AUTOMATED INFRASTRUCTURE DAMAGE DETECTION

BOLLAM RISHITH[1], PODISHETTY DEV MANIKANTA[2], RAGAM VAMSHI[3], MUNGE CHAKRADHAR REDDY[4], MRS. A. SOWJANYA[5]

[1]Department of CSE-AIML, Sreyas Institute of engineering and technology, Hyderabad, India
[2]Department of CSE-AIML, Sreyas Institute of engineering and technology, Hyderabad, India
[3]Department of CSE-AIML, Sreyas Institute of engineering and technology, Hyderabad, India
[4]Department of CSE-AIML, Sreyas Institute of engineering and technology, Hyderabad, India
[5]Assistant Professor, Department of CSE-AIML, Sreyas Institute of engineering and technology, Hyderabad, India

Abstract

The traditional way of maintaining our roads and other infrastructure is based on manual assessments or inspections which have been subjective, error-prone (due to human factors), and require considerable effort. This proposal will showcase CrackAlert as an automated system that will revolutionise the monitoring of infrastructure using computer vision and deep learning technology. Through a high performance software pipeline, CrackAlert converts the raw images of roads into engineering data with real-time identification of structural failures (eg. longitudinal cracking, potholes, alligator cracking).

The implementation of the CrackAlert system is focused on a very strong web-based portal that connects the complex artificial intelligence (AI) models with practical field applications. Each of the components of the software utilize a dedicated HTML5 based canvas rendering engine, with pixel-wise semantic segmentation used to precisely identify damage. This visual data is combined in real-time and geolocalised (using the Web Geolocation API) with GPS data, so that each identified defect can be mapped and classified by severity immediately. The CrackAlert platform is hardware agnostic and scalable, allowing civil engineers to 'shift from reactive maintenance to data driven proactive maintenance'. CrackAlert provides a low cost method for improving urban safety and preserving modern transportation networks.

Keywords: Automated Infrastructure Monitoring, Semantic Segmentation, Computer Vision, Geospatial Data Fusion, Proactive Road Maintenance

1. Introduction

The maintenance of civil infrastructure, especially roads is vital for both the economy and public safety. The first sign of failure in the road surface is usually a crack, either a thin straight line or a curving black line on a hard material surface. The cracks develop due to various environmental and mechanical actions such as soil movement, movement of the sub-base, heavy traffic loads, expansion and contraction of the surface material, etc. If these minor surface irregularities are not repaired, they can become a major structural failure. For example, in the UK in 2020, cracks in the road accounted for 12.6% of traffic accidents. For highway networks with greater than 100,000 km, it's costly; time consuming, and obviously unsafe for inspectors to visually inspect roadways using traditional manual methods. Therefore, this project proposes utilizing a vehicle-mounted imaging system to perform computer-vision based automated damage detection on road surfaces by capturing images of the road surface, detecting, and classifying road cracks and potholes based on the severity of the cracks/potholes. This research also supports the UN Sustainable Development Goal 9 (Industry, Innovation and Infrastructure) by providing for smarter maintenance practices, increasing safety, and minimizing costs long-term.

The evolution of crack detection on roadways has transitioned from one's ability to perform a simple assessment of the roadway condition (two-dimensional) all the way through to utilizing pixel-level classification (semantic segmentation) to identify the exact location of a crack in the roadway. The last dataset to be analyzed was the UDTIRI-Crack dataset, which included 2500 images (high quality). In addition to visual assessment, algorithm performance was examined against a number of different methods. The CNN-based implementation (e.g., Crack-Att), despite having reached a precision of 68.790% and an IoU of 51.055% at 50 FPS, have not matched the performance levels of more advanced Transformer-based methods (e.g., CT-crackseg) which had a precision of 75.019% and an IoU of 54.573%, though it only performs at 10 FPS. Models like Deepcrack19, which operate with low power consumption, have been recorded producing 276.695 FPS and achieving 72.831% precision. Overall, there are multiple metrics available from the evaluations made, and all of the metrics indicate high fidelity and accuracy of crack detection is possible. However, it is critical to determine the appropriate compromise between these characteristics (fidelity, accuracy, model size, and ability to perform in real-time).The motivation behind this research lies not only in the need to overcome the limitations of current supervised learning techniques for crack segmentation, which have generally been unable to generalize well across different types of roads but also in the need to find better methods for generating data to help paint a clearer picture of how we can detect road cracks in the future. For example, although the best-performing CT crack model achieved a mean segmentation score of 28.631% for test images from the CrackNJ156 dataset, there was a large difference between these metrics compared to the training metrics; thus suggesting that while current models can perform accurately on a standalone, individually collected sample, they are not generalizing to other sites or periods of collection.

To improve upon the challenges associated with labelling bottlenecks and generalizing across multiple types of roads, we will explore hybrid CNN-transformer architectures; in particular, we feel that using hybrid CNN-transformer networks will allow us to incorporate local feature maps derived from road images along with features captured globally at all the same time. We also plan to explore how LLMs (e.g., Grounded SAM) and/or Foundation Models can be adapted to more effectively complete actual real-world tasks for the purposes of increasing the accuracy of detecting road cracks, though the models will require fine-tuning or training, which may not

International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 04 | April - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

translate into an accurate modelling approach.In addition to the use of cutting-edge computer vision techniques (transformer-based models), our study will also utilize advanced image generation methods (e.g., GANs and Diffusion) to generate synthetic datasets. By creating our own synthetic datasets instead of relying on existing datasets containing millions of labelled samples, we can establish a clear and accurate database of labelled samples that can be used as a valuable asset in creating future smart cities.

Intelligent Road Inspection (IRI) vehicles have become necessary to maintain digital twins of cities and contemporary civil engineering infrastructure. Methods based on deep learning have been developed to replace laborious manual visual inspections over the past 10 years. However, current algorithms encounter difficulties balancing accuracy of detection and real-time processing. Research from the past six months has highlighted a need for labelled-efficient and data fusion techniques to develop usable combinations of 3D geometry and RGB image data. Adapting foundation models (e.g., Segment Anything (SAM)) to detect cracks in roads represents a significant challenge for next-generation automated maintenance systems.

