out-of-range loss of package or other unexpected communication failures.

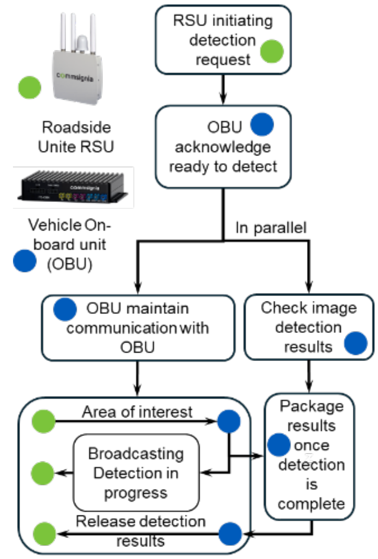

Figure 7 Updated CV2X Communication protocol

Once the communication is finished, the OBU sends images as WSMP (SAE standard communication unit) chunks, followed by a flag indicating completion.

Subsequently, RSU receives image data (PNG/TXT chunks), reconstructs them, and saves all the necessary detection information.

Protocol Test Results

We have conducted investigations on communication protocol design to evaluate the feasibility of communicating inspection requests and results between infrastructure (i.e., bridges) and vehicles. The performance of a set of standalone tests is summarized in Figure #. In the test, a set of 7 crack detection results, each of a different size, is packaged and transmitted between the OBU and the RSU. Each image is sent 10 times consecutively. The bidirectional test was conducted to investigate heterogeneity in communication performance between RSU and OBU, as RSU is, in general, more powerful in computation and processing. Time profiling is done to differentiate the encoding (chunking using base64), decoding, processing, and transmission processes. The images are indexed in ascending order according to size. The histogram shows time profiling of the first six sets of images. The majority of time is spent sending messages, while the other parts, including logging (which includes the chunk time) and the initialization phase, are negligible. While it is not surprising that the time depends on the image size, since the number of chunks is proportional to it, it is interesting that both directions (RSU to OBU vs OBU to RSU) take roughly the same time. It is important to know that even images of ~24 kB, which represent a large crack detection result of pure masks (e.g. Figure bottom row), takes roughly 500 ms to transmit in one direction, despite that 100 % success rate is experienced.

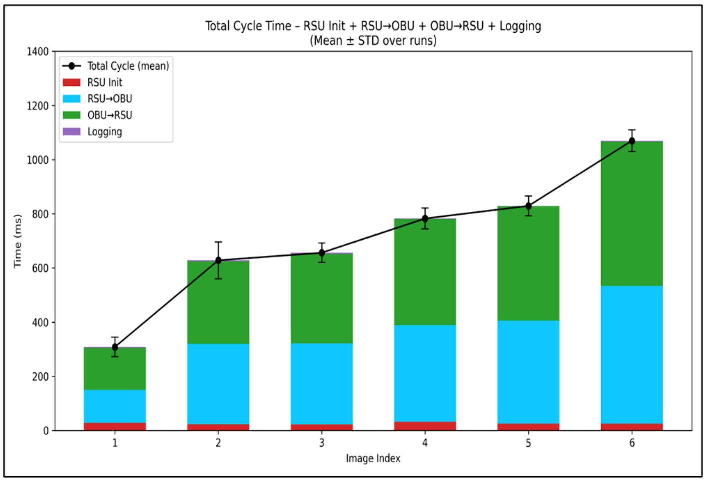

Image #	Stats	Total Cycle Time	Size (IN kB)	Chunk size	Success rate
1	MEAN	308.4	5.2	7	100%
	STD	36.4			
2	MEAN	682.2	11.9	16	100%
	STD	68.1			
3	MEAN	656.1	14	19	100%
	STD	35.8			
4	MEAN	782.6	16.5	23	100%
	STD	38.7			
5	MEAN	829.9	18.2	25	100%
	STD	36.4			
6	MEAN	1069.8	23.8	32	100%
	STD	39.7			
7	MEAN	13807.5	329	439	85%
	STD	485.4			

Figure 8 Communication Protocol Testing Statistics

The 7th set of images corresponds to raw images of the crack area of interest taken by the vehicle's forward-facing camera. According to the statistics shown in the table on the right, the success rate has dropped significantly due to an increase in the number of packages per image. More importantly, the transmission time increases to 5 seconds. While these results meet expectations, they imply that the full image detection results (4 MB / image) would not be feasible.

