Open Access and Peer Reviewed Journal ISSN 2394-2231 https://www.ijctjournal.org/

RESEARCH ARTICLE OPEN ACCESS

An AI and GIS-Based Framework for Pothole Severity Assessment, Repair Priority Ranking, and Cost Prediction

Vaishnavi Ghogare, Ashwini Bhavar, Rohan Shelke

School of Information Technology, Shree Chanakya Education Society's Indira University, Pune, Maharashtra, India

vaishnavi.ghogare@indirauniversity.edu.in | ashwni.bhavar@indirauniversity.edu.in | Rohan.shelke@indirauniversity.edu.in

Under the Guidance of Mr. Nitin Joshi and Dr. Manisha Patil

Abstract—Potholes are a major form of road deterioration that can affect road safety, vehicle performance, traffic movement, and transportation infrastructure. Traditional road inspection mainly depends on manual surveys and complaint-based reporting, which can be time-consuming, costly, and difficult to apply across large road networks. This paper proposes an integrated Artificial Intelligence (AI) and Geographic Information System (GIS)-based framework for automated pothole detection, severity assessment, repair-priority ranking, and approximate cost prediction. The framework uses YOLO-based computer vision to detect potholes from road images and extract visual characteristics including bounding-box width, height, area, aspect ratio, and relative image coverage. These features are intended for machine-learning models such as Random Forest and XGBoost to classify potholes into Low, Medium, and High severity categories. GIS-based analysis then incorporates contextual information such as road class, traffic conditions, proximity to critical services, and pothole distribution. Multi-Criteria Decision Analysis (MCDA) combines severity and spatial risk factors to generate repair-priority rankings. Regression models are proposed for approximate repair-cost prediction using defect characteristics, repair category, and road type. The current cleaned dataset contains 1,740 pothole records with image and bounding-box information. Because actual repair-cost labels are not present in the current dataset, cost prediction is defined as a subsequent stage requiring supplementary maintenance-cost records. The main contribution is a unified decision-support framework connecting visual detection, severity assessment, geographic context, repair prioritization, and cost planning.

Keywords—Pothole Detection, Artificial Intelligence, Computer Vision, YOLO, Machine Learning, GIS, Random Forest, XGBoost, MCDA, Road Maintenance.

1. Introduction

Road transportation infrastructure is fundamental to mobility, economic activity, and urban development. Potholes are among the most visible and disruptive forms of pavement deterioration. They can reduce ride quality, damage vehicles, interrupt traffic movement, and create safety risks. Effective maintenance therefore depends not only on identifying damaged locations but also on understanding their severity, geographical context, urgency, and expected repair requirements.

require considerable human effort and may be d frequently across extensive road networks. Computer vision and machine learning provide an opportunity to automate the identification of visible pavement defects from road images. Recent road-damage research has demonstrated the use of YOLO-family object detectors and benchmark datasets such as RDD2022 for automated road-damage detection.

However, detection is only the first stage of a maintenance decision. A detected pothole still needs to be assessed for severity and maintenance priority. Location also matters: a defect on a high-traffic road or near a critical public facility may present a different risk from a similar defect on a low-traffic road. GIS can provide this spatial context, while Multi-Criteria Decision Analysis (MCDA) can combine heterogeneous factors into an interpretable priority score.

This research therefore proposes an integrated AI and GIS framework that connects pothole detection, visual feature extraction, severity classification, spatial risk analysis, repair prioritization, and approximate repair-cost prediction. The framework is designed as a decision-support approach rather than a replacement for engineering inspection.

2. Literature Review

Research on automated road-condition monitoring has increasingly moved from manual inspection and conventional image processing toward deep-learning-based computer vision. Traditional image-processing approaches may use edges, texture, colour, and intensity, but their performance can be sensitive to lighting, shadows, weather, pavement texture, and image quality. Deep-learning detectors learn visual representations from annotated training images and can directly localize road defects.

YOLO is a single-stage object-detection family widely used when object localization and computational efficiency are important. Road-damage studies have adapted YOLO architectures to detect potholes and other pavement defects. The RDD2022 dataset is a major benchmark resource, containing 47,420 road images from six countries and more than 55,000 annotated road-damage instances across four damage categories.

Severity assessment is distinct from detection. Detection answers whether and where a pothole is present, whereas severity assessment estimates how serious the defect is. The proposed study considers Low, Medium, and High severity categories and identifies Random Forest and XGBoost as candidate classifiers. Classification can be evaluated using accuracy, precision, recall, F1-score, and a confusion matrix.

GIS adds geographical context to image-derived information. Relevant contextual variables include latitude and longitude, road type, road importance, traffic level, nearby intersections, schools, hospitals, and other critical public locations. A multi-factor priority score can combine these variables with severity to support maintenance planning.

