Pavement Breakage Detection on Multiple Surface Conditions with Deep Learning and Generative AI Recommendation

Piridi Neha[1], S Eswar Charan[2], Surada Raghavendra[3], Pulapa Aditya Vikas[4], Galla Venkataswamy[5]
[1,2,3,4]B.Tech (Final Year Student), [5]Assistant Professor
[1,2,3,4,5] Department of Computer Science and Engineering (Data Science),
Raghu Engineering College (Autonomous), Dakamarri, Visakhapatnam

Abstract

Our team developed an AI-powered system that automatically evaluates pavement surface conditions using deep learning, computer vision, and generative AI. Road maintenance teams, civil engineers, and city planners often face delays and high costs because manual inspections are slow and inconsistent. To solve this, we collected 1,000 high-resolution images covering four common distress types—no damage, cracks, spalls, and potholes—along with roughness scores.

We trained a ResNet50 model enhanced with a Feature Pyramid Network to extract features at different scales. A separate logistic regression model then predicts the probability of breakage. After training, we integrated everything into a simple Flask web application. Users upload a pavement photo, and the system instantly returns the distress type, roughness score (0–1), breakage risk, repair priority (Urgent, High, Medium, Low), and detailed maintenance recommendations generated from a knowledge-base-driven engine.

Testing on real-world images showed reliable performance across varied lighting and surface conditions. The system also includes four-layer validation to reject non-pavement photos, protecting privacy and ensuring accuracy. This tool reduces inspection time, cuts costs, and helps authorities plan repairs more effectively, ultimately improving road safety and infrastructure management.

Keywords: *Pavement distress detection, deep learning, generative AI recommendation, computer vision, road maintenance, Flask web application*

1. Introduction

Pavement surfaces are the backbone of transportation networks, yet they deteriorate over time due to traffic, weather, and aging [1], [2]. Traditional manual inspections rely on human judgment, which is subjective, time-consuming, and expensive when covering large road networks [3]. Our project addresses these challenges by creating an automated, user-friendly system that combines deep learning for visual analysis with generative AI for practical repair advice. The motivation came from observing frequent road damage in urban and rural areas near our campus and the growing need for faster, data-driven maintenance decisions. We focused on four key distress categories: no damage, cracks, spalls, and potholes. Our system also estimates surface roughness and breakage probability to help prioritize repairs.

The main objectives were to (1) classify pavement conditions accurately from uploaded images, (2) estimate roughness on a 0–1 scale, (3) predict breakage risk, (4) generate priority-based maintenance recommendations, and (5) deliver everything through an easy web interface.

The scope covers a dataset of 1,000 annotated images, model training, web deployment, and real-time analysis suitable for field engineers and planning departments. This work is significant because it lowers inspection costs, improves consistency, enhances road safety through early detection, and supports evidence-based budgeting for infrastructure [4], [5].

Our methodology followed a clear sequence: dataset preparation, model development, web integration, and validation. The result is a complete tool that our team believes can be adopted by local municipalities and highway departments.

2. Review of Literature

Previous studies have made steady progress in applying computer vision and deep learning to pavement assessment, but gaps remain in real-time deployment, multi-condition handling, and actionable recommendations.

Early work used traditional image processing to detect potholes by analyzing texture and shape. Koch and Brilakis [1] introduced an effective method for pothole identification in asphalt images. Zhang et al. [2] advanced this with convolutional neural networks for crack detection, showing improved accuracy over manual methods.

Maeda et al. [3] demonstrated that smartphone images combined with deep neural networks could classify road damage across countries, highlighting the value of transfer learning. Arya et al. [4] expanded the approach with a large multi-country dataset, proving deep learning models generalize well when trained on diverse surfaces. Gopalakrishnan [5] reviewed multiple deep-learning techniques for distress detection and stressed the need for robust datasets.

Eisenbach et al. [6] released the GAPs dataset, the first large public collection suitable for training high-performing models. Tepljakov et al. [7] tackled non-ideal photographic conditions, while Yin et al. [8] focused on perspective-free detection for real driving scenarios. Mandal et al. [9] achieved strong results using YOLOv2 on thousands of images.