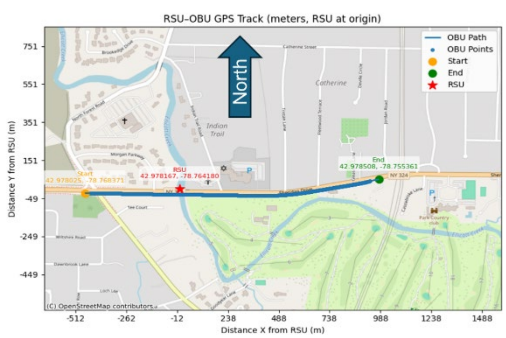

Figure 9 Field Test Communication

The team is also validating the communication protocol. The major focus is on identifying a feasible transmission time window and assessing the communication quality. In Figure , one test run is demonstrated where the vehicle was traveling eastbound with OBU and was expected to inspect the bridge surface. The blue points on the blue curve correspond to instances where basic safety messages (BSMs) were recorded. It started communicating with the RSU around 500 meters before passing the bridge. After passing the bridge, communication between the RSU and the OBU remains possible up to 1km. This means that the transmission of a raw, cropped window image (300 kb or ~6 seconds) is feasible. Nonetheless, this still underscores the risk of transmitting raw detection results from the vehicle to the infrastructure. Either cropping is needed to show the focus area (which may be larger than the area of interest), or other communication techniques (e.g., cellular network).

As increasing the communication bandwidth is beyond the scope of this year’s project, the team has decided to preserve raw images locally and only report detection measurements and cropped images from the vehicle to the infrastructure.

5.4. Cooperative Vehicle Control

The research team has implemented the CV2X-based vehicle motion control on the Lincoln connected automated vehicle (CAV) Platform at UB [13, 14]. The team successfully implemented and tested

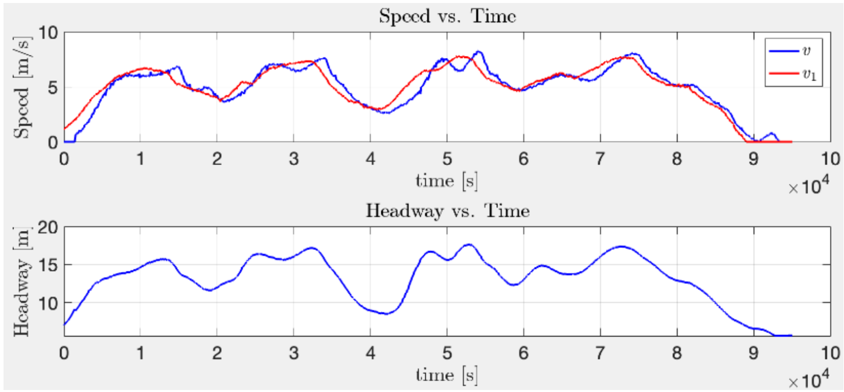

Figure 10 CV2X-based motion control performance: longitudinal speed tracking

longitudinal control algorithms on CAVs, enabling them to regulate their speed in accordance with the requested set speed while maintaining safety in traffic. The time profiles corresponding to an example run are illustrated in Figure 10. As shown in the figure, the CAV (blue) can maintain the same speed as the vehicle ahead (red) while staying a safe distance (Headway). A demonstration video can be found **here**. Part of the longitudinal control is included in the accepted paper [14].

5.5. Detection Pipeline

Architecture

The team has developed a streamlined detection pipeline based on ROS2. In Figure , the overall pipeline architecture and its corresponding data flow have been presented. There are 4 carefully designed ROS2 nodes supporting the real-time detection process: the OBU node, the Vehicle node, the Camera node, and the Detection node. At the beginning of each test, all 4 nodes, along with the required hardware driver packages, are launched together using a single launch command. The operation logic of the pipeline is outlined below:

OBU Node: The OBU node first sends the target crack location to the vehicle node; meanwhile, it recursively publishes the vehicle information read from the CV2X device to the vehicle node.

