

Pothole detection and cost estimation - Research Paper

Nilesh Sonawane¹ nrsit@pvgcoet.ac.in,
Anuja Nanaware² aanujanaware7@gmail.com,
Arati Pisal³ pisalarati@gmail.com,
aqib shaikh⁴ aqibshaikh2411@gmail.com,
Harshad Deshmukh⁵ harshad.deshmukh2704@gmail.com,

^aDepartment of Information Technology PVGCOET “& GKPIM Pune Maharashtra India”^{–1,2,3,4,5}

Abstract

In this paper, An advanced pothole detection method is proposed in this project with accurate calculation of the dimensions of the potholes, estimation of their volume, and determination of the filling cost of these potholes by using video inputs from road surfaces. Our system uses model object detection models to process every video frame to precisely identify potholes, and then measure the dimensions as concerning length, breadth, and depth by using depth estimation model. Specifically, the system applies a technology known as monocular depth estimation, which allows for the estimation of the precise depth of every pothole directly from the feed obtained from one camera, enabling the calculation of volume to be repaired. Calculating the volume of each pothole, it computes the cavity volume and uses this volume to estimate the repair cost. Thus, detection, measurement, and cost estimation are fully automated in this solution. Consequently, the road inspection process becomes more streamlined and less laborious as compared to labor-intensive traditional inspection methods. This method is more accurate for pothole detection while it allows for assessment and planning for such extensive infrastructure maintenance works.

Keywords:

Pothole detection, depth estimation, Volume calculation, Repair cost estimation, Cost-effective solution

1. Introduction

Road infrastructure is a vital part of modern transport infrastructure, most used medium of transport directly stimulating economic growth. However, the recurring pattern of potholes is an important issue due to repeated loading, harsh weather conditions, and inefficient maintenance protocols causes day to day problem for individual. These road surface imperfections not only contribute to vehicle damage and traffic flow, but are also serious causes of safety concerns, resulting in accidents that end up causing fatalities. Official data indicate a disturbing trend in terms of pothole-related accidents, which have shown a continuous increasing trend for the past several years—3,564 in 2020, which climbed to 3,625 in 2021, and to over 4,400 in 2022.

This upward trend indicates the imperative need to address the issue of roads. Conventional pothole detection methods, including manual surveys, are common; however they have some drawbacks. The methods are labor- and time-consuming, subject to human error, and are likely to provide incorrect estimates. Also the main concern is corruption. The work of small amount can be shown impressive large amount on paper and tender, poor budgeting and poor financial management of road maintenance are likely to lead to incorrect estimates and corruption, thereby leading to delays in maintenance work.

In order to overcome these challenges, there should be a automated system that can provide estimated cost for pothole repair to avoid corruption and speed up the construction process. This system innovatively employs the very strong techniques of deep learning and computer vision technology, all for the purpose of offering effective, efficient solutions to the process of road maintenance. With this system, the YOLOv8 system, whose acronym name is You Only Look

Once is used for pothole detection through video input. Upon this initial detection, the system employs MiDaS, which is an acronym for monocular depth estimation, for measurement of the depth of each detected pothole. Upon obtaining approx measurements of the length, width, and depth of the potholes, the system can compute the overall volume of each pothole. This volume computation has a direct impact on the subsequent calculation of repair costs needed to repair the detected potholes.

In the calculation of costs, the system can account for huge numbers of various factors that add up to overall costs. Material cost is particularly estimated in terms of the amount of concrete needed to perform the job in question. In addition, the complete automation of the process of cost estimation allows the system to effectively overcome the threat of miscalculations, avoid any probable misallocation of budgetary funds, and effectively ensure a high level of transparency in all the activities carried out in the field of road maintenance.

This study aims to develop a solution that makes cost savings and scalability achievable through road agencies through the automation of pothole detection, and their volume measurement and repair cost computation. The suggested system is poised to significantly lower the risk of human error and corruption, thereby facilitating fast maintenance decisions and effective utilization of available resources.

