vehicles

MDPI

Article

AI Meets ADAS: Intelligent Pothole Detection for Safer AV Navigation

Ibrahim Almasri, Dmitry Manasreh and Munir D. Nazzal *

Center for Smart, Sustainable& Resilient Infrastructure (CSSRI), Department of Civil and Architectural Engineering and Construction Management, College of Engineering and Applied Science, University of Cincinnati, Cincinnati, OH 45221-0071, USA
* Correspondence: munir.nazzal@uc.edu

Abstract

Potholes threaten public safety and automated vehicles (AVs) safe navigation by increasing accident risks and maintenance costs. Traditional pavement inspection methods, which rely on human assessment, are inefficient for rapid pothole detection and reporting due to potholes' random and sudden occurring. Advancements in Artificial Intelligence (AI) now enable automated pothole detection using image-based object recognition, providing innovative solutions to enhance road safety and assist agencies in prioritizing maintenance. This paper proposes a novel approach that evaluates the integration of 3 state-of-the-art AI models (YOLOv8n, YOLOv11n, and YOLOv12n) with an ADAS-like camera, GNSS receiver, and Robot Operating System (ROS) to detect potholes in uncontrolled real-life scenarios, including different weather/lighting conditions and different route types, and generate ready-to-use data in a real-time manner. Tested on real-world road data, the algorithm achieved an average precision of 84% and 84% in recall, demonstrating its effectiveness, stable, and high performance for real-life applications. The results highlight its potential to improve road safety, allow vehicles to detect potholes through ADAS, support infrastructure maintenance, and optimize resource allocation.

Keywords: autonomous vehicles; computer vision; pothole detection; crowdsourcing; real-time detection; sensors

check for updates

Academic Editors: Eva Michelaraki and George Yannis

Received: 30 August 2025
Revised: 23 September 2025
Accepted: 26 September 2025
Published: 28 September 2025

Citation: Almasri, I.; Manasreh, D.; Nazzal, M.D. AI Meets ADAS: Intelligent Pothole Detection for Safer AV Navigation. *Vehicles* **2025**, *7*, 109. https://doi.org/10.3390/vehicles7040109

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

The United States has approximately 4 million miles of paved roads, yet according to the American Society of Civil Engineers (ASCE) 2025 Report Card [1], about 39% of public roads are in poor to mediocre conditions. This deterioration leads to billions of dollars in pavement and vehicle maintenance costs annually. Pavement distress is a major contributor to the rising maintenance costs of both pavements and vehicles. Among the most critical pavement distresses are potholes, which contribute to approximately 500 fatal accidents each year in the U.S. Furthermore, in 2021 [2] alone, pothole-related vehicle damage cost Americans over $26.5 billion.

Government agencies are striving to keep up with the increasing number of pothole reports and perform the necessary maintenance. However, with the continued rise in the number of vehicles and the expanding roadway network, addressing these issues effectively remains a challenge. A 2022 AAA survey [2] found that 1 in 10 drivers required repairs due to pothole-related damage. One of the key challenges in pothole mitigation is their unpredictable formation, as they result from a combination of factors such as weather conditions, traffic loads, and underlying pavement weaknesses.

Traditional pothole inspection methods rely on trained professionals manually surveying road sections to identify and report pavement deficiencies. While this approach has been standard for years, it is labor-intensive, time-consuming, and increasingly inefficient given the growing scale of road maintenance needs.

Another critical need for the efficient detection and localization of potholes is the safe deployment of Connected and Autonomous Vehicles (CAVs), especially as transportation agencies increasingly explore the implications of CAV adoption on existing infrastructure [3]. Autonomous Vehicles (AVs) rely primarily on two integrated systems: the *perception system*, responsible for self-localization and environmental understanding, and the *decision-making system*, which governs navigation and motion planning [4]. Both systems are fundamental to ensuring the safe and effective operation of AVs.

The perception system is equipped with a suite of advanced sensors, including LiDAR, cameras, GPS, and radar, which collectively enable the vehicle to detect and classify environmental elements such as other vehicles, pedestrians, traffic signs, and roadway features. To enhance navigation and environmental awareness, AVs commonly utilize high-definition (HD) maps that offer rich semantic and geometric information, such as lane boundaries, road markings, roadside infrastructure, and speed limits, allowing for centimeter-level localization accuracy. These HD maps, such as those provided by companies like HERE, can be updated in near real-time through internet connectivity.

However, AVs cannot rely solely on HD maps, as they are limited by update latency and require constant internet access. In scenarios where connectivity is lost or road conditions change rapidly, such as the sudden appearance of a pothole, reliance on predownloaded maps becomes insufficient. Therefore, AVs must be capable of autonomously detecting road surface defects in real time to ensure safe navigation and minimize the risk of damage or accidents.

A key enabler of such real-time perception is the Advanced Driver-Assistance System (ADAS), which leverages sensors like cameras and LiDAR to support functions such as cruise control, lane keeping, and collision avoidance. When integrated with artificial intelligence (AI) and computer vision algorithms, ADAS can be extended to detect and localize road surface defects, such as potholes, effectively. This integration not only enhances driving safety and comfort but also facilitates continuous monitoring and maintenance of road infrastructure by enabling a crowdsourced approach to pothole detection and reporting.

2. Related Work

In recent years, there has been a significant increase in the application of artificial intelligence and sensor-driven systems for pothole detection and reporting. The advances in deep learning, computer vision, and embedded hardware technologies allowed researchers to propose a wide range of solutions, with many leveraging convolutional neural network (CNN)-based object detectors, such as YOLO variants, Single Shot Multibox Detector (SSD), and others, to analyze street-level imagery or aerial footage for the identification and localization of road potholes.

Multiple studies implemented low-cost, edge-based systems using a Raspberry Pi computer and OAK-D. For instance, Jeffreys et al. [5] mounted Raspberry Pi 4+ units with EfficientDet-Lite0 and TL-VGG models on garbage trucks to detect potholes; the results showed 85% and 65% in precision and recall, respectively. Moreover, Asad et al. [6] deployed Tiny-YOLOv4 on Raspberry Pi with OAK-D for real-time detection at 31 frames per second (FPS) and achieved 90% accuracy. Heo et al. [7] also tested multiple YOLO versions for pothole detection. The results showed that SPFPN-YOLOv4 tiny achieved the highest mean average precision (mAP) of 79.6%, while estimating pothole distance and size using monocular vision.