Vehicle Node (Core coordination): The vehicle node subscribes to the target crack location and to multiple resources regarding vehicle status, including real-time speed, GPS location, heading angle, and satellite status from the OBU node, RTK GPS, and the Vehicle CAN Bus. In real time, the vehicle node also publishes two indicators: the cropping indicator and the process indicator, which are binary signals indicating when image recording and subsequent processing should start.

When the node is first initiated, both indicators will be set to 0. Once the vehicle node receives the crack location, it continuously computes the real-time distance between the CAV and the target crack to determine when cropping should begin, ensuring the best visual resolution. Once the CAV reaches the appropriate range, the vehicle node generates the ideal cropping boundary and switches the cropping indicator to 1 to tell the camera node to log raw image asynchronously. After the CAV passes the crack location, the cropping indicator will be switched back to 0, and the processing indicator will be switched from 0 to 1.

Camera Node (Asynchronously logging raw images): The camera node subscribes to the cropping indicator. When it detects the cropping indicator changes from 0 to 1, it will start logging the raw images asynchronously with the cropping info generated from the vehicle node.

Detection Node (Synchronize cropping, inference and compute crack length): The detection node subscribes to the processing indicator. When the process indicator changes to 1, the detection node reads the image logs that contain the raw image and its generation timestamps, and the cropping logs that contain the cropping boundary coordinates in the image pixel frame and its corresponding timestamp. For each logged raw image, the node syncs the corresponding cropping boundary, crops the raw image to a 512 by 512 window size, and further feeds the cropped window into the inference step to generate the mask image and eventually compute the crack length.

CV2X to RSU transmission: To close the detection loop, the OBU node continuously checks the directory where mask results are saved. Whenever new results are generated, it will package the latest generated mask image and computed crack length and send it to the RSU to close the loop.