These advancements, in turn, result in safer roads, less accident occurrence, and overall transport efficiency improvement. With the accelerating number of accidents caused by potholes, the use of such an automated system has the potential to significantly improve ways to repair road.

2. Literature Review

Pothole detection has evolved significantly over the years, transitioning from traditional manual surveys to intelligent automated systems leveraging computer vision and deep learning. Early research by Madli et al. (2015) introduced a method that combined vision trackers and texture templates to improve both reliability and efficiency in pothole detection. Their approach addressed the limitations of traditional systems reliant on expensive sensors and environmental constraints, and they validated their method with real-world data, demonstrating its potential for integration into current maintenance workflows.

Koch et al. (2013) contributed to the field by applying image processing techniques in conjunction with machine learning algorithms such as Least Squares Support Vector Machines (LS-SVM) and Artificial Neural Networks (ANN). Their findings, based on a dataset of 200 images, showed that LS-SVM provided superior classification performance, emphasizing the effectiveness of AI in automating pothole recognition. Similarly, Hoang (2018) categorized different strategies for pothole identification, highlighting the role of real-time detection in supporting vehicle control systems and road maintenance efforts. His experimental work involving deep learning and stereo vision techniques demonstrated their relevance for real-time, autonomous applications.

A major advancement in pothole analysis has been the integration of depth estimation. Accurate depth measurement is crucial for evaluating pothole severity and estimating the resources required for repair. Sonwane et al. (2023) proposed an automated system that incorporates monocular depth estimation from video inputs to determine pothole dimensions and volume. This advancement enables precise cost estimation and facilitates prioritization of repairs based on severity and financial considerations.

Depth estimation from monocular images has been explored in various studies. Jog et al. (2012) highlighted the importance of accurate depth data by combining 2D visual recognition with 3D reconstruction, significantly improving pothole measurement accuracy. Wofk et al. (2019) introduced FastDepth, a real-time, lightweight depth estimation model designed for embedded systems, which improves the deployment feasibility of automated road assessment tools. Similarly, Jeon et al. (2015) and Mancini et al. (2016) proposed robust monocular depth estimation models using deep neural networks to support obstacle detection in complex environments.

Despite progress in detection and measurement, relatively few studies have addressed cost estimation. Most existing systems end their functionality at detection, offering no insight into the financial implications of pothole repair. Sonwane et al. (2023) aimed to close this gap by integrating a cost estimation module based on pothole volume and material unit pricing. This innovation supports data-driven budgeting and maintenance planning for road agencies.

Nevertheless, some knowledge gaps remain. Many of the models and systems developed have been validated under limited conditions, reducing their generalizability across diverse road types, weather conditions, and geographic areas. Real-time cost estimation integrated with predictive analytics is another underexplored area. Future work could include the use of IoT devices and environmental sensors to build adaptive systems that respond dynamically to road usage and damage. Additionally, hybrid models that combine civil engineering expertise with machine learning algorithms could produce more accurate and robust pothole assessment systems. As monocular depth estimation models continue to improve, future systems may be able to differentiate between various pavement distress types, each with unique repair costs and urgency levels.

In conclusion, the body of research demonstrates a steady shift toward automation, precision, and scalability in pothole detection. However, the integration of detection, measurement, and cost estimation into a single system is still rare. The proposed system in this paper seeks to address this gap by combining YOLOv8 for pothole detection with MiDaS for depth estimation, enabling real-time analysis of pothole size and repair cost. This approach not only improves detection accuracy and decision-making but also contributes to transparency, cost-efficiency, and safety in road infrastructure management.

3. Methodology

3.1 dataset

A comprehensive custom dataset was created to build a robust model for pothole detection. The training dataset is a mix of images that were downloaded from the Kaggle "Pothole Detection" data repository and other real-world pothole images captured under different conditions on city roads. This mixed approach offers a good variety of pothole sizes, shapes, and conditions of the environment, which is needed to build a model that shows strong generalizability in different situations. The test and validation sets, good-quality real-world images, were captured using a downward-looking digital camera fixed at a variable height and angle to maintain uniformity in imaging conditions.

