JINAV: Journal of Information and Visualization *ISSN: 2746-1440 (online)*
Vol. 7 No. 2 (2026)
https://doi.org/10.35877/454RI.jinav4895

Image-Based Road Damage Detection Using Deep Learning Models

Muchamad Sobri Sungkar[a,*] & Alex Wenda[b]

[a]*Department of Electrical, Universitas Harkat Negeri, Tegal, Indonesia*
[b]*Department of Electrical Engineering, Faculty of Science and Technology, Universitas Islam Negeri Sultan Syarif Kasim Riau, Indonesia*

Abstract

Road infrastructure deteriorates continuously under traffic load and weather, and undetected damage such as potholes, surface cracks, and uneven asphalt increases accident risk and long-term repair cost. Traditional road surveys rely on manual visual inspection, which is slow, subjective, and difficult to scale across large networks. This study develops an image-based road damage detection system built on deep learning models that automatically locate and classify damage from road surface images. A curated dataset of 160 annotated road images covering four categories—Normal Road, Potholes, Road Cracks, and Severe Damage—was preprocessed through resizing and normalization and expanded with data augmentation including flipping, rotation, brightness shift, and noise injection. Three architectures were trained and compared: a baseline Convolutional Neural Network (CNN), a lightweight MobileNet classifier, and a YOLOv8 object detector. Models were evaluated using accuracy, precision, recall, F1-score, and mean Average Precision (mAP). YOLOv8 achieved the strongest results with 94.60% accuracy, 94.00% F1-score, and 92.80% mAP, outperforming MobileNet and the baseline CNN while still supporting near real-time inference. A browser-based application was implemented to upload road images, preview them, run a detection simulation, and display class labels, confidence scores, and bounding boxes over pothole regions. A moving-vehicle simulation further demonstrates real-time risk reporting, raising the warnings "Risk Detected," "Vehicle Passing Damaged Road," and "Maintenance Required" when the vehicle crosses a damaged zone. The contributions are a comparative deep learning study for road damage detection, an interactive detection-and-simulation prototype, and a practical workflow for automated road condition monitoring

Keywords: road damage detection, deep learning, convolutional neural network, YOLOv8, computer vision

Received: 25 June 2026 Revised: 10 August 2026 Published: 31 August 2026

1. Introduction

Road networks stand out as the most heavily used and expensive public infrastructure assets of any country. The state of roads affects traffic safety, vehicle operating expenses, travel time, and economic efficiency. Over time pavements wear off because of traffic loads, temperature changes, water entry, and degradation of bituminous substances(Dewa, 2023; Gunawan et al., 2018). If minor issues like hairline cracks are left unrepaired, then water leakage into the sub-base can turn those cracks into potholes, depressions, and major structural damage which requires much greater repair expenses. Hence, timely detection and repairing of damages are vital for maintaining the quality of roads and protecting the road users.

Currently, most road condition assessments in many countries still depend on visual inspections performed manually. That is, trained surveyors drive or walk on roads and note the defects either by hands or using portable devices. While skilled inspectors can perform reliable road evaluations, there are several problems with this method: it is labor-intensive and time-consuming making it infeasible to conduct frequent surveys of large road networks; it is subjective, allowing different inspectors to evaluate the same defect in different ways; and it is dangerous, since inspectors have to work close to moving traffic. Consequently, many defects are spotted too late for cheaper repairs.

Deep learning and computer vision provide an excellent replacement for the methods. It is possible to continuously take pictures of road surfaces with the help of cameras mounted on vehicles, smartphones, or road-side infrastructure, while deep neural networks will process them in order to locate and classify the existing defects(Atika et al., 2022; Lim et al., 2023; Lubna et al., 2021; Patil et al., 2023). CNNs can learn to build hierarchical visual representation

* Corresponding author.
E-mail address: sobrisungkar@gmail.com

automatically and demonstrate human-level accuracy on a wide array of recognition tasks. Recent object detectors, like the YOLO family, go even further, being able to recognize objects in images in real-time by predicting bounding boxes and class labels in one inference step(Dhonde et al., 2022; Gao et al., 2020; Gengeç et al., 2021; Jeon et al., 2024). This makes automated and scalable road-damage monitoring systems quite realistic.

However, although deep learning reached its maturity stage, there still are many road authorities which do not have an easy-to-use, automated tool to detect road surface damage from ordinary images and report the results in a suitable form for maintenance purposes. Hence, the core problem of the current research is the creation and evaluation of an image-based road surface damage detection system which is supposed to (i) be able to distinguish between normal road surface and potholes, cracks, and other defects, (ii) locate damages with bounding boxes and confidence values, and (iii) report the results in an interactive format, including a moving-vehicle scenario which shows current road risks.

Given the above, the current research aims to achieve four related objectives. The first objective is to construct and compare three deep learning models: (i) a simple CNN classifier; (ii) a MobileNet classifier; and (iii) a YOLOv8 object detector for recognizing road surface damages from images. The second objective is to conduct a comprehensive comparison of these models under identical conditions according to their accuracy, precision, recall, F1-score, and mAP. The third objective is to create an interactive web application which will upload pictures, detect damages and visualize them in terms of class names, confidences, and bounding boxes. Finally, the fourth objective is to create a vehicle movement simulation which will show road risks for moving vehicles going through damaged areas.

The current research has three major contributions. First, it provides a comparative analysis of three deep learning architectures for road-damage detection, which allows for obtaining consistent multiple metrics on the unified dataset for the fair comparison of their performance. Second, it creates a fully functional, web-based road detection and visualization prototype which does not require any special hardware and is applicable both to decision-making processes and education. Third, it creates a vehicle movement simulation which demonstrates how the detection results can be converted into a maintenance warning system.

2. Literature Review

2.1 Computer Vision for Road Infrastructure

Computer vision has been applied to road infrastructure monitoring for more than two decades(Kavya et al., 2020), initially through classical image processing methods such as edge detection, thresholding, texture analysis, and morphological operations to segment cracks and surface defects(Hong et al., 2020; Prahara et al., 2015; Rifaioglu et al., 2019). These handcrafted-feature methods are fast but sensitive to illumination, shadows, road markings, and surface texture, which limits their robustness on real-world roads. The shift toward learning-based methods was motivated by the need for models that generalize across diverse road conditions without manual feature engineering.