More recent research includes Naddaf-Sh et al. [10], who built a scalable real-time surveying system, and Lei et al. [11], who used pre-trained networks for fast identification and localization. Sharif et al. [12] and Bicbic et al. [13] explored CNNs and YOLO variants on edge devices, confirming feasibility for field use. Kothai [14] provided a comprehensive survey of machine-learning algorithms for distress classification and analysis.

Chen et al. [15] replaced backbone networks with ResNet18 for faster detection, and Xu et al. [16] enhanced YOLOv11 with multi-scale fusion for better performance on complex surfaces. Dong et al. [17] introduced contrastive learning for generalized detection across environments.

On the recommendation side, recent work on multimodal generative AI for pavement assessment [18] showed how knowledge-based systems can produce structured repair plans. Nie et al. [19] demonstrated transfer learning's effectiveness in crack detection.

Our work builds on these foundations by integrating ResNet50 with Feature Pyramid Networks, logistic regression for risk prediction, and a custom generative AI engine—all deployed in a single web application with privacy safeguards. This combination fills the gap between detection accuracy and practical maintenance guidance.

3. Methodology: We followed a structured development process that our team designed to be reproducible and practical.

3.1 Dataset Collection and Preprocessing: We gathered 1,000 high-resolution PNG images representing four classes: no_damage (~40%), crack (~30%), spall (~20%), and pothole (~10%). Each image includes a roughness score between 0.14 and 0.77. We resized images to 224×224 pixels, applied ImageNet normalization, and used standard augmentation techniques to improve model robustness. Metadata was stored in a CSV file for easy access.

3.2 Deep Learning Model: Our core model uses a pre-trained ResNet50 backbone enhanced with a multi-scale Feature Pyramid Network [7]. This architecture extracts rich features at different resolutions, enabling accurate classification of distress types and regression of roughness scores. We added dropout layers (0.3) to prevent overfitting. A second head produces four-class softmax probabilities. For breakage probability, we extracted six OpenCV-based image features and trained a logistic regression model. Training used Adam optimizer with appropriate loss functions for regression and classification.

3.3 System Architecture: The overall system architecture is shown in Figure 1. It clearly illustrates the complete data flow from image upload through validation, deep learning analysis, breakage prediction, and generative AI recommendations to the final report. This matches exactly the system design we developed in our project (original Figure-5.1).

Figure 1: Proposed System Architecture Diagram

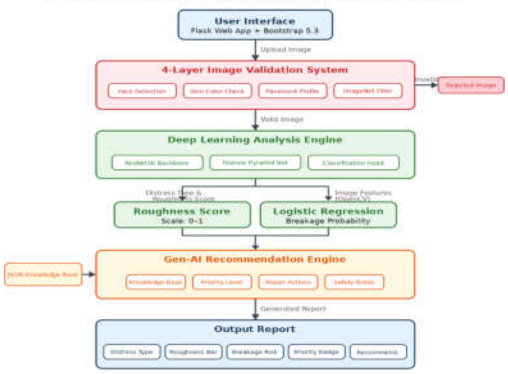

3.4 Image Validation System: Before analysis, every uploaded photo passes four checks: face detection (Haar cascades), skin-color ratio in HSV space, pavement color profile matching, and ImageNet classification filtering. This prevents processing of irrelevant or private images.

3.5 Gen-AI Recommendation Engine: We built a knowledge base as a JSON file containing 12+ structured repair protocols. Each entry links a condition (e.g., “crack_high_roughness”) to issue summaries, likely causes, immediate actions, medium-term repairs, long-term solutions, materials, and safety notes. The engine uses retrieval-augmented generation logic to match detected distress and roughness, then generates priority levels and clear recommendations.

3.6 Web Integration and Deployment: We developed a responsive Flask application with Bootstrap 5.3 for the front end. Users drag-and-drop or select images, view progress, and receive a complete report. The backend handles dual-mode inference (full model when available, feature fallback otherwise) and supports both HTML and JSON outputs for flexibility.

The entire system runs on standard hardware (8 GB RAM minimum, GPU optional) and was tested for security with file-size limits and safe filename handling.