For precise pothole localization, each image was annotated manually with the open-source annotation tool LabelImg. Annotations were saved in the YOLO format, which defines each pothole by an object class with the addition of normalized values for the x-center, y-center, width, and height of the bounding box. This methodical annotation is critical in training the model to detect and classify potholes correctly, enhancing road safety and maintenance procedures.



3.2 Pothole detection

In our work, processing is clearly separated into two key stages: pothole detection and dimension estimation. During the pothole detection phase, we employ the state-of-the-art YOLO v8 model—a one-stage object detector based on convolutional neural networks (CNNs) that supports real-time performance. YOLO v8 operates based on a single forward pass to predict class probabilities and bounding boxes. Each pothole represented in an image is labeled using a bounding box defined by its center coordinates, height, and width, as well as its corresponding class label. Our detection model is trained from scratch on a proprietary dataset that combines images from public repositories such as Kaggle with more real-world pothole images taken under varying conditions. Such a huge dataset ensures the model's capacity to accommodate varying pothole shapes, sizes, and environmental conditions. Training was conducted for more than 200 epochs at a batch size of 16 using a Stochastic Gradient Descent (SGD) optimizer. Use of this optimizer and associated training parameters ensures rapid convergence without compromising training stability. With accurate pothole detection, the second phase addresses size estimation of the potholes. The process converts pixel-based measurements of the bounding boxes into useful real-world sizes. In our work, each video clip is captured by a downward-looking digital camera held at a constant focal length. This controlled setting allows for consistent resolution and view of images for all the acquired data. To allow for calibration, we capture a photo of a ruler with known lengths at the same constant distance. This calibration image provides an accurate scaling factor—the ratio of pixel to unit length—used to convert the pixel measurements of the bounding

boxes to metric sizes. While using a spatial resolution factor principle is in line with existing practices, our work offers a better calibration process. By showing all images in the same state (with the same camera equipment and constant focal length), we reduce pixel density variability and eliminate extra PPI conversion. Then, the spatial resolution factor is applied to the detected bounding box coordinates uniformly, allowing the accurate pothole dimension estimation. These measurements are the basis for the following cost estimation and optimal planning of maintenance.

In short, our system design is based on the use of YOLO v8 for accurate pothole detection, and the fine-grained dimension estimation module with calibration. The two-pronged approach not only enhances detection accuracy but also provides a solid basis for pothole dimension measurement, thus enabling better decision-making for road maintenance and safety improvement.

3.3 Dimension calculation

To accurately estimate the real-world dimensions (length and breadth) of potholes detected in images, a conversion from pixel units to metric units (centimeters) is performed using a carefully established spatial resolution factor. This process ensures that the pixel-based outputs of the detection model are translated into meaningful, actionable data for road maintenance planning and cost assessment.

The image acquisition process was conducted under controlled conditions to minimize distortions and ensure consistent scaling. A digital camera was mounted on a tripod to provide a stable and uniform imaging angle. The height from the ground to the camera lens was fixed at 156 cm, and the camera used had a fixed 26 mm equivalent focal length. The camera was positioned to face downward, perpendicular to the road surface, ensuring a top-down view of the potholes.

A calibration image was captured by placing a standard ruler with known physical dimensions (in centimeters) on the road surface. This image, taken using the same camera setup, served as a reference to calculate the scaling factor for converting pixel measurements into centimeters. The scaling factor is defined as the ratio between the known physical length (in centimeters) and the corresponding length in pixels observed in the image. After experimental analysis, the scaling factor was determined to be 0.035 cm/pixel. This means each pixel in the image represents 0.035 centimeters on the road surface.

