Industrial Engineering Journal
ISSN: 0970-2555
Volume : 53, Issue 8, August : 2024

Efficient Road Damage Monitoring: Advanced UAV and YOLO-Based Detection Techniques

[1]SUMIYA SULTANA, [2] KUDUMULA SHILPA REDDY
[1]Assistant Professor,[2]M.Tech student
[1,2]Department of CSE
[1,2] Sri Sai Institute of Technology and Science, Rayachoty, Annamaya (Dist) Andhra Pradesh, India.
shaiksumiyasultana@gmail.com, silpareddy.44@gmail.com

ABSTRACT

This paper introduces an innovative method for automatically detecting road damage using UAV images and deep learning. Road infrastructure maintenance is crucial for safe transportation, but manual data collection is time-consuming and hazardous. To address this, we propose a solution that leverages UAVs and AI to enhance detection accuracy and efficiency. Our approach utilizes YOLOv5 and YOLOv7, and the enhanced YOLOv8 algorithms for object detection in UAV images, trained and tested with datasets from China and Spain. Results show impressive performance with up to 85% accuracy. This study showcases the potential of UAVs and deep learning in automated road damage detection, opening avenues for future research.

Index Terms: Road Damage, YOLOV8

INTRODUCTION:

Effective road maintenance is crucial for a country's economic development. Regular assessments are essential to ensure road longevity and safety. Traditionally, this task has been performed manually by state or private agencies using sensor-equipped vehicles to detect road damage. However, this method is often time-consuming, costly, and hazardous for workers. To overcome these challenges, researchers and engineers are increasingly leveraging Unmanned Aerial Vehicles (UAVs) and Artificial Intelligence (AI) to automate road damage detection. UAVs, equipped with high-resolution cameras and sensors, offer a comprehensive view of road conditions from various angles and heights. They can swiftly cover large areas, minimizing the need for dangerous manual inspections. The integration of UAVs with deep learning techniques has shown promise in developing efficient and cost-effective solutions for road damage detection, gaining significant attention for urban inspections and various applications like monitoring swimming pools, rooftops, and vegetation.

LITERATURE SURVEY:

J. Redmon *et al*

We introduce YOLO9000, a real-time object detection system capable of recognizing over 9000 categories. YOLOv2, an improved model, achieves state-of-the-art results on PASCAL VOC and COCO tasks with multi-scale training. It balances speed and accuracy, reaching 76.8 mAP at 67 FPS. YOLO9000 jointly trains on detection and classification, performing well even on unlabelled classes, validating its versatility.

Y.-J. Cha et al

Various image processing techniques (IPTs) have been employed to detect civil infrastructure defects, but real-world conditions pose challenges. This article proposes a vision-based method using convolutional neural networks (CNNs) for detecting concrete cracks without calculating defect features. Trained on 40K images, the CNN achieved 98% accuracy. Tested on diverse conditions, it outperformed traditional methods like Canny and Sobel edge detection, proving effective for real-world applications.

PROBLEM STATEMENT:

The current manual collection of road damage data is labor-intensive, unsafe, and inefficient, hindering effective road infrastructure maintenance. This process is not only costly but also poses significant safety risks to inspectors, resulting in delayed detection of defects. These delays can lead to increased risks and financial losses, underscoring the need for a more efficient and accurate approach to road damage inspection and maintenance.

PROPOSED METHOD:

The proposed model employs Unmanned Aerial Vehicles (UAVs) with high-resolution cameras and deep learning algorithms YOLOv5 and YOLOv7 to automate road damage detection. This approach ensures efficient, accurate, and cost-effective road maintenance, addressing the limitations of traditional manual inspections.

ARCHITECTURE:

