CTIPS

Center for Transformative Infrastructure Preservation and Sustainability

USDOT Region 8 University Transportation Center

CTIPS-25-002

B. Kuang
J. Chen

ARTIFICIAL INTELLIGENCE AND MOBILE PHONE-BASED PAVEMENT MARKING CONDITION ASSESSMENT AND LITTER IDENTIFICATION

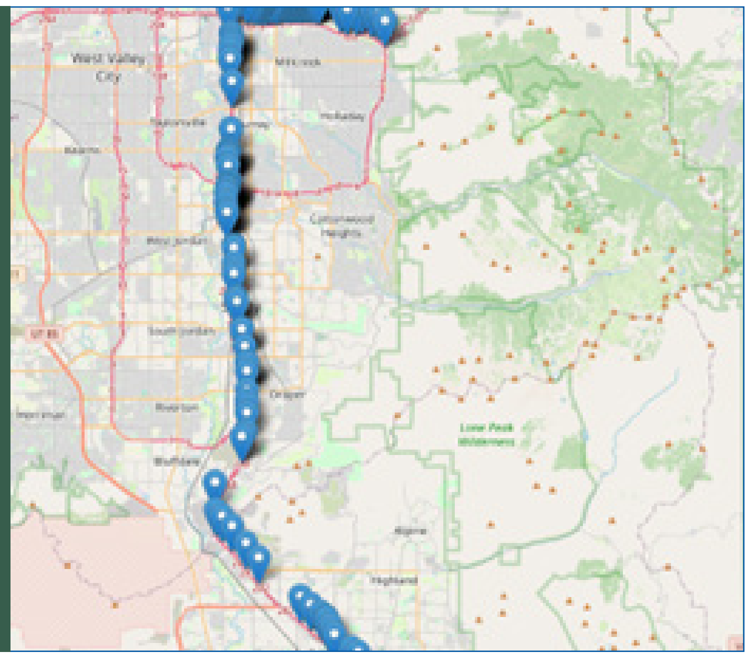

A University Transportation Center sponsored by the U.S. Department of Transportation serving the Center for Transformative Infrastructure Preservation and Sustainability members:

Colorado State University
Fort Lewis College
North Dakota State University
South Dakota State University
United Tribes Technical College
University of Colorado Denver
University of Denver
University of North Dakota
University of Utah
University of Wyoming
Utah State University

Technical Report Documentation Page

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
CTIPS-007		

4. Title and Subtitle	5. Report Date
Artificial Intelligence and Mobile Phone-Based Pavement Marking Condition Assessment and Litter Identification	November 2025
	6. Performing Organization Code CTIPS
7. Author(s) Biao Kuang https://orcid.org/0000-0001-5787-9480 Jianli Chen https://orcid.org/0000-0001-7106-7833	8. Performing Organization Report No. CTIPS-25-002
9. Performing Organization Name and Address University of Utah Floyd and Jeri Meldrum Civil Engineering Building, 110 S Central Campus Dr, Salt Lake City, UT 84112	10. Work Unit No.
	11. Contract or Grant No. 69A3552348308
12. Sponsoring Agency Name and Address Department of Transportation Research and Innovative Technology Administration 400 7th Street, SW Washington, DC 20590	13. Type of Report and Period Covered Final Report
	14. Sponsoring Agency Code

15. Supplementary Notes

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

16. Abstract

Regular inspection of transportation assets is essential to ensure pavement markings and pavements remain in good, clean, and safe condition. Our previous MPC-funded project (Report No. MPC-668) demonstrated the strong potential of using artificial intelligence (AI) and mobile phone imagery to identify various transportation assets. However, that initial effort was limited in scale, using only ~1,000 images for training and validation. Building upon this foundation, the present project focuses on two targeted assets: pavement marking issues and roadside litter, while expanding the capability of the previously developed AI packages. In this study, the dataset increases to over 6,000 images for each asset type. Using the You Only Look Once (YOLO) deep learning architecture, two detection models are trained and achieved strong accuracy metrics, with F1-scores of 0.88 for pavement marking issues and 0.84 for roadside litter. In addition, counting and geolocation models are developed to quantify detected objects within a road section or video clip and to determine their precise locations by integrating data from a phone-based Global Positioning System (GPS) tracker. The geolocation model demonstrates high spatial accuracy, achieving an average positional error of only 0.27 meters. To facilitate practical application, an interactive mapping interface is implemented to visualize the geolocation, object class, inspection time, and cropped image of each identified object. This interface enables clear and intuitive assessment of pavement conditions, specifically faded markings and roadside litter. Overall, this project enhances our prior work by extending capabilities in detection, counting, geolocation, and visualization, which supports regular asset inspection, informs maintenance planning, and ultimately improves roadway safety.

17. Key Words	18. Distribution Statement
Artificial intelligence (AI), Pavement marking; Roadside litter and trash; Object detection; Object Counting; Object geolocation	No restrictions. This document is available through the National Technical Information Service, Springfield, VA 22161.

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	44	n/a

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

Artificial Intelligence and Mobile Phone-Based Pavement Marking Condition Assessment and Litter Identification

Biao Kuang

Dr. Jianli Chen

University of Utah
Department of Civil & Environmental Engineering
110 Central Campus Dr., RM 2000
Salt Lake City, UT 84112

November 2025

About the Center for Transformative Infrastructure Preservation and Sustainability

The Center for Transformative Infrastructure Preservation and Sustainability (CTIPS) is the US DOT Region 8 University Transportation Center (UTC) located at North Dakota State University. CTIPS will focus on "Preserving the Existing Transportation System," along with consortia members Colorado State University, Fort Lewis College, Fort Lewis College, South Dakota State University, United Tribes Technical College, University of Colorado Denver, University of Denver, University of North Dakota, University of Utah, University of Wyoming and Utah State University.

Acknowledgements

The authors sincerely thank Utah Department of Transportation (UDOT) and The Center for Transformative Infrastructure Preservation and Sustainability (CTIPS) for the invaluable support and assistance.

ABSTRACT

Regular inspection of transportation assets is essential to ensure pavement markings and pavements remain in good, clean, and safe condition. Our previous MPC-funded project (Report No. MPC-668) demonstrated the strong potential of using artificial intelligence (AI) and mobile phone imagery to identify various transportation assets. However, that initial effort was limited in scale, using only ~1,000 images for training and validation. Building upon this foundation, the present project focuses on two targeted assets: pavement marking issues and roadside litter, while expanding the capability of the previously developed AI packages. In this study, the dataset increases to over 6,000 images for each asset type. Using the You Only Look Once (YOLO) deep learning architecture, two detection models were trained and achieved strong accuracy metrics, with F1 scores of 0.88 for pavement marking issues and 0.84 for roadside litter. In addition, counting and geolocation models are developed to quantify detected objects within a road section or video clip and to determine their precise locations by integrating data from a phone-based global positioning system (GPS) tracker. The geolocation model demonstrates high spatial accuracy, achieving an average positional error of only 0.27 meters. To facilitate practical application, an interactive mapping interface is implemented to visualize the geolocation, object class, inspection time, and cropped image of each identified object. This interface enables clear and intuitive assessments of pavement conditions, specifically faded markings and roadside litter. Overall, this project enhances our prior work by extending capabilities in detection, counting, geolocation, and visualization, which supports regular asset inspection, informs maintenance planning, and ultimately improves roadway safety.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF KEY TERMS

AAA	American Automobile Association
AI	Artificial Intelligence
ASTM	American Society for Testing and Materials
CNN	Convolutional Neural Networks
FN	False Negative
FP	False Positives
FPS	Frames Per Second
GIS	Geographic Information System
GPS	Global Positioning System
IoU	Intersection over Union
mAP@50	Mean Average Precision at 50% IoU
MPC	Mountain-Plains Consortium
TP	True Positive
YOLO	You Only Look Once

EXECUTIVE SUMMARY

Pavement markings are critical transportation assets that directly affect roadway safety and mobility but have a relatively short service life. Faded markings can significantly impair drivers' ability to navigate safely. Also, roadside litter (e.g., vehicle debris) poses hazards that may contribute to traffic accidents. Regular inspection and maintenance, such as repainting faded markings and removing litter, are essential to keep roadways in good, clean, and safe condition. However, traditional inspection methods rely heavily on manual labor, which is subjective, time-consuming, and unsuitable for large-scale and frequent assessments. Advances in artificial intelligence (AI), particularly deep learning and computer vision, now offer promising alternatives for automating transportation infrastructure inspections.

In our previous MPC project (Report No. MPC-668), we developed an AI-based package for detecting transportation assets and related issues. However, that effort was limited in scope, using only ~1,000 images per task for training and validation, which restricted the generalizability of the developed models. In this new project, we focus specifically on pavement markings and roadside litter, expanding both the dataset and the capabilities of the developed AI packages.

First, the dataset is expanded to over 6,000 self-collected images for each asset type. Faded pavement markings are classified into white faded and yellow faded markings, while roadside litter is categorized into four classes: white litter, black litter, leaves, and dirt. Based on these annotated images and using the You Only Look Once (YOLO) deep learning architecture, we trained two AI detection models that achieved strong performance, with F1 scores of 0.88 for faded pavement markings and 0.84 for litter.

Beyond detection, we extend our AI models to quantify the number of identified objects within a road segment or video clip. Based on object tracking, we develop a specialized counting model, which has been successfully deployed in a case study of a highway route in Utah. Moreover, we enhance the AI packages to geolocate each identified object. By integrating a phone-based GPS tracker (GPX tracker) and employing time-based interpolations, we propose a geolocation model to match the video frames with GPS recorders. This geolocation model is designed to estimate the latitude and longitude of the identified objects accurately; it exhibits high accuracy in pinpointing object positions, with an average distance error of only 0.27 meters, as validated through a case study on a test highway route (I-15) in Utah.

Finally, we implement an interactive mapping interface using Folium, a Python package, to visualize the geolocation, object class, inspection time, and cropped image of each identified object. This tool offers transportation agencies an intuitive, data-rich platform for monitoring roadway conditions, prioritizing maintenance, and improving asset management efficiency. Overall, this project builds significantly on our earlier work, extending AI capabilities from simple detection to comprehensive counting, geolocation, and visualization. These advancements enable more frequent, accurate, and scalable inspections of transportation assets, supporting proactive maintenance planning and contributing to improved roadway safety.

1. INTRODUCTION AND BACKGROUND

1.1 Project Motivation

Pavement markings are critical components of roadway infrastructure, providing essential visual guidance to drivers by indicating regulations, navigation instructions, and safety warnings (FHWA 2009). However, due to continuous exposure to traffic wear, adverse weather conditions, and mechanical actions like snowplowing, these markings often deteriorate over time, becoming faded or entirely worn out (Sassani et al. 2021; Xu et al. 2021). Studies have shown that pavement markings generally have a short service life, ranging from as little as 0.5 years to a maximum of about three years (Alzraiee et al. 2021). Degraded or missing markings can lead to confusion, misjudgment, and accidents among road users (Kawano et al. 2017; Sassani et al. 2021), emphasizing the importance of regular inspections and timely maintenance to uphold road safety. Traditionally, pavement marking assessments are carried out manually or with specialized equipment, e.g., retroreflectometers (Pike et al. 2011; Wei et al. 2021). While effective, these methods are labor-intensive, time-consuming, expensive, and infeasible for frequent or large-scale inspections (Kuang and Chen 2024b). The manual process also introduces subjectivity and variability in results. Therefore, there is a growing need for a more scalable, objective, and cost-effective solution to monitor pavement markings efficiently and accurately.

Another persistent challenge in highway maintenance is roadside litter (Chamberlin et al. 2021). Common forms of litter, such as construction debris, vehicle debris, food containers/wrappers, plastics, and leaves, are strewn across roads, which easily lead to blockage of the roadside drainage system (Karimi and Faghri 2021; Mullaney and Lucke 2014) and pose significant hazards to road users by causing traffic accidents. A research report by the American Automobile Association (AAA) Foundation for Traffic Safety indicates that, over the years 2011–2014, an estimated average of 50,658 police-reported crashes in the U.S. were attributed to road debris, resulting in approximately 9,800 injuries and 125 fatalities annually (Tefft 2016). Therefore, maintaining a clean and hygienic pavement is an important task for traffic safety and municipal works (Ramalingam et al. 2021). Similar to pavement marking issues, current practices for identifying litter on the roadside rely on the manual process, which is repetitive and laborious (Gómez et al. 2022). Hence, there is a pressing demand for an intelligent and efficient way to identify litter on the roadways.

One way to develop an efficient and automated approach to inspect pavement marking and identify litter is the utilization of artificial intelligence (AI), such as computer vision and deep learning. The advancement of AI has enabled the automatic inspection of transportation assets through object detection and classification with high accuracy and at a nearly real-time speed (Krizhevsky et al. 2017; Redmon et al. 2016). In transportation applications, deep learning has been employed for pavement distress detection (Farhadmanesh et al. 2021; Peng et al. 2024; Wen et al. 2022; Zakeri et al. 2017), traffic sign recognition identification (Karsten et al. 2021), and even segmentation-based pavement marking evaluations (Xu et al. 2021). These AI tools offer not only improved accuracy but also the scalability required for systematic infrastructure monitoring.