Aerial-based pothole detection has seen growing adoption through the integration of unmanned aerial vehicles (UAVs) equipped with computer vision models. For example, Parmar et al. [8] and Alzamzami et al. [9] employed UAVs to capture high-resolution road surface imagery, leveraging YOLOv5 through YOLOv8 for object detection. Among these, YOLOv8 demonstrated superior performance, achieving precision and recall rates of up to 96% and 92%, respectively [8]. Additionally, Pehere et al. [10] highlighted the effectiveness of UAVs using traditional image processing techniques, such as edge detection and morphological filtering, reporting 80% precision and 74.4% recall.

Hybrid systems that fuse visual and non-visual data streams have demonstrated promising results in pothole detection. Silvester et al. [11] combined SSD-based image detection with inertial measurements from accelerometers and gyroscopes to enhance detection reliability. Similarly, Xin et al. [12] employed a crowdsourced approach by leveraging smartphone GPS and video data, achieving a 6% improvement in performance compared to conventional methods. In another study, Matouq et al. [13] developed a YOLOv8-based detection framework that integrates GPS data and camera calibration for real-time pothole detection and area estimation, reporting a mAP of 93%. In addition, P. and B.K. [14] proposed a self-deployed model to detect potholes in various weather conditions using thermal imaging. The model was validated on 300 images in various weather conditions and achieved outstanding results with about 98% in precision.

LiDAR-based approaches have demonstrated high accuracy in pothole detection, particularly for detailed 3D surface profiling. Talha et al. [15] utilized mobile LiDAR synchronized with GNSS to generate cross-sectional road images, which were then analyzed using YOLOv5n and YOLOv5s, achieving a detection accuracy of 98%. In another study, Salcedo et al. [16] integrated UNet-based semantic segmentation with YOLOv5 and EfficientDet for object detection on annotated road datasets such as IDD and RDD2020. Their segmentation model attained an average precision (AP) of 94%, while object detection performance ranged from 63% to 74% AP, depending on the detector used.

In addition, other research has focused on comparative evaluations of CNN architectures for object detection. For example, Ping et al. [17] benchmarked YOLOv3, SSD, Histogram of Oriented Gradients (HOG) with Support Vector Machine (SVM), and Faster R-CNN on 2036 smartphone-captured images, reporting YOLOv3 as the most accurate with 82% accuracy. Moreover, Yik et al. [18] integrated YOLOv3 with the Google Maps API to enhance geolocation capability, achieving 90% precision and an mAP of 65%.

Earlier approaches, such as the method proposed by Koch and Brilakis [19], relied on classical image processing techniques to segment defective pavement regions using histogram thresholding and geometric modeling, achieving an accuracy of 86%. Another research by Bharadwaj et al. [20] proposed an approach to detect potholes using software designed through MATLAB that processes images from a camera placed on top of an AV. The results showed a limitation related to lighting conditions, where the approach did not perform well under inconsistent lighting conditions.

Although the use of Advanced Driver Assistance Systems (ADAS) for pothole detection has gained some attention, only a limited number of studies have explored this approach in depth. For example, studies [21,22] proposed methodologies that analyze grayscale images to detect the features of a pothole, such as width and length. In their approach, the camera view must be mounted at the front bottom of a vehicle. Similarly, Martins et al. [23] presented an approach that evaluated multiple AI models that analyze images from a front-facing camera, achieving a mAP of approximately 72%.

Across these studies, YOLOv5 and YOLOv8 have consistently demonstrated strong performance in terms of both detection speed and accuracy. The use of UAVs and edge devices further enables flexible and scalable deployment in diverse environments. Moreover,

the integration of GNSS and IMU data significantly improves localization precision and system reliability. Recent developments increasingly focus on real-time operation and the feasibility of large-scale deployment, reflecting a shift toward practical implementation in intelligent transportation systems.

With the rapid emergence of new technologies, it is important to acknowledge that solutions like LiDAR and UAVs come with notable drawbacks. While LiDAR technology is becoming more affordable due to technical advancements, it remains a costly addition to most systems compared to conventional cameras. Similarly, drones, despite their ability to cover large areas efficiently, are limited by their short flight duration and the need for frequent recharging. In contrast, vehicle-mounted cameras present a more cost-effective alternative, providing a continuous stream of data as they operate on roads where vehicles naturally travel.

This paper contributes to the ongoing efforts for wider deployment of autonomous vehicles by ensuring safer and more efficient navigation of AVs on roads with potholes. It also establishes a foundation for leveraging crowdsourced pothole data to assist transportation agencies in more effective maintenance planning. This contribution is achieved by evaluating the integration of artificial intelligence, object detection techniques, a GNSS navigation system, and an ADAS-like camera, which is already deployed in production-level vehicles, for real-time pothole detection, enabling AVs to accurately identify and avoid road potholes while supporting data-driven infrastructure management.

3. Methodology

3.1. Hardware Overview

Figure 1 illustrates the hardware system used in this study. The primary sensors include a Lucid Triton TRI054S camera (LUCID Vision Labs, Richmond, BC, Canada) equipped with a 12.5 mm lens Figure 1a and a GNSS receiver (u-blox, Thalwil, Switzerland) Figure 1b, which is the navigation system that is being used in production-level vehicles. These sensors were connected to a Lenovo P16 Gen2 laptop (Lenovo, Whitsett, NC, USA) featuring an Intel i9 processor and an NVIDIA RTX 4070 GPU. As shown in Figure 1, the camera location, which was mounted on the windshield of an SUV near the rearview mirror at a 5° downward tilt from level, is the same location as the camera utilized in the Advanced Driver Assistance System (ADAS), providing a clear view of the road ahead. The camera was configured to operate at 10 frames per second (fps), while the GPS system was calibrated to run at 50 Hz, enabling high-precision pothole localization.