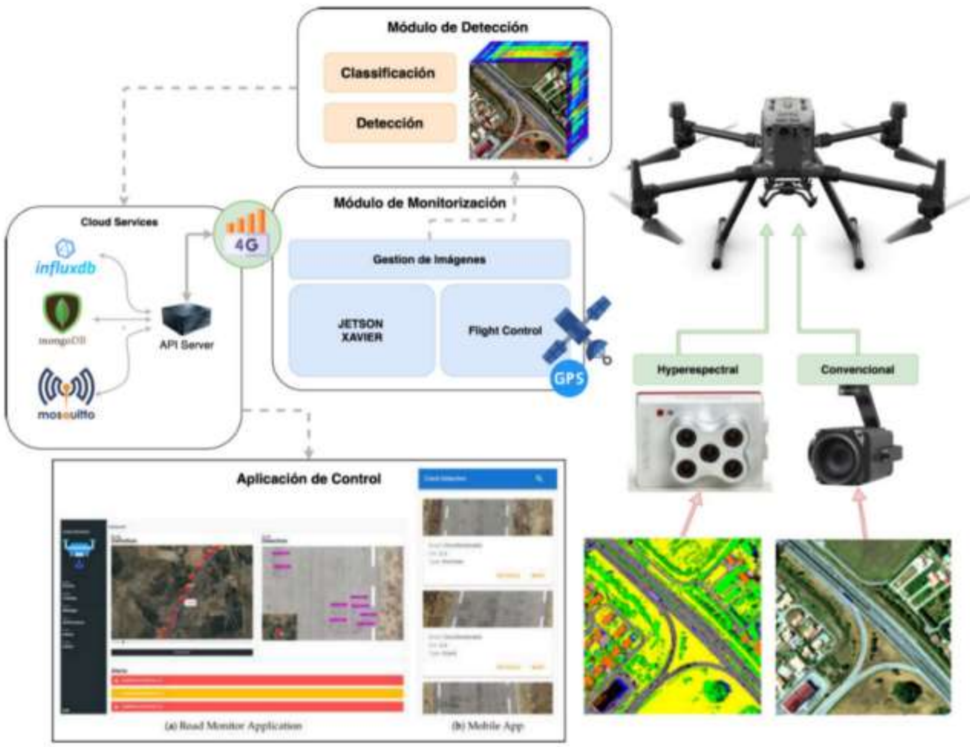

ROAD DAMAGE DATASET:

Different Labels images in dataset

METHODOLOGY:

Data Collection and Preprocessing

Industrial Engineering Journal
ISSN: 0970-2555
Volume : 53, Issue 8, August : 2024

Dataset

Our analysis leverages the RDD2022 China Drone dataset, an extensive collection of drone-captured images paired with detailed annotations that highlight areas of road damage. This dataset offers a rich source of visual data specifically designed to train and evaluate models for detecting and analyzing damaged road segments.

Data Preprocessing

To ensure consistency and maximize the effectiveness of our model training, several preprocessing steps were employed. First, all images were resized to a uniform resolution. This resizing is crucial for standardizing the input data, which helps in the efficient training of machine learning models. Next, the bounding boxes within these images, which indicate the regions of interest (i.e., damaged road areas), were normalized. Normalization of bounding boxes is a key step in simplifying the model's task of learning to predict these regions accurately. Additionally, the labels associated with these images were encoded, converting categorical data into a numerical format suitable for classification and bounding box regression tasks.

Model Selection and Development

Model Architecture

For the task of detecting damaged roads, we selected several variations of the YOLO (You Only Look Once) architecture, known for its efficiency in object detection tasks. Our choices included YOLOv5, YOLOv7, and a customized version of YOLOv8.

- **YOLOv5:** This version of YOLO is noted for its balance between speed and accuracy. It utilizes a custom architecture incorporating convolutional layers that extract features from the images, followed by dense layers that handle classification and bounding box regression. This setup allows YOLOv5 to detect and localize damaged road areas effectively.
- **YOLOv7:** Building on the advancements of YOLOv5, YOLOv7 introduces additional layers designed to enhance performance further. These modifications aim to improve the model’s ability to detect and classify road damage with higher precision and recall.

- **YOLOv8 Extension:** For our advanced approach, we employed an extended version of YOLOv8, which utilizes pre-trained models from Ultralytics. This extension includes further refinements and additional modifications to boost accuracy and model robustness.

Training and Evaluation

Data Splitting

To ensure that our models are rigorously evaluated, we divided the dataset into training and testing subsets. Specifically, 80% of the data was allocated for training, while the remaining 20% was reserved for testing. This split allows us to train the models on a large portion of the data while evaluating their performance on a separate, unseen portion.

Model Training

The training process for each model was tailored to its specific architecture:

- **YOLOv5 and YOLOv7:** Both of these models were trained using custom training loops. This approach involved selecting appropriate loss functions and optimizers to fine-tune the models effectively. For YOLOv5, this meant adapting the training to its convolutional and dense layers, while YOLOv7's additional layers required adjustments in the training process to leverage their enhanced capabilities.
- **YOLOv8 Extension:** This model was fine-tuned using pre-trained weights, which provided a strong starting point and reduced the training time required. Customized training parameters were applied to further refine the model's performance.

