



International Journal of Advanced Research in Education and Technology (IJARETY)

Volume 13, Issue 3, May - June 2026

Impact Factor: 9.971



Pothole and Speed Breaker Detection System Using Deep Learning

H N Sumantha

Department of Master of Computer Applications, CMR Institute of Technology, Visvesvaraya Technological University, Bangalore, India

ABSTRACT: Broken roads create serious problems, impacting how safely people travel while raising expenses for car repairs. Bumps like potholes or badly built speed bumps disrupt smooth driving, sometimes holding up ambulances or fire trucks. Keeping track of pavement quality nonstop isn't easy, particularly where cities are expanding fast. Old ways of checking streets depend on workers walking routes or tools measuring shakes in vehicles; these approaches struggle to grow, react poorly to weather shifts, and fail to deliver steady live updates. This study introduces a live method to spot potholes and speed bumps using smart image analysis. Built around a slimmed-down neural network derived from MobileNetV2, it uses prior learning patterns to sort roads into one of three types: undamaged, cracked with holes, or featuring raised platforms. Tests show it spots things right most times while reacting quickly, shaping up into something cities could roll out widely without draining budgets for traffic oversight networks.

KEYWORDS: Deep Learning, MobileNetV2, Transfer Learning, Road Hazard Detection, Image Processing, Computer Vision.

I. INTRODUCTION

Getting around by road keeps economies moving and people linked in today's world. When roads are kept in good shape, things move easier — trucks deliver faster, people reach places quicker, safety goes up too. Yet holes in the pavement or badly built speed bumps without signs create serious problems. Bumpy surfaces lead to broken cars, more fuel burned, jams on busy routes, sometimes even crashes with harm or loss of life.

Water sneaks into roads, then vehicles keep driving over weak spots, so damage grows slowly but surely. Tiny splits appear first, yet they grow wider when rain gets underneath and breaks things apart. Bumps meant to slow cars sometimes cause trouble if built wrong, creating jolts rather than control. Roads start failing from these flaws, which is why noticing problems fast makes a difference later on.

Old methods such as workers walking miles scanning pavement cracks or vibration-based sensors suffer from inconsistency, poor scalability, and inability to distinguish damage types without visual support. Progress in computer vision and convolutional neural networks (CNNs) has enabled more reliable image-based detection. Models like VGGNet, ResNet, and Inception excel at classification but require heavy processing power. MobileNetV2 addresses this with depthwise separable convolutions, cutting computation while maintaining accuracy. This study introduces a live pothole and speed bump detector built with transfer learning based on MobileNetV2, suitable for deployment on regular computers in real-time road surveillance tasks.

II. LITERATURE REVIEW

Spotting damage on roads draws much scientific interest because it affects both travel safety and upkeep of public works. Methods meant to catch these flaws have evolved from manual inspection all the way to deep neural networks. Manual inspection, while straightforward, is time-consuming, inconsistent, and cannot scale to large urban networks.

Researchers subsequently developed accelerometer and vibration-based systems placed directly on cars. Though more automated, these systems depend heavily on vehicle speed, shock system design, and sensor calibration. Without image support, distinguishing damage types remains a challenge. Later, image-based approaches using filters, contrast checks, and pixel grouping improved detection but struggled under uneven light, painted lines, or messy backgrounds. Their accuracy dropped outside controlled settings.

Recent advances in deep learning, especially CNNs, have boosted road damage detection. Transfer learning using pre-trained ImageNet weights enables faster training and better performance with limited data. Starting from recent progress, this study applies MobileNetV2 alongside transfer learning to build a system that detects potholes and speed breakers instantly, supporting deployment across diverse environments where resources are limited.

III. METHODOLOGY

A sequence of connected steps forms the core of the suggested pothole and speed breaker identification method. Built around real-time analysis, it splits tasks into separate units that work together without tight coupling. Vision techniques feed data into trained models which interpret road surface patterns. Design choices support growth, yet keep computation demands within practical limits.

A. Overall System Workflow

The overall workflow consists of the following sequential stages: (1) Live Image Acquisition, (2) Image Preprocessing, (3) Feature Extraction using Convolutional Neural Network, (4) Classification using MobileNetV2, (5) Confidence Threshold Filtering, (6) Multi-frame Validation, and (7) Visual Output Display. Each stage processes the output of the previous stage, ensuring systematic data flow and reduced computational redundancy.




Fig. 1. System Perspective Diagram of the Proposed Detection System

B. Image Acquisition

Using a regular webcam paired with OpenCV, images are taken on demand. While the system runs, new frames stream in without pause. One by one, these snapshots feed into the classifier as separate inputs. Not stopping, the camera keeps grabbing frames so motion can be tracked instead of checking single snapshots, improving performance for outdoor, on-the-move conditions.