Building on this potential, our prior project, funded by the Mountain-Plains Consortium (MPC) (report No. MPC-668) and titled “Mobile Phone-Based Artificial Intelligence Development for Maintenance Asset Management” (Kuang and Chen 2024a; Kuang et al. 2024), successfully demonstrated AI-based models for detecting faded pavement markings and roadside litter using mobile phone video data. However, that effort was small-scale in scope, with a relatively small dataset of ~1,000 images per task for model training and validation. Additionally, the prior work is limited to identifying objects of interest and lacks the ability to count the number of detected objects within a specific road section or video clip. Object detection in the previous work is a static task focused on locating regions of interest in individual frames (Koch et al. 2013; Wang et al. 2023). However, the same object can appear across consecutive frames in a video sequence (Wei et al. 2021), leading to repeated detection of the same object in multiple frames. Furthermore, the precise geolocation of identified objects is another gap in our previous project. Accurate geolocation is essential for planning effective maintenance work. A few prior studies have employed high-resolution global positioning system (GPS) devices to record geolocations during data collection (Ma et al. 2022; Okpe and Idachaba 2023). Based on these GPS records, the location of detected objects can be inferred. However, the sampling rates of GPS devices and video recordings are heterogeneous, and how to accurately synchronize GPS records with video frames remains an unsolved issue in existing studies. This lack of precise location information for identified objects (e.g., transportation assets or asset issues) presents a significant research gap, which hinders maintenance teams from effectively understanding the locations of these issues for timely intervention.

As a result, significant gaps remain in the development of an end-to-end AI-based inspection package that can not only detect pavement-related issues and identify roadside litter but also count and geolocate them accurately. Addressing these gaps through an enhanced, scalable, and automated AI solution would offer transportation agencies a powerful tool for improving roadway maintenance, enhancing safety, and reducing operational costs.

1.2 Research, Objectives, and Tasks

To address the limitations in our previous project, this project aims to expand upon prior work through a large-scale investigation focused on the detection, counting, and geolocation of pavement markings and roadside litter. Specifically, we seek to significantly increase dataset size and diversity, improve AI model accuracy, and enhance the practical utility of object detection systems for transportation asset management. By advancing both the technical and operational capabilities of the AI package, we aim to support more effective and efficient maintenance planning. The following research tasks are proposed to achieve these objectives:

Task 1: Data Collection

Expand the dataset by collecting a significantly larger volume of images and videos of transportation assets and issues, including pavement markings and guardrails, across a broader geographic area and under diverse environmental conditions (e.g., varying lighting, weather, and time of day).

Task 2: Advancement of Object Detection Models

Leverage the expanded dataset to enhance and refine AI models for object detection. This includes improving the reliability and accuracy of assessing pavement marking conditions and identifying roadside litter under various operational conditions.

Task 3: Development of Object Counting Algorithms

Design and implement algorithms capable of accurately counting distinct objects detected within a given road section or video segment. These algorithms will address the challenge of repeated detections by tracking objects across consecutive frames, thereby reducing overcounting.

Task 4: Geolocation of Identified Objects

Develop a method to synchronize GPS data with video frames to accurately geolocate each detected object. This task addresses a major limitation in previous work, ensuring that maintenance teams can pinpoint the exact locations of identified issues for timely and effective intervention.

Task 5: Large-Scale Model Validation and Testing

Conduct extensive validation and testing of the enhanced AI models using the expanded dataset. This step will assess the performance of models in terms of accuracy, robustness, and scalability across different road environments and operational conditions.

1.3 Report Overview

This report is organized to provide a clear understanding of the background, methodology, and outcomes. The remaining report is structured as follows: Section 2 presents a comprehensive review of current practices in transportation asset data collection, discussing both their advantages and limitations, as well as recent applications of AI in the maintenance of transportation infrastructure. Section 3 outlines the methods used in this study, including the data collection process, model development, and the accuracy metrics employed to evaluate performance. Section 4 presents the results of the project, showcasing the effectiveness of the developed AI models in object detection, counting, geolocation, and visualization. Finally, Section 5 summarizes the key findings and contributions of the study and provides recommendations for future research and practical implementation.

2. LITERATURE REVIEW

2.1 AI Applications in Identifying Transportation Assets

Various AI algorithms have been applied to transportation asset maintenance, with pavement condition assessment emerging as a prominent area of research. Deep learning techniques such as convolutional neural networks (CNN) (Gopalakrishnan et al. 2017), faster R-CNN (Majidifard et al. 2020), and You Only Look Once (YOLO) (Mandal et al. 2020). Numerous public datasets related to transportation assets have been widely adopted for developing automated detection models. Publicly available datasets used in these studies have been collected through smartphones, mounted cameras, and Google Street View imagery (Majidifard et al. 2020; Mandal et al. 2020). Using these models and data sources, researchers have developed AI models for detecting pavement distresses, identifying traffic signs, assessing pavement markings, and identifying roadside litter (Du et al. 2020; Ghosh & Smadi 2021; Majidifard et al. 2020).

2.1.1 Pavement Marking Detection

Compared with pavement distress detection, relatively few studies have focused on assessing the condition of pavement markings. Among existing approaches, a common method involves quantifying the remaining proportion of pavement markings. For instance, Lee & Cho (2023) employed Mask R-CNN to automatically calculate pavement marking defect ratios by analyzing retroreflectivity data collected via a vehicle-mounted retroreflectometer. In addition to sensor-based approaches, visual imagery has been widely used. Zhang & Ge (2012) utilized traditional image processing techniques, such as the Hough transform and feature recognition, to estimate the percentage of remaining pavement markings. Similarly, Xu et al. (2021) applied edge extraction and object segmentation to quantify the severity of fading at the pixel level. Kang et al. (2020), using a YOLOv3-based model, focused on identifying arrow markings and evaluating their visibility by measuring intensity contrast between the markings and the surrounding pavement. Wei et al. (2021) estimated damage ratios by segmenting regions corresponding to both original and faded markings through hierarchical semantic segmentation.

Vision-based deep learning approaches have also gained attention for assessing pavement marking conditions. Iparraguirre et al. (2022) applied several deep learning architectures, including Faster R-CNN, SSD, and EfficientDet, to detect faded markings. However, their study was limited to a single class (white faded lane markings) and lacked clear data labeling criteria. Alzraiee et al. (2021) categorized faded markings into nine classes and trained a Faster R-CNN model to detect them; however, the resulting accuracy was relatively low (less than 0.25). Similarly, YOLO-based models were adopted by Kawano et al. (2017) and Bronuela-Ambrocio & Antes (2023), but these studies also achieved modest precision scores (below 60%). In contrast, Sun et al. (2024) evaluated multiple architectures: YOLOv5, YOLOv7, Faster R-CNN, and ResNet on a dataset of 2,000 self-labeled images, achieving accuracies between 0.86 and 0.94. Maeda et al. (2018) developed a CNN-based model capable of identifying two pavement marking issues—faded crosswalks and faded white lane markings—with accuracies ranging from 0.62 to 0.95. Additionally, Kong et al. (2022) implemented the DeepLab V3+ semantic segmentation model to assess various pavement marking types, such as line markings, arrow markings, and evenly spaced lines, and reported F1 scores between 0.81 and 0.85.

Despite these advancements, many existing studies focus only on specific types or shapes of pavement markings (most often lane lines), while neglecting other forms of markings that may also experience fading. This narrow scope highlights the need for more comprehensive models capable of assessing a wider variety of pavement marking conditions under diverse real-world scenarios.

2.1.2 Litter and Trash Detection

In contrast to research on pavement condition assessment, significantly fewer studies have focused on the detection and classification of litter and trash on roadways (Wu et al. 2023). Wu et al. (2023) conducted a comprehensive review and found that only 10.4% of existing studies in this area addressed urban trash detection, indicating a substantial research gap. Among the early contributions, Rad et al. (2017) developed a model to detect and quantify various types of litter, such as cigarette butts and fallen leaves. Liu et al. (2018) applied the YOLOv2 architecture for garbage detection on pavements but limited the scope to a single trash class. Similarly, Sayyad et al. (2020) developed an AI model focused on detecting large-sized litter without classifying different garbage types, which constrained the applicability of the developed model in real-world scenarios. Zhang et al. (2019) adopted Faster R-CNN to automatically identify and count multiple categories of litter, including organic waste, inorganic materials, general trash, and tree leaves. However, their dataset was exclusively collected from walking pavements, limiting the model's generalizability to highway or street environments. More recently, Mandhati et al. (2024) evaluated and compared the performance of several deep learning models, such as Faster R-CNN, RetinaNet, YOLOv3, and YOLOv5, for detecting plastic litter on roadsides. Their study reported an average precision of 40%, highlighting the challenges in accurately detecting lightweight and visually diverse roadside litter. V et al. (2022) developed a YOLOv4-based model trained on 343 images to detect plastic bottles and trash cans, achieving a precision of 0.71 and a recall of 0.81.

Overall, deep learning and computer vision techniques have demonstrated strong capabilities in object detection and image classification, with widespread applications in transportation asset monitoring and inspection. While pavement condition assessment remains a relatively extensive domain, litter detection has received less attention. In particular, the limited focus on detecting trash and debris, especially across varied environments like highways, represents a significant gap. Expanding research in these underexplored areas could greatly enhance the automation, efficiency, and comprehensiveness of transportation asset management systems.

2.2 AI Applications in Object Counting

Quantifying the number of identified objects, such as defects or issues in transportation assets, is essential for effective infrastructure inspection and maintenance planning. For example, determining the number of faded pavement markings within a specific road segment is critical for assessing the extent of deterioration and evaluating the need for repainting. However, in video-based detection, objects often appear across consecutive frames, leading to duplicate detections of the same item. Despite this, relatively few studies have addressed the challenge of accurately counting unique objects in such contexts.

Okpe & Idachaba (2023), for instance, claimed that their model could count pavement potholes; however, their approach only tallied the number of detections without accounting for duplicate appearances of the same object across multiple frames. As a result, the reported count is likely to overestimate the actual number of distinct potholes. Wei et al. (2021) recognized this issue and applied kernelized correlation filters to track pavement markings across video frames. While their method addressed duplicate detections to some extent, it did not use object tracking explicitly for the purpose of counting the actual number of faded markings within a defined road section.

Object tracking offers a promising solution for eliminating duplicate detections and accurately counting individual objects across a sequence of frames. Multi-object tracking and counting algorithms have already been widely applied in transportation domains such as vehicle and pedestrian counting (Fernández-Sanjurjo et al. 2019; Wickramasinghe & Ganegoda 2020). However, their application in counting transportation infrastructure assets, such as degraded pavement markings or other asset-related issues, still remains limited.

A few notable studies have begun to explore this area. Koch et al. (2013) employed a kernel-based tracking approach to count potholes in video sequences, demonstrating the feasibility of tracking-based counting for infrastructure defects. Similarly, Ma et al. (2022) utilized YOLOv3 for crack detection, followed by the median flow tracking algorithm to monitor and count pavement cracks within road sections. These studies illustrate the potential of integrating object detection with tracking mechanisms to improve the accuracy of object counting in transportation asset management. However, further research is needed to adapt and refine these methods for broader and more complex use cases.

2.3 AI Applications in Geolocating Objects

The geolocation of identified transportation asset issues is a critical component for effective maintenance planning and resource allocation. Despite its importance, geospatial information is often lacking in many AI-based studies focused on detecting and assessing transportation infrastructure. Without precise location data, it becomes challenging for maintenance teams to accurately locate and address issues in a timely and efficient manner.

Some studies have leveraged geo-tagged or GIS-integrated imagery. For example, Kong et al. (2022) assessed pavement marking conditions using street-view images from open-source mapping platforms that contained embedded geographic information system (GIS) data. Likewise, Kargah-Ostadi et al. (2020) integrated GPS coordinates into their dataset, enabling precise geolocation of transportation assets captured in each image. While these methods are effective, they rely on pre-existing street-view inventories, making them less suitable for applications requiring frequent data updates or self-collected imagery.

Other research efforts have incorporated dedicated GPS hardware to capture location data alongside images or videos of transportation assets. Okpe & Idachaba (2023) integrated a GPS module with a Raspberry Pi-based system to detect potholes and transmit the corresponding location data to a remote monitoring platform via email. Similarly, Ma et al. (2022) developed a custom detection device equipped with a GPS unit and 4G wireless module to geolocate pavement cracks in real time. Both approaches, however, depend on specialized hardware, which may increase system complexity and deployment costs.

Mobile phones have emerged as a more accessible alternative for data collection, combining image capture with built-in GPS functionality. Souza et al. (2018), for instance, developed a mobile app that recorded GPS data at a frequency of 1 Hz to evaluate pavement conditions along road segments. However, the study did not aim to geolocate individual pavement defects but rather focused on broader section-level assessments using relatively low-resolution spatial data.