(**a**) (**b**)

Figure 1. (**a**) Sony camera used for image collection. (**b**) GNSS receiver utilized for positioning and localization.

3.2. Algorithm Overview

Figure 2 below illustrates the overall workflow of the proposed algorithm. Data collection, pothole detection, and post-processing were all implemented within the Robot Operating System (ROS) framework. ROS is a flexible middleware platform that offers a suite of software libraries and tools for building robotic applications. A fundamental component of ROS is the node, which operates either as a publisher (sending data) or a subscriber (receiving data). In this algorithm, a custom ROS node, entirely developed in Python version 3.8.10 programming language, was created to collect images and GPS data from sensors. Captured images are analyzed for pothole detection using the YoloV8n model. After a pothole detection, a dedicated ROS topic is published containing the pothole-related information, which are (1) the number of potholes in the image, (2) GPS coordinates, (3) pothole's relative location with respect to the front of the vehicle (left of vehicle, center of vehicle, right of vehicle). Finally, detection results were stored in a ROS bag file (ros bag) in 30 min intervals of recording, which enables real-time offline post-processing and analysis. As outlined in the introduction, the reporting functionality in the algorithm is designed to serve infrastructure maintenance agencies for long-term maintenance planning. Details on the reporting mechanism will be discussed in a later section.

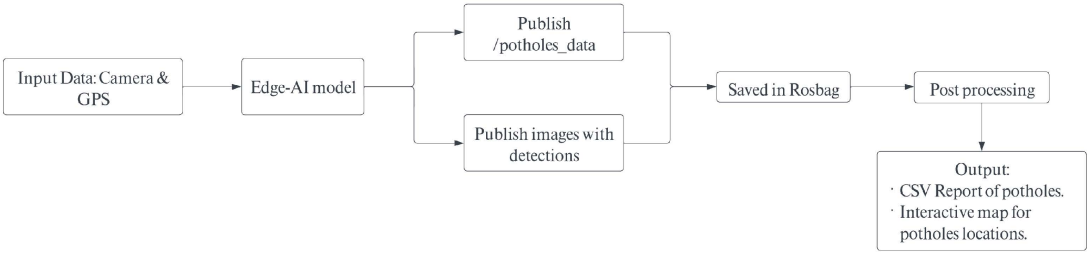

Figure 2. Workflow of the proposed algorithm outlining the sequential processing steps from data acquisition to final output.

3.3. Detection Model

The primary objective of this paper is to detect potholes in real time. To achieve this, the Ultralytics YOLO object detection model, considered a state-of-the-art solution in the field, was selected. Ultralytics YOLO represents the latest advancement in the YOLO (You Only Look Once) series, specifically designed for real-time object detection and image segmentation tasks. YOLO models use an end-to-end convolutional neural network that simultaneously predicts bounding boxes and class probabilities in a single forward pass, making them highly efficient for real-time applications.

In this study, YOLOv8n, YOLOv11n, and YOLOv12n were trained and evaluated to be employed to detect potholes from camera images. All versions architecture is available in multiple sizes: nano, small, medium, large, and x-large, each offering a trade-off between accuracy and computational demand. Larger models tend to yield higher detection accuracy but require more processing power, which may compromise real-time performance. Since this application focuses solely on detecting potholes and emphasizes real-time performance on embedded hardware, the nano variant was selected because it offers a balanced combination of detection accuracy and high inference speed, making it well-suited for deployment in real-time pavement inspection scenarios.

Data collected was split into 80% for training and 20% for validation. Each image that contained the objects of interest was annotated with a 2D bounding box using the LabelMe version 5.4.1 annotation software.

The evaluation metrics chosen for assessing model performance are precision and recall. Precision evaluates the model's ability to correctly identify potholes while mini-

mizing false positives (incorrectly classifying non-potholes as potholes). Precision can be calculated using the following equation:

$$\textbf{\textit{Precision}} = \frac{\textbf{\textit{TP}}}{\textbf{\textit{TP}} + \textbf{\textit{FP}}}$$

where

- True Positives (*TP*) are the number of instances that are correctly classified as positive by the model.
- False Positives (*FP*) are the number of instances that are incorrectly classified as positive by the model.

On the other hand, the recall metric indicates the model's ability to identify all relevant instances in a dataset. It is calculated using the following equation:

$$\textbf{\textit{Recall}} = \frac{\textbf{\textit{TP}}}{\textbf{\textit{TP}} + \textbf{\textit{FN}}}$$

where

- False Negatives (*FN*) are the number of instances that are positive, but are classified as negative by the model.

During training and evaluation, it was observed that the model struggled with false positives, particularly due to the challenging front-facing perspective of potholes from the vehicle's viewpoint. To address this limitation, several solutions were explored and detailed in the following sections.

3.3.1. Training the Model on Different Classes

The model frequently misclassified various objects, such as manhole covers, tire markings, pavement sealing, patches, and other random artifacts, as potholes. To address this issue and enhance the model's precision, the dataset annotations were expanded to include four distinct classes: potholes, patches, manhole covers, and tire markings. Those objects were chosen since they can be distinguished clearly from the images and have a consistent pattern. Nevertheless, due to the irregular and unpredictable nature of pavement sealing patterns and miscellaneous objects, consistent annotation was not feasible. Therefore, an alternative approach was employed to mitigate their impact on detection performance.

3.3.2. Adding Background Images

Several studies [3,13] recommend that background images, images that do not contain any objects of interest, should comprise approximately 10% of the total dataset. This inclusion is intended to help the model distinguish between meaningful objects and irrelevant content, thereby reducing the likelihood of false positives. However, the standard 10% background image ratio proved to be insufficient in the pothole detection scenario. This limitation became particularly evident due to the challenging nature of pothole detection, which often involves subtle surface variations and inconsistent lighting conditions.

To address this issue more effectively, we adopted a targeted approach by introducing selective background images. Instead of using arbitrary background images, we carefully curated images that had previously caused the model to trigger false positive detections during earlier rounds of inference. These selectively chosen backgrounds provided more realistic and challenging negative examples for the model to learn from.