Cost prediction is another decision-support component. Candidate regression methods include Linear Regression, Random Forest Regressor, XGBoost Regressor, and Gradient Boosting Regressor. Appropriate metrics include Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and R-squared (R^2). Reliable cost prediction requires actual repair-cost records; such labels are not present in the current 1,740-record pothole dataset.

automatic detection and localization. Although these systems can identify potholes, detection alone does not provide a complete maintenance decision. The practical workflow also requires severity assessment, geographic interpretation, prioritization, and budget planning.

The research gap addressed here is the lack of a unified framework connecting: (1) pothole detection, (2) severity assessment, (3) geographic context, (4) repair-priority ranking, and (5) cost prediction. The proposed framework integrates these stages into a common pipeline and uses the current 1,740-record pothole dataset as a foundation for visual-feature analysis and machine-learning experiments.

4. Objectives

1) Automatically detect potholes from road images using a YOLO-based object detector.
2) Extract quantitative visual features from detected potholes.
3) Classify potholes into Low, Medium, and High severity categories using machine-learning models.
4) Integrate pothole records with GIS-based road and location context.
5) Generate repair-priority rankings using MCDA.
6) Develop an approximate repair-cost prediction stage when suitable historical cost labels are available.
7) Provide an integrated framework to support data-driven road-maintenance planning.

5. Proposed Methodology

The proposed framework consists of five connected stages: road-image input, pothole detection and feature extraction, severity classification, GIS-based risk analysis and MCDA prioritization, and cost prediction. The output is a structured pothole record containing visual, geographic, severity, priority, and—when cost data are available—estimated cost attributes.

5.1 Data Collection and Preparation

The current study uses a cleaned pothole dataset containing 1,740 records. The available information includes image and bounding-box measurements. The data should be checked for duplicate records, invalid annotations, missing values, and inconsistent measurements before model development. For object detection, images and annotations are divided into training, validation, and testing subsets. Augmentation may be applied to improve robustness to realistic variations in illumination, scale, and viewpoint.

5.2 YOLO-Based Pothole Detection

A YOLO-based detector is used to identify potholes and return bounding boxes with confidence scores. The detector output provides the spatial location of each pothole in the image. The bounding-box measurements form the basis for downstream feature extraction.

5.3 Visual Feature Extraction

For each detection, width, height, bounding-box area, aspect ratio, relative image area, confidence, and number of potholes can be recorded. Bounding-box area is $A = W \times H$. Relative image coverage is $C = A_p / A_i$, where A_p is pothole bounding-box area and A_i is total image area. Aspect ratio can be represented as $R = W / H$.

XGBoost classifiers. The target variable is the severity category $S \in$ {Low, Medium, High}. Severity labels must be defined using an explicit annotation protocol or validated engineering criteria before supervised training. A rule-based threshold system may be used as a baseline, but it should not be presented as ground truth unless supported by expert annotation.

5.5 GIS Integration

Each pothole is linked to geographic information where coordinates are available. GIS attributes can include latitude, longitude, road type, road importance, traffic level, proximity to intersections, schools, hospitals, and other critical public locations. OpenStreetMap can be used as a source of road-network and point-of-interest information when appropriate.

5.6 MCDA Repair-Priority Ranking

A normalized priority score is calculated from multiple criteria. A general formulation is $P = \Sigma(w_i x_i)$, where P is the priority score, w_i is the weight of criterion i, and x_i is its normalized value. Candidate criteria are pothole severity, traffic level, road importance, critical-location risk, number of reports, and report age. The resulting score can be mapped to Immediate, High, Moderate, and Low priority categories after threshold selection and validation.

5.7 Cost Prediction

When historical repair-cost records become available, regression models can estimate approximate cost from pothole size, severity, repair category, road type, and other relevant variables. Candidate models are Linear Regression, Random Forest Regressor, XGBoost Regressor, and Gradient Boosting Regressor. Cost estimates are intended for planning and should not be treated as final contractor quotations.

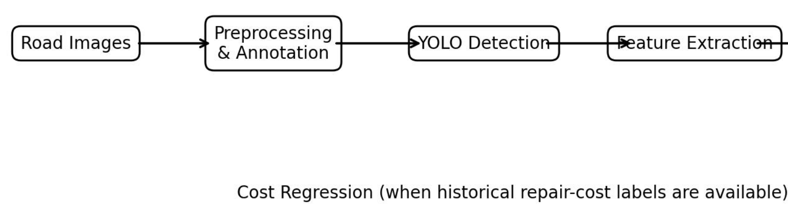

Fig. 1. Proposed integrated AI and GIS framework for pothole maintenance decision support.

6. Dataset Description

The current cleaned dataset contains 1,740 pothole records. According to the supplied research material, the dataset provides image and bounding-box information, including measurements such as width, height, and area. Aspect ratio and relative image area can be derived from these measurements.