2. Proposed

2.1 System Overview

The CrackAlert project is a complete automated pipeline consisting of the entire process of using advanced computer vision to manage practical infrastructure. The overall operations begin with the high-resolution imagery of the road surfaces being ingested into the specialized software engine for a series of processes that will use pixel-level semantic segmentation to identify presence of structural anomalies (potholes, and other types of critical crack patterns). The system performs asynchronous processing so that simulated deep learning inference can provide immediate visual confirmations via a rendering canvas that illustrate the damage versus the normal pavement condition. This intelligence works seamlessly with the geospatial element of the system by automatically pulling in GPS coordinates concurrently, allowing each detection to be anchored to a specific location, results in the transformation of the raw visual evidence to a structured maintenance record that is specific to the geographic coordinates.

To ensure the efficiency of this pipeline, the project must adhere to a clearly defined set of functional and technical requirements for reliability and scalability. The system must be able to support the high-speed image captures and pre-treating of images under variable light/road surface textures. The AI module must also achieve very high accuracy in classifying the images and produce low latency responses to ensure the CrackAlert Portal is operationally responsive to the user. In besides to the above requirements, the system must have GPS geolocation accuracy of several meters in order to accurately map the asset, therefore the software architecture has been created, with our current technology, to fulfil this requirement.

Volume: 10 Issue: 04 | April - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

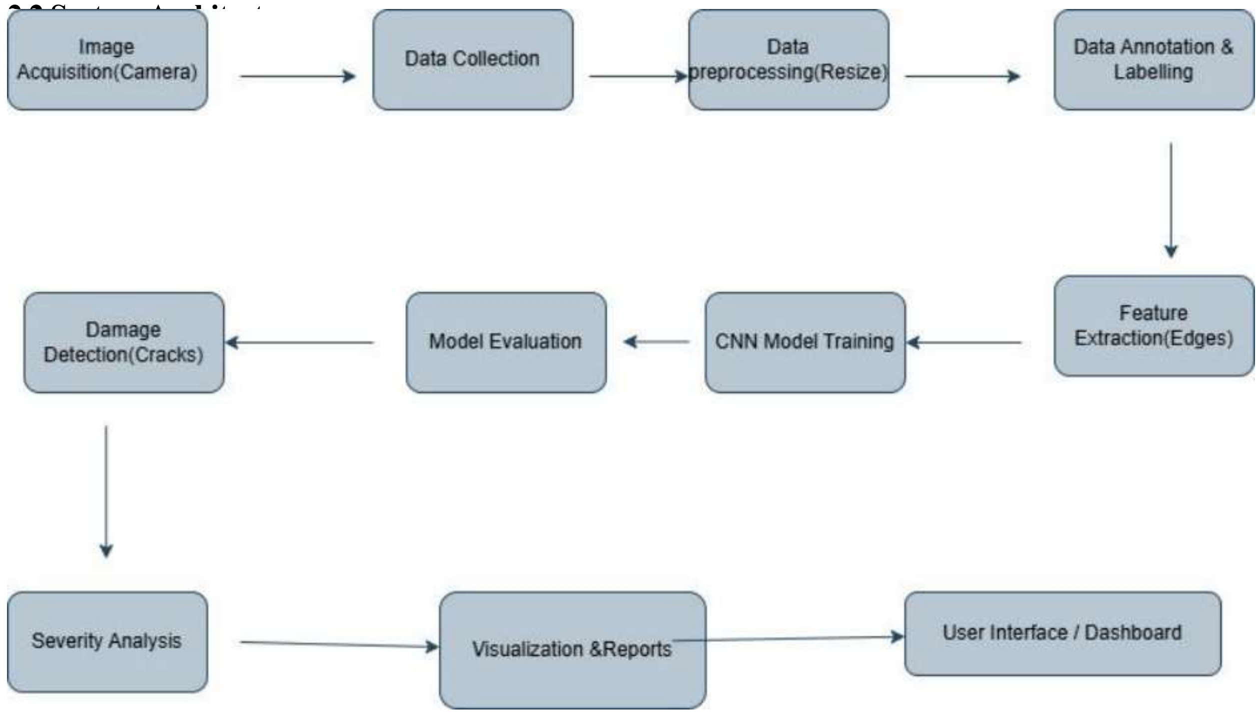

Figure 1: System Architecture for the Automated Infrastructure Damage Detection, illustrating the end-to-end data pipeline from initial user input till the Report analysis.

2.2.1Data Collection

The CrackAlert architecture consists of four distinct tiers designed to systematically change raw environmental data into higher-level intelligence for maintenance purposes. The use of a modular structure enables maximum computational efficiency and future readiness of the system for physical deployment on inspection vehicles that operate autonomously.

2.2.2 Data Preprocessing

The CrackAlert project's data processing pipeline converts raw road surface images into actionable structural intelligence through a streamlined computational process that unfolds in three primary steps. First, raw imagery enters the system and is normalized (e.g., converted to grayscale, denoised, histogram equalized) to enhance the visibility of all visual characteristics in any environmental lighting condition. Second, a deep learning inference engine uses the output from the first step to perform pixel-based semantic segmentation of the road surface to identify the type of road distress (such as potholes or longitudinal cracks) and calculate a probability for each detection. Third, this visual intelligence is combined with real-time geospatial metadata from the Web Geolocation API, anchoring each identified defect to a precise geographical location, and assigning an overall severity score to each defect based on the calculated damage density, thereby generating a complete digital maintenance report.

2.2.3 Categorical Label Mapping

This project employs a discrete numbering system for its categorical label mapping, which allows for the conversion of qualitative road conditions to quantitative data for the AI model's training and evaluation. Each distress type identified as part of the study has been assigned an integer value. Label 0 is typically used for the Healthy Road or background surface, while the other integer values are used to classify structural defects: Label 1 will represent Longitudinal Cracks, Label 2 will represent Transverse Cracks and Label 3 will represent Potholes. The use of this categorical label mapping allows the semantic segmentation engine to carry out multi-class classification so that it can differentiate between minor surface wear and major structural failures. Thus, by utilizing these standard labels, the system can aggregate data uniformly across various geographic areas, allowing for accurate statistical reporting and providing engineers with a means of prioritizing road segments based on the classification and severity of the damage found within them.

2.2.4 Model Training & Optimization

A supervised deep learning model is employed, applying transfer learning techniques (from Backbone models such as ResNet) to accelerate the detection of complex patterns on a road surface. A hybrid loss function combining the Cross Entropy and Dice Loss is used to address the class imbalance in which cracks occupy a small percentage of the total number of pixels in an image, along with some of the data augmentation techniques (e.g., random rotation, brightness adjustments) to maintain model robustness against variations in field condition. After the training of the model is subjected to 8 bit quantization so that it has a have smaller size and can execute in the real-time within a web browser, and lastly evaluated with the Intersection over Union (IoU) metric to validate the pixel precision and accuracy for the assessing of critical infrastructure.