In summary, while some progress has been made in integrating geolocation with AI-driven detection of transportation asset issues, challenges remain, particularly in achieving high-resolution, frame-level geolocation for self-collected video data. Accurate synchronization of GPS data with image frames, especially when using heterogeneous sampling rates, is an area that requires further investigation to improve the spatial precision and utility of AI-assisted asset monitoring systems.

3. METHODOLOGY

3.1 Research Process Overview

The overall workflow for developing the AI package is illustrated in Figure 1.

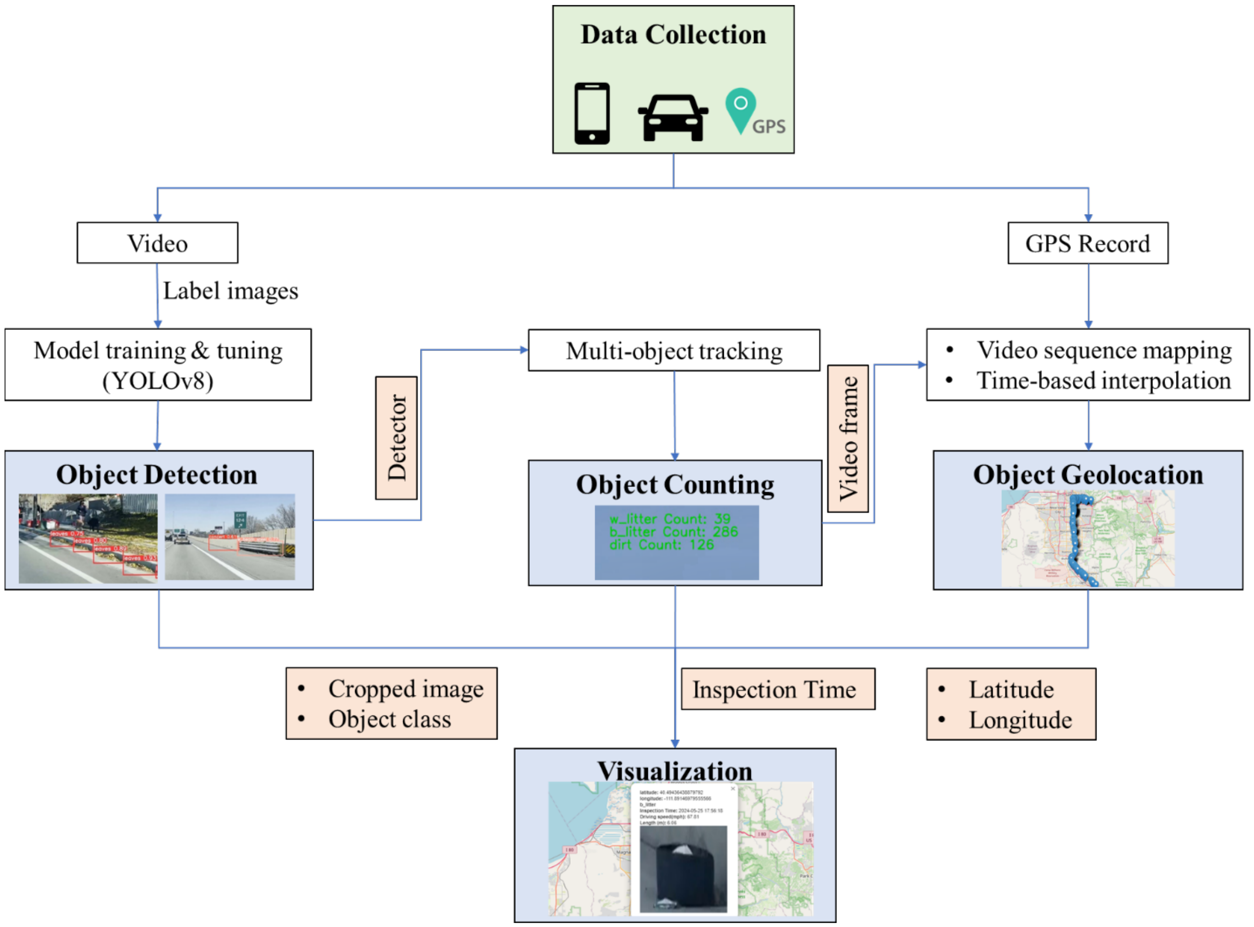

Figure 1. Research Framework

The process begins with data collection using a mobile phone mounted on the front windshield of a vehicle. This setup captures forward-facing videos along with GPS driving data. The recorded videos are then converted into individual image frames, which are manually annotated for training and validation of object detection models. YOLO architecture is employed to detect various targeted transportation assets, including pavement markings and roadside litter.

Following object detection, a multi-object tracking approach is applied to develop a counting model, which quantifies the number of unique objects within a defined road segment or video clip. The tracking process links detections across consecutive video frames, reducing redundancy from repeated identifications of the same object.

Based on the tracking output, a geolocation model is developed to estimate the precise locations (latitude and longitude) of each detected object. This is achieved by aligning GPS records with video frames using time-based interpolation. The resulting geolocation estimates are integrated into a visualization interface that displays detected objects on an interactive map. This map includes cropped images of the detected assets, classification labels, geolocation coordinates, and inspection timestamps, offering a user-friendly platform for infrastructure monitoring and decision-making.

3.2 Data Collection and Processing

Consistent with Phase I of this project, a mobile phone mounted on the interior of the vehicle's front windshield (passenger side) is used to collect video data of transportation assets from a forward-facing perspective. The data collection setup is shown in Figure 2.

(a) Exterior View (b) Interior View

Figure 2. Setup of Data Collection

The videos are recorded at 30 frames per second (fps) and cover a variety of road environments across Utah, including both freeways and local streets. Data are collected under diverse conditions, such as different road types and varying illumination levels, to ensure model robustness. Targeted transportation assets include pavement markings and roadside litter.

During data processing, individual frames are extracted from the recorded videos. To build a high-quality and balanced dataset, only frames that clearly capture objects of interest are manually selected. To reduce redundancy, no more than three images of the same object are included in the training set. This strategy ensures diversity in the dataset while avoiding overrepresentation of any single asset instance.

3.3 Data Annotations

For annotation, the open-source tool LabelImg[1] is used to manually label objects with bounding boxes. LabelImg provides a user-friendly interface for accurate annotation, and it supports the generation of labels compatible with common deep learning frameworks. Separate datasets are constructed for each detection task: assessing pavement marking conditions and identifying roadside litter.

For pavement markings, we categorize them into two color-based classes, white and yellow, with further classification based on their condition. Following ASTM (2020) guidelines, pavement markings that are

[1] https://github.com/heartexlabs/labelImg

more than 50% faded or missing are labeled as “faded,” excluding those that are completely worn and unidentifiable. Faded markings are further divided into two categories:

- “y_faded”: Faded yellow markings, including single or double yellow curb and lane lines.
- “w_faded”: Faded white markings, including longitudinal lane lines, horizontal markings (e.g., crosswalks, stop lines), arrow symbols, and word-based markings.

For trash and litter, objects are annotated into four distinct classes:

- “leaves”: Vegetation and leaves accumulated along the roadside.
- “dirt”: Accumulations of dirt or soil on or near the pavement.
- “w_litter”: Litter in white or light colors (e.g., plastic, foam).
- “b_litter”: Litter in black or dark colors (e.g., used tires, rubber pieces, tree branches).

3.4 Model Development for Object Detection

YOLO is a state-of-the-art object detection algorithm known for its high accuracy and real-time performance, achieved through the simultaneous execution of object localization and classification (Redmon et al. 2016). In this study, YOLOv8 was selected as the deep learning architecture to train detection models for identifying a range of targeted transportation assets: faded pavement markings and roadside litter.

YOLO operates by predicting bounding boxes and their associated class probabilities directly from full images in a single pass, which contributes to its computational efficiency. In our implementation, detections are retained only if the associated confidence score exceeds a defined threshold of 0.35. Furthermore, YOLO provides the capability to extract cropped images of detected objects, enabling additional downstream processing such as object tracking, counting, and geolocation.

To evaluate the performance of the trained detection models, we employed three standard metrics: precision, recall, and F1-score, as defined in Equations (1)–(3). A predicted bounding box is considered a true positive (TP) if its Intersection over Union (IoU) with the corresponding ground truth box exceeds a predefined threshold of 0.65. Predictions that do not meet this IoU criterion are treated as false positives (FP), including boxes incorrectly identifying background objects. A false negative (FN) refers to an instance where a ground truth object (e.g., a faded white or yellow marking) was not detected by any prediction box.

- Precision indicates the reliability of the model in classifying objects as positive (i.e., TP / [TP + FP]).
$$precision = \frac{TP}{TP+FP} \quad (1)$$
- Recall measures the model’s ability to detect all relevant objects (i.e., TP / [TP + FN]).
$$recall = \frac{TP}{TP+FN} \quad (2)$$

- F1-score provides a balanced measure by calculating the harmonic mean of precision and recall, ensuring neither metric dominates the evaluation.
- $$F1 = \frac{2*precision*recall}{precision+recall} \quad (3)$$

To further assess overall detection performance, we report mean average precision at 50% IoU (mAP@50). This metric captures the mean of average precision values across all classes at a fixed IoU threshold of 0.50, offering a holistic view of detection effectiveness. Higher values in precision, recall, F1-score, and mAP@50 indicate better model performance.

3.5 Model Development for Object Counting

The core principle of object counting in video sequences lies in object tracking, which involves the continuous identification and monitoring of detected objects across consecutive frames by leveraging both spatial and temporal information. This process, illustrated in Figure 3, ensures that each object is consistently tracked throughout its appearance in the video by assigning and maintaining a unique identifier. Object tracking integrates two key components: object detection (identifying objects in individual frames) and motion prediction, which uses data association techniques to link objects across frames based on their location and movement (Wojke et al. 2017).

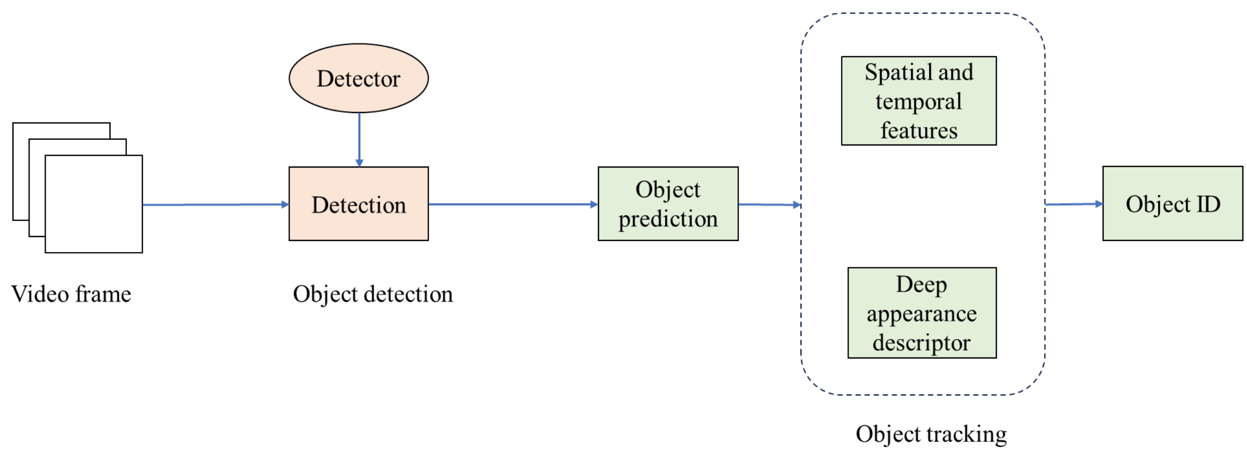

Figure 3. Illustration of Object Tracking

In this study, we adopt ByteTrack, which is an advanced tracking algorithm integrated into the YOLOv8 framework, as the core method for object tracking and counting. ByteTrack is known for its effective “tracking-by-detection” strategy, where object identities are preserved across frames even under challenging conditions such as occlusion, motion blur, or variable lighting. It achieves robust performance by refining associations between low-confidence and high-confidence detections, allowing it to maintain accurate object IDs across a range of real-world scenarios.

Given its reliability and compatibility with YOLO-based detection outputs, ByteTrack is selected as the foundational tracking method in this project. Once objects are detected in each frame, ByteTrack assigns consistent identifiers to each object, enabling the tracking of unique instances across video frames. By counting the number of unique object IDs observed over the duration of a video clip or within a specific road segment, we accurately estimate the total number of distinct objects, avoiding the duplicate counting that typically occurs with per-frame detection results.

3.6 Model Development for Object Geolocating

To accurately determine the geographic location of identified transportation asset issues (faded pavement markings or roadside litter), a geolocation model was developed to integrate video data with GPS coordinates. During data collection, an open-source mobile application called GPX Tracker was used to continuously record GPS data in parallel with video capture: latitude (*lat*), longitude (*lon*), elevation (*ele*), and the corresponding timestamp (*time_gpx*). While the video was recorded at 30 fps, the GPS data were sampled at a variable rate, depending on the moving speed, with a maximum of two samples per second.