To understand the impact of different background image ratios on model performance, various proportions of selective backgrounds were evaluated, specifically 10, 30, and 50%

of the total training dataset. For each configuration, the model was retrained, and evaluated the outcomes were evaluated using the standard performance metrics previously discussed, such as precision, recall, and F1-score. This approach supported the assessment of whether increasing the presence of challenging negative examples could enhance the model's ability to differentiate potholes from look-alike objects and reduce the false positive rate.

3.3.3. Two-Step Solution with an AI Image Classifier

To further address the persistent misclassification problem, an additional verification system was introduced. YOLOvn_cls classifier was introduced as a secondary filtering step. The model is built upon the YOLO architecture but repurposed as a multi-class image classifier. This model was explicitly trained to distinguish between four categories: potholes, manhole covers, surface patches, and tire markings. These categories were chosen based on an analysis of the most common sources of false positives identified during model inference.

The integration of the classifier into the detection pipeline serves as a filtering step. After the primary object detection model identified potential potholes, each detected region was cropped and subsequently passed through the classifier. A detection was only retained and labeled as a pothole if the classifier also confirmed it as such with high confidence. If the classifier assigned the region to one of the other three categories, it was discarded as a false positive.

The classifier was trained using a balanced dataset composed of cropped image samples representing each class, with careful attention given to ensuring variability in lighting, texture, and angle to encourage robust generalization. Once trained, the classifier's performance was thoroughly evaluated using the same set of performance metrics adopted in earlier stages. These evaluations were conducted on a hold-out validation set specifically curated to include challenging and ambiguous samples to test the classifier's robustness under realistic conditions.

3.3.4. Data Augmentation and Data Collection Under Various Weather Conditions

To enhance the model's generalization capabilities and improve performance under diverse real-world conditions, a variety of augmentation techniques were applied specifically to the pothole images. These augmentations included random rotations, horizontal flipping (mirroring), as well as modifications to brightness and contrast levels to simulate different lighting environments. The purpose of these transformations was to expose the model to a wider range of visual variations, helping it become more resilient to changes in orientation, lighting, and appearance of potholes during inference. In addition, the potholes data collection was conducted in different weather conditions to provide the model with real-life examples of lighting conditions that might not be created through the augmentation techniques.

3.3.5. Limiting Detection Range

Since the front-facing dashcam captures a narrow field of view directly ahead of the vehicle, it was observed that the model's performance tended to degrade with increasing distance from the camera. Specifically, objects located farther away appeared smaller and less detailed, which made accurate detection more difficult. This limitation often resulted in a higher number of false positives, particularly for distant features that resemble potholes in shape or texture.

To mitigate this issue, a distance-based filtration step was introduced into the detection pipeline. This filter automatically discards any pothole detections located approximately 3 m from the front of the vehicle. The threshold was selected based on empirical observations and performance analysis during testing. By focusing on the near-field region,

where the camera provides more detail and the detections are more reliable, the algorithm was able to significantly reduce false positives without compromising detection of critical road defects. This filtering approach is visually illustrated in Figure 3, which shows how detections outside the defined range are excluded from the final output.



Figure 3. Region of potholes to be considered in the detection process.

3.4. Output Postprocessing

As described in the algorithm overview, once a pothole is detected in an image, both the annotated image and corresponding pothole metadata are published and stored in a ROS bag. Before generating final reports, two critical post-processing steps were required: (1) removing duplicate pothole detections using geolocation-based filtering, and (2) generating a CSV report and interactive map for agency use.

To eliminate duplicates, we utilized GPS coordinates associated with each detection to filter out redundant detections within a two-meter radius. In cases where consecutive images contained overlapping detections, the image with the greater number of unique potholes was retained for reporting. For the second task, pandas and simpleKML libraries were utilized in Python to generate both a CSV report and an interactive map, as illustrated in Figure 4 and Table 1. These outputs provide a structured summary and visual representation of the detected potholes for agency review.

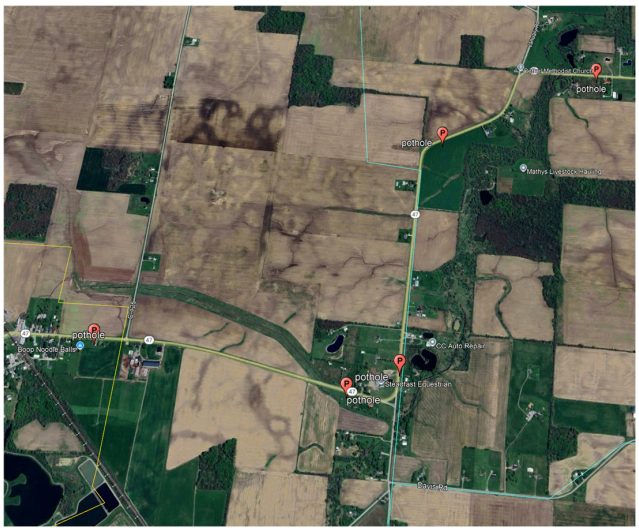

Figure 4. Interactive map for the location of detected potholes.

Table 1. Generated potholes report.

ID	Detections	GPS Time	Latitude	Longitude	Confidence	Location
1	1	1.74×10^9	40.44644	−83.4494	0.6	Right
2	1	1.74×10^9	40.48333	−83.4801	0.74	Right
3	1	1.74×10^9	40.49242	−83.4876	0.76	Left
4	1	1.74×10^9	40.49699	−83.4913	0.62	Center
5	1	1.74×10^9	40.49255	−83.4877	0.70	Left

The post-processing workflow is outlined below in Figure 5.