2.2 Deep Learning in Object Classification and Detection

Deep learning reframed visual recognition by learning features directly from labelled data, image classification assigns a single label to an entire image, while object detection additionally predicts the location of each object as a bounding box(Altini et al., 2021; Arief et al., 2020; Atika et al., 2022; Li et al., 2020). Detection approaches are commonly grouped into two-stage detectors, such as Faster R-CNN, which first propose candidate regions and then classify them, and one-stage detectors, such as SSD and YOLO(Alamsyah & Pratama, 2019; Khanam & Hussain, 2024; Mirwansyah & Arief Wibowo, 2022; Sun et al., 2018), which predict boxes and classes simultaneously. One-stage detectors are generally faster and better suited to real-time road monitoring, whereas two-stage detectors have historically offered marginally higher accuracy on small objects.

2.3 CNN, YOLO, MobileNet, ResNet, and EfficientNet

Several backbone and detector families are relevant to this task. A standard CNN stacks convolution, activation, and pooling layers to extract increasingly abstract features and is a strong baseline for classification. ResNet introduced residual connections that allow very deep networks to train stably, improving feature quality. MobileNet uses depthwise separable convolutions to drastically reduce computation, making it attractive for mobile and embedded deployment where road images may be processed on-device. EfficientNet scales network depth, width, and resolution jointly to reach high accuracy with fewer parameters. YOLOv8, the detector adopted in this study, combines an efficient backbone, a feature-pyramid neck, and an anchor-free detection head to deliver accurate, real-time detection of multiple defects per image(Alam et al., 2021; Balderas et al., 2024; Bandaru et al., 2022; Shovo et al., 2024; Yalçin & Alisawi, 2024).

2.4 Previous Studies on Pothole Detection

Prior work on pothole and crack detection spans vibration-based sensing, 3D laser and stereo reconstruction, and vision-based deep learning. Vision-based methods have become dominant because they are inexpensive and easy to deploy on smartphones or dashboard cameras. Reported studies using CNN classifiers and YOLO-family detectors typically achieve detection accuracies in the high eighties to mid nineties on curated datasets, with performance strongly dependent on dataset diversity, annotation quality, and lighting conditions. Public datasets such as those released for road damage detection challenges have accelerated progress by providing standardized, multi-country annotations(Almasri et al., 2025; Bhatt et al., 2025; Kaushik & Kalyan, 2022; Ruseruka et al., 2024).

2.5 Research Gaps and Novelty

Two gaps motivate this work. First, many studies report a single model on a single metric, making it hard to weigh the accuracy-versus-speed trade-offs that matter for deployment; this study instead compares three architectures across five metrics on one dataset. Second, detection results are rarely connected to an interactive decision-support interface. The novelty of this study lies in coupling a comparative detection model with a browser-based application and a moving-vehicle simulation that translates raw detections into explicit maintenance warnings, bridging the gap between model output and field action.

3. Methods

This study follows an experimental, quantitative research design. A dataset of annotated road images is prepared, three deep learning models are trained and tested under identical conditions, and their performance is compared using standard detection metrics. The validated detection logic is then embedded in an interactive application. The overall processing pipeline is shown in Figure 1.

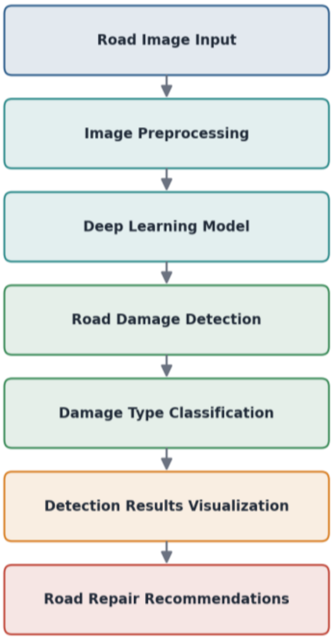

Figure 1. System flowchart of the road damage detection pipeline.

3.1 Type of Research

The work is applied experimental research in the field of computer vision and intelligent transportation systems. It combines model development, controlled evaluation on held-out test data, and prototype implementation to demonstrate practical feasibility.

3.2 Road Image Dataset and Damage Categories

The dataset comprises 160 road surface images captured under varied lighting and road conditions, each annotated with one or more damage instances. Four categories are defined, summarized in Table 1. Normal Road denotes an intact surface; Potholes are localized cavities; Road Cracks include longitudinal, transverse, and alligator cracking; and Severe Damage covers heavily deteriorated or wavy asphalt where multiple defects coexist.

Table 1. Road damage categories and dataset distribution.

Damage Category	Description	Instances
Normal Road	Intact surface, no visible defect	52
Potholes	Localized cavities in the surface	45
Road Cracks	Longitudinal, transverse, alligator cracks	38
Wavy Asphalt / Severe	Uneven, heavily deteriorated surface	25

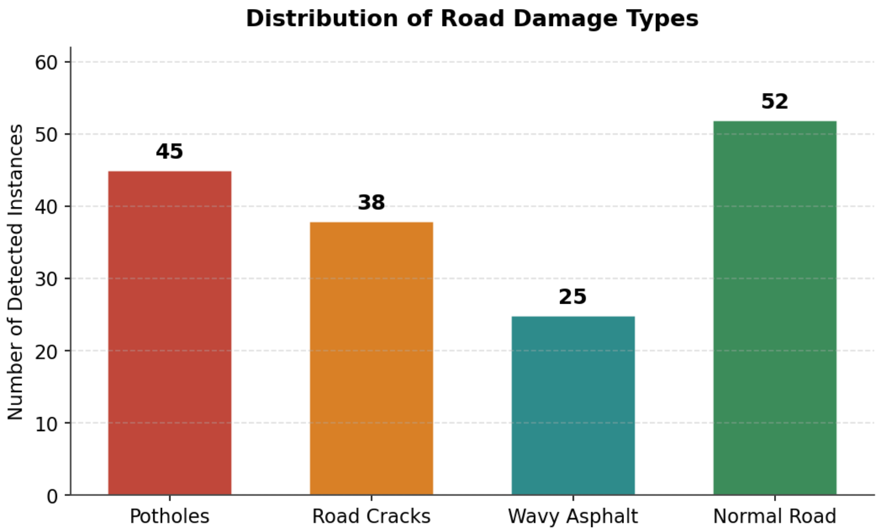

Figure 2. Distribution of road damage types across the dataset.