This discrepancy in sampling frequency between the video and GPS data introduces challenges in directly aligning each video frame with a unique geolocation. The GPX timestamps have a precision limited to seconds. In contrast, the video frame timestamps are accurate to milliseconds, often leading to multiple GPS records sharing identical timestamps and a mismatch between video frames and location points.

To address this issue, we developed a time-based interpolation model to estimate the precise geolocation of each identified faded pavement marking using nearby GPS data points. The process includes the following steps:

(1) The issues of pavement markings are detected by the detection model, with the output of the class (*cls*) and cropping image (*img*) of the faded markings. The timestamp of identified marking issues (*time_img*) can be inferred based on the number of frames and the video start recording time.

(2) Identified marking issues are matched with their nearest GPX records based on their timestamps (i.e., *time_img* and *time_gpx*, respectively).

(3) Due to the short distance within a few seconds, the location (latitude and longitude) of each identified marking issue is estimated based on the second closest records before and after the corresponding matched GPX record using linear interpolations based on the time differences, as illustrated in Figure 4 and equations (4) and (5). For example, $time_{img}$ =10:01:27.96 (accuracy up to milliseconds) matches with $time_{gpx}$ at $time_{idx}$ =10:01:27 (accuracy up to seconds). The second closest points before and after the $time_{idx}$ are $time_{idx-2}$ and $time_{idx+2}$, i.e., 10:01:26 (*time1*) and 10:01:28 (*time2*) in the example. Based on the time differences between $time_{img}$, $time_{idx+2}$, and $time_{idx-2}$, the location of the identified marking issue at $time_{img}$ (lat_{img} and lon_{img}) can be estimated.

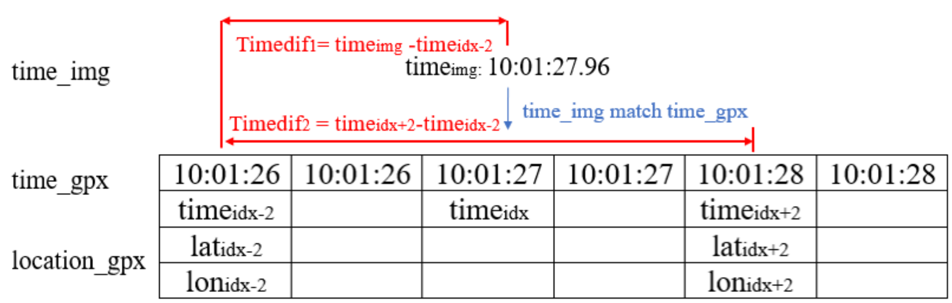

Figure 4. Illustration of Geolocation Estimation

$$lat_{img} = lat_{idx-2} + (lat_{idx+2} - lat_{idx-2}) * \frac{(time_{img} - time_{idx-2})}{time_{idx+2} - time_{idx-2}} \quad (4)$$

$$lon_{img} = lon_{idx-2} + (lon_{idx+2} - lon_{idx-2}) * \frac{(time_{img} - time_{idx-2})}{time_{idx+2} - time_{idx-2}} \quad (5)$$

Where $time_i$ is indicated in Figure 4, lon_i and lat_i are the longitude and latitude of the location at $time_i$.

4. RESULTS AND DISCUSSION

4.1 Results of Object Detection

Our detection models were trained individually for each task using a desktop system equipped with an 8-Core CPU, an NVIDIA GeForce RTX 3070 GPU, and 64 GB of memory. The training was conducted based on the YOLOv8 architecture to ensure optimal performance and accuracy for the tasks in this project.

4.1.1 Identifying Pavement Marking Issues

To develop the object detection model for identifying faded pavement markings, 7,000 labeled images were prepared and divided into training and validation datasets, with 5,600 images allocated for training and 1,400 for validation, following an 8:2 split. The model was trained using the YOLOv8 architecture. The training process reached optimal performance at epoch 173, where both loss and accuracy metrics stabilized.

Figure 5 displays the training and validation curves, illustrating steady convergence and consistent improvements in model performance over time. As shown in Table 1, the final trained model achieved an average precision of 0.88, a recall of 0.89, an F1-score of 0.89, and a mean average precision at an IoU threshold of 0.50 (mAP@50) of 0.93. These evaluation metrics indicate that the model performs reliably in detecting and classifying faded pavement markings.

Table 1. Accuracy Metrics of Identifying Faded Pavement Markings

Class	Precision	Recall	F1-score	mAP50
all	0.88	0.89	0.88	0.93
w_faded	0.86	0.86	0.86	0.91
y_faded	0.91	0.91	0.91	0.95

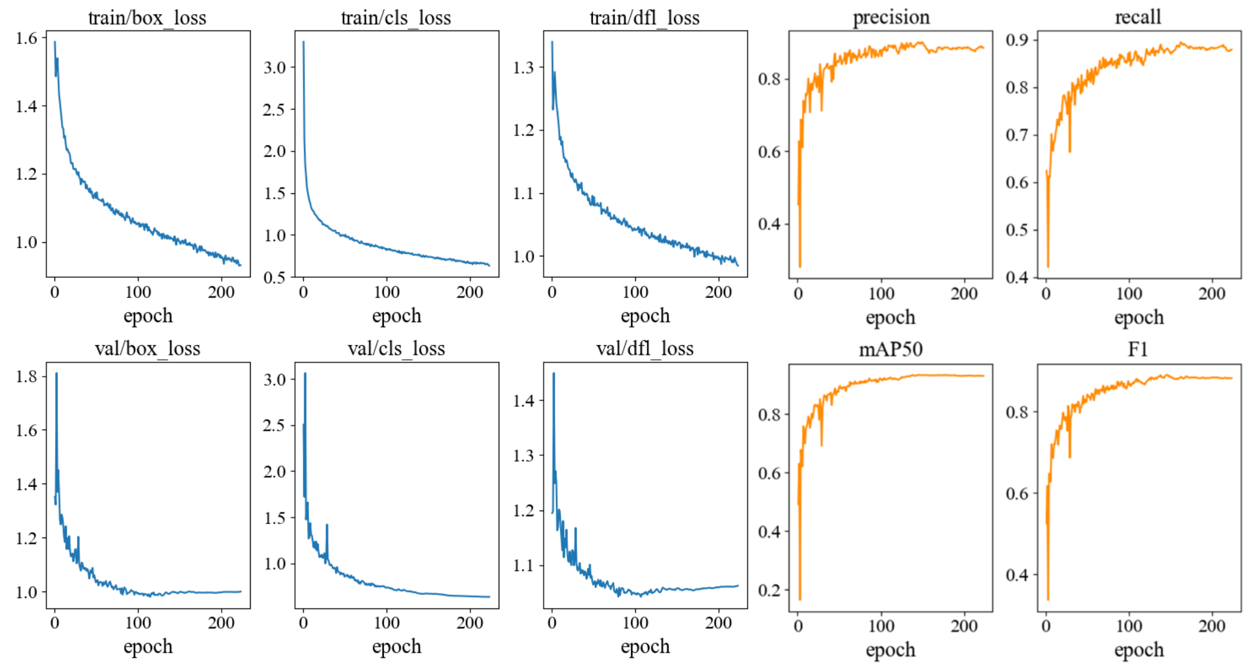

(a) Training Loss (b) Accuracy Metrics

Figure 5. Training Loss and Accuracy Metrics of Model for Pavement Marking Issue
Note: epoch refers to one complete pass through the entire training dataset by the model.

Figure 6 presents several example cases that illustrate the model's effectiveness in detecting both white and yellow faded pavement markings under various conditions. The white faded markings identified by the model include single-lane lines, word-based road messages, and directional arrow markings, while the yellow faded markings comprise both single- and double-lane lines. The model demonstrated consistent accuracy across these different types of markings, highlighting its versatility. In addition, the model's robustness was evaluated using tiger markings, which are road markings characterized by alternating white and black striping to enhance visibility. These high-contrast designs posed no challenge to the model, which successfully identified the white faded portions without being misled by the adjacent black elements.

(a) Faded White Lane Marking-1

(b) Faded White Lane Marking-2

(c) Faded White Word Marking-1

(d) Faded White Word Marking-2

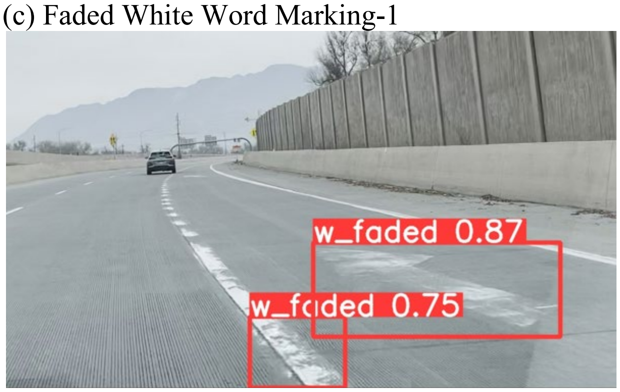

(e) Faded White Arrow and Lane Markings-1

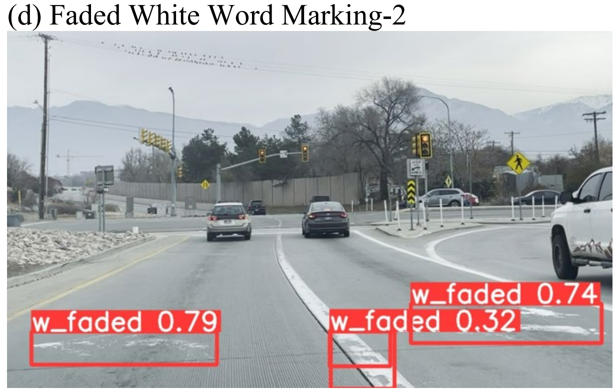

(f) Faded White Arrow and Lane Markings-2

(g) Faded Yellow Single Lane Marking

(h) Faded Yellow Double Lane Marking

Figure 6. Examples of Detection Results of Pavement Marking Issues

However, despite the overall strong performance, the developed model for pavement marking issue detection exhibits several noteworthy limitations. A primary challenge lies in the model's inability to differentiate between genuinely faded markings and previously removed markings. For instance, as illustrated in Figure 7 (a), the model incorrectly classifies a previously removed marking, with only faint residue remaining on the pavement surface, as a faded marking, which indicates a lack of contextual sensitivity in distinguishing historical markings from currently degraded ones.

Another limitation is the model's difficulty in detecting markings that are severely faded or nearly imperceptible. These markings often lack sufficient contrast with the pavement background, making them particularly challenging for the model to identify. Figure 7 (b) provides an example of such a failure, where the model overlooks a nearly invisible marking due to its extremely low visual salience.

Furthermore, the model is prone to false positives when encountering objects that share visual similarities with faded markings. As shown in Figure 7 (c) through Figure 7 (e), elements such as snow patches, pavement discolorations, or other unrelated white artifacts are occasionally misclassified as faded markings. These misidentifications underscore the difficulty of differentiating between relevant and irrelevant features in complex roadway environments, particularly under variable lighting, shadow, and weather conditions.

(a) Previously Removed Marking

(b) Almost Fully Faded Marking

(c) Other Marking

(d) Other/Previous Marking

(e) Other Object with White Piece-1

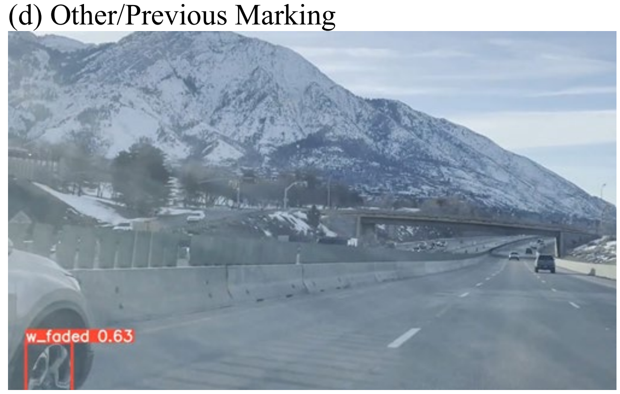

(f) Other Object with White Piece-2

Figure 7. Examples of Wrong Detection of Pavement Marking Issues

4.1.2 Identifying Litter & Trash

To train the model for litter and trash detection, 6,250 annotated images were compiled, with 5,000 images allocated for training and 1,250 for validation. The dataset has a diverse array of roadside debris, including both organic materials (e.g., leaves and dirt) and inorganic items (e.g., plastic wrappers, foam, and used vehicle components). The model was trained using the YOLOv8 architecture, with performance monitored throughout the training process via loss and accuracy curves. As shown in Figure 8, the training reached optimal performance at epoch 219, indicating convergence and stability in learning.

The resulting model demonstrated consistently strong performance across multiple evaluation metrics. It achieved a precision of 0.86, a recall of 0.82, an F1-score of 0.84, and a mean average precision at 50% IoU (mAP50) of 0.88, as detailed in Table 2. These results underscore model's high reliability in detecting various forms of roadside litter across diverse visual contexts.