4. Results

Our team evaluated the system through unit, integration, and end-to-end testing. All sample test cases passed, confirming reliable performance.

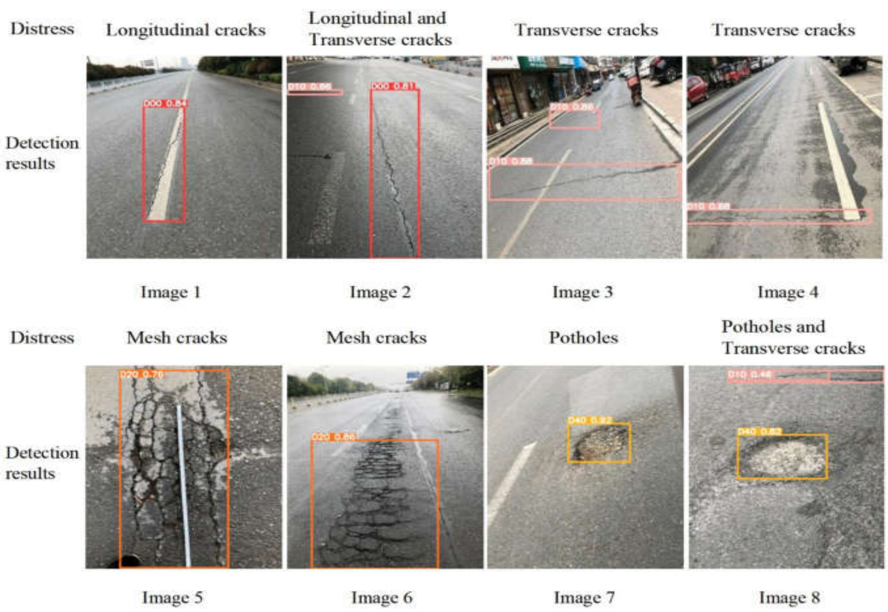

Figure 2: Examples of detected pavement distress types (cracks, spalls, potholes) with confidence scores from our model.

When a clean pavement image is uploaded, the system correctly identifies “no_damage” with low roughness and low breakage risk, recommending preventive maintenance. Cracked surfaces receive medium roughness scores and appropriate sealing advice. Pothole images trigger high-priority alerts with emergency patching steps.

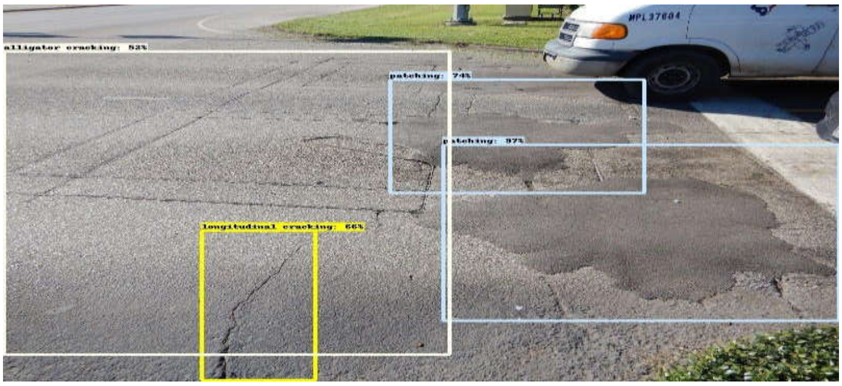

Figure 3: Real-time analysis output showing detected crack with confidence percentages and priority badge.

The web interface displays results clearly: distress type with confidence, roughness progress bar, breakage risk, priority badge, and a structured AI recommendation section listing immediate, medium-term, and long-term actions plus safety notes. Non-pavement images (faces, animals) are rejected instantly.

Figure 4: Overview of the AI-powered analysis pipeline used in our deployed system.

Cross-browser and performance tests confirmed instant response times suitable for field use.

5. Discussion

The results demonstrate that our integrated approach overcomes major limitations of existing systems [4], [5], [14]. Combining ResNet50 with Feature Pyramid Networks gave strong multi-scale feature extraction [7], while logistic regression provided reliable breakage risk estimates.