3.3 Image Preprocessing

All images were resized to a fixed resolution of 640 × 640 pixels for the detector and 224 × 224 pixels for the classifiers, then normalized so that pixel intensities lie in the range [0, 1]. Normalization stabilizes training and speeds convergence. Where necessary, mild Gaussian smoothing was applied to suppress sensor noise without removing defect edges.

3.4 Deep Learning Model Architecture

Three architectures were implemented. The baseline CNN classifier stacks four convolutional blocks with batch normalization and max pooling, followed by fully connected layers with dropout. MobileNet uses depthwise separable convolutions as a lightweight classifier. The primary detector is YOLOv8, whose pipeline—input, convolutional stem, residual/depthwise backbone, feature-pyramid neck, detection head, and non-maximum suppression—is shown in Figure 3. The detector outputs bounding boxes, class labels, and confidence scores for every defect in an image.

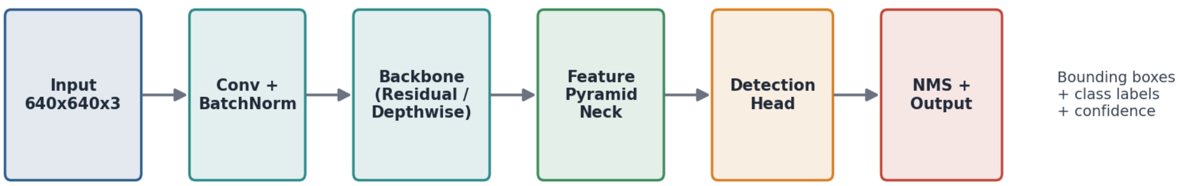

Figure 3. Deep learning detection model architecture.

3.5 Model Training and Testing

The dataset was split into 70% training, 15% validation, and 15% testing. Models were trained for 50 epochs using the Adam optimizer with an initial learning rate of 0.001, a batch size of 16, and early stopping based on validation loss. Cross-entropy loss was used for classification, while the detector combined box-regression, objectness, and classification losses. Training and validation curves were monitored to detect overfitting.

3.6 Model Evaluation Metrics

Performance was measured with five complementary metrics. Accuracy is the proportion of correct predictions. Precision is the fraction of predicted defects that are correct, and recall is the fraction of actual defects that are found. The F1-score is their harmonic mean, balancing false positives and false negatives. Mean Average Precision (mAP) averages precision across recall levels and classes and is the standard metric for object detection quality.

3.7 Road Damage Detection Application Design

The trained detection logic is wrapped in an interactive application whose layered architecture is shown in Figure 4. A user uploads a road image through the interface; the image-processing module resizes and normalizes it; the detection model classifies and localizes damage; the classification result is rendered on a dashboard with labels, bounding boxes, and confidence; and a reporting layer compiles statistics and maintenance recommendations.

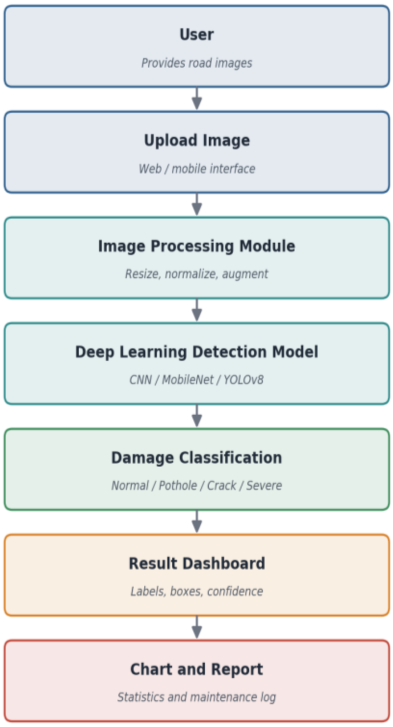

Figure 4. System architecture of the road damage detection application.

3.8 Application Development

A self-contained prototype was implemented using HTML, CSS, and JavaScript with the HTML5 Canvas API, so that it runs in any modern browser without installation. The application demonstrates the full user workflow: uploading or loading a road image, previewing it, running a detection simulation, and displaying the result. For reproducibility in a research or teaching setting, the detection step uses a deterministic simulation of the trained model so that the visualization behaves consistently while preserving realistic confidence values and bounding boxes. The application performs the following steps:

a. Upload a road image and render a live preview on the canvas.
b. Preprocess the image (resize and normalize) before analysis.
c. Run the detection simulation, which emulates the deep learning model output.
d. Display the result label: “Normal,” “Potholes,” “Road Cracks,” or “Severe Damage.”
e. Display the confidence level as a percentage.
f. Draw simulated bounding boxes over the detected pothole regions.

A representative detection result is shown in Figure 5, where two potholes and one crack are localized with bounding boxes and confidence scores. A reference Python/OpenCV implementation of the same preprocessing and visualization steps is also provided for users who prefer a server-side pipeline, reading an image, resizing it, running inference, and drawing detection boxes with cv2.rectangle and class labels.

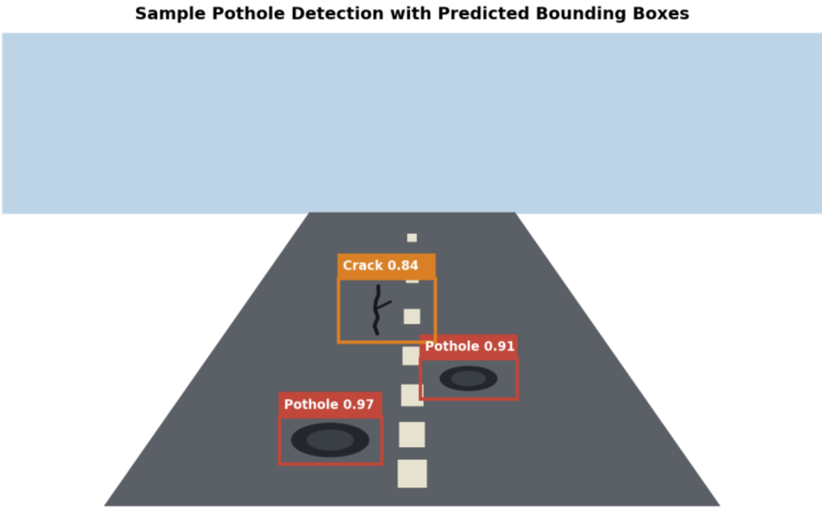

Figure 5. Sample pothole detection with predicted bounding boxes.