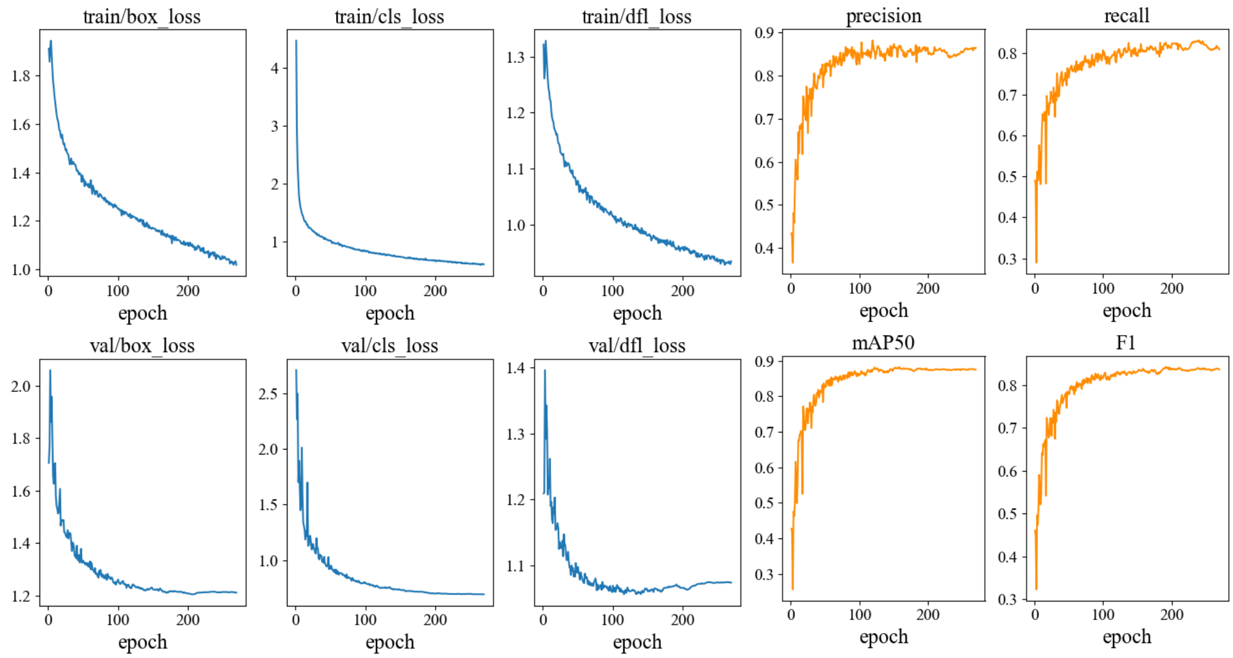

(a) Training Loss (b) Accuracy Metrics

Figure 8. Training Loss and Accuracy Metrics of Model for Litter & Trash

Table 2. Accuracy Metrics of Trash & Litter

Class	**Precision**	**Recall**	**F1 score**	**mAP@50**
all	0.86	0.82	0.84	0.88
leaves	0.82	0.80	0.81	0.84
dirt	0.86	0.82	0.84	0.88
w_litter	0.87	0.82	0.84	0.90
b_litter	0.88	0.83	0.85	0.90

Figure 9 presents a set of representative examples showcasing the detection capabilities in real-world scenarios. These examples include successful identification of black litter (e.g., discarded tires), white litter (e.g., plastic debris and foam objects), as well as natural waste materials such as dried leaves and soil. The diversity of detected object types reflects the robustness and adaptability of the model in recognizing distinctive textures, colors, and shapes associated with roadside waste, even under fluctuating lighting and environmental conditions.

(a) Black Litter on the Highway-1

(b) Black Litter on the Highway-2

(c) White Litter on the Highway-1

(d) White Litter on the Highway-2

(e) Leaves on the Street Road-1

(f) Leaves on the Street Road-2

(g) Dirt on the Highway-1

(h) Dirt on the Highway-2

Figure 9. Examples of Detection Results of Trash & Litter

Despite its overall strong performance, the litter detection model exhibits several limitations that affect its classification accuracy in certain scenarios. A key issue involves the occurrence of false positives, where non-litter pavement features are incorrectly identified as litter. Specifically, elements such as asphalt patches, oil stains, and pavement cracks are occasionally misclassified as black litter ("b_litter"), as illustrated in Figure 10 (a)–(d). Likewise, white pavement marks or gray-colored objects, such as small stones or concrete chips, can be erroneously detected as white litter ("w_litter"), as shown in Figure 10 (e) and (f). These misclassifications typically stem from visual similarities in color, shape, and surface texture with litter.

(a) Asphalt Patches-1

(b) Asphalt Patches-2

(c) Shadow

(d) Black Taint

(e) Other Object with White Piece-1

(f) Other Object with White Piece-2

Figure 10. Examples of Wrong Detection of Trash & Litter

4.2 Results of Object Counting

Based on the developed object counting algorithm, a case study was conducted using a 35-minute dashcam video recorded along a highway corridor stretching from Salt Lake City to Lehi, covering approximately 35 miles (56.4 kilometers [km]), as illustrated in Figure 11. The video served as a comprehensive test dataset and captured a range of key transportation assets relevant to this study, including pavement markings and roadside litter.

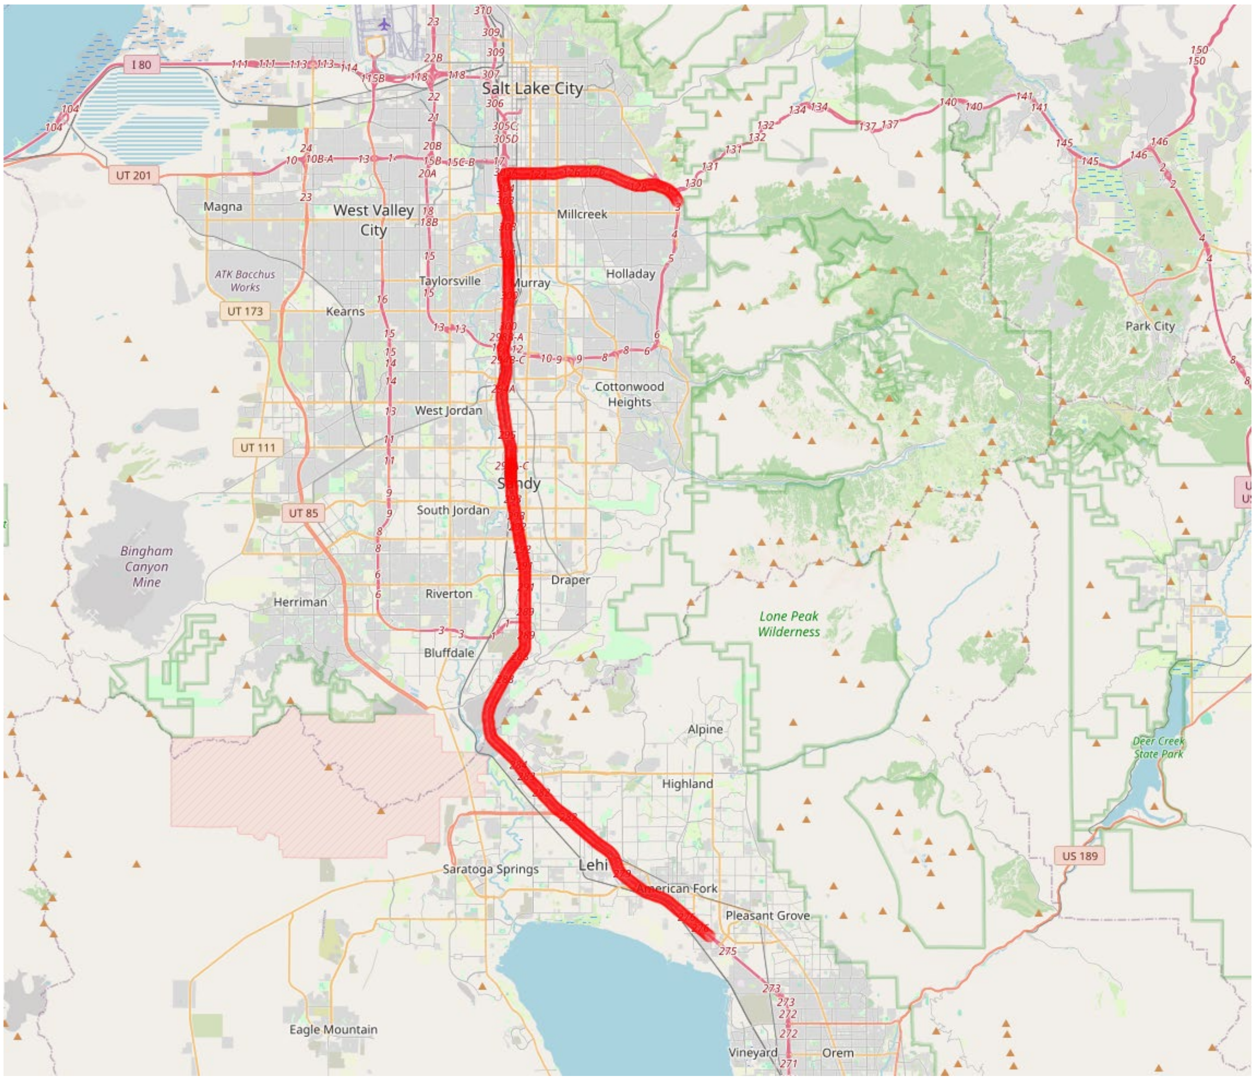

Figure 11. Case Study Area

The counting results are dynamically overlaid and updated in the top left corner of each video frame throughout the case study, as illustrated in Figure 12. Across the test highway route, the model successfully identifies 518 white faded pavement markings. It also detects 33 instances of white litter, 263 instances of black litter, and 102 instances of dirt. These comprehensive results underscore the capability of the developed AI-based system to accurately detect, classify, and quantify different maintenance-related concerns (i.e., faded markings and various types of roadside litter) across extended roadway segments.

(a) Counting Results of Pavement Marking Issues

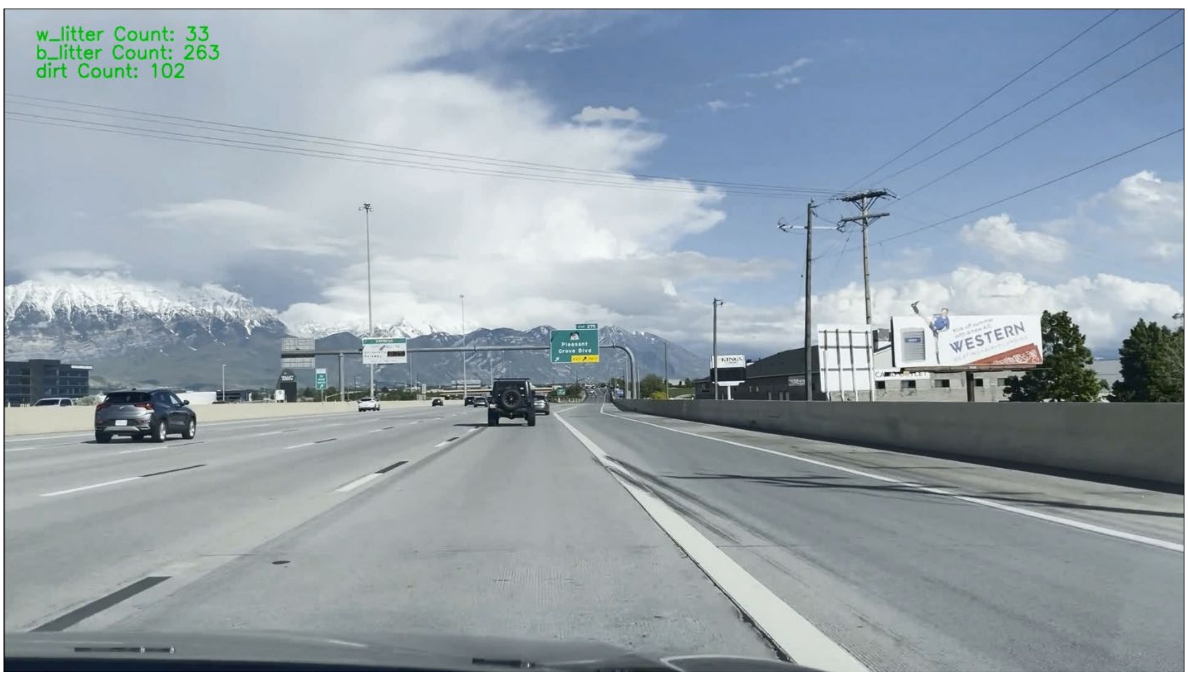

(b) Counting Results of Road Litter

Figure 12. Results of Object Counting

4.3 Results of Object Geolocation

4.3.1 Accuracy of Geolocation Model

A geolocation model was developed to accurately determine the positions of identified objects by utilizing data obtained from a GPX tracker. Given the challenges associated with acquiring precise ground truth data for the exact locations of these objects in real-world settings, the performance of the model was evaluated through a cross-section validation method. Specifically, from a dataset consisting of 3,500 continuous GPS records collected over approximately 24 miles (39 km) along I-15, a random sample of 175 points, representing 5% of the total dataset, was selected as validation points, as illustrated in Figure 13 (a). The remaining GPS records were used to estimate the positions of the validation points, and the predicted coordinates were then compared to their corresponding true locations obtained from the GPX tracker.