Model Evaluation

To assess the effectiveness of our models, several performance metrics were calculated:

- **Accuracy:** Measures the overall correctness of the model in detecting and classifying damaged road areas.
- **Precision:** Indicates the proportion of true positive detections among all positive predictions made by the model.
- **Recall:** Reflects the proportion of true positive detections out of all actual damaged areas present in the images.
- **F1 Score:** Provides a balanced measure of precision and recall, offering a single metric to evaluate model performance.

Industrial Engineering Journal
ISSN: 0970-2555
Volume : 53, Issue 8, August : 2024

Additionally, confusion matrices were generated to visualize the classification results. These matrices help in understanding how well the models distinguish between different classes of road damage.

Performance Comparison and Analysis

Performance Metrics

For each model, we computed accuracy, precision, recall, and F1 score to gauge their performance. These metrics provide a comprehensive view of how well each model performs in detecting and classifying road damage.

Comparative Analysis

In our comparative analysis, YOLOv7 and the YOLOv8 extension demonstrated superior performance relative to YOLOv5. Specifically, YOLOv7 and the YOLOv8 extension achieved higher precision, recall, and F1 scores, indicating better accuracy and robustness in detecting damaged road areas. The enhanced layers and modifications in these models contributed significantly to their improved performance.

In summary, while YOLOv5 provided a solid baseline, the advancements seen in YOLOv7 and the YOLOv8 extension highlight their superior capabilities in the domain of road damage detection. These findings underscore the importance of using advanced architectures and fine-tuning strategies to achieve optimal performance in real-world applications.

Industrial Engineering Journal
ISSN: 0970-2555
Volume : 53, Issue 8, August : 2024

RESULTS:

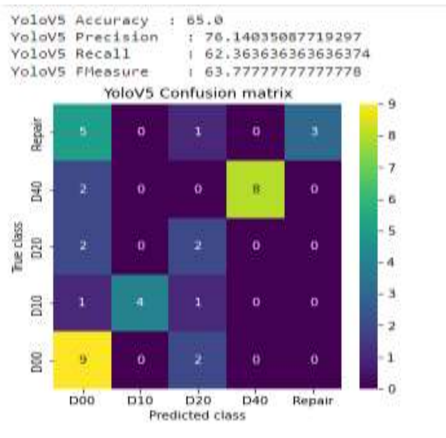

Training YoloV5 got 65% accuracy

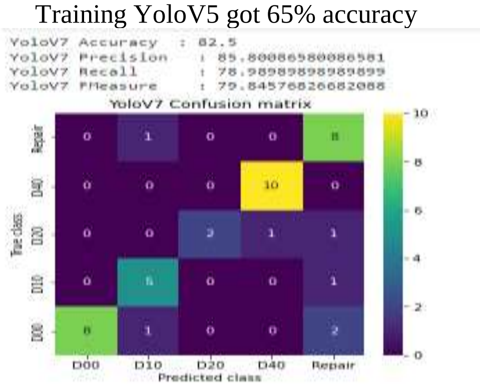

Training Yolov7 got 82% accuracy

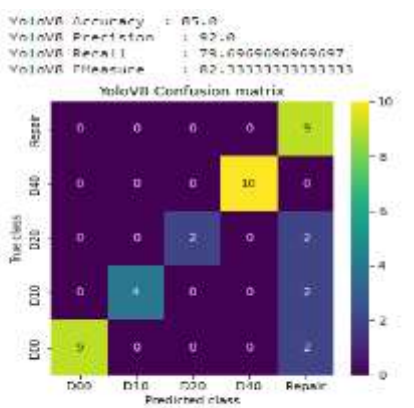

Training Yolov8 got 85% accuracy

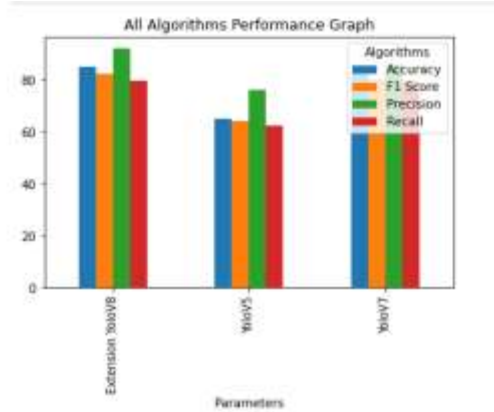

x-axis represents algorithm names and y-axis represents accuracy and other metrics in different color bars

Prediction:

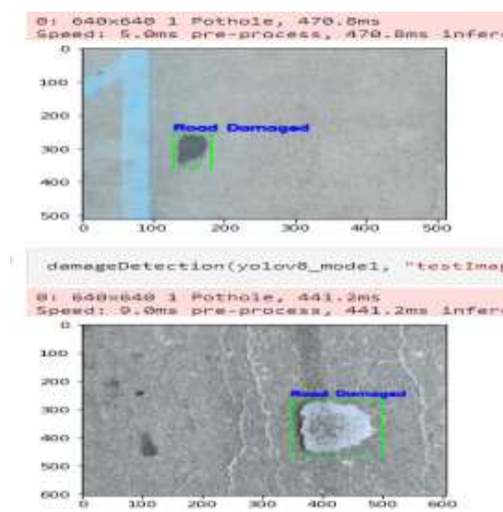

Damage road detected from test input images

Industrial Engineering Journal
ISSN: 0970-2555
Volume : 53, Issue 8, August : 2024

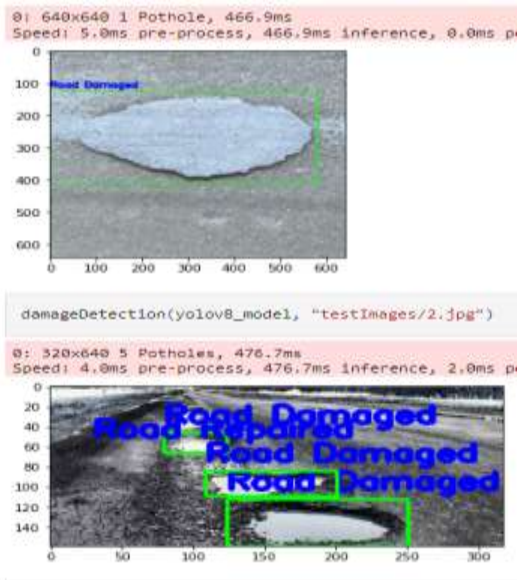

Damage road detected from test input images

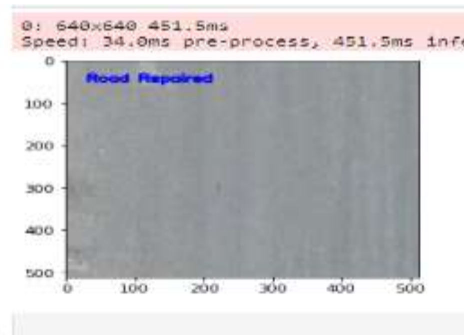

Road repaired is predicted without damage

CONCLUSION

This paper investigates the effectiveness of YOLOv5 and YOLOv7 algorithms in detecting road damage from UAV images. Utilizing the RDD2022 dataset, YOLOv7 emerged as the most accurate model. Despite limited training resources, both YOLOv5 and YOLOv7 performed well. The recent introduction of YOLOv8 by Ultralytics marks a notable improvement in object detection accuracy. This study underscores the value of advanced algorithms and datasets in improving object detection for essential applications such as infrastructure maintenance and monitoring.

REFERENCES:

[1] H. S. S. Blas, A. C. Balea, A. S. Mendes, L. A. Silva, and G. V. González, ‘‘A platform for swimming pool detection and legal verification using a multi-agent system and remote image sensing,’’ Int. J. Interact. Multimedia Artif. Intell., vol. 2023, pp. 1–13, Jan. 2023.

[2] V. J. Hodge, R. Hawkins, and R. Alexander, ‘‘Deep reinforcement learning for drone navigation using sensor data,’’ Neural Comput. Appl., vol. 33, no. 6, pp. 2015–2033, Jun. 2020, doi: 10.1007/s00521-020-05097-x.

[3] A. Safonova, Y. Hamad, A. Alekhina, and D. Kaplun, ‘‘Detection of Norway spruce trees (Piceaabies) infested by bark beetle in UAV images using YOLOs architectures,’’ IEEE Access, vol. 10, pp. 10384–10392, 2022.

[4] D. Gallacher, ‘‘Drones to manage the urban environment: Risks, rewards, alternatives,’’ J. Unmanned Vehicle Syst., vol. 4, no. 2, pp. 115–124, Jun. 2016.

[5] L. A. Silva, A. S. Mendes, H. S. S. Blas, L. C. Bastos, A. L. Gonçalves, and A. F. de Moraes, ‘‘Active actions in the extraction of urban objects for information quality and knowledge recommendation with machine learning,’’ Sensors, vol. 23, no. 1, p. 138, Dec. 2022, doi: 10.3390/ s23010138.