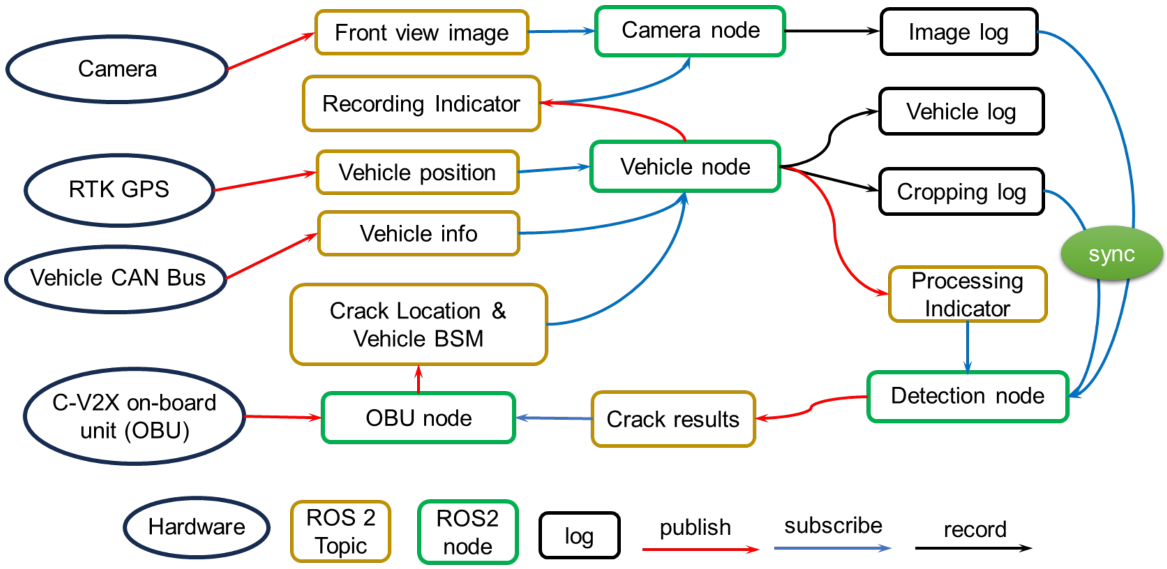

Figure 11 Detection Pipeline Schematics

Dynamic Cropping Window

The team has developed a dynamic cropping window determination intended to enhance crack detection performance. The area of interest transmitted by the Bridges will be mapped to a cropping window within the camera's field of view. While such translation may be straightforward using a pinhole camera model, the challenges reside in the uncertainties related to camera extrinsic parameters (e.g., yaw, pitch, and roll angles), the vehicle's GPS measurement accuracy, vehicle motion, and other factors. The goal is to acquire cropping windows that 1) include the target crack in the cropping window with high confidence and 2) acquire images where the cropping windows appear with high resolution of pixels for distance translation. In this quarter, the team has implemented these algorithms and achieved satisfactory performance.

$$\text{Optimal parameter} = \underset{\text{Allsamples}}{\arg\max} \sum \text{DIOU(YOLO Box, Dynamic Cropping window)}$$

Statistics	Mean	STD
DIOU score for overlapping	0.86	0.10

Dynamic Cropping window

Yolo bounding box on checkerboard

Figure 12 Dynamic Cropping Windows

To obtain the ground truth of the area of interest, a checkerboard is used to mark the potential area. The GPS location could be acquired using a portable RTK GPS antenna. A YOLOv8 model is trained to detect a checkboard on the ground and mark a bounding box. Distance-IOU [15] is used to evaluate the performance of the cropping window. Furthermore, it is also used to optimize camera extrinsic parameter calibrations, allowing for robust run-time performance of the dynamic cropping window algorithms. In Figure the performance metrics with samples are presented. The reported DIOU score ideally = 1, which indicates that the target area of interest (center of checkboard) should coincide with that of the cropping window.

Crack Size Estimation

The team has also deployed a crack size estimation based on crack detection results. The estimation is based on a standard pinhole camera model, which is illustrated using a simulated view in Figure . The simulated crack is represented by a black line in the area of interest specified (purple square on the road). With proper cropping window logic, the area of interest will be covered, and the crack size can be correctly inferred. When implemented on the real vehicle, it requires proper camera calibration. We leverage the calibration process detailed in the above cropping window logic using the checkboard to fine-tune the camera extrinsic parameters. However, real cracks are not perfectly aligned or uniformly shaped like a checkerboard. To more accurately recover the crack size, we further developed an edge corner detection algorithm. The algorithm uses the mask image as input, extracts every highlighted pixel covered in the crack region, and reconstructs every highlighted pixel from the image frame to the vehicle frame. It then divides all reconstructed points into four quadrants and uses the quadrant information, along with the correlation matrix, to estimate edge locations by locating the pixel with the largest Euclidean distance from the center point in each quadrant. We are able to achieve satisfactory diagonal results when lighting conditions are favorable, and the crack is close to the camera.

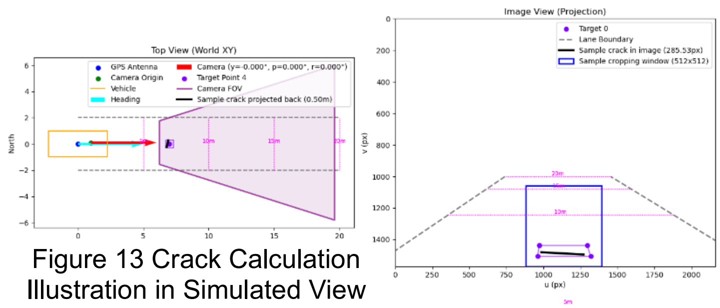

Figure 13 Crack Calculation Illustration in Simulated View

Real-time Crack Detection Performance

With all the techniques developed (modeling, communication and the image preprocessing + crack calculation, we integrate them all and implement them on the CAV platform, as summarized in Figure 14. Acknowledging the limitation in communication bandwidth, the detected results