Model accuracy was quantified by calculating distance error, defined as the distance between the predicted and actual locations. The distribution of these distance errors is presented in Figure 13 (b). The results indicate that the prediction error ranged from less than 0.01 meters to a maximum of 8.16 meters, with an average error of only 0.27 meters. These outcomes demonstrate that the geolocation model delivers a high level of geolocation accuracy, making it a reliable tool for field applications involving the localization of pavement marking issues and detected roadside litter.

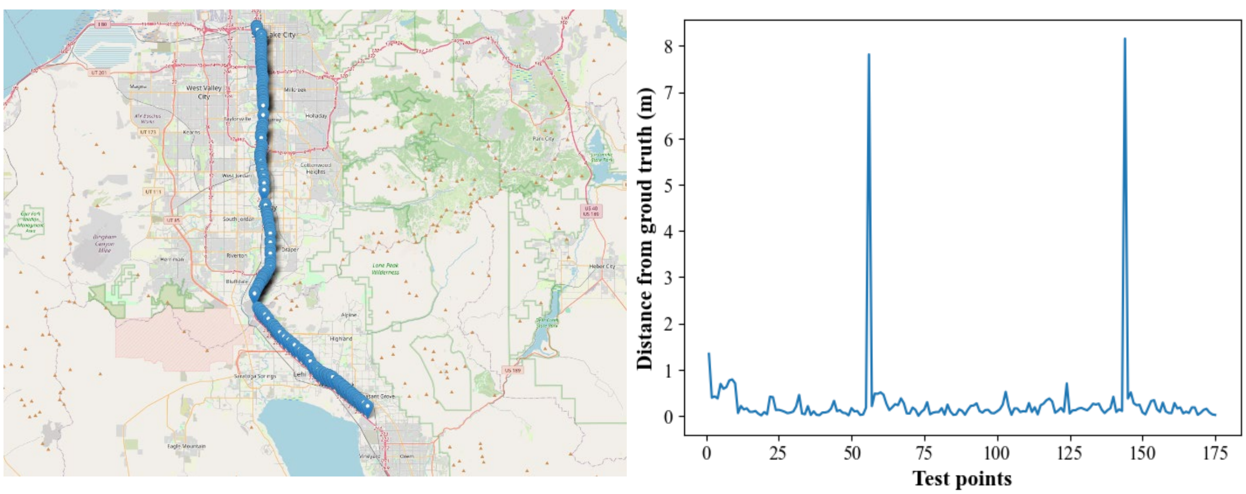

(a) Validation Points of Geolocation Model (b) Locating Error of Geolocation Model

Figure 13. Validation Points and Locating Error

To effectively visualize the information associated with identified objects, such as their geolocation, classification, and inspection time, an interactive mapping interface was developed using the Python package Folium. This interface is designed to present data in a clear and intuitive format, which enables users to quickly interpret and analyze results. As shown in Figure 14, the generated map includes a pop-up window for each identified object. Each pop-up displays the object's latitude and longitude coordinates, classification label, inspection timestamp, and the estimated driving speed, which is derived from the geolocation model. To provide a visual reference, the largest available cropped image of the object is also embedded in the pop-up.

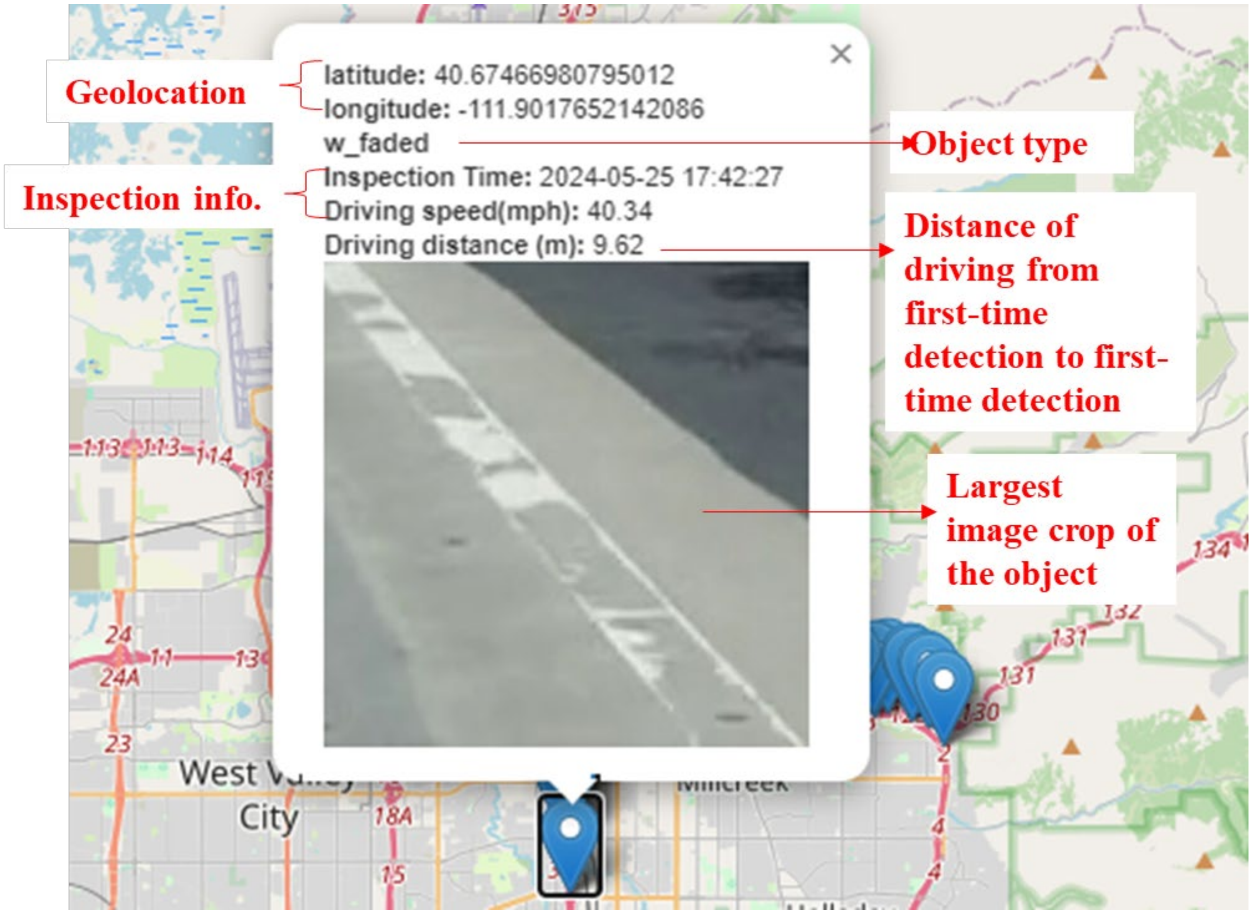

Figure 14. Example of Visualization

For continuous objects such as faded pavement issues, the interface also calculates the driving distance between the first and last detection points. While this measurement does not reflect the actual physical length of the object, it serves as a valuable comparative metric for evaluating the relative extents of different objects along the route. This visualization tool offers a practical and user-friendly means of reviewing inspection results, supporting more informed decision-making in transportation asset management.

4.4 Case Studies of Object Geolocation and Visualization

The geolocations of pavement marking issues detected along the test route are presented in Figure 15 (a), with each marker indicating the inferred location of an identified issue. This spatial visualization enables quick identification of problem areas along the corridor. Figures 15 (b)–(d) provide illustrative examples of individual white faded marking detections, showcasing the associated details such as location, inspection time, and cropped imagery for enhanced clarity.

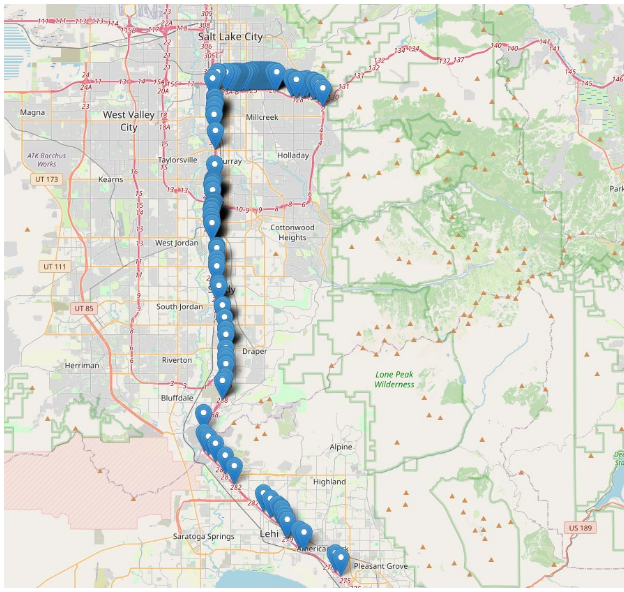

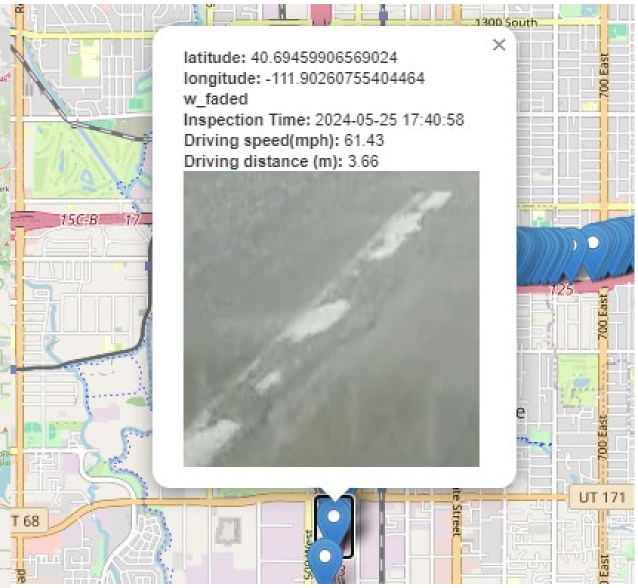

(a) Geolocation of Marking Issues

(b) Visualization of Marking Issue-1

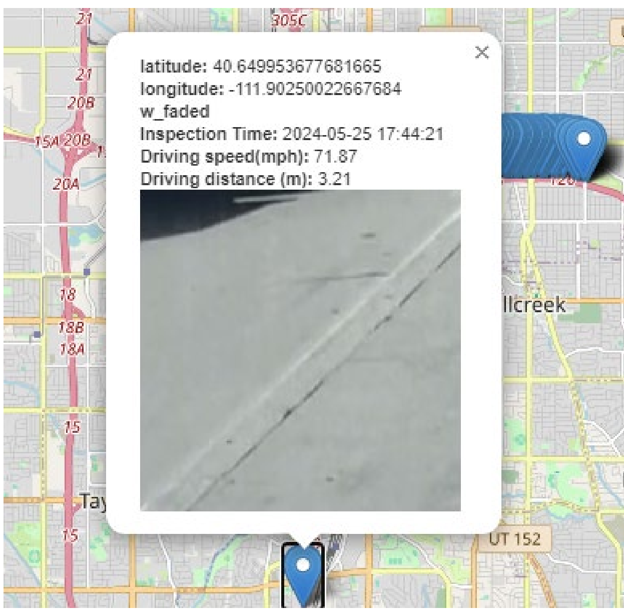

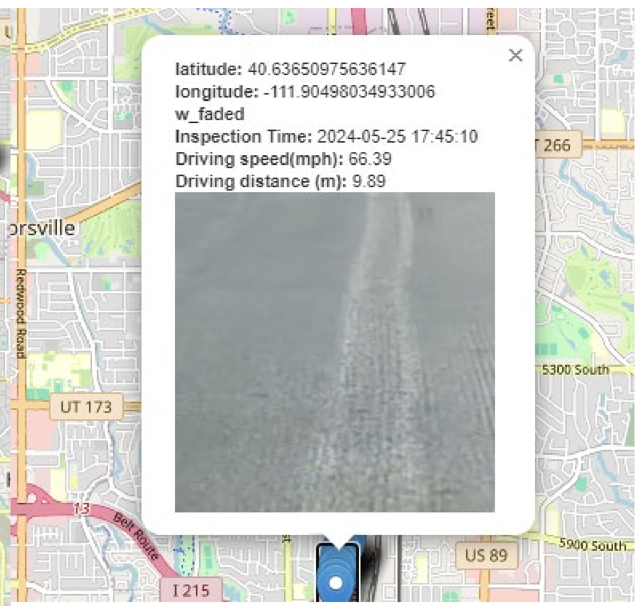

(c) Visualization of Marking Issue-2

(d) Visualization of Marking Issue-3

Figure 15. Geolocation and Visualization of Pavement Marking Issues

Similarly, Figure 16 (a) depicts the geolocations of identified litter and trash along the same route, offering a comprehensive spatial overview of debris distribution. Figure 16 (b)–(d) present representative examples of different litter categories, such as white litter, black litter, and organic debris, complete with relevant details—including geographic coordinates, classification type, inspection timestamp, and high-resolution cropped images—to assist with field verification and decision-making.