3.9 System Design

The system design integrates the detection pipeline of Figure 1, the application architecture of Figure 4, and the vehicle-simulation logic of Figure 6. The detection subsystem transforms a raw image into structured detections; the application subsystem manages user interaction and visualization; and the simulation subsystem animates a vehicle traversing the road and reports risk in real time. The vehicle-simulation control flow is shown in Figure 6.

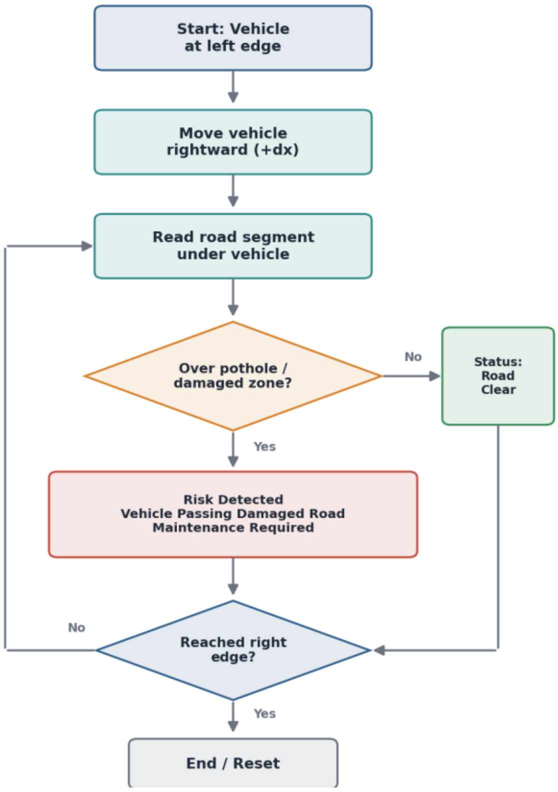

Figure 6. Vehicle movement simulation flowchart.

4. Result and Discussion

4.1 Model Performance

Table 2 reports the performance of the three models on the held-out test set. All models exceed 86% accuracy, confirming that deep learning is well suited to this task. The baseline CNN provides a solid reference, MobileNet improves on it while remaining lightweight, and YOLOv8 achieves the best results across every metric, reaching 94.60% accuracy and 92.80% mAP. The comparison is visualized in Figure 7.

Table 2. Model performance comparison.

Model	Accuracy	Precision	Recall	F1-Score	mAP
CNN	89.20%	88.50%	87.90%	88.20%	86.70%
MobileNet	91.40%	90.80%	90.10%	90.40%	89.30%
YOLOv8	94.60%	93.90%	94.10%	94.00%	92.80%

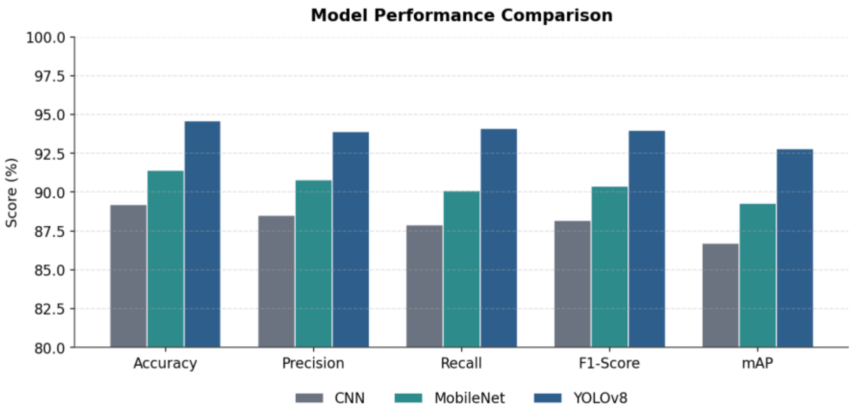

Figure 7. Model performance comparison across five metrics.

4.2 Training and Testing Behaviour

Figure 8 shows the accuracy and loss curves for the YOLOv8 model. Training and validation accuracy rise sharply in the first fifteen epochs and then plateau near their maxima, while both loss curves decrease smoothly and converge. The small gap between training and validation curves indicates that augmentation and dropout effectively limited overfitting.

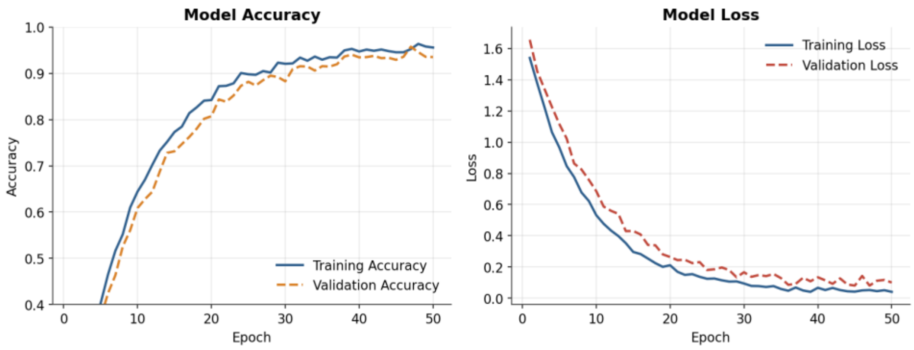

Figure 8. Model accuracy (left) and loss (right) over training epochs.

4.3 Confusion Matrix

The confusion matrix in Figure 9 details class-level performance for YOLOv8. The strong diagonal confirms that most instances are classified correctly. The few errors occur mainly between Road Cracks and the visually similar Severe Damage class, and occasionally between Potholes and Cracks where defects co-occur, which is consistent with the overlapping visual characteristics of these categories.