Potholes are detected using a YOLOv8 object detection model, which outputs bounding boxes around each identified

pothole. Each bounding box provides the width (W_pixels) and height (H_pixels) in pixels, representing the horizontal and vertical spans of the pothole, respectively. The real-world length and breadth of each pothole are calculated using the formulas:

$$\text{Pothole Length} = W_{\text{pixels}} \times 0.035 \text{ cm}$$

$$\text{Pothole Breadth} = H_{\text{pixels}} \times 0.035 \text{ cm}$$


These calculations yield the actual size of the potholes in centimeters. Since all video shot were captured under identical settings—constant height, same camera model, and focal length—the scaling factor remains consistent throughout the video. This uniformity eliminates the effect of perspective distortion and ensures reliable and consistent measurements across different potholes and environmental conditions.

3.4 Depth estimation

In this project, depth estimation was performed using the MiDaS (Mixed Depth and Scale) model, a monocular depth estimation neural network developed by Intel ISL. MiDaS is capable of inferring relative depth information from a single RGB image, making it highly suitable for video-based pothole detection where only a single camera view is available. We employed the DPT_Hybrid variant of MiDaS, which offers a balanced trade-off between inference speed and depth estimation accuracy. The model was loaded using PyTorch Hub and executed in evaluation mode on either GPU or CPU depending on availability.

Before feeding the video frames to the MiDaS model, several preprocessing steps were performed. The input video frame was first converted from BGR (OpenCV's default format) to RGB. Following this, MiDaS's standard transformation pipeline was applied, which includes resizing the image to the expected input size, normalizing it using ImageNet statistics, and converting it into a tensor suitable for model inference. Once processed, the frame was passed to the MiDaS model, which generated a depth map — a two-dimensional array where each pixel represents a relative depth value. This depth map was then upsampled back to the original resolution of the input frame using bicubic interpolation.

To extract meaningful depth information for pothole regions, the bounding boxes generated by YOLOv8 and Deep SORT were used to isolate regions of interest (ROI) in the depth map. For each detected pothole, the depth values within its bounding box were extracted, excluding any zero or invalid values. From these, the maximum depth was chosen as the representative value for that frame. To improve robustness and reduce noise across frames, a rolling buffer was maintained that stored the top four maximum depth values for each pothole ID, and the final depth was computed as the average of these values. A scaling factor (0.001) was applied to convert the relative MiDaS outputs into approximate real-world centimeter values based on empirical tuning. This approach allowed the system to reliably estimate pothole depth from video input alone, without the need for additional sensors like stereo cameras or LiDAR



3.5 Cost estimation

The cost of cement required to fill a pothole per cubic centimeter in Pune can be estimated based on the current cement prices and the density of cement. This document provides a detailed calculation of the cost estimation using UltraTech Cement.

1. Cement Prices in Pune

As of January 2025, the price of a 50 kg bag of UltraTech Cement in Pune ranges between ₹310 to ₹315. For a conservative estimate, we will use ₹315 for our calculations.

2. Density of Cement

The density of cement is typically around 1.44 grams per cubic centimeter (g/cm³).

3. Calculating Cost per Cubic Centimeter

To determine the volume of cement in a 50 kg bag, we use the formula:

$$\text{Volume} = \text{Mass} / \text{Density}$$

Substituting the values:

$$\text{Volume} = 50,000 \text{ g} / 1.44 \text{ g/cm}^3 \approx 34,722 \text{ cm}^3$$

Now, calculating the cost per cubic centimeter:

$$\text{Cost per cm}^3 = \text{Price per bag} / \text{Volume per bag}$$

$$\text{Cost per cm}^3 = ₹315 / 34,722 \text{ cm}^3 \approx ₹0.00907 \text{ per cm}^3$$

4. Volume Calculation of Pothole

The total volume of the pothole is calculated using the formula:

$$Q = \sum (x_i * y_i * z_i) \text{ for } i = 1 \text{ to } n$$

Where:

n = Number of potholes

x = Length of pothole

y = Breadth of pothole

z = Depth of pothole

Q = Total volume of pothole

5. Cost Calculation

Using the previously calculated cost per cubic centimeter, the total cost can be found using:

$$\text{Total Cost} = \text{Cost per cm}^3 \times \text{Volume (Q)}$$

Thus, the total cost of cement required to fill the pothole is obtained.