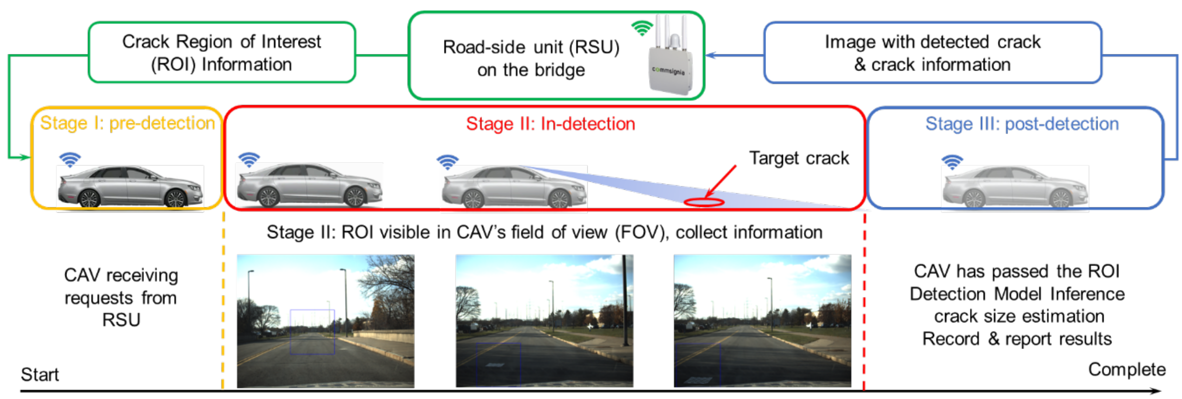

Figure 14 Overall Real-time Pipeline

The team first compares the reconstructed checkerboard diagonal length with the measured ground truth. This step aims to justify the robustness of our camera calibration and reconstruction accuracy under the assumption that we have the perfect detected mask, and what is the difference between the diagonal length computed from our own algorithm and the ground truth from direct measurement.

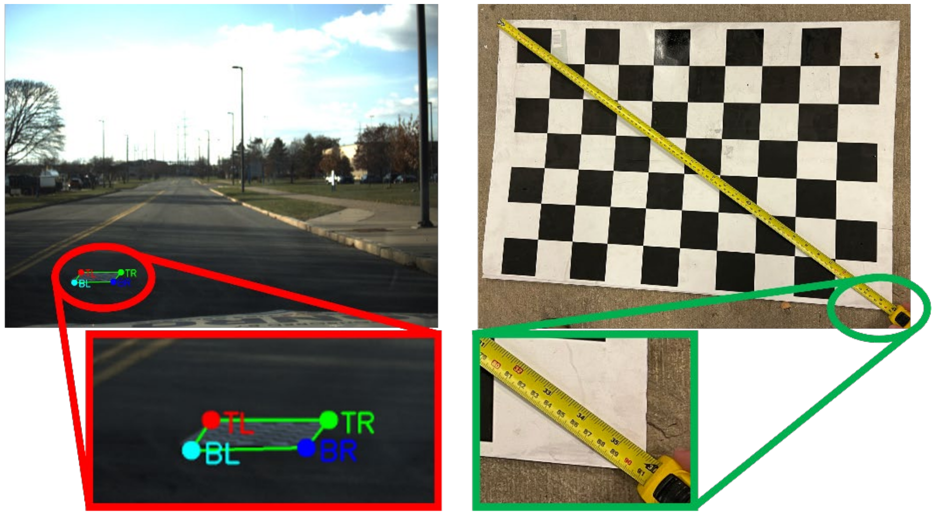

Figure 15 Crack length algorithm applied to checkerboard with a known size

Diagonal Edge	TL-BR(from detection)	TR-BL(from detection)	Ground Truth
Length [cm]	90.8	89.6	91.0 ± 0.1

Table 2 Comparison between computed checkerboard length and the ground truth

Figure 16 Test on UB campus

The team then used the calibrated camera parameters, conducted a field test to test the performance of the dynamic cropping. In the field test, a designated crack was selected on the service center road bridge at the University at Buffalo North Campus; the detailed location is shown in Figure 16. A video showcasing the current dynamic cropping status can be found in the video link.

We acknowledge that the limitations of the crack size, light condition, and the reflection on the windshield may affect the inference results. To better simulate the crack detection on a public bridge, we selected a longer and thicker crack on the service center road to evaluate our model inference performance. Based on the accuracy shown in Table 2, we used the diagonal length computed from a hand annotation mask as the ground truth and further compared all the results from our trained models so far. To better visualize the detected crack corners, we labeled edge points on the raw image from both the ground truth and our trained models. A detailed comparison of the computed diagonal length is provided in Table 3.