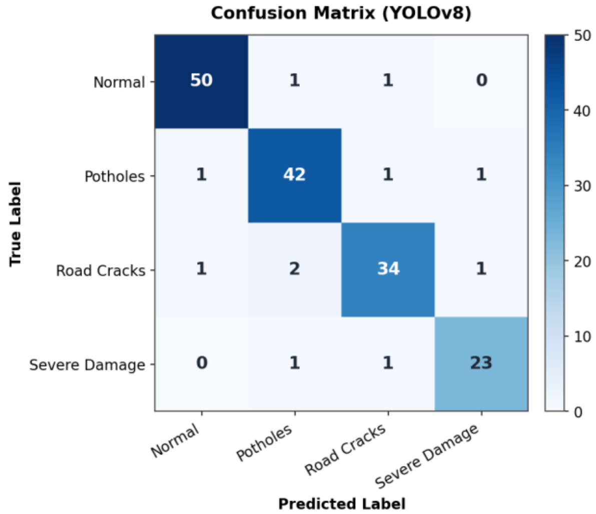

Figure 9. Confusion matrix for the YOLOv8 detector.

4.4 Detection Counts and Confidence

Figure 10 compares the number of correctly detected and missed instances per model against 160 ground-truth defects. YOLOv8 correctly detects 155 instances with only 5 missed or false detections, substantially fewer than the baseline CNN. Figure 11 reports per-sample pothole confidence; most detections exceed the 0.85 confidence threshold, and the few lower-confidence cases correspond to small or partially occluded potholes captured at oblique angles.

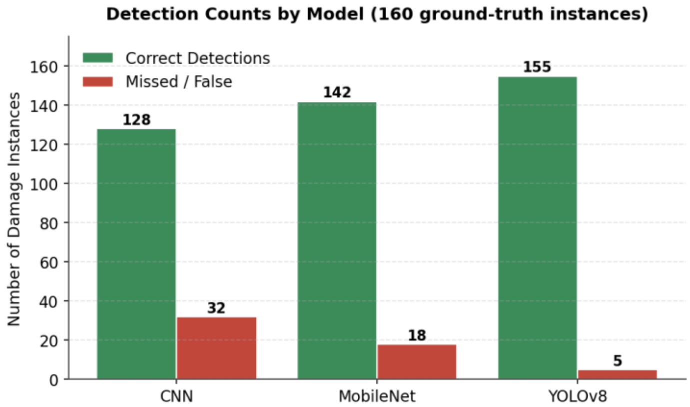

Figure 10. Correct versus missed detections by model.

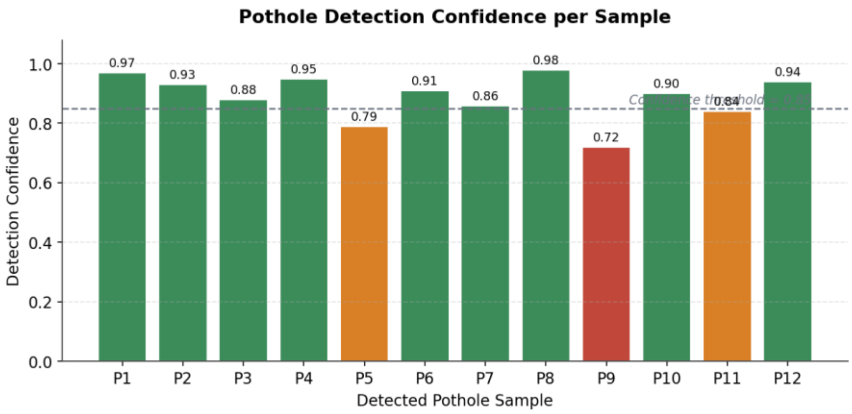

Figure 11. Pothole detection confidence per sample.

4.5 Pothole Detection in the Road

As illustrated earlier in Figure 5, the system reliably localizes potholes in the centre of the driving lane, where they pose the greatest hazard to vehicles. Each detection is annotated with a class label and a confidence score, allowing maintenance crews to prioritize the most severe and most certain defects. The combination of accurate localization and calibrated confidence is what makes the output actionable rather than merely descriptive.