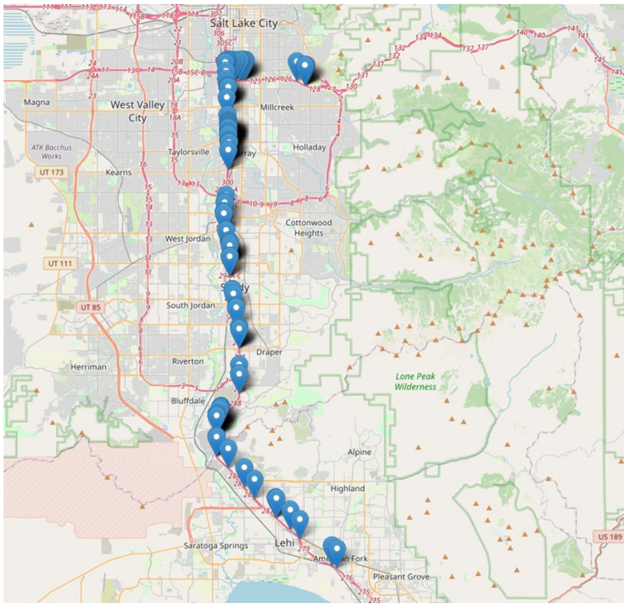

(a) Geolocation of Litter & Trash

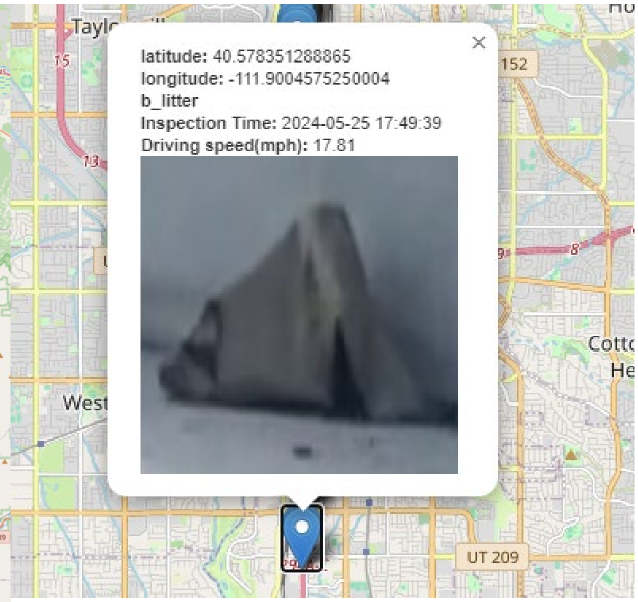

(b) Visualization of Black Litter

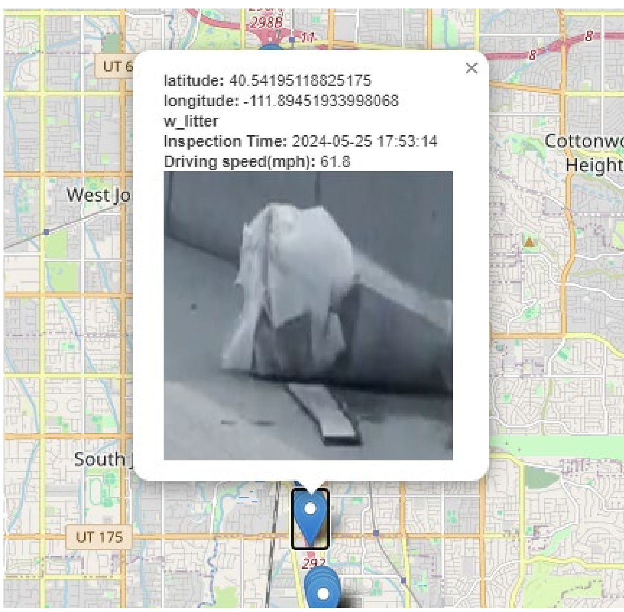

(c) Visualization of White Litter

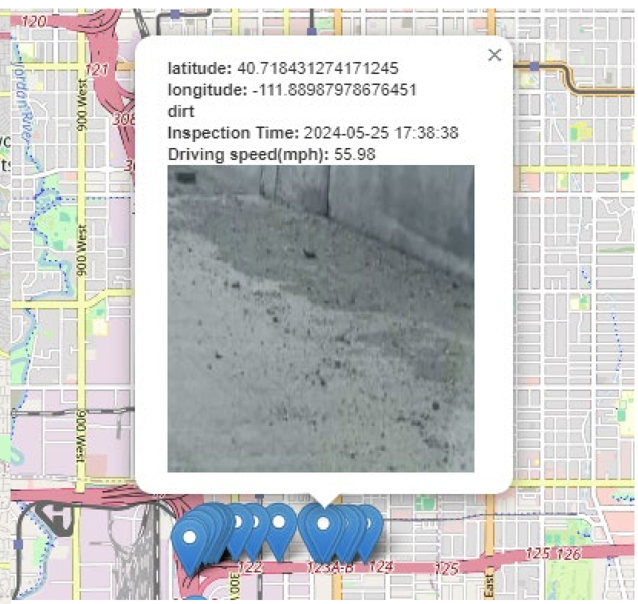

(d) Visualization of Dirt

Figure 16 Geolocation and Visualization of Litter & Trash

5. CONCLUSIONS

5.1 Summary of Findings

This project builds upon the AI packages developed in our prior MPC Project (Report No. MPC-668) by focusing on two targeted transportation assets: pavement marking issues and roadside litter. The training datasets are expanded to over 6,000 images for each asset type, and the models are validated on a larger scale to improve robustness and generalizability. In addition, new functionalities are implemented, including object counting within defined road segments or video clips, geolocation estimation using a mobile phone-based GPS tracker, and interactive visualization of detected object details through a Folium-based mapping interface. Specifically, we increase the training datasets for each project task, as follows:

(1) Pavement Marking Issues: 7,000 images are used to train the AI model for detecting pavement marking issues. These issues are classified into two classes by color: faded yellow markings (“y_faded”) and faded white markings (“w_faded”). The AI model achieves a precision of 0.88, a recall of 0.89, an F1 score of 0.88, and mAP@50 of 0.93.

(2) Trash and Litter Identification: A dataset of 6,250 images is utilized to develop the AI model for trash and litter identification. The model categorizes litter into four major classes: leaves, dirt, white litter (“w_litter”), and black litter (“b_litter”). The model attains a precision of 0.86, recall of 0.82, F1 score of 0.84, and a mAP@50 of 0.88.

Also, building upon object tracking, a specialized counting model is developed to quantify the identified objects within video clips or road sections. This model has been effectively employed to count the number of identified objects, specifically focusing on the objects of interest within this project. Notably, this model has been applied successfully on a test highway route in Utah, showcasing its practical application and reliability in real-world scenarios.

Moreover, expanding the capabilities further, a model is developed to estimate the locations of identified objects using time-based interpolations with the assistance of a mobile phone-based GPS tracker (GPX). This geolocation model, integrated with the detection model, demonstrates remarkable accuracy with an average distance error of only 0.27 meters, which is validated through cross-section analysis. This precision underscores the reliability and effectiveness of the geolocation model in pinpointing the positions of identified objects with high fidelity, as demonstrated in a case study of a test highway route in Utah.

Finally, utilizing Folium (a Python package), an interactive interface is developed to visually represent the detailed information of these identified objects on a map. The display information consists of geolocation coordinates (i.e., latitude and longitude), object class, inspection time, and cropped images. The developed interface serves as a valuable tool for transportation authorities to comprehensively understand and assess the conditions of transportation assets through intuitive visualizations.

This project expands upon our previous work to improve the capabilities of models in identifying targeted transportation assets while introducing additional functionalities, such as counting, geolocating, and visualizing these detected objects. These advancements facilitate more frequent inspections of transportation assets, contributing to the continual enhancement of road user safety.

5.2 Implementation Plan

The delivered AI package consists of two trained detection models tailored to different transportation asset issues: faded pavement markings and roadside litter. It also includes models for object counting, geolocation, and visualization. All models were developed in Python and require a configured computational environment compatible with the YOLOv8 architecture.

To implement the AI package in practice, users should follow a sequence of key steps. First, during data collection, a mobile device equipped with a GPS tracking application (e.g., GPX Tracker) should be mounted on the vehicle's windshield. While driving, the device should simultaneously record videos and GPS data. Care must be taken to adjust the camera angle, device position, and lane alignment to ensure clear visibility of the target objects in the footage.

Next, in the data analysis phase, users must export both the recorded videos and their associated GPS logs. The appropriate trained detection model should then be selected and executed in conjunction with the object counting model. This process will detect and count relevant objects, save cropped images, and annotate video frames, with object counts displayed in the top-left corner of each frame. After detection, the geolocation model should be applied to estimate the precise locations of each identified object, followed by interactive visualization on a map.

By following this implementation process, users can deploy the AI package effectively to detect, count, geolocate, and visualize transportation assets. This workflow provides a powerful and scalable tool to support infrastructure evaluation and enhance the quality and efficiency of transportation asset management.

5.3 Limitations

While the developed system shows strong performance, several limitations remain. The object counting model does not include objects detected only once in a video sequence due to tracking constraints. In the visualization interface, pop-up windows can display details for only a single object per frame; if multiple objects share the largest cropped image in the same frame, only one is shown on the map. For continuous objects, the reported driving distance reflects the span between first and last detections rather than the true physical length. In addition, classification errors may still occur in challenging cases, such as distinguishing faded markings from previously removed markings or differentiating litter from pavement surface artifacts.

Future improvements could address these issues by:

- Refining tracking methods to capture single-frame detections without inflating false positives
- Enabling the visualization interface to display multiple objects from the same frame
- Incorporating offset distance adjustments to estimate true object lengths for continuous assets
- Expanding datasets to include challenging and visually similar scenarios to reduce misclassification rates

6. REFERENCES

Alzraiee, Hani, Andrea Leal Ruiz, and Robert Sprotte. 2021. "Detecting of Pavement Marking Defects Using Faster R-CNN." *Journal of Performance of Constructed Facilities* 35 (4): 04021035. https://doi.org/10.1061/(ASCE)CF.1943-5509.0001606.

ASTM. 2020. *Standard Practice for Evaluating Degree of Pavement Marking Line Wear*. Nos. D913-20. https://compass.astm.org/document/?contentcode=ASTM%7CD0913-20%7Cen-US.

Bronuela-Ambrocio, Lea B., and Emilson Ryan D. Antes. 2023. "Assessment of Longitudinal Pavement Marking Conditions Based on Deep Learning." *Philippine Engineering Journal* 44 (1): 65–88.

Chamberlin, Robert, Troy Lincoln, Steve Stein, "Chamberlin Analytics, and ER Planning." 2021. *Sources of Highway Litter in Utah*. UT-21.01. https://rosap.ntl.bts.gov/view/dot/60313.

Du, Yuchuan, Ning Pan, Zihao Xu, Fuwen Deng, Yu Shen, and Hua Kang. 2020. "Pavement Distress Detection and Classification Based on YOLO Network." *International Journal of Pavement Engineering*, January 24, 1–14. https://doi.org/10.1080/10298436.2020.1714047.

Farhadmanesh, Mohammad, Chandler Cross, Ali H. Mashhadi, Abbas Rashidi, and Jessica Wempen. 2021. "Highway Asset and Pavement Condition Management Using Mobile Photogrammetry." *Transportation Research Record* 2675 (9): 296–307. https://doi.org/10.1177/03611981211001855.

Fernández-Sanjurjo, Mauro, Brais Bosquet, Manuel Mucientes, and Víctor M. Brea. 2019. "Real-Time Visual Detection and Tracking System for Traffic Monitoring." *Engineering Applications of Artificial Intelligence* 85 (October): 410–20. https://doi.org/10.1016/j.engappai.2019.07.005.

FHWA. 2009. *Manual on Uniform Traffic Control Devices (MUTCD) 2009 Edition*. https://mutcd.fhwa.dot.gov/pdfs/2009/pdf_index.htm.

Ghosh, Rohit, and Omar Smadi. 2021. "Automated Detection and Classification of Pavement Distresses Using 3D Pavement Surface Images and Deep Learning." *Transportation Research Record*, April 28, 03611981211007481. https://doi.org/10.1177/03611981211007481.

Gómez, Braulio Félix, Lim Yi, Balakrishnan Ramalingam, et al. 2022. "Deep Learning Based Litter Identification and Adaptive Cleaning Using Self-Reconfigurable Pavement Sweeping Robot." *2022 IEEE 18th International Conference on Automation Science and Engineering (CASE)*, August 2301–6. https://doi.org/10.1109/CASE49997.2022.9926489.

Gopalakrishnan, Kasthurirangan, Siddhartha K. Khaitan, Alok Choudhary, and Ankit Agrawal. 2017. "Deep Convolutional Neural Networks with Transfer Learning for Computer Vision-Based Data-Driven Pavement Distress Detection." *Construction and Building Materials* 157 (December): 322–30. https://doi.org/10.1016/j.conbuildmat.2017.09.110.