Because the supplied dataset does not contain validated severity labels, GIS attributes, traffic information, report history, or actual repair-cost labels, those fields should be added or derived before the corresponding supervised models can be trained and evaluated. This distinction is important to avoid reporting unsupported experimental results.

1	YOLO-based detection	Pothole bounding boxes and confidence
2	Feature extraction	Width, height, area, aspect ratio, coverage
3	Random Forest / XGBoost	Low / Medium / High severity
4	GIS analysis	Road and spatial context
5	MCDA	Repair-priority score
6	Regression models	Approximate repair cost when labels exist

7. Model Evaluation

Object detection should be evaluated using Intersection over Union (IoU), precision, recall, and mean Average Precision (mAP). IoU is defined as IoU = Area(Bpred ∩ Btrue) / Area(Bpred ∪ Btrue).

For severity classification, Accuracy = correctly classified samples / total samples. Precision, Recall, and F1-score should be reported per class and as appropriate macro/weighted averages. A confusion matrix should be included to identify confusion among Low, Medium, and High classes.

For cost regression, MAE, RMSE, and R^2 should be reported. These values should only be calculated after genuine cost labels are collected and a train/test protocol is defined.

8. System Architecture

The proposed architecture follows the sequence: Road Images → Preprocessing → YOLO Pothole Detection → Bounding-Box Feature Extraction → Random Forest/XGBoost Severity Classification → GIS Spatial Context → MCDA Priority Ranking → Cost Regression → Maintenance Decision Dashboard.

The dashboard can present each pothole's image, geographic location, severity, priority level, risk factors, and estimated cost when available. This structure transforms raw image detections into information that can support maintenance planning.

9. Expected Results

The study is expected to demonstrate that automated image analysis can identify potholes and generate quantitative visual features at scale. Machine-learning classifiers are expected to provide a consistent mechanism for severity assessment once reliable severity labels are available. GIS integration is expected to add contextual information that is not available from images alone, while MCDA can combine severity and spatial factors into a transparent repair-priority score.

No numerical accuracy, F1-score, mAP, MAE, RMSE, or R^2 values are reported in this paper because the supplied materials do not contain completed experimental results for these stages. Reporting invented values would not constitute a valid research result.

10. Discussion

The central contribution of the proposed work is integration. A computer-vision detector answers where a pothole appears in an image, but maintenance agencies need additional information to decide what to do next. Severity classification adds an estimate of defect seriousness. GIS provides information about the environment in which the defect occurs. MCDA provides an explicit mechanism for combining these factors, and cost prediction can support budget planning when appropriate records calculated directly from detections. Severity labels require a defensible annotation process. GIS variables require geographic data. Cost prediction requires historical repair costs. Keeping these dependencies explicit improves the scientific validity and reproducibility of the study.

11. Limitations

- Image-based measurements may be affected by camera angle, lighting, shadows, weather, occlusion, and image quality.
- Bounding-box area is only a two-dimensional proxy and does not directly measure pothole depth.
- Severity classification depends on the availability and consistency of ground-truth severity labels.
- GIS prioritization depends on the accuracy and availability of traffic, road, and point-of-interest information.
- Cost prediction cannot be validated without actual historical repair-cost records.

12. Future Scope

Future work can incorporate depth estimation using stereo cameras, LiDAR, depth sensors, or validated monocular-depth techniques. Real-time inference could be deployed on smartphones, vehicle-mounted systems, or edge devices. GIS analysis can be expanded with accident history, drainage conditions, rainfall, congestion, population density, and road-age information. Priority scores can be updated dynamically as new observations arrive. Historical municipal work orders can be incorporated to improve repair-cost estimation. A mobile application and maintenance dashboard can provide an operational interface for authorities.

13. Conclusion

This paper presents an integrated AI and GIS-based framework for pothole severity assessment, repair-priority ranking, and approximate cost prediction. The proposed pipeline combines YOLO-based visual detection, bounding-box feature extraction, Random Forest and XGBoost severity classification, GIS contextual analysis, MCDA-based prioritization, and regression-based cost estimation.

The current 1,740-record pothole dataset provides a foundation for the visual-feature and machine-learning stages. However, severity labels, GIS attributes, and actual repair-cost records must be incorporated before those components can be experimentally validated. The proposed framework therefore provides a practical and research-oriented foundation for intelligent road-maintenance decision support while clearly distinguishing the current evidence from future experimental outcomes.

14. References

[1] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, “You Only Look Once: Unified, Real-Time Object Detection,” Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), pp. 779–788, 2016.

[2] L. Breiman, “Random Forests,” Machine Learning, vol. 45, pp. 5–32, 2001.

[3] T. Chen and C. Guestrin, “XGBoost: A Scalable Tree Boosting System,” Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, pp. 785–794, 2016.

[4] D. Arya et al., “RDD2022: A multi-national image dataset for automatic Road Damage Detection,” arXiv:2209.08538, 2022.

[5] D. Arya et al., “RDD2020: An annotated image dataset for automatic