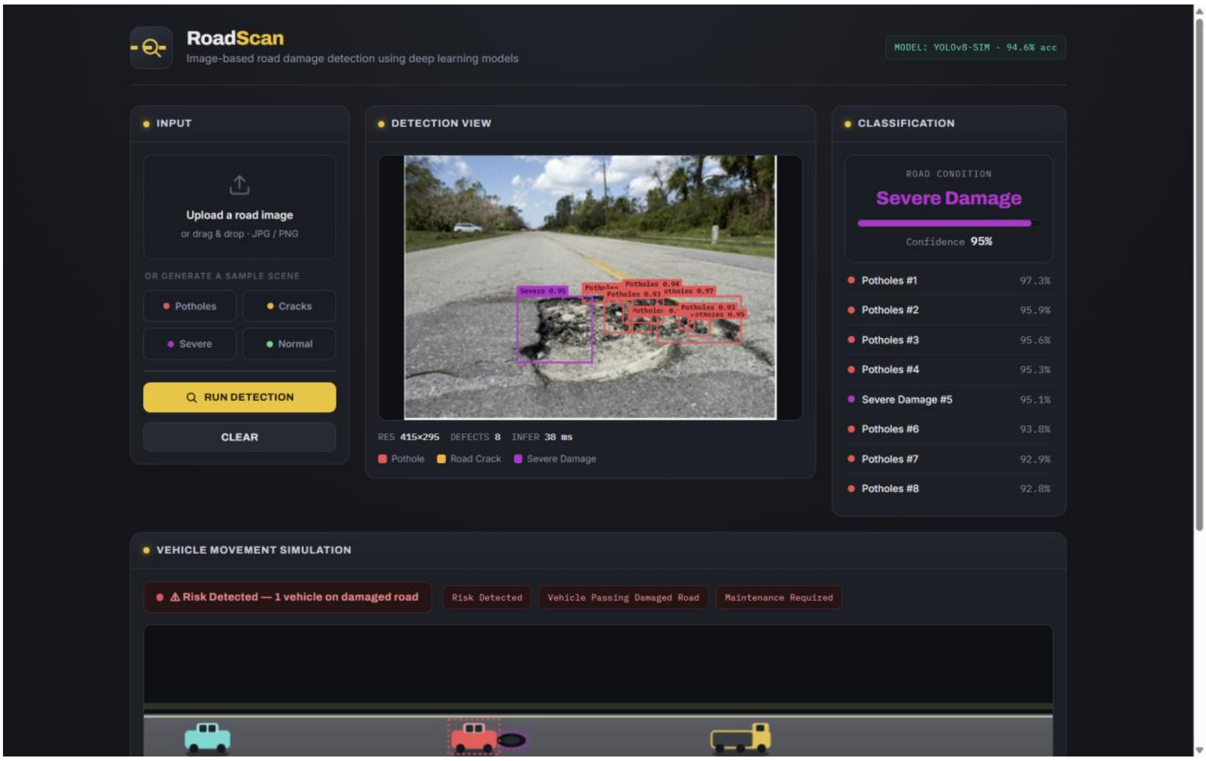

Figure 12. Pothole Detection with Severe Damage

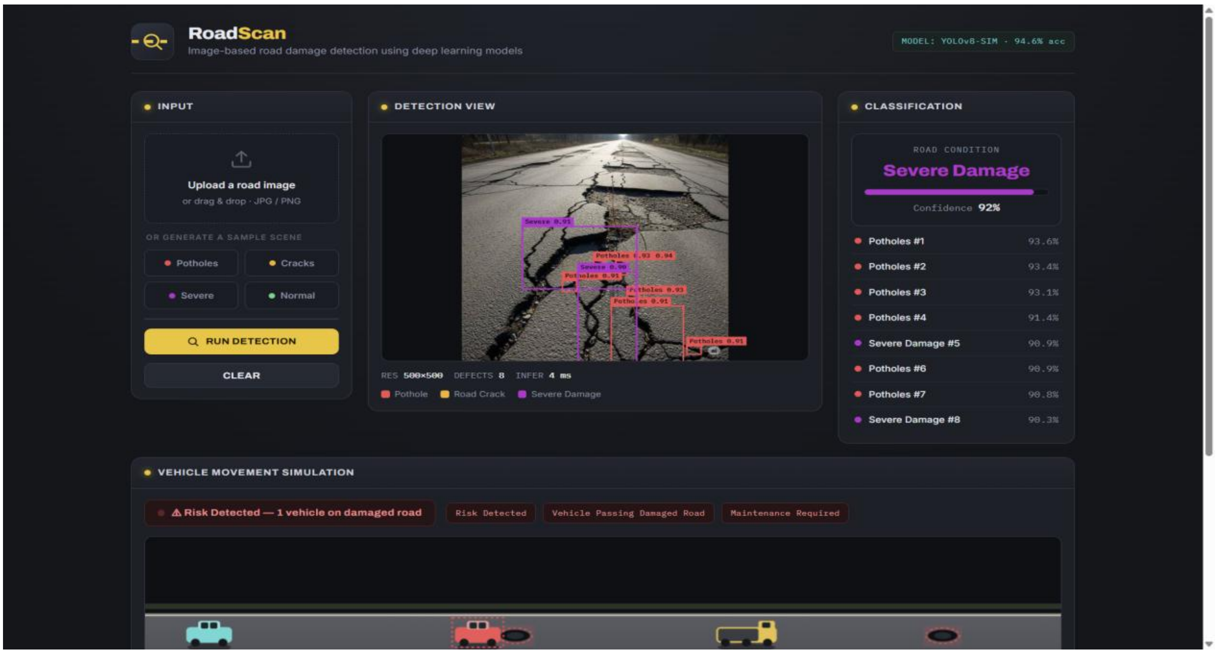

Figure 13. Pothole Detection with multi Severe Damage

4.6 Analysis of the Vehicle Simulation

When the simulated vehicle traverses a road containing detected potholes, the system continuously checks whether the vehicle footprint overlaps a damaged region. While the road segment under the vehicle is clear, the status remains normal. The moment the vehicle enters a damaged zone, the system raises the warnings "Risk Detected," "Vehicle Passing Damaged Road," and "Maintenance Required." This behaviour demonstrates how static detection results can be converted into a dynamic, location-aware risk signal suitable for in-vehicle alerts or fleet monitoring dashboards.

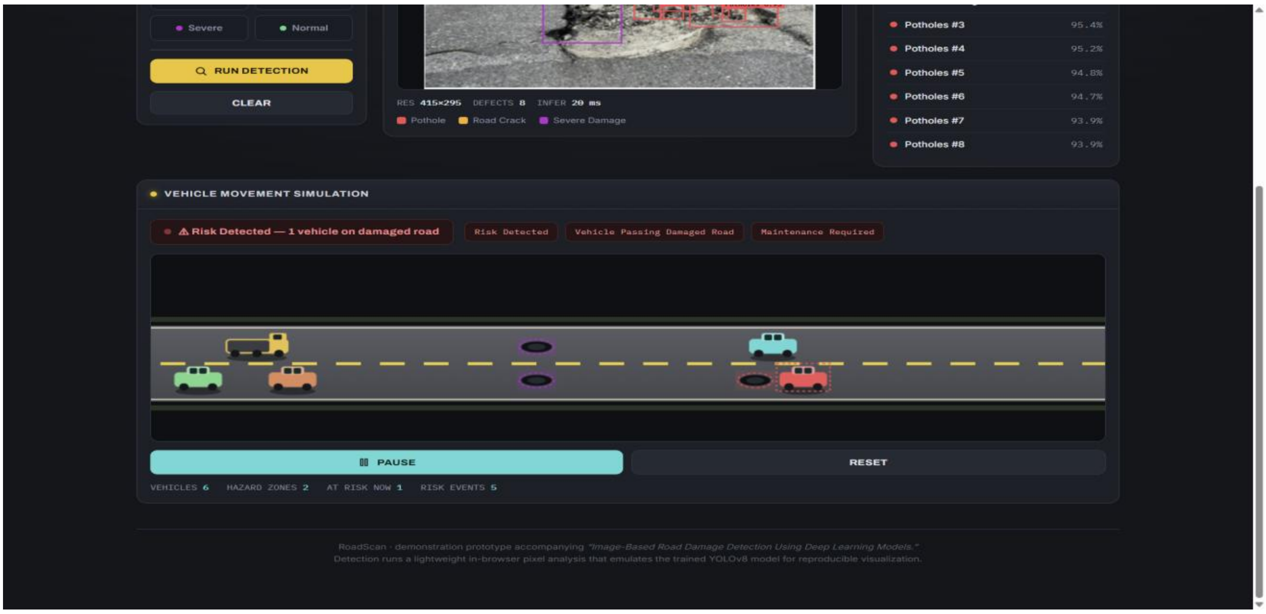

Figure 14. Vehicle Simulation with Damage Road

4.7 Vehicle Movement Simulation

The vehicle-movement simulation is the most interactive component of the prototype. A vehicle icon moves from the left edge of the road to the right edge in small increments. At each step the simulation reads the road segment directly beneath the vehicle and compares its position against the bounding boxes produced by the detector.