[6] L. Melendy, S. C. Hagen, F. B. Sullivan, T. R. H. Pearson, S. M. Walker, P. Ellis, A. K. Sambodo, O. Roswintiarti, M. A. Hanson, A. W. Klassen, M. W. Palace, B. H. Braswell, and G. M. Delgado, ‘‘Automated method for measuring the extent of selective logging damage with airborne LiDAR data,’’ ISPRS J. Photogramm. Remote Sens., vol. 139, pp. 228–240, May 2018, doi: 10.1016/j.isprsjprs.2018.02.022.

[7] L. A. Silva, H. S. S. Blas, D. P. García, A. S. Mendes, and G. V. González, ‘‘An architectural multi-agent system for a pavement monitoring system with pothole recognition in UAV images,’’ Sensors, vol. 20, no. 21, p. 6205, Oct. 2020, doi: 10.3390/s20216205.

[8] M. Guerrieri and G. Parla, ‘‘Flexible and stone pavements distress detection and measurement by deep learning and low-cost detection devices,’’ Eng. Failure Anal., vol. 141, Nov. 2022, Art. no. 106714, doi: 10.1016/j.engfailanal.2022.106714.

[9] D. Jeong, ‘‘Road damage detection using YOLO with smartphone images,’’ in Proc. IEEE Int. Conf. Big Data (Big Data), Dec. 2020, pp. 5559–5562, doi: 10.1109/BIGDATA50022.2020.9377847.

[10] M. Izadi, A. Mohammadzadeh, and A. Haghighattalab, ‘‘A new neuro-fuzzy approach for post-earthquake road damage assessment using GA and SVM classification from QuickBird satellite images,’’ J. Indian Soc. Remote Sens., vol. 45, no. 6, pp. 965–977, Mar. 2017.

[11] Y. Bhatia, R. Rai, V. Gupta, N. Aggarwal, and A. Akula, ‘‘Convolutional neural networks based potholes detection using thermal imaging,’’ J. King Saud Univ.-Comput. Inf. Sci., vol. 34, no. 3, pp. 578–588, Mar. 2022, doi: 10.1016/j.jksuci.2019.02.004.

[12] J. Guan, X. Yang, L. Ding, X. Cheng, V. C. Lee, and C. Jin, ‘‘Automated pixel-level pavement distress detection based on stereo vision and deep learning,’’ Automat. Constr., vol. 129, p. 103788, Sep. 2021, doi: 10.1016/j.autcon.2021.103788.

[13] D. Arya, H. Maeda, S. K. Ghosh, D. Toshniwal, and Y. Sekimoto, ‘‘RDD2022: A multi-national image dataset for automatic road damage detection,’’ 2022, arXiv:2209.08538.

[14] J. Redmon and A. Farhadi, ‘‘YOLO9000: Better, faster, stronger,’’ in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), Honolulu, HI, USA, 2017, pp. 6517–6525, doi: 10.1109/CVPR.2017.690.

[15] J. Redmon and A. Farhadi. YOLOv3: An Incremental Improvement. [Online]. Available: https://pjreddie.com/yolo/

[16] A. Bochkovskiy, C.-Y. Wang, and H.-Y. Mark Liao, ‘‘YOLOv4: Optimal speed and accuracy of object detection,’’ 2020, arXiv:2004.10934.

[17] G. Jocher, A. Chaurasia, A. Stoken, J. Borovec, Y. Kwon, K. Michael, J. Fang, C. Wong, D. Montes, Z. Wang, C. Fati, J. Nadar, V. Sonck, P. Skalski, A. Hogan, D. Nair, M. Strobel, and M. Jain, ‘‘Ultralytics/YOLOv5: V7.0—YOLOv5 SOTA realtime instance segmentation,’’ Zenodo, Tech. Rep., Nov. 2022. [Online]. Available: https://zenodo.org/record/7347926

[18] C.-Y. Wang, A. Bochkovskiy, and H.-Y. Mark Liao, ‘‘YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors,’’ 2022, arXiv:2207.02696.

[19] R. Ali, D. Kang, G. Suh, and Y.-J. Cha, ‘‘Real-time multiple damage mapping using autonomous UAV and deep faster region-based neural networks for GPS-denied structures,’’ Autom. Construct., vol. 130, Oct. 2021, Art. no. 103831. [Online]. Available: https://www.sciencedirect.com/ science/article/pii/S092658052100282X

[20] D. Kang and Y.-J. Cha, ‘‘Autonomous UAVs for structural health monitoring using deep learning and an ultrasonic beacon system with geo-tagging: Autonomous UAVs for SHM,’’ Comput.-Aided Civil Infrastruct. Eng., vol. 33, no. 10, pp. 885–902, Oct. 2018.