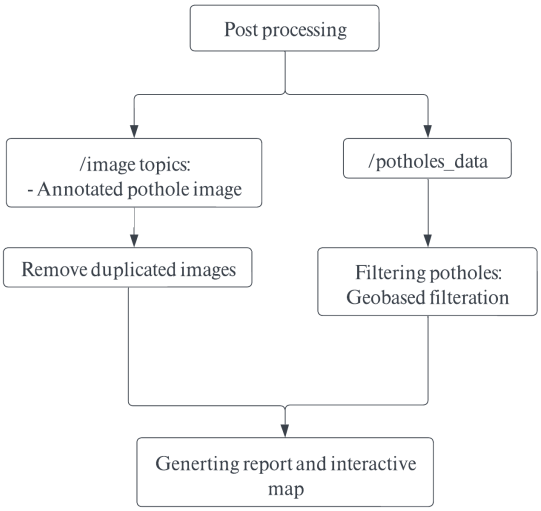

Figure 5. Workflow of the post-processing step.

4. Results and Discussion

4.1. Dataset and Data Collection

For this research, data was collected from over 500 miles of roadway across the state of Ohio, resulting in a dataset of more than 4800 annotated images consisting of 2700 potholes, 1550 manholes, 1620 patches, and approximately 1780 tire markings. The dataset was split randomly into 80% for training and 20% for validation. Data information is illustrated in Figure 6.

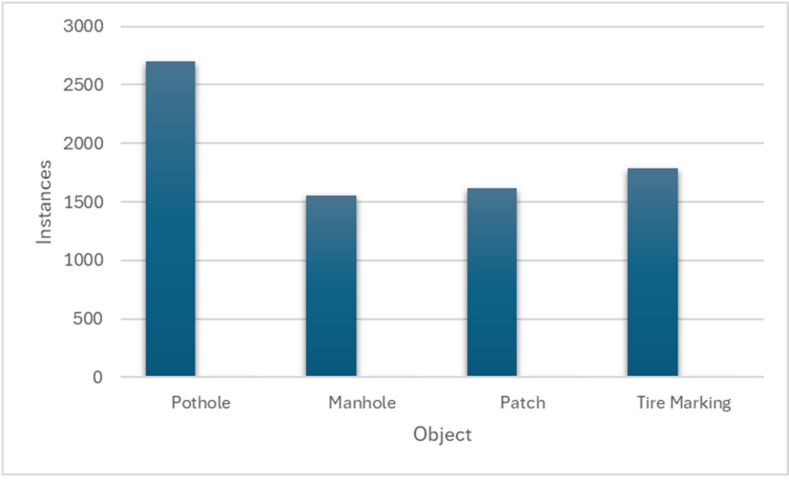

Figure 6. Collected data classes and quantities for model training.

4.2. Model Performance Evaluation

As previously discussed, the front-facing view posed challenges in accurately detecting potholes, often resulting in false positives. To address this, multiple mitigation techniques were applied. The following section presents these techniques and their corresponding results.

4.2.1. Selective Background

Initially, YoloV8 was trained using different percentages of selected background images to determine the best background ratio that the models would stabilize at while increasing the overall precision metric. 10%, 30%, and 50% percentages were assessed. The addition of more selected background images that included the random elements, which were mistakenly considered as potholes, enhanced the model's performance.

The results of the proposed solution are demonstrated in Figure 7.

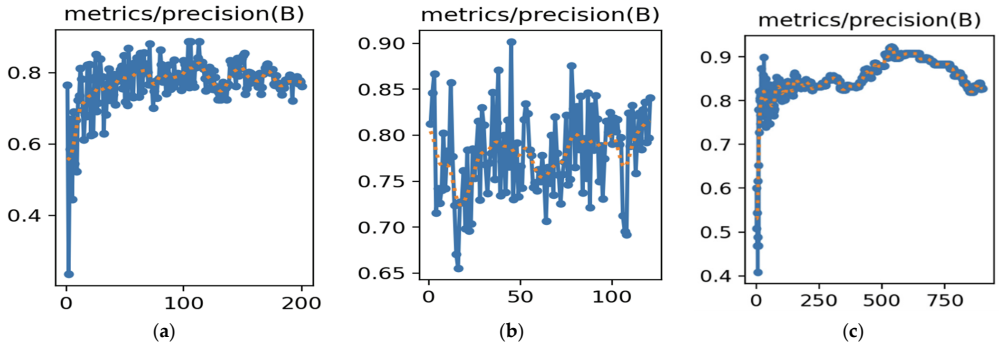

Figure 7. Performance of the models on the validation dataset under varying background ratios: (**a**) results with 10% background, (**b**) results with 30% background, and (**c**) results with 50% background.

As demonstrated, the model achieved approximately 78% precision on unseen data when trained with 10% background images. Increasing the background proportion to 30% introduced some fluctuations in performance, but it stabilized around 80%. Training with 50% background images resulted in more consistent performance, reaching approximately 85% precision.

4.2.2. Evaluating YOLO Versions 11 and 12, Based on the Chosen Background Percentage

YOLOv11n and YOLOv12n were trained and evaluated based on the 50% background ratio. The results of the three trained models used in this study are illustrated in Figures 8 and 9, and Table 2.

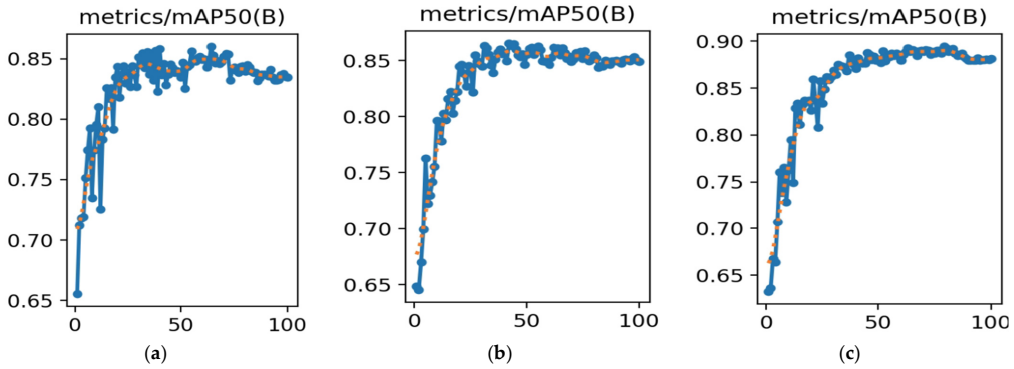

Figure 8. Mean average precision results for (**a**) Yolov8n, (**b**) YOLOv11n, and (**c**) YOLOv12n.