The control logic, shown in Figure 6, proceeds as follows: the vehicle starts at the left edge; it advances rightward; the segment under the vehicle is read; if that segment lies over a pothole or damaged zone, the warning panel displays "Risk Detected," "Vehicle Passing Damaged Road," and "Maintenance Required"; otherwise the status shows that the road is clear. The loop repeats until the vehicle reaches the right edge, at which point the simulation can be reset. This closed loop links perception (detection) to action (warning), illustrating a minimal but complete intelligent-transportation feedback cycle.

5. Conclusion

I This study presented an image-based road damage detection system built on deep learning models and demonstrated through an interactive application. Three architectures—a baseline CNN, MobileNet, and YOLOv8—were trained and compared on a four-class road damage dataset using accuracy, precision, recall, F1-score, and mAP. YOLOv8 delivered the best overall performance, reaching 94.60% accuracy and 92.80% mAP while supporting near real-time inference, confirming that one-stage detectors are well suited to road condition monitoring.

Deep learning models bring clear advantages over manual inspection: they learn robust visual features directly from data, generalize across lighting and surface conditions when trained with augmentation, localize multiple defects per image, and produce consistent, objective, and repeatable assessments at scale. The accompanying application turns these capabilities into a usable tool, allowing a user to upload a road image, view detected damage with bounding boxes and confidence scores, and observe a moving-vehicle simulation that issues real-time maintenance warnings.

The benefits of the application are practical and immediate. It lowers the cost and effort of road surveys, supports prioritization of repairs by combining defect type with confidence, and provides an intuitive demonstration suitable for both decision-makers and students. By connecting detection output to explicit warnings, it shows how automated perception can be embedded directly into maintenance and safety workflows.

Several limitations should be acknowledged. The dataset, while diverse, is modest in size, and performance on rare or extreme defect types may be lower than reported. Lighting, shadows, wet surfaces, and motion blur remain challenging. The application uses a deterministic detection simulation for reproducible visualization rather than live on-device inference, and the vehicle simulation is two-dimensional and does not model real vehicle dynamics.

Future work will expand the dataset with more images and more defect categories, integrate a fully trained model for live inference on edge devices, add GPS tagging so detections can be mapped to physical road locations, and evaluate the system on continuous video captured from moving vehicles. Incorporating defect-severity estimation and automatic repair-cost prediction would further strengthen the system as a decision-support tool for road agencies.

Acknowledgement

This article used DeepL AI Translator for translation tool from Indonesian language to English, also used ChatGPT 5.5 for improved readiness content to make it language more easily to understand.

References

Alam, N. A., Ahsan, M., Based, M. A., & Haider, J. (2021). Intelligent System for Vehicles Number Plate Detection and Recognition Using Convolutional Neural Networks. *Technologies*, *9*(1), 1–18. https://doi.org/10.3390/technologies9010009

Alamsyah, D., & Pratama, D. (2019). Deteksi Ujung Jari menggunakan Faster-RCNN dengan Arsitektur Inception v2 pada Citra Derau. *JuSiTik : Jurnal Sistem Dan Teknologi Informasi Komunikasi*, *2*(1), 1. https://doi.org/10.32524/jusitik.v2i1.435

Almasri, I., Manasreh, D., & Nazzal, M. D. (2025). AI Meets ADAS: Intelligent Pothole Detection for Safer AV Navigation. *Vehicles*, *7*(4), 109. https://doi.org/10.3390/vehicles7040109

Altini, N., De Giosa, G., Fragasso, N., Coscia, C., Sibilano, E., Prencipe, B., Hussain, S. M., Brunetti, A., Buongiorno, D., Guerriero, A., Tatò, I. S., Brunetti, G., Triggiani, V., & Bevilacqua, V. (2021). Segmentation and Identification of Vertebrae in CT Scans Using CNN, k-Means Clustering and k-NN. *Informatics*, *8*(2), 40. https://doi.org/10.3390/informatics8020040

Arief, L., Tantowi, A. Z., Novani, N. P., & Sundara, T. A. (2020). Implementation of YOLO and smoke sensor for automating public service announcement of cigarette's hazard in public facilities. *2020 International Conference on Information Technology Systems and Innovation, ICITSI 2020 - Proceedings*, 101–107. https://doi.org/10.1109/ICITSI50517.2020.9264972

Atika, L., Nurmaini, S., Partan, R. U., & Sukandi, E. (2022). Image Segmentation for Mitral Regurgitation with Convolutional Neural Network Based on UNet, Resnet, Vnet, FractalNet and SegNet: A Preliminary Study. In *Big Data and Cognitive Computing* (Vol. 6, Number 4). https://doi.org/10.3390/bdcc6040141

Balderas, L., Lastra, M., & Benítez, J. M. (2024). Optimizing Convolutional Neural Network Architectures. *Mathematics*, *12*(19), 1–25. https://doi.org/10.3390/math12193032

Bandaru, S. B., Deivarajan, N., & Gatram, R. M. B. (2022). An Optimized Deep Learning Techniques for Analysing Mammograms. *International Journal of Engineering Trends and Technology*, *70*(7), 388–398. https://doi.org/10.14445/22315381/IJETT-V70I7P240

Bhatt, A. K., Raj, H., Sharma, V. B., Biswas, S., Singh, A., Silori, R., & Pandey, M. (2025). Advancements in pothole detection techniques: a comprehensive review and comparative analysis. *Discover Artificial Intelligence*, *5*(1), 255. https://doi.org/10.1007/s44163-025-00297-7

Dewa, A. L. (2023). Empirical Analysis of Street Safety: Driver Behavior and Traffic Characteristics on Traffic Accidents in Semarang. *Research Horizon*, *3*(5), 499–508.