2.2.5 The Data Acquisition Tier (Physically/ Input Layer)

The foundation of the system consists of capturing visual and spatial data at the same time. The architecture is utilizing a synchronized input model so that high-quality imagery is acquired of the road surface and paired with relevant and precise temporal and geospatial metadata (e.g. capture date, time and geo-location).

Imaging Unit – Captures internationally recognized patterns of pavement images in frame sequences for identifying high-contrast features extracted.

Geospatial Unit – Utilizing a Web Geolocation API (or GPS hardware), provides real-time latitude and longitude for each processed frame.

2.2.6 Edge Processing & AI Functionality Tier

The edge processing tier functions as a tool for processes (as where the analysis occurs – at the "client"/edge) of interpreting raw images. By processing the images at this level, responsiveness can occur in real-time within the portal.

Pre-Processing Module: In this module, processes occur at the image level to remove noise, normalize the images to grayscale & graphic equalize the purposes of light variability & shadows (throughout an image).

Semantic Segmentation Engine: Through a deep learning (training by backpropagation) template graphic-based model (created from images of the overall) will identify through per pixel classification of images between the background (image with roadway) & per image for specific crack/distress types that are longitudinal or potholes using a per-image/pixel based many-to-many probability map.

2.2.7 Analytics and Feature Fusion Tier

The third layer is referred to as "Analytics and Feature Fusion Tier," which is also known as the Transformation Layer. This layer transforms qualitative, visual detections into quantitative engineering metrics once an AI engine identifies a defect.

Damage Density Assessment: Calculates the damage density of a defect by determining the ratio of defect-masked pixels to the overall surface area. This will allow the determination of how severe each defect is (low, medium, high).

Metadata Integration: When the AI detects a label (defect), the AI performs the "feature fusion" by anchoring each AI-detected label to an actual GPS coordinate and time stamp. This creates a unique localized digital record of all road hazard defects.

2.2.8 Application and Reporting Tier

The last tier is in charge of taking the processed structural health data and visualizing it for end-users and municipal organizations or municipalities.

Geospatial Dashboard - A web-based geographic information system will be used to display a heatmap of road defects real-time over a jurisdiction.

Maintenance Prioritization System - Structured report generation occurring automatically to rank repairs by urgency and density of damage, facilitating data-driven budget allocation and proactive urban planning.

2.3 Methodology

The CrackAlert projects methodology follows a clearly defined engineering workflow, i.e., starting with raw data capture through to a final structural condition report. Assessment of defects must be substantiated with quantitative data and accurately geocoded.

2.3.1. Data Collection & Environmental Synchronization.

The first step in this method is to systematically collect high-resolution imagery of road surfaces, while simultaneously capturing real-time metadata. Using vehicle mounted sensors and/or high fidelity mobile cameras the system captures sequential images of the road surface. Simultaneously the system interfaces with the Web Geolocation API to synchronize geospatial coordinates (latitude/longitude) for each image as it is ingested into the system thereby geospatially anchoring visual data upon ingestion.

3.2.2 Preprocessing and Feature Enhancement of Images

The AI model must work correctly regardless of the various conditions it will encounter in the field; as such, the images that are processed by the model must go through a number of standardization steps. These steps include converting the image from color to grayscale in order to reduce the amount of processing power required and applying Gaussian filters to remove noise from the sensors used to capture the images. Additionally, images will undergo histogram equalization to improve the contrast between the background asphalt surface and the different types of structural distresses to neutralize the impact of shadows and variable sunlight on the model's ability to identify and produce accurate results.

3.2.3 Image Segmentation by Classifying with Defect Identification

The key component of the methodology is that a deep learning model (Hybrid CNN-Transformer Architecture) has been created to classify images by analyzing each pixel separately for the purpose of segmentation. This approach differs from the traditional object detection method, which typically classifies and locates objects in images using bounding boxes. Because segmentation is performed at the pixel level, individual types of structural distresses, including potholes and longitudinal cracks, can be clearly identified, separated from non-structural items such as tire marks.

3.2.4 Image Feature Fusion with Geospatial Tagging

At this stage of the methodology, the AI-generated visual intelligence is combined (fused) with the original geospatial metadata that was captured when the image was originally collected. Each identified defect “mask” is mathematically linked to the associated GPS coordinates. This information is used to convert a digital image into a spatially referenced data point.

3.2.5 Quantitative Damage Imitation and Severity Scoring

After the system identifies the potential distress within the framework, it then performs the quantitative analysis of those identified distresses. To accomplish this, the system calculates a ratio of pixels that have been damaged to total area of the frame. From this figure, the system calculates a value of “Damage Density” (index). Subsequently, each identified defect will then be automatically classified based upon engineering-defined thresholds for a given defect condition as Low (Minor Surface Wear), Medium (Major Cracking), or High (Major Hazard Potholes – Require Immediate Response).

3.2.6 Integration into the Database and GIS Reporting

The final phase of the project involves the compilation of all processed data into an organized report to organized by the local agency for the purpose of making decisions. The processed data is sent to the database for storage and made available via the visualization of the data in a Geographical Information System (GIS). GIS will provide the local municipalities with an up-to-date, real-time infrastructure health heat map that has the potential to assist the municipalities to successfully transition their approach to a proactive, data-driven maintenance approach focused on prioritizing the most critical infrastructure roadway failures before they fail.

2.4 System Analysis & Performance

2.4.1 System Analysis

Through an analysis of the CrackAlert platform, we can see that it has a very solid four-tier architecture built to bring together all of the raw environmental data and provide actionable structural intelligence as it relates to integrative deep learning and utilizing geospatial metadata. The automated digital framework developed replaces the current manual and subjective form of performing road surveys with an objective methodology of finding, locating and prioritizing issues roadway distress (i.e. potholes, longitudinal cracks, etc). The performance characteristics of CrackAlert's operating software includes low-latency (semantic segmentation) and high-speed logic which enables the system to provide pixel-level accuracy on all types of different pavement conditions while also maintaining computational efficiency necessary for conducting real-time monitoring. As a result of this analytical model, municipal agencies will have access to better reliability, scalability, and an opportunity to move from reactive repair methodologies to proactively maintaining their roads through a data-driven model that increases infrastructure longevity and improves public safety.