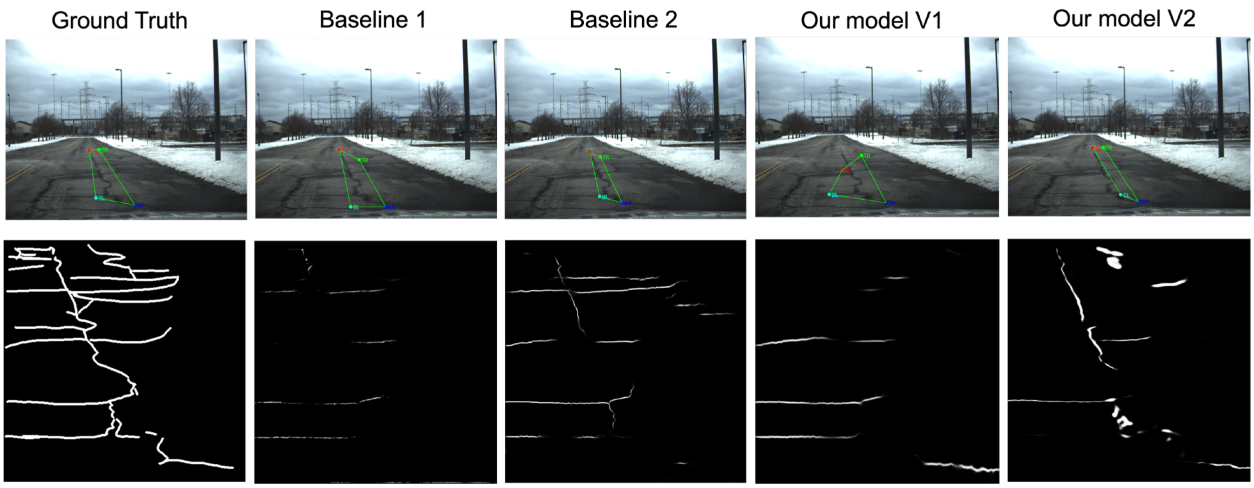

Figure 17 Model inference results

Models	Diagonal (TL-BR)	Unit
Ground Truth	2021.7	cm
Baseline1	1982.7	cm
Baseline2	1605.8	cm
Our model V1	561.4	cm
Our model V2	1894.4	cm

Table 3 Crack length estimation comparison

We observed that the LECSFormer models perform better at visualizing the vertical edge of the crack mask; however, they also miss more details along the horizontal crack edges compared with the DeepCrack model. Based on this observation, we further trained LECSFormer models on more comprehensive datasets and achieved closer alignment with the ground truth in model V2.

6. Conclusions and recommendations

In this project, we established the foundation of a continuous and low-cost precast bridge inspection system using connected automated vehicles. We established a prototype pipeline that consist of three pillars. 1) a communication protocol based on vehicle-to-infrastructure (V2I) communication that enables inspection requests and results transmission; 2) cooperative vehicle motion that allows optimal detection motion by the vehicle, and 3) a crack detection algorithm using a vehicle front-facing camera assisted with V2I communication and vehicle motion.
We thoroughly presented the results, setting up baselines. We can conclude that such targeted inspection systems can be potentially achieved by properly designed algorithms on CAV and the infrastructure. The major recommendations for future works include

a. The need to extend a proper training dataset for the detection model is crucial.
b. The communication bandwidth with existing CV2X communication protocols is the bottleneck for real-time sharing of inspection results.

We are excited to continue this development to achieve the full target inspection systems.

7. Practical application/impact on transportation infrastructure:

This work builds on increasing levels of autonomy and connectivity and aims to involve road users in maintaining the longevity of transportation infrastructure. It has the potential to unlock a large number of data probes with high-quality data. The methods and techniques described in this can be used to broaden cooperative sensing applications that involve vehicles and connectivity and apply to broad transportation infrastructures inspection and early hazard detection. This creates a huge economic impact by reducing inspection and maintenance costs, and allows more just-in-time, efficient maintenance and repair activities.