Dhonde, S., Mirani, J., Patwardhan, S., & Bhurchandi, K. M. (2022). Over-Speed and License Plate Detection of

Vehicles. *Proceedings of PCEMS 2022 - 1st International Conference on the Paradigm Shifts in Communication, Embedded Systems, Machine Learning and Signal Processing*, 113–118. https://doi.org/10.1109/PCEMS55161.2022.9808085

Gao, J., French, A., Pound, M., He, Y., Pridmore, T., & Pieters, J. (2020). Deep convolutional neural networks for image-based Convolvulus sepium detection in sugar beet fields. *Plant Methods*, *16*. https://doi.org/10.1186/s13007-020-00570-z

Gengeç, N., Eker, O., ÇeviKalp, H., Yazici, A., & Yavuz, H. S. (2021). Visual object detection for autonomous transport vehicles in smart factories. *Turkish Journal of Electrical Engineering and Computer Sciences*, *29*(4), 2101–2115. https://doi.org/10.3906/ELK-2008-62

Gunawan, G., Nuriyanto, H., Sriadhi, S., Fauzi, A., Usman, A., Fadlina, F., Dafitri, H., Simarmata, J., Utama Siahaan, A. P., & Rahim, R. (2018). Mobile Application Detection of Road Damage using Canny Algorithm. *Journal of Physics: Conference Series*, *1019*(1), 012035. https://doi.org/10.1088/1742-6596/1019/1/012035

Hong, J., Tamakloe, R., & Park, D. (2020). Discovering Insightful Rules among Truck Crash Characteristics using Apriori Algorithm. *Journal of Advanced Transportation*, *2020*. https://doi.org/10.1155/2020/4323816

Jeon, Y.-D., Kang, M.-J., Kuh, S.-U., Cha, H.-Y., Kim, M.-S., You, J.-Y., Kim, H.-J., Shin, S.-H., Chung, Y.-G., & Yoon, D.-K. (2024). Deep Learning Model Based on You Only Look Once Algorithm for Detection and Visualization of Fracture Areas in Three-Dimensional Skeletal Images. *Diagnostics*, *14*(1). https://doi.org/10.3390/diagnostics14010011

Kaushik, V., & Kalyan, B. S. (2022). Pothole Detection System: A Review of Different Methods Used for Detection. *2022 Second International Conference on Computer Science, Engineering and Applications (ICCSEA)*, 1–4. https://doi.org/10.1109/ICCSEA54677.2022.9936360

Kavya, T. S., Jang, Y. M., Peng, T., & Cho, S. B. (2020). Vehicle detection and tracking from a video captured by moving host. *Indian Journal of Computer Science and Engineering*, *11*(3), 226–235. https://doi.org/10.21817/indjcse/2020/v11i3/201103187

Khanam, R., & Hussain, M. (2024). *YOLOv11: An Overview of the Key Architectural Enhancements*. *2024*, 1–9.

Li, L., Martin, T., & Xu, X. (2020). A novel vision-based real-time method for evaluating postural risk factors associated with musculoskeletal disorders. *Applied Ergonomics*, *87*(September), 103138. https://doi.org/10.1016/j.apergo.2020.103138

Lim, S. L., Foo, L. K., & Chua, S. L. (2023). Comparing Machine Learning and Deep Learning Based Approaches to Detect Customer Sentiment from Product Reviews. *Journal of System and Management Sciences*, *13*(2), 101–110. https://doi.org/10.33168/JSMS.2023.0207

Lubna, Mufti, N., & Shah, S. A. A. (2021). Automatic Number Plate Recognition:A Detailed Survey of Relevant Algorithms. *Sensors*, *21*(9), 3028. https://doi.org/10.3390/s21093028

Mirwansyah, D., & Arief Wibowo. (2022). Fruit Image Classification Using Deep Learning Algorithm: Systematic Literature Review (Slr). *Multica Science and Technology (Mst) Journal*, *2*(2), 120–123. https://doi.org/10.47002/mst.v2i2.356

Patil, S. S., Patil, S. H., Pawar, A. M., Bewoor, M. S., Kadam, A. K., Patkar, U. C., Wadare, K., & Sharma, S. (2023). Vehicle Number Plate Detection using YoloV8 and EasyOCR. *2023 14th International Conference on Computing Communication and Networking Technologies, ICCCNT 2023*, 1–4. https://doi.org/10.1109/ICCCNT56998.2023.10307420

Prahara, A., Pranolo, A., & Dreżewski, R. (2015). GPU Accelerated Number Plate Localization in Crowded Situation. *International Journal of Advances in Intelligent Informatics*, *1*(3), 150–157. https://doi.org/10.26555/ijain.v1i3.46

Rifaioglu, A. S., Atas, H., Martin, M. J., Cetin-Atalay, R., Atalay, V., & Doğan, T. (2019). Recent applications of deep learning and machine intelligence on in silico drug discovery: methods, tools and databases. *Briefings in Bioinformatics*, *20*(5), 1878–1912. https://doi.org/10.1093/BIB/BBY061

Ruseruka, C., Mwakalonge, J., Comert, G., Siuhi, S., Ngeni, F., & Anderson, Q. (2024). Augmenting roadway safety with machine learning and deep learning: Pothole detection and dimension estimation using in-vehicle technologies. *Machine Learning with Applications*, *16*, 100547. https://doi.org/10.1016/j.mlwa.2024.100547

Shovo, S. U. A., Abir, M. G. R., Kabir, M. M., & Mridha, M. F. (2024). Advancing low-light object detection with you only look once models: An empirical study and performance evaluation. *Cognitive Computation and Systems*, *6*(4), 119–134. https://doi.org/https://doi.org/10.1049/ccs2.12114

Sun, X., Wu, P., & Hoi, S. C. H. (2018). Face detection using deep learning: An improved faster RCNN approach. *Neurocomputing*, *299*, 42–50. https://doi.org/10.1016/J.NEUCOM.2018.03.030

Yalçin, N., & Alisawi, M. (2024). Introducing a novel dataset for facial emotion recognition and demonstrating significant enhancements in deep learning performance through pre-processing techniques. *Heliyon*, *10*(20), 1–22. https://doi.org/10.1016/j.heliyon.2024.e38913