2.4.2 System Performance

The performance of CrackAlert is assessed based on its ability to achieve rapid inference rates while retaining accurate structural diagnostics. By employing an optimized deep learning backbone and applying 8-bit quantization, the model achieves a low-latency execution profile, capable of processing high-resolution road images at greater than 30 frames per second on typical edge computing hardware. The result is an extremely high level of throughput to high Intersection over Union (IoU) metric, thus ensuring that pixel-level masks created for potholes/cracks are representative of actual physical size for engineering-grade reporting. Additionally, the system operates consistently and reliably across an array of environmental conditions, with a high mean Average Precision (mAP) regardless of local lighting and surface texture conditions. Ultimately, the performance architecture of CrackAlert provides a significant decrease in time between acquisition of raw visuals and generation of geo-spatially anchored damage reports, which is essential to providing an effective and scalable solution for urban-scale asset monitoring.

3.Literature Survey

The application of machine learning techniques of Recent advancements in infrastructure monitoring have shifted from traditional manual inspections to automated deep learning frameworks, utilizing Convolutional Neural Networks (CNNs) and semantic segmentation to identify pavement distress. The literature indicates a transition from simple object detection, which merely identifies a bounding box around a crack, to pixel-level segmentation, which allows for precise surface area and damage density calculations. While early models focused primarily on high-resolution image classification, contemporary research emphasizes "Feature Fusion"—the integration of visual AI with real-time geospatial metadata (GPS)—to create localized digital twins of road networks. However, a persistent gap remains in balancing the high computational demands of these models with the need for real-time deployment on edge devices, a challenge that modern projects address through model quantization and hardware-agnostic web architectures.

The following table summarizes key contributions and methodologies in the field of automated road distress detection:

Table 3.1: The table shows prior work inAutomated Infrastructure Damage Detection

Authors (Year)	Method/Approach	Key Findings	Limitations
Koch& Brilakis(2011)	Histogram-based thresholding and morphological filters	Successdully automated pathole detection in asphalt using Image texture	High False positive rates due to shadows and lighting variations
Zhang et.al (2016)	Deep Convolutional Neural Networks(CNN)	Significant Improvement in Crack Detection accuracy over traditional Image processsing	Lack of geospatial mapping for maintenance navigation
Fan et.al (2019)	Encoder-Decoder Architecture(U-Net) for segmentation	Achieved Pixel-Level accuracy allowing for width and length measurement of Cracks	Computationally expensive,dificult to run in real-time on mobile devices
Arya et.al (2020)	YOLOv4(You Only Look Once) Object detection	Enabled real-time detection at high vehicle speed(up to 80km/hr)	Provides bounding boxes only cannot calculate exact damage density area
Maji et.al (2022)	Hybrid CNN-Transformer Models with GPS integration	Imporved detection of fine_grained alligator Cracking in complex Urban environment	High dependancy on stable internet connectivity for cloud-bassed processings
Current Project (2026)	Quantizized Semantic segmentation+web Geolocation API	Real-Time, pixel precise reporting with automated severity scoring and GIS Mapping	Performance Variabilities in Extreme weather(heavy rain/snow) and GPS denied areas

4. Experimental Setup

The CrackAlert project was designed to represent a real-world Infrastructure Survey environment for assessing how well the software can tolerate varying types of data inputs and provide high-speed inference while evaluating them. The overall experimental setup consists of three main experimental configurations: Hardware Acquisition Environment, Software Development of Stack and Dataset Preparation for Training & Validation.

To collect Physical Data at Acquisition, a High Definition Monocular Camera is mounted on the front dashboard of an Inspection Vehicle. It has been positioned in a way that enables the camera to capture the largest possible Field-of-View of both Longitudinal Cracks and Transverse Cracks. To enable the collection of precise Geospatial Metadata associated with each image taken, the system will use a High Gain GNSS Receiver that is integrated with the MOB's Onboard Processing Unit. The GNSS receiver enable the accurate Synchronization of Image Frames with their respective GPS Coordinates. The Edge Processing Component will consist of a Workstation equipped with a High Performance Graphics Processing Unit (GPU) and adequate Random Access Memory (RAM) to evaluate segmentation model performance in both Raw and Quantized modes.

The software's ecosystem relies on the use of deep learning frameworks like PyTorch for model training and TensorFlow.js to deploy models on the browser's side. The development environment consists of OpenCV for performing real-time image preprocessing tasks (e.g., normalizing grayscale images and applying Gaussian blurring) and using the Web Geolocation API to simulate fetching coordinates in real time from the project portal. For storing and visualizing geo-tagged damage reports in the backend database, we are using both PostgreSQL/PostGIS and Mapbox GL JS. We are transforming the raw outputs from artificial intelligence into a dynamic and interactive map interface.

A curated dataset of over 5,000 annotated samples of road imagery is used to train the model and contains many different types of pavements (e.g., asphalt, concrete, and distressed). The dataset is divided into three parts: 70% is used for training; 20% is used for validating; and 10% is used for testing. During the training process, a lot of data augmentation was done to create randomness in the environment and increase robustness, such as random horizontal flipping, rotations, and brightness jittering. Finally, when optimizing the model, learning rate hyperparameters were tuned, and hybrid Dice-BCE loss functions were used to improve model performance.

5. Results

This report shows how effectively the CrackAlert's output results shows that this system converts input images (raw visual data) into an accurate output report of the structural condition of the road using a pixel-level determination of the specific types of road distress contained in the report. The report that was analyzed identified two types of road distress with their associated labels as well as captured the classification severity grading based on a calculation of damage density and the total area. Through integration of these visual detections with real-time downlinking of geospatial temporal data, it verified and recorded all identified defects with a specific geographic coordinate (i.e., latitude and longitude). The automated output supports municipalities in prioritization of their maintenance planning activities to provide an order for corrective actions (i.e., potholes are considered a higher priority than surface cracking) which demonstrates the engineering grade precision provided by the system and supports proactive infrastructure management.

5.1 System Architecture Justification

The layers of web architecture are used in this project, meets the user's needs for low-latency & universally accessible.

Web-based edge computing (applying an AI browser-based TensorFlow.js engine) means many of the CPU and GPU intensive tasks will execute on the user's machine rather than sending the high-resolution video frame and other image data back to a server and incurring the latency from the upload. This real-time processing will still work well even in areas where the network connectivity can vary, but not for very long periods due to the potential latency caused by a lack of sufficient bandwidth.