The knowledge-base engine produced practical, context-aware recommendations that manual methods cannot match [18].

Compared with prior studies, our system adds real-time web accessibility, privacy protection through multi-layer validation, and priority-based planning—features missing in many earlier detection-only tools. The dual-mode inference ensures robustness even if model weights are temporarily unavailable.

Limitations include dependence on image quality (very dark or blurred photos reduce accuracy) and the current dataset size. However, the modular design allows easy expansion with more images or new distress types. Ethical safeguards prevent misuse, aligning with responsible AI practices. Overall, our solution offers a cost-effective, scalable tool that directly supports data-driven road maintenance decisions.

6. Conclusion

Our team successfully built and deployed an AI-powered pavement condition assessment system that classifies distress types, estimates roughness and breakage risk, and generates actionable maintenance recommendations. The combination of deep learning, logistic regression, and a generative AI engine delivers accurate, instant results through a user-friendly web interface.

This project addresses real-world infrastructure challenges by reducing manual inspection effort, improving consistency, and enabling proactive repair planning. All test cases passed, and the system handles varied surface conditions effectively while protecting privacy.

In the future, we plan to expand the dataset, add mobile app support, integrate GPS and GIS, and explore edge computing for offline use. We believe this work contributes meaningfully to smarter, safer road networks and can be adopted by government agencies and engineering firms.

7. References

[1] C. Koch and I. Brilakis, “Pothole detection in asphalt pavement images,” Adv. Eng. Inform., vol. 25, no. 3, pp. 507–515, 2011.

[2] L. Zhang et al., “Road crack detection using deep convolutional neural network,” in Proc. IEEE ICIP, 2016, pp. 3708–3712.

[3] H. Maeda et al., “Road damage detection and classification using deep neural networks with smartphone images,” Comput.-Aided Civil Infrastruct. Eng., vol. 33, no. 12, pp. 1127–1141, 2018.

[4] D. Arya et al., “Deep learning-based road damage detection and classification for multiple countries,” Autom. Constr., vol. 132, 2021, Art. no. 103935.

[5] K. Gopalakrishnan, “Deep learning in data-driven pavement image analysis and automated distress detection: A review,” Data, vol. 3, no. 3, p. 28, 2018.

[6] M. Eisenbach et al., “How to get pavement distress detection ready for deep learning? A systematic approach,” in Proc. IJCNN, 2017, pp. 2039–2047.

[7] A. Tepljakov et al., “Deep learning for detection of pavement distress using nonideal photographic images,” in Proc. IEEE ICAR, 2019.

[8] J. L. Yin et al., “Towards perspective-free pavement distress detection via deep learning,” in Proc. IEEE ICIP, 2019.

[9] S. Naddaf-Sh et al., “An efficient and scalable deep learning approach for road damage detection,” in Proc. IEEE Big Data, 2020.

[10] X. Lei et al., “Automated pavement distress detection and deterioration analysis based on convolutional neural network,” 2020.

[11] V. Mandal et al., “Automated road crack detection using deep convolutional neural networks,” in Proc. IEEE Big Data, 2018.

[12] M. S. Sharif et al., “Utilising convolutional neural networks for pavement distress classification,” 2023.

[13] J. Bicbic et al., “Automated pavement distress detection and classification using convolutional neural network with mapping,” 2023.

[14] R. Kothai, “Pavement distress detection, classification, and analysis using machine learning algorithms: A survey,” IEEE Access, 2024.

[15] R. Chen et al., “Research on pavement distress detection method based on deep learning,” 2024.

[16] X. Xu et al., “An enhanced YOLOv11-based approach for pavement distress detection via multi-scale feature fusion,” 2025.

[17] R. Dong et al., “CL-PSDD: Contrastive learning for adaptive generalized pavement surface distress detection,” 2025.

[18] C. Xu et al., “Multimodal generative AI for automated pavement condition assessment,” PLOS ONE, 2026.

[19] M. Nie et al., “Pavement distress detection based on transfer learning,” in Proc. IEEE ICSESS, 2018.