Iparraguirre, Olatz, Nagore Iturbe-Olleta, Alfonso Brazalez, and Diego Borro. 2022. "Road Marking Damage Detection Based on Deep Learning for Infrastructure Evaluation in Emerging Autonomous Driving." *IEEE Transactions on Intelligent Transportation Systems* 23 (11): 22378–85. https://doi.org/10.1109/TITS.2022.3192916.

Kang, Kyubyung, Donghui Chen, Cheng Peng, Dan Koo, Taewook Kang, and Jonghoon Kim. 2020. "Development of an Automated Visibility Analysis Framework for Pavement Markings Based on the Deep Learning Approach." *Remote Sensing* 12 (22): 22. https://doi.org/10.3390/rs12223837.

Kargah-Ostadi, Nima, Ammar Waqar, and Adil Hanif. 2020. "Automated Real-Time Roadway Asset Inventory Using Artificial Intelligence." *Transportation Research Record* 2674 (11): 220–34. https://doi.org/10.1177/0361198120944926.

Karimi, Kasra, and Ardeshir Faghri. 2021. "The Issues of Roadside Litter: A Review Paper." *Current Urban Studies* 9 (4): 4. https://doi.org/10.4236/cus.2021.94046.

Karsten, Lloyd, Suliman Gargoum, Mohamed Saleh, and Karim El-Basyouny. 2021. "Automated Framework to Audit Traffic Signs Using Remote Sensing Data." *Journal of Infrastructure Systems* 27 (3): (ASCE)IS.1943-555X.0000618, 04021014. https://doi.org/10.1061/(ASCE)IS.1943-555X.0000618.

Kawano, Makoto, Kazuhiro Mikami, Satoshi Yokoyama, Takuro Yonezawa, and Jin Nakazawa. 2017. "Road Marking Blur Detection with Drive Recorder." *2017 IEEE International Conference on Big Data (Big Data)*, December, 4092–97. https://doi.org/10.1109/BigData.2017.8258427.

Koch, Christian, Gauri M. Jog, and Ioannis Brilakis. 2013. "Automated Pothole Distress Assessment Using Asphalt Pavement Video Data." *Journal of Computing in Civil Engineering* 27 (4): 370–78. https://doi.org/10.1061/(ASCE)CP.1943-5487.0000232.

Kong, Wanyue, Teng Zhong, Xin Mai, Shuliang Zhang, Min Chen, and Guonian Lv. 2022. "Automatic Detection and Assessment of Pavement Marking Defects with Street View Imagery at the City Scale." *Remote Sensing* 14 (16): 16. https://doi.org/10.3390/rs14164037.

Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. 2017. "ImageNet Classification with Deep Convolutional Neural Networks." *Communications of the ACM* 60 (6): 84–90. https://doi.org/10.1145/3065386.

Kuang, Biao, and Jianli Chen. 2024a. "Mobile Phone-Based Artificial Intelligence Development for Maintenance Asset Management." Mountain-Plains Consortium. https://rosap.ntl.bts.gov/view/dot/77215.

Kuang, Biao, and Jianli Chen. 2024b. "Mobile-Phone-Based Artificial Intelligence Package Development and Validation in Large Scale for Maintenance Asset Management." Utah Department of Transportation. https://rosap.ntl.bts.gov/view/dot/84576.

Kuang, Biao, Jianli Chen, and Yuqing Hu. 2024. "Identification and Geolocation of Pavement Marking Issues Based on Artificial Intelligence and Mobile Phone." *Computing in Civil Engineering 2024*.

Lee, Sangyum, and Byoung Hooi Cho. 2023. "Evaluating Pavement Lane Markings in Metropolitan Road Networks with a Vehicle-Mounted Retroreflectometer and AI-Based Image Processing Techniques." *Remote Sensing* 15 (7): 7. https://doi.org/10.3390/rs15071812.

Liu, Ying, Zhishan Ge, Guoyun Lv, and Shikai Wang. 2018. "Research on Automatic Garbage Detection System Based on Deep Learning and Narrowband Internet of Things." *Journal of Physics: Conference Series* 1069 (August): 012032. https://doi.org/10.1088/1742-6596/1069/1/012032.

Ma, Duo, Hongyuan Fang, Niannian Wang, Chao Zhang, Jiaxiu Dong, and Haobang Hu. 2022. "Automatic Detection and Counting System for Pavement Cracks Based on PCGAN and YOLO-MF." *IEEE Transactions on Intelligent Transportation Systems* 23 (11): 22166–78. https://doi.org/10.1109/TITS.2022.3161960.

Maeda, Hiroya, Yoshihide Sekimoto, Toshikazu Seto, Takehiro Kashiyama, and Hiroshi Omata. 2018. "Road Damage Detection and Classification Using Deep Neural Networks with Smartphone Images." *Computer-Aided Civil and Infrastructure Engineering* 33 (12): 1127–41. https://doi.org/10.1111/mice.12387.

Majidifard, Hamed, Peng Jin, Yaw Adu-Gyamfi, and William G. Buttlar. 2020. "Pavement Image Datasets: A New Benchmark Dataset to Classify and Densify Pavement Distresses." *Transportation Research Record* 2674 (2): 328–39. https://doi.org/10.1177/0361198120907283.

Mandal, Vishal, Abdul Rashid Mussah, and Yaw Adu-Gyamfi. 2020. "Deep Learning Frameworks for Pavement Distress Classification: A Comparative Analysis." *2020 IEEE International Conference on Big Data (Big Data)*, December, 5577–83. https://doi.org/10.1109/BigData50022.2020.9378047.

Mandhati, Sriram Reddy, N. Lakmal Deshapriya, Chatura Lavanga Mendis, et al. 2024. "pLitterStreet: Street Level Plastic Litter Detection and Mapping." arXiv:2401.14719. Preprint, arXiv, January 26. https://doi.org/10.48550/arXiv.2401.14719.

Mullaney, Jennifer, and Terry Lucke. 2014. "Practical Review of Pervious Pavement Designs." *CLEAN – Soil, Air, Water* 42 (2): 111–24. https://doi.org/10.1002/clen.201300118.

Okpe, JONAH Bameyi, and Francis Enejo Idachaba. 2023. "Development of an Iot Model for Pothole Detection , Counting and Geolocation." SSRN Scholarly Paper No. 4671532. Rochester, NY, December 21. https://doi.org/10.2139/ssrn.4671532.

Peng, Jun, Weidong Wang, Wenbo Hu, et al. 2024. "Automated Detection and Quantification of Pavement Cracking around Manhole." *Engineering Applications of Artificial Intelligence* 130 (April): 107778. https://doi.org/10.1016/j.engappai.2023.107778.

Pike, Adam M., Lance D. Ballard, and Paul J. Carlson. 2011. "Evaluation of Retroreflectivity Measurement Techniques for Profiled and Rumble Stripe Pavement Markings." *Transportation Research Record* 2258 (1): 80–87. https://doi.org/10.3141/2258-10.

Rad, Mohammad Saeed, Andreas Von Kaenel, Andre Droux, et al. 2017. "A Computer Vision System to Localize and Classify Wastes on the Streets." In *Computer Vision Systems*, edited by Ming Liu, Haoyao Chen, and Markus Vincze, vol. 10528. Lecture Notes in Computer Science. Springer International Publishing. https://doi.org/10.1007/978-3-319-68345-4_18.

Ramalingam, Balakrishnan, Abdullah Aamir Hayat, Mohan Rajesh Elara, et al. 2021. "Deep Learning Based Pavement Inspection Using Self-Reconfigurable Robot." *Sensors* 21 (8): 8. https://doi.org/10.3390/s21082595.

Redmon, Joseph, Santosh Divvala, Ross Girshick, and Ali Farhadi. 2016. "You Only Look Once: Unified, Real-Time Object Detection." *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June, 779–88. https://doi.org/10.1109/CVPR.2016.91.

Sassani, Alireza, Omar Smadi, and Neal Hawkins. 2021. “Developing Pavement Marking Management Systems: A Theoretical Model Framework Based on the Experiences of the US Transportation Agencies.” *Infrastructures* 6 (2): 18. https://doi.org/10.3390/infrastructures6020018.

Sayyad, Shahzia, Shaily Parmar, Mrinalini Jadhav, and Karan Khadayate. 2020. “Real-Time Garbage, Potholes and Manholes Monitoring System Using Deep Learning Techniques.” *2020 4th International Conference on Trends in Electronics and Informatics (ICOEI)(48184)*, June, 826–31. https://doi.org/10.1109/ICOEI48184.2020.9143030.

Souza, Vinicius M.A., Rafael Giusti, and Antônio J.L. Batista. 2018. “Asfault: A Low-Cost System to Evaluate Pavement Conditions in Real-Time Using Smartphones and Machine Learning.” *Pervasive and Mobile Computing* 51 (December): 121–37. https://doi.org/10.1016/j.pmcj.2018.10.008.

Sun, Yazhen, Haixiang Tang, and Huaizhi Zhang. 2024. “Automatic Detection of Pavement Marking Defects in Road Inspection Images Using Deep Learning.” *Journal of Performance of Constructed Facilities* 38 (2): 04024002. https://doi.org/10.1061/JPCFEV.CFENG-4619.

Tefft, B.C. 2016. “The Prevalence of Motor Vehicle Crashes Involving Road Debris, United States, 2011-2014” (Technical Report). Washington, D.C.: AAA Foundation for Traffic Safety.

V, Shalini, Shrikant Tangade, Prajna P.K., Sangeetha J.P., Farooque Azam, and Anoop G.L. 2022. “Real-Time Litter Recognition Using Improved YOLOv4 Tiny Algorithm.” *2022 IEEE 2nd Mysore Sub Section International Conference (MysuruCon)*, October 16, 1–5. https://doi.org/10.1109/MysuruCon55714.2022.9972356.

Wang, Niannian, Lihang Shang, and Xiaotian Song. 2023. “A Transformer-Optimized Deep Learning Network for Road Damage Detection and Tracking.” *Sensors* 23 (17): 17. https://doi.org/10.3390/s23177395.

Wei, Chong, Shurong Li, Kai Wu, Zijian Zhang, and Ying Wang. 2021. “Damage Inspection for Road Markings Based on Images with Hierarchical Semantic Segmentation Strategy and Dynamic Homography Estimation.” *Automation in Construction* 131 (November): 103876. https://doi.org/10.1016/j.autcon.2021.103876.

Wen, Tian, Shuo Ding, Hong Lang, et al. 2022. “Automated Pavement Distress Segmentation on Asphalt Surfaces Using a Deep Learning Network.” *International Journal of Pavement Engineering* 0 (0): 1–14. https://doi.org/10.1080/10298436.2022.2027414.

Wickramasinghe, K.S., and G.U. Ganegoda. 2020. “Pedestrian Detection, Tracking, Counting, Waiting Time Calculation and Trajectory Detection for Pedestrian Crossings Traffic Light Systems.” *2020 20th International Conference on Advances in ICT for Emerging Regions (ICTer)*, November, 172–77. https://doi.org/10.1109/ICTer51097.2020.9325479.

Wojke, Nicolai, Alex Bewley, and Dietrich Paulus. 2017. “Simple Online and Realtime Tracking with a Deep Association Metric.” *2017 IEEE International Conference on Image Processing (ICIP)*, September, 3645–49. https://doi.org/10.1109/ICIP.2017.8296962.

Wu, Ting-Wei, Hua Zhang, Wei Peng, Fan Lü, and Pin-Jing He. 2023. “Applications of Convolutional Neural Networks for Intelligent Waste Identification and Recycling: A Review.” *Resources, Conservation and Recycling* 190 (March): 106813. https://doi.org/10.1016/j.resconrec.2022.106813.

Xu, Shuyuan, Jun Wang, Peng Wu, Wenchi Shou, Xiangyu Wang, and Mengcheng Chen. 2021. “Vision-Based Pavement Marking Detection and Condition Assessment—A Case Study.” *Applied Sciences* 11 (7): 7. https://doi.org/10.3390/app11073152.

Zakeri, H., Fereidoon Moghadas Nejad, and Ahmad Fahimifar. 2017. “Image Based Techniques for Crack Detection, Classification and Quantification in Asphalt Pavement: A Review.” *Archives of Computational Methods in Engineering* 24 (4): 935–77. https://doi.org/10.1007/s11831-016-9194-z.

Zhang, Pengcheng, Qi Zhao, Jerry Gao, Wenrui Li, and Jiamin Lu. 2019. “Urban Street Cleanliness Assessment Using Mobile Edge Computing and Deep Learning.” *IEEE Access* 7: 63550–63. https://doi.org/10.1109/ACCESS.2019.2914270.

Zhang, Yunlong, and Hancheng Ge. 2012. “Assessment of Presence Conditions of Pavement Markings with Image Processing.” *Transportation Research Record* 2272 (1): 94–102. https://doi.org/10.3141/2272-11.

NDSU Dept 2880 • P.O. Box 6050
Fargo, ND 58108-6050

(701)231-7767

ndsu.ugpti@ndsu.edu