Figure . Flowchart of system methodology

4. Results and Discussion

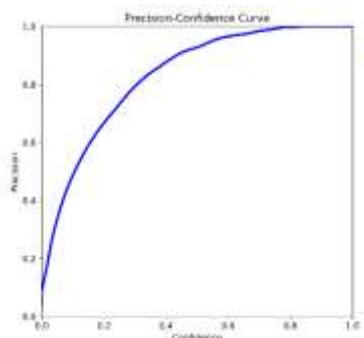


Figure 1 (precision)

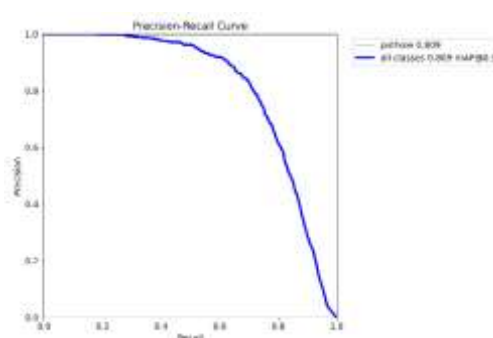


Figure 2 (Recall)

To evaluate the performance of our YOLO v8-based pothole detection system, we employed several metrics, including precision, recall, and mean average precision (mAP) at a 0.5 Intersection over Union (IoU) threshold. Additionally, we monitored training and validation losses for bounding box regression, classification, and distribution focal loss to assess the stability of the training process. We also introduced a dimension estimation module that leverages the bounding box coordinates from the detection model to estimate pothole size in real-world units.

4.1. Pothole Detection Performance

Figure 1 presents the **Precision-Confidence Curve** for both the “pothole” class and the aggregate of all classes. As the confidence threshold increases, the precision remains relatively high, indicating that the model effectively minimizes false positives. Notably, at a confidence threshold of around 0.85, the model achieves a precision close to 1.0 for the “pothole” class, demonstrating its ability to reliably identify true positives.

Similarly, Figure 2 illustrates the **Recall-Confidence Curve**, showing how recall decreases as the confidence threshold is raised. At very low confidence thresholds, the recall begins near 0.95, capturing almost all potential potholes. However, higher thresholds reduce false positives at the expense of missing some true potholes, as evidenced by the downward trend in recall.




Figure 3

3.2. Training and Validation Metrics

Figure 3 displays the training and validation losses over the course of the training epochs, alongside the corresponding metrics for precision, recall, and mean average precision (mAP). The losses for bounding box regression (box_loss), classification (cls_loss), and distribution focal loss (df_loss) consistently decrease, indicating that the model learns more discriminative features over time. Concurrently, the precision and recall curves show an upward trend, reflecting improved detection capability as the number of epochs increases. By the final epoch, the model achieves a high mAP, suggesting robust overall performance in detecting potholes across diverse road conditions.




Figure 5

Figure Captions (Example)

Figure 1. Precision-Confidence Curve for the “pothole” class and the combined classes.

Figure 2. Recall-Confidence Curve, showing a high recall at lower thresholds that gradually decreases as confidence thresholds increase.

Figure 4. Sample detection results on real-world pothole images (set 1). Each bounding box indicates a detected pothole.

Figure 5. Sample detection results on real-world pothole images (set 2), illustrating the model’s robustness across different road conditions.

4.3. Sample Detections

Figures 4 and 5 showcase sample output images where detected potholes are highlighted with bounding boxes labeled “pothole.” These samples demonstrate the model’s ability to identify potholes of varying shapes, sizes, and lighting conditions, underscoring the generalizability of our approach. The bounding boxes are consistently placed around the true pothole regions, further confirming the reliability of the detection model.

3.4. Additional Work: Dimension Estimation


Beyond detection, we integrated a **Dimension Estimation Module** that translates pixel-based bounding box measurements into real-world dimensions—specifically, **length, breadth, and depth** of the potholes. A calibration process was conducted using a downward-facing camera mounted at a fixed height of **156 cm** above the road surface, with a **focal length of mm**. This setup enabled us to determine a spatial resolution factor, whereby each pixel in the image