8. References:

[1] U.S. Department of Transportation, "USDOT Releases National Deployment Plan for Vehicle-to-Everything (V2X) Technologies to Reduce Death and Serious Injuries on America's Roadways," ed, 2024.

[2] "New York State Highway Bridge Data." https://www.dot.ny.gov/main/bridgedata (accessed 08/31/2024.

[3] "The Evolution of Drones Part 3: Building Highways in the Sky," American Association of State Highway and Transportation Officials, 2019. [Online]. Available: https://archive.transportation.org/the-evolution-of-drones-part-3/

[4] NEXCO - WEST USA. "Deck Top Scanning System (DTSS)." https://www.w-nexco-usa.com/assets/documents/DTSS.pdf (accessed 08/31/2024.

[5] Alynix. "Decker: Take the guesswork out of bridge deck assessments." https://alynix.com/decker/ (accessed 08/31/2024.

[6] Mercedes-Benz. "Trafikverket uses vehicle data to identify road surface conditions in Sweden." https://developer.mercedes-benz.com/success-stories/trafikverket (accessed 08/31/2024, 2024).

[7] Nexteer Automotive. "Road Surface Detection." https://www.nexteer.com/software/road-surface-detection/ (accessed 08/31/2024, 2024).

[8] A. Alan, C. R. He, T. G. Molnar, J. C. Mathew, A. H. Bell, and G. Orosz, "Integrating Safety With Performance in Connected Automated Truck Control: Experimental Validation," *IEEE Transactions on Intelligent Vehicles,* vol. 9, no. 1, pp. 3075-3088, 2024, doi: 10.1109/tiv.2023.3305204.

[9] *Requirements for Probe Data Collection Applications*, J2945/C_202206, S. I. T. Standard, 2022.

[10] N. Ma *et al.*, "Vehicular Road Crack Detection with Deep Learning: A New Online Benchmark for Comprehensive Evaluation of Existing Algorithms," *arXiv preprint arXiv:2503.18082,* 2025.

[11] Q. Zou, Z. Zhang, Q. Li, X. Qi, Q. Wang, and S. Wang, "DeepCrack: Learning Hierarchical Convolutional Features for Crack Detection," *IEEE Trans Image Process,* Oct 31 2018, doi: 10.1109/TIP.2018.2878966.

[12] J. Chen, N. Zhao, R. Zhang, L. Chen, K. Huang, and Z. Qiu, "Refined Crack Detection via LECSFormer for Autonomous Road Inspection Vehicles," *IEEE Transactions on Intelligent Vehicles,* vol. 8, no. 3, pp. 2049-2061, 2023, doi: 10.1109/tiv.2022.3204583.

[13] S. Beregi, S. S. Avedisov, C. R. He, D. Takacs, and G. Orosz, "Connectivity-Based Delay-Tolerant Control of Automated Vehicles: Theory and Experiments," *IEEE Transactions on Intelligent Vehicles,* vol. 8, no. 1, pp. 275-289, 2023, doi: 10.1109/tiv.2021.3131957.

[14] H. Xiao and C. R. He, "Safe and Efficient Data-driven Connected Cruise Control," presented at the The 5th Modeling, Estimation and Control Conference (MECC 2025), Pittsburgh, PA, 2025.

[15] Z. Zheng, P. Wang, W. Liu, J. Li, R. Ye, and D. Ren, "Distance-IoU Loss: Faster and Better Learning for Bounding Box Regression," *Proceedings of the AAAI Conference on Artificial Intelligence,* vol. 34, no. 07, pp. 12993-13000, 2020, doi: 10.1609/aaai.v34i07.6999.

9. Acknowledgements

We would like to acknowledge tremendous support and valuable inputs from the UB team at the Civil Engineering department, Prof. Pinar Okumus, Prof. Ravi Ranade; at the Computer Science and Engineering department, Prof. Qiao. We really appreciate the technical support from the NYS DOT, especially Mr. Gary Jakubowski, and from the Commsignia tech support team. Finally, we would like to express our gratitude to the USDOT TRANS-IPIC center leadership team, particularly Prof. Bassem Andrawes, for their continued unwavering support.