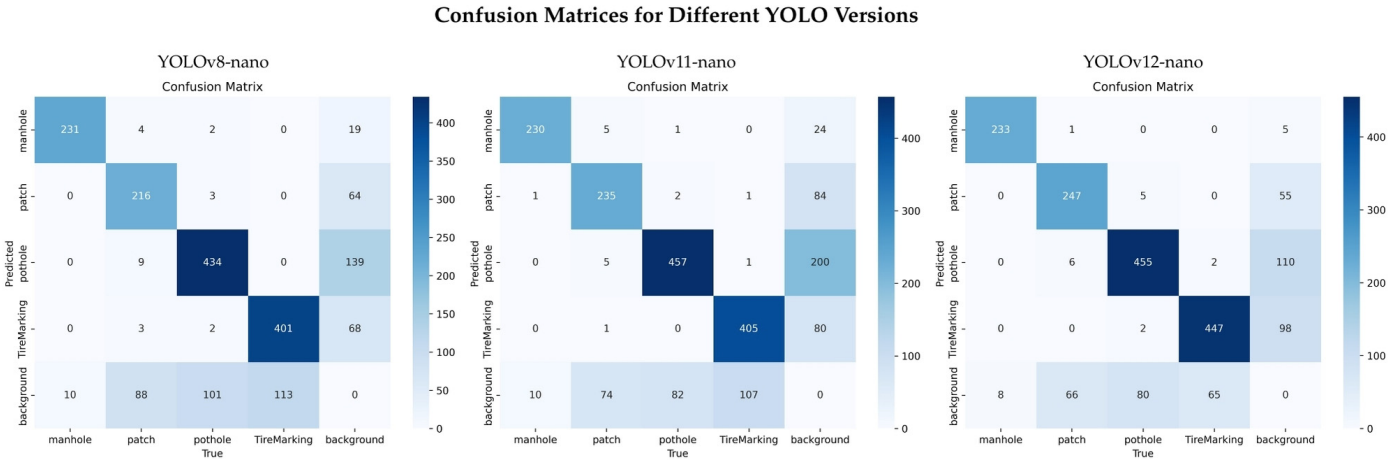

Figure 9. Confusion matrices for the trained YOLO Versions.

Table 2. Precision and recall analysis for the pothole class.

Model Version	True Potholes	False Potholes	False Negatives	Precision	Recall
V8	434	148	108	75%	80%
V11	457	206	85	69%	84%
V12	455	118	87	79%	84%

As demonstrated in Figure 8, YOLO v8n and v11n have about the same mAP with a value of 86%, for YOLO v12n, the mAP was 89% which outperforms the older versions. Nevertheless, since our focus is on the pothole class particularly, Figure 9 shows the confusion matrix for the models' results for each class, and Table 2 provides a calculation of the precision and recall for each model based on the pothole class.

Based on the results, YOLOv12 outperforms versions 8 and 11 in precision with a value of 79% and achieved a high recall of 84% proving its ability to detect most potholes.

A representative sample of detection results from the model trained with 50% background images on the validation dataset is presented below in Figure 10.



Figure 10. (**a**,**b**) Examples of detected potholes on validation dataset.

4.2.3. Adding a Classifier for Filtration

A YOLOv8n classifier was introduced as a second-stage filtering step to further reduce false positives. However, due to the challenging visual characteristics of potholes, particularly when images are tightly focused on the defect, the classifier consistently overfitted and failed to achieve the desired performance.

4.2.4. Data Augmentation and Data Collection Under Various Weather Conditions

Data augmentation techniques involving adjustments to image lighting and contrast failed to improve model performance. This is likely because such synthetic modifications do not accurately replicate the complex lighting conditions encountered in real-world scenarios. As a result, these augmentations degraded the model's effectiveness, leading to a noticeable decline in both precision and recall. On the other hand, data collection under various weather conditions enhanced the model's performance and made it capable of detecting potholes under various lighting conditions while also limiting the occurrence of false positives.

4.2.5. Limiting Detection Region

Restricting pothole detection to a range of approximately 3 m ahead of the vehicle significantly reduced the false positive rate. This improvement is attributed to minimizing the influence of distant objects, whose textures, distorted by distance, often resemble potholes triggers the model to produce more false positives.

4.3. Field Evaluation of the Developed Model

Although the validation results of the developed AI detection model appeared sufficient, they do not necessarily reflect real-world performance. Therefore, the YOLOv12n, which was the best-performing model based on the validation dataset, was deployed for real-time inference to evaluate its effectiveness in practical scenarios.

The evaluation was conducted along different types of roads and under different weather conditions. Evaluation sections are shown in Figure 11. During evaluation, 1134 potholes were detected, 1015 of which were confirmed as actual potholes, resulting in an observed precision of approximately 89.5%. A sample of the detected potholes is demonstrated in Figure 12.

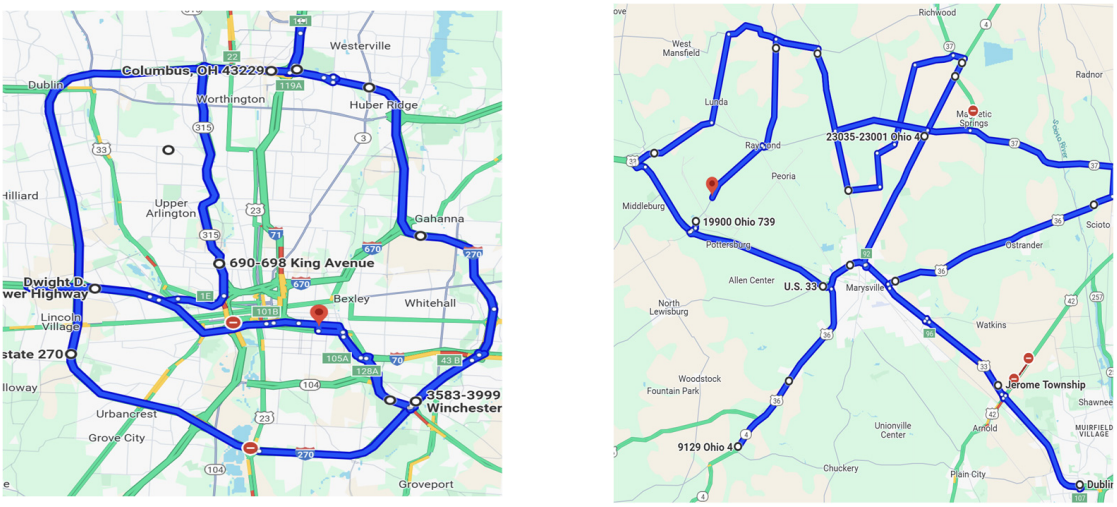

(**a**) (**b**)

Figure 11. (**a**,**b**) evaluation sections (in blue) in Cincinnati, OH, USA.