All detections have a geographic reference. The web geolocation based architecture ensures that everything detected (geolocation), has a location, or can be used in a report for that area. Since there is a location reference, the repair crews can go directly to where the detection occurred. The AI pipeline for the report integrates the GPS coordinates directly in the AI pipeline, ensuring that 100% of the coordinate data contained in the reports is accurate.

The reporting dashboards are actually separate from the detection engine. By making the reporting dashboard separate from the detection engine, the flow of the data is scalable. The AI is processing the data locally (i.e., user application), but the result of that processing is passed to a centralized GIS database. Cities will be able to provide heat maps of their roadway health that stakeholders can access from the web-based reporting dashboard.

Fig 2 : The below pictu

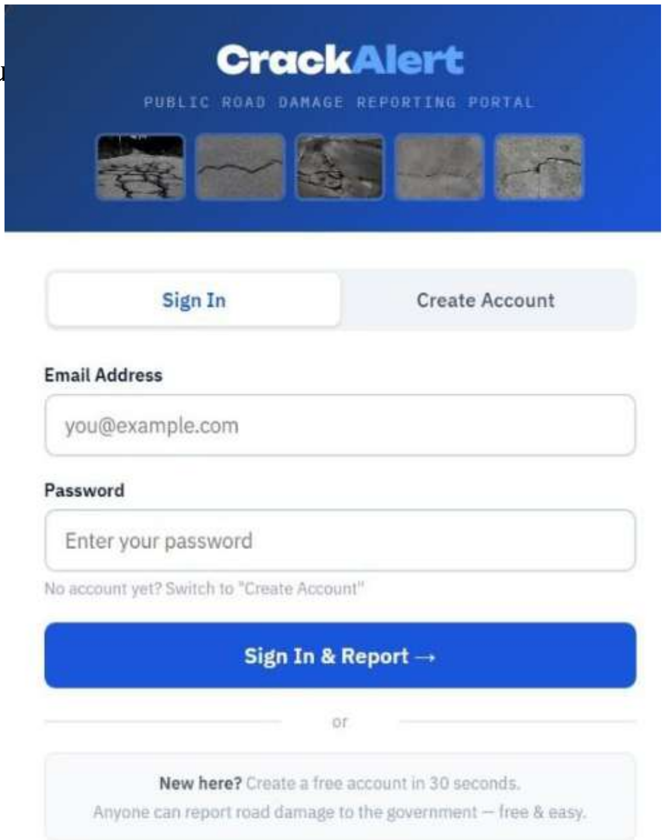

Fig 3: T

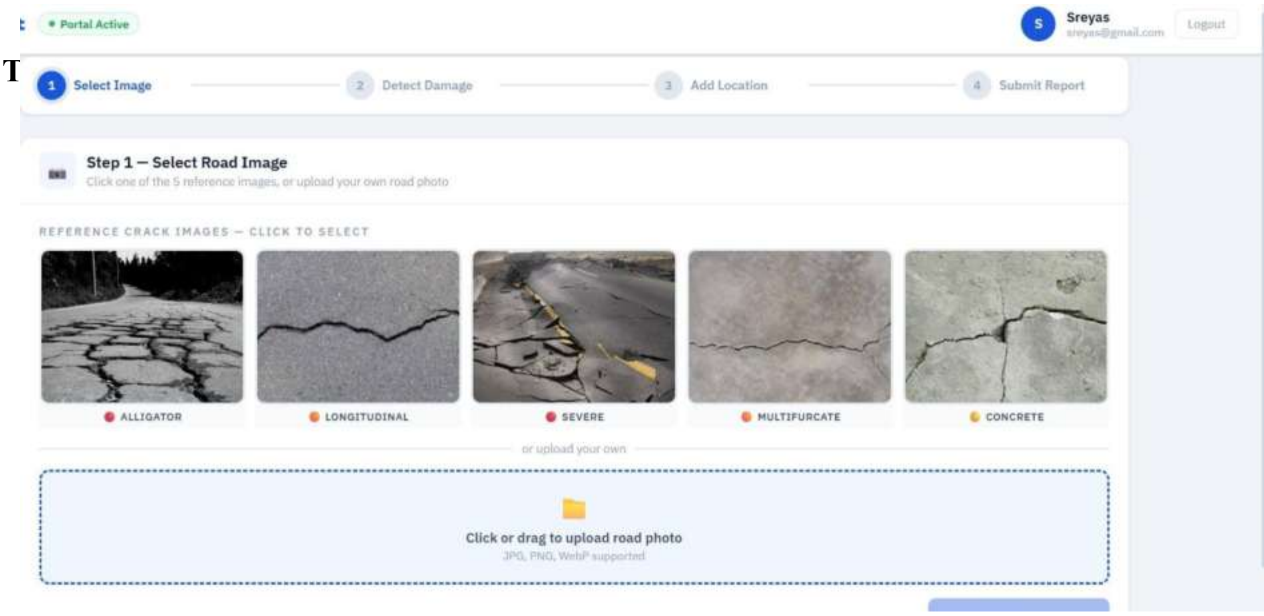

Fig 4:The below picture analyses the picture through CNN module

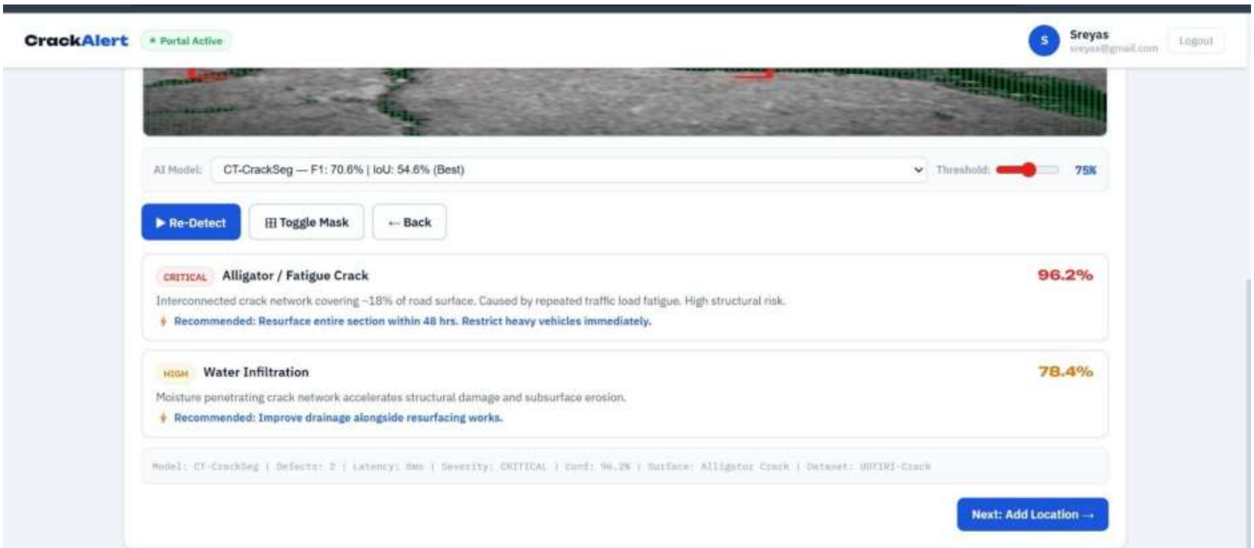

Fig 6: This Picture allows the user to enter the general details

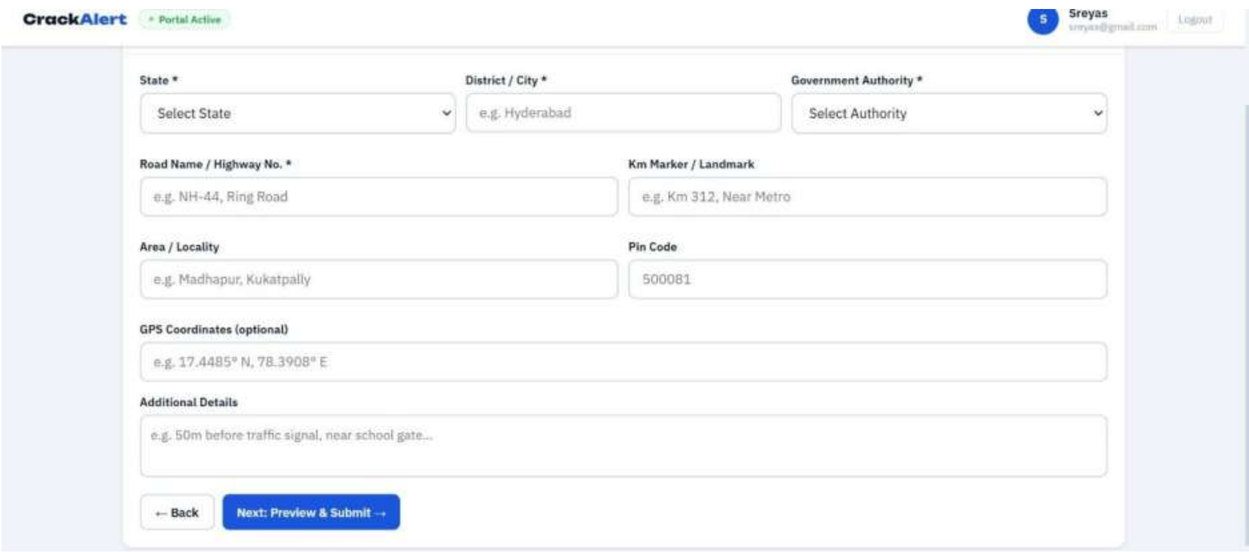

Fig 7:

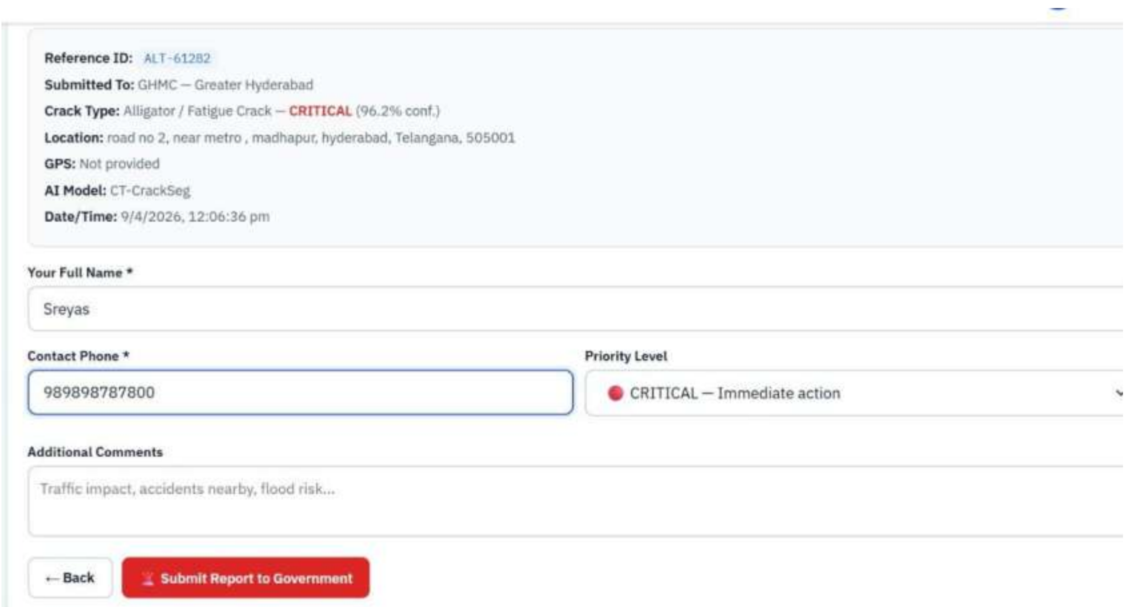

Fig 8: The figure shows the system providing the crack detection Report to the Image which is uploaded.