C. Image Preprocessing

Before classification, each frame undergoes preprocessing to standardize input format and improve model inference accuracy. The preprocessing pipeline includes: resizing image to 224×224 pixels, converting color format from BGR to RGB, normalizing pixel intensity values ($I_{no} \leftarrow_m(x,y) = I(x,y)/255$), and expanding dimensions to match batch input format. Normalization improves convergence stability and prediction reliability.

D. Deep Learning Model Architecture

A lightweight deep learning model processes image data during classification, leveraging MobileNetV2 architecture fine-tuned for performance. MobileNetV2 applies depthwise separable convolutions to reduce computational load while maintaining accuracy. Transfer learning is applied by loading pre-trained ImageNet weights. The final fully connected layer is replaced with a custom classification layer containing three output neurons: Normal Road, Pothole, and Speed Breaker. The Softmax activation function is applied to produce class probabilities.

E. Training Strategy

Training employs categorical cross-entropy loss. The Adam optimizer is used alongside data augmentation methods including image rotation, brightness adjustment, zoom, and horizontal flipping to boost generalization. The loss

function is defined as: $L = -\sum y_i \log(\hat{y}_i)$, where y_i represents the true label and $\hat{y}_i$ represents the predicted probability. Training is conducted over multiple epochs until convergence is achieved.

F. Confidence Thresholding & Multi-Frame Validation

When confidence scores fall below a set level, the system ignores them, preventing weak signals from causing unnecessary notifications. For multi-frame validation, the system checks several frames in a row instead of just one — only when the same outcome shows up repeatedly does it accept the finding as real, improving robustness and ensuring reliable real-time detection.

TABLE I. Major System Components and Their Functions

Component	Function
Camera Module	Captures live road images
Preprocessing Module	Standardizes image input
MobileNetV2 Model	Extracts features and classifies images
Validation Module	Applies threshold and multi-frame checks
Alert Interface	Displays detection result to user

IV. RESULTS AND DISCUSSION

The system was tested under controlled and real-time conditions. The model demonstrated reliable classification across three categories: Normal Road, Pothole, and Speed Breaker. The system achieved near real-time performance with minimal latency. Multi-frame validation significantly improved reliability by filtering out glare-induced false detections.




Fig. 2. Application Home Window



Fig. 3. Normal Road Detection



Fig. 4. Speed Breaker Detection



Fig. 5. Pothole Detection

TABLE II. Performance Evaluation

Parameter	Observation
Detection Accuracy	High
Response Time	1–2 seconds
False Detection Rate	Low
System Stability	Stable

V. CONCLUSION

The proposed Pothole and Speed Breaker Detection System effectively classifies road hazards using MobileNetV2-based transfer learning. The system provides real-time detection with high computational efficiency.

A. Advantages

- Real-time detection
- Lightweight model suitable for regular computers
- Cost-effective implementation
- Easy deployment without specialized hardware

B. Disadvantages

- Sensitive to extreme lighting conditions
- Limited dataset for training
- No GPS integration for location tagging

C. Future Scope

Future enhancements include GPS tagging for geo-referenced hazard reporting, cloud integration for centralized monitoring, mobile deployment for wider access, and adoption of advanced object detection models such as YOLO or Faster R-CNN for improved detection granularity.

REFERENCES

- [1] A. Howard et al., "MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications," arXiv:1704.04861, 2017.
- [2] A. Krizhevsky, I. Sutskever, and G. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," in Proc. NIPS, 2012.
- [3] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in Proc. CVPR, 2016.
- [4] C. Szegedy et al., "Going Deeper with Convolutions," in Proc. CVPR, 2015.
- [5] M. Sandler et al., "MobileNetV2: Inverted Residuals and Linear Bottlenecks," in Proc. CVPR, 2018.
- [6] R. Girshick, "Fast R-CNN," in Proc. ICCV, 2015.
- [7] J. Redmon et al., "You Only Look Once: Unified, Real-Time Object Detection," in Proc. CVPR, 2016.
- [8] W. Liu et al., "SSD: Single Shot MultiBox Detector," in Proc. ECCV, 2016.
- [9] D. Ciresan et al., "Multi-column Deep Neural Network for Traffic Sign Classification," Neural Networks, 2012.
- [10] S. Ren et al., "Faster R-CNN: Towards Real-Time Object Detection," IEEE TPAMI, 2017.
- [11] Y. LeCun, Y. Bengio, and G. Hinton, "Deep Learning," Nature, vol. 521, pp. 436–444, 2015.
- [12] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. MIT Press, 2016.
- [13] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," in Proc. ICLR, 2015.
- [14] M. Everingham et al., "The Pascal Visual Object Classes Challenge," IJCV, 2010.
- [15] S. Maeda et al., "Road Damage Detection Using Deep Neural Networks," in Proc. ICIP, 2018.

International Journal of Advanced Research in Education and Technology

ISSN: 2394-2975

Impact Factor: 9.971