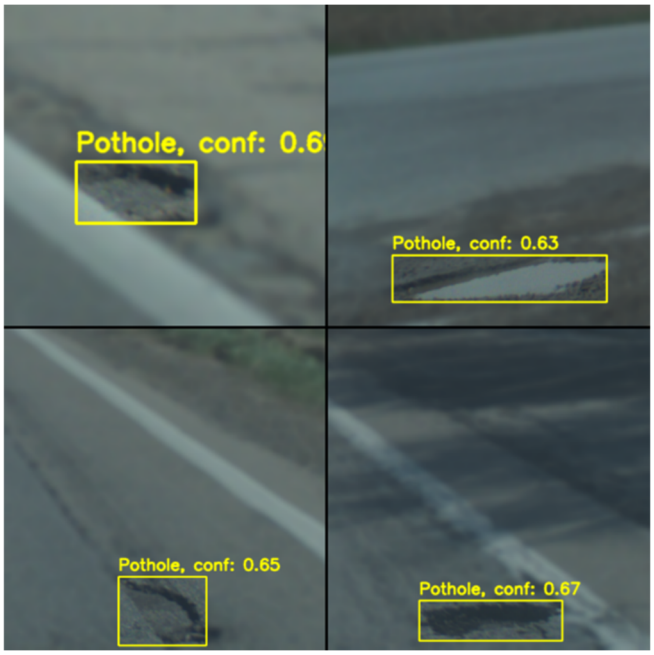

Figure 12. Sample of pothole detection from real-life application.

The results of the validation process indicate a strong and stable performance under different route types, weather, and lighting conditions.

5. Conclusions

This study introduces a practical AI-based algorithm that validates the integration of an ADAS-like camera with a GNSS receiver and state-of-the-art AI models to enable real-time pothole detection and avoidance in autonomous vehicles. The algorithm also supports a crowdsourced data collection approach, allowing transportation agencies to gather accurate pothole location data from vehicles and improve maintenance planning.

Multiple state-of-the-art YOLO models were tested, specifically the nano variant of versions 8.11, and the most recent YOLOv12 model. YOLOv12 showed the best results compared to older versions, with a mean average precision of 89%. But since in this paper we are trying to validate the integration and application of AI in ADAS, we validated the model on multiple miles under variant conditions, which in turn showed a high stable performance in rainy, cloudy, and sunny days. Compared to the cited references, this approach provides a more realistic and validated performance for pothole detection since it does not rely on limited road sections or static validation images. In addition, results show that with the development of machine vision models, a high and stable performance is achievable with balanced data, allowing researchers to start focusing on the potential use cases from the output rather than spending more time collecting more data, trying to stabilize the model's performance.

Despite the good performance of the algorithm, it still encounters false positives that are triggered by random objects that might appear on roads; for example, the most common false positive triggers are pavement sealing, water pots, tire shreds, and dead animals. Thus, in order to improve the model's resilience, we recommend that fine-tuning the model by adding those rare cases will help it be more stable and reduce the false positives.

Looking ahead, we are planning to expand our studies to incorporate more data into the model output, such as area and volume, which in turn will help DOTs in the quantification of materials. In addition, since the algorithm was able to detect and produce data in real-time, we are planning to investigate the possibility of using this data to (1) inform other vehicles using vehicle-to-vehicle communication about pothole locations to take

corresponding action before hitting a pothole, (2) allow vehicles to avoid potholes on the lane through their ADAS camera.

Finally, this AI-driven approach holds significant potential for enhancing infrastructure resilience and supporting transportation agencies in delivering timely and cost-effective maintenance. Moreover, it supports the overarching goal of ensuring autonomous vehicles are safely and reliably deployed in real-world environments. In addition, the algorithm can be deployed in passenger vehicles, enabling an efficient crowdsourcing mechanism for pothole data collection across diverse road networks.