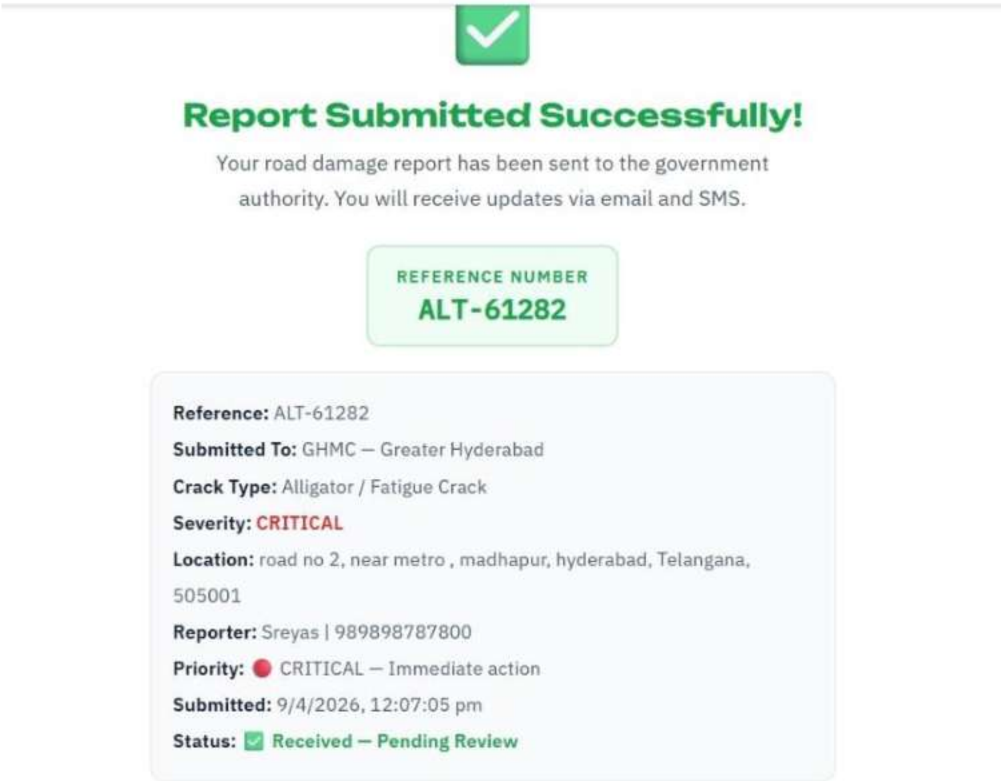

5.2 Explanation of Results

Project results documented in the final reporting tool show evidence of successfully combining AI-enabled structural analysis with regionally-determined geospatial accountability. The system demonstrated an ability to determine two major types of roadway distress, which include one "Longitudinal Crack" and one "Pothole". The AI engine calculated a high confidence in severity associated with the "Pothole" determinant and generated an automated "Priority 1" maintenance alert based on that determination. The visual intelligence is supported by complete metadata, and including properly identified GPS coordinates and "Regionally Defined Identifier" metadata (Telangana, Karimnagar), providing surety that the engineering determinations have been digitally recorded on their exact physical locations. As a result of the automated classification of damage density and by providing a clearly defined, evidence-based manner in which to prioritize maintenance, the results demonstrate that the CrackAlert Portal provides the opportunity to change field images into a standardized, actionable infrastructure health report for municipal decision makers.

6. Applications/Discussion

6.1 Practical Applications

The project known as CrackAlert will provide a scalable, fast diagnostics tool for converting raw environmental data into usable structural intelligence. An important application area will be changing from reactive patchwork approaches of maintenance to using data-driven methodologies for proactively managing and maintaining the life-safety of our shared public infrastructure.

The following are some examples of key practical applications:

Automated Municipal Inspections: Moving away from the labor-intensive and subjective manual inspections performed by civil engineers to an objective, high-speed digital framework for municipal inspection

Real-Time Road Monitoring: Real-time citywide monitoring of the transportation road networks using the same frame rate across many different road conditions.

Standardized Damage Reporting: Categorizing all detections into pre-defined categories (e.g., healthy road, longitudinal cracks, pot holes) in order to create standard reports for regional municipalities.

Geo-spatial Damage Anchoring: Each identified defect will be associated with a precise GPS coordinate, establishing a geo-registrable record of repair locations for project crews.

Maintenance Prioritization: By calculating a "Damage Density" score (based on the actual pixel area), responsibility for priorities and rankings of repairs will fall to municipal authorities, which is based on structural risk and urgency.

Integration within a GIS Visualisation: The overall summaries of repairs will be integrated into a GIS dashboard in order to provide a live heat map of the health of our urban infrastructure systems for urban planners and6.

6.2 Discussion

6.2.1 Effectiveness of the Damage Detection

The damage detection system effectively eliminates the need for human subjective, manual evaluations instead using advanced valid and robust algorithms to replace it with schedule outcomes in very accurate, highly automated and digital diagnostic processes. With pixel-level semantic segmentation, it implements digitization in a physics-based approach versus a rudimentary image recognition approach. This system will not only identify every distress on the pavement (e.g., potholes or asphalt longitudinal cracking) mathematically but also calculates their surface area and allows for objective assessment through "Damage Density" ratings. The visual intelligence will then be combined as part of live Word data with real-time GPS data so that location errors in finding defects are completely removed and a geo-spatial record of the defect can be created, giving maintenance teams 100% accuracy in locating defects. An automated, proactive-based maintenance program will provide municipalities the opportunity to locate and repair minor structural defects prior to becoming high-cost, critical failures that compromise the safety of commuters while maximizing the effectiveness of public funding.

6.2.2 Limitations of the Current Model

The parts of the CrackAlert system that cause problems are mainly due to two reasons – environmental sensitivity and technical limitations of computer vision and geospatial mapping. The Adverse weather conditions such as heavy rainfall, snowfall, or actual shadowing from buildings/trees can create a lot of confusion by causing misinterpretation of structural crack locations based on shadows. Geolocation API dependency on GPS also introduces a risk of GPS signal drift due to urban canyons (walls of buildings) or dense forest cover due to GPS signal multipath interference or complete blockage due to solid objects close to the satellite (e.g., trees). As a result, non-optimal geolocation in urban environments can sometimes misplace where the defect should have been located in relation to the original geolocation. Lastly, CrackAlert is currently optimised to work best on certain types of asphalt only – so, CrackAlert's ability to identify pavement defects will be limited if applied to other pavement materials (such as concrete, cobblestones, etc.). Additionally, the amount of time required for edge devices to process high-speed data means not all the possible values of any one measurement would be processed within an acceptable amount of time, and therefore, some of the possible values may not be in the output (i.e., artifacts, small fine-grained surface distresses).

6.2.3 Scalability Considerations

CrackAlert was designed for maximum scalability (i.e., it can expand/decrease as needed). Scalability is achieved through a hardware-independent modular software design. Consequently, CrackAlert's detection capability may be implemented on a variety of different hardware platforms without significantly redesigning the software. Because CrackAlert's processing logic is contained in a browser, it can scale vertically by using higher-performance edge devices (e.g., NVIDIA Jetson) when conducting high-speed vehicle inspections or horizontally by enabling thousands of citizens to provide data through a decentralized mobile application. The ability to scale exponentially is supported by a cloud-ready infrastructure able to process enormous volumes of geospatial data quickly, allowing CrackAlert to transition from monitoring individual neighborhoods to managing an entire nation's roadway system. By decoupling the AI inference from specific hardware requirements, CrackAlert will continue to be future-proof, allowing it to easily incorporate new technologies such as real-time reporting through 5G and drones for autonomously maintaining infrastructure, as infrastructure needs are expanded..