Author Contributions: Conceptualization, M.D.N., D.M. and I.A.; methodology, M.D.N., D.M. and I.A.; validation, I.A., D.M. and M.D.N.; formal analysis, I.A., D.M. and M.D.N.; investigation, I.A., D.M. and M.D.N.; resources, M.D.N.; data curation, I.A., D.M. and M.D.N.; writing—I.A., D.M. and M.D.N.; writing—review and editing, I.A., D.M. and M.D.N.; supervision, M.D.N. and D.M.; project administration, M.D.N. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The images used in this study are not publicly available due to ethical and privacy considerations, as they contain identifiable information, including vehicle license plates and individuals. Sharing these data could compromise privacy, and therefore, the data cannot be made available.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. American Society of Civil Engineers. 2025 Report Card for America's Infrastructure. Available online: https://infrastructurereportcard.org/ (accessed on 28 August 2025).
2. AAA. AAA: Potholes Pack a Punch as Drivers Pay $26.5 Billion in Related Vehicle Repairs. Available online: https://newsroom.aaa.com/2022/03/aaa-potholes-pack-a-punch-as-drivers-pay-26-5-billion-in-related-vehicle-repairs/ (accessed on 28 August 2025).
3. Manasreh, D.; Nazzal, M.D.; Talha, S.A.; Khanapuri, E.; Sharma, R.; Kim, D. Application of Autonomous Vehicles for Automated Roadside Safety Assessment. *Transp. Res. Rec.* **2022**, *2676*, 255–266. [CrossRef]
4. Kortmann, F.; Fassmeyer, P.; Funk, B.; Drews, P. Watch Out, Pothole! Featuring Road Damage Detection in an End-to-End System for Autonomous Driving. *Data Knowl. Eng.* **2022**, *142*, 102091. [CrossRef]
5. Jeffreys, Z.; Kumar, K.; Xie, Z.; Bae, W.; Alkobaisi, S.; Narayanapppa, S. PotholeVision: An Automated Pothole Detection and Reporting System Using Computer Vision. In Proceedings of the ACM Conference on Computer-Human Interaction, San Francisco, CA, USA, 15–18 May 2024; pp. 695–697. [CrossRef]
6. Asad, M.; Khaliq, S.; Yousaf, M.H.; Ullah, M.; Ahmad, A. Pothole Detection Using Deep Learning: A Real-Time and AI-on-the-Edge Perspective. *Adv. Civ. Eng.* **2022**, *2022*, 9221211. [CrossRef]
7. Heo, D.-H.; Choi, J.-Y.; Kim, S.-B.; Tak, T.-O.; Zhang, S.-P. Image-Based Pothole Detection Using Multi-Scale Feature Network and Risk Assessment. *Electronics* **2023**, *12*, 826. [CrossRef]
8. Parmar, A.; Gajjar, R.; Gajjar, N. Drone-Based Potholes Detection Using Machine Learning on Various Edge AI Devices in Real-Time. In Proceedings of the 2023 IEEE International Symposium on Smart Electronic Systems (iSES), Jaipur, India, 18–20 December 2023; pp. 22–26. [CrossRef]
9. Alzamzami, O.; Babour, A.; Baalawi, W.; Al Khuzayem, L. PDS-UAV: A Deep Learning-Based Pothole Detection System Using Unmanned Aerial Vehicle Images. *Sustainability* **2024**, *16*, 9168. [CrossRef]
10. Pehere, S.; Sanganwar, P.; Pawar, S.; Shinde, A. Detection of Pothole by Image Processing Using UAV. *J. Sci. Technol.* **2020**, *5*, 101–110. [CrossRef]
11. Silvister, S.; Komandur, D.; Kokate, S.; Khochare, A.; More, U.; Musale, V.; Joshi, A. Deep Learning Approach to Detect Potholes in Real-Time Using Smartphone. In Proceedings of the 2019 IEEE PuneCon, Pune, India, 19–21 December 2019; pp. 1–4. [CrossRef]
12. Xin, H.; Ye, Y.; Na, X.; Hu, H.; Wang, G.; Wu, C.; Hu, S. Sustainable Road Pothole Detection: A Crowdsourcing-Based Multi-Sensors Fusion Approach. *Sustainability* **2023**, *15*, 6610. [CrossRef]
13. Matouq, Y.; Manasreh, D.; Nazzal, M.D. AI-Driven Approach for Automated Real-Time Pothole Detection, Localization, and Area Estimation. *Transp. Res. Rec.* **2024**, *2678*, 2018–2031. [CrossRef]

14. Pathmanaban, P.; Gnanavel, B.K. Robust Pothole Detection in Adverse Weather Conditions Using Thermal Imaging and Image Processing. In Proceedings of the 2024 23rd IEEE Intersociety Conference onThermal and Thermomechanical Phenomena in Electronic Systems (ITherm), Orlando, FL, USA, 28–31 May 2024; pp. 1–6. [CrossRef]
15. Talha, S.A.; Manasreh, D.; Nazzal, M.D. The Use of Lidar and Artificial Intelligence Algorithms for Detection and Size Estimation of Potholes. *Buildings* **2024**, *14*, 1078. [CrossRef]
16. Salcedo, E.; Jaber, M.; Requena Carrión, J. A Novel Road Maintenance Prioritisation System Based on Computer Vision and Crowdsourced Reporting. *J. Sens. Actuator Netw.* **2022**, *11*, 15. [CrossRef]
17. Ping, P.; Yang, X.; Gao, Z. A Deep Learning Approach for Street Pothole Detection. In Proceedings of the 2020 IEEE 6th International Conference on Big Data Computing Service and Applications (BigDataService), Oxford, UK, 3–6 August 2020; IEEE: Piscataway, NJ, USA, 2020; pp. 198–204. [CrossRef]
18. Yeoh, K.Y.; Alias, N.E.; Yusof, Y.; Isaak, S. A Real-time Pothole Detection Based on Deep Learning Approach. In Proceedings of the International Symposium on Automation, Information and Computing (ISAIC 2020), Beijing, China, 2–4 December 2020. [CrossRef]
19. Koch, C.; Brilakis, I. Pothole Detection in Asphalt Pavement Images. *Adv. Eng. Inf.* **2011**, *25*, 507–515. [CrossRef]
20. Bharadwaj, S.; Sundra Murthy, G.; Varaprasad, G. Detection of Potholes in Autonomous Vehicle. *IET Intell. Transp. Syst.* **2014**, *8*, 543–549. [CrossRef]
21. Kumar, B.N.; Kulkarni, S.G.; Kakkar, S.; Vipin, K. ADAS Pothole Detection System Using FPGA. In Proceedings of the 2023 IEEE Asia Pacific Conference on Postgraduate Research in Microelectronics and Electronics (PRIMEAsia), Hyderabad, India, 13–15 December 2023; pp. 56–57. [CrossRef]
22. Lakmal, H.K.I.S.; Dissanayake, M.B. Pothole Detection with Image Segmentation for Advanced Driver Assisted Systems. In Proceedings of the 2020 IEEE International WIE Conference on Electrical and Computer Engineering (WIECON-ECE), Bhubaneswar, India, 26–27 December 2020; pp. 308–311. [CrossRef]
23. Martins, W.D.; Osorio, F.S.; Bruno, D.R. Real-Time Deep Learning Based Pothole Detection on Low-Cost Embedded Devices for ADAS. In Proceedings of the 2025 Brazilian Conference on Robotics (CROS), Belo Horizonte, Brazil, 19–21 May 2025; pp. 1–6. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.