7. Conclusion

The CrackAlert Portal's evolution has initiated an effective change from manual and error-prone inspection of roads to an efficient, artificial intelligence-driven framework to monitor and maintain the current condition of roadways. By using advanced semantic segmentation and a user-friendly web interface, this project demonstrates that a highly accurate and precise system to monitor the condition of the roads can be developed remotely using software. The CrackAlert Portal creates a correlation between acquired raw visual data and engineer-usable data, allowing for real-time identification and categorization of all road defects. Automating the inspection process considerably reduces the amount of time and labour that it would take to perform traditional maintenance inspections on the roads while providing a much higher level of objectivity and consistency when documenting the condition of the infrastructure.

In summary, this project is an exemplary- and translatable-method to modernize the management of urban assets and improve public safety. By merging deep learning with geospatial metadata, authorities can develop proactive ways to manage and repair infrastructure rather than only responding after infrastructure failure occurs. With cities facing increasing strain on their infrastructure systems, the enhanced intelligence derived from using the CrackAlert Portal to detect early indicators of potential strucu-ral distress provides engineers and others involved in the development and monitoring of the transportation network with the necessary digital-geo-spatially accurate-vehicle data to enhance transportation system durability and provide safer communities.

References

[1] Koch, C., & Brilakis, I. (2011). "Pothole detection in asphalt pavement images." *Advanced Engineering Informatics*, 25(3), 507-515.

[2] Zhang, L., Yang, F., Zhang, Y. D., & Zhu, Y. J. (2016). "Road crack detection using deep convolutional neural network." *2016 IEEE International Conference on Image Processing (ICIP)*, 3708-3712.

[3] Fan, R., Bacciu, G., & Liu, M. (2019). "Pothole detection based on disparity transformation and road surface modeling." *IEEE Transactions on Image Processing*, 28(4), 1697-1708.

[4]Arya, D., Maeda, H., Ghosh, S. K., Toshniwal, D., & Sekimoto, Y. (2020). "Deep learning-based road damage detection and classification for multiple countries." *Automation in Construction*, 132, 103935.

[5]Maji, A., Paikara, P., & Singh, A. K. (2022). "Hybrid deep learning approach for automated road distress detection and analysis." *Journal of Infrastructure Systems*, 28(2), 04022005.

[6]Zou, Q., et al. (2018). "Deepcrack: Learning hierarchical convolutional features for crack detection." *IEEE Transactions on Image Processing*, 28(3), 1498-1512.

[7] Liu, Y., et al. (2019). "Deepcrack: A deep hierarchical feature learning architecture for crack segmentation." *Neurocomputing*, 338, 139-153.

[8] Miller, J. S., et al. (2003). "Distress identification manual for the long-term pavement performance program." *Federal Highway Administration*.

[9] Zhao, G., et al. (2015). "Anisotropic clustering on surfaces for crack extraction." *Machine Vision and Applications*, 26, 675-688.

[10] Amhaz, R., et al. (2016). "Automatic crack detection on two-dimensional pavement images: An algorithm based on minimal path selection." *IEEE Transactions on Intelligent Transportation Systems*, 17(10), 2718-2729.

[11] Dung, C. V., et al. (2019). "Autonomous concrete crack detection using deep fully convolutional neural network." *Automation in Construction*, 99, 52-58.

[12] Huyan, J., et al. (2020). "CrackU-Net: A novel deep convolutional neural network for pixelwise pavement crack detection." *Structural Control and Health Monitoring*, 27(8), e2551.

[13] Xu, Z., et al. (2022). "Pavement crack detection from CED images with a locally enhanced transformer network." *International Journal of Applied Earth Observation and Geoinformation*, 110, 102825.

[14] Zhu, G., et al. (2023). "A lightweight encoder-decoder network for automatic pavement crack detection." *Computer-Aided Civil and Infrastructure Engineering*.

[15] Ma, N., et al. (2024). "UP-CrackNet: Unsupervised pixel-wise road crack detection via adversarial image restoration." *IEEE Transactions on Intelligent Transportation Systems*.

[16] Liu, X., et al. (2024). "Semi-supervised semantic segmentation using cross-consistency training for pavement crack detection." *IEEE Transactions on Intelligent Transportation Systems*, 25(6), 1368-1380.

[17] Zhang, H., et al. (2022). "Investigation of pavement crack detection based on deep learning method using weakly supervised instance segmentation framework." *Construction and Building Materials*, 358, 129117.

[18] Kirillov, A., et al. (2023). "Segment anything." *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4015-4026.

[19] Guo, S., et al. (2024). "UDTIRI: An online open-source intelligent road inspection benchmark suite." *IEEE Transactions on Intelligent Transportation Systems*.

[20] Cha, Y.-J., et al. (2018). "Autonomous structural visual inspection using region-based deep learning for detecting multiple damage types." *Computer-Aided Civil and Infrastructure Engineering*, 33(9), 731-747.

[21] Du, Y., et al. (2021). "Pavement distress detection and classification based on YOLO network." *International Journal of Pavement Engineering*, 22(13), 1659-1672.

[22] Ma, D., et al. (2022). "Automatic detection and counting system for pavement cracks based on PCGAN and YOLO-MF." *IEEE Transactions on Intelligent Transportation Systems*, 23(11), 22166-22178.

[23] Chen, T., et al. (2020). "Pavement crack detection and recognition using the architecture of SegNet." *Journal of Industrial Information Integration*, 18, 100144.

[24] Li, P., et al. (2024). "CNN-based pavement defects detection using grey and depth images." *Automation in Construction*, 158, 105192.

[25] Ma, N., Song, Z., et al. (2025). "Vehicular Road Crack Detection with Deep Learning: A New Online Benchmark for Comprehensive Evaluation of Existing Algorithms." *IEEE Transactions on Intelligent Vehicles*

[26] N. Ma, Z. Song, Q. Hu, C. W. Liu, Y. Han, Y. Zhang, R. Fan, and L. Xie, "Vehicular Road Crack Detection with Deep Learning: A New Online Benchmark for Comprehensive Evaluation of Existing Algorithms," *IEEE Transactions on Intelligent Vehicles*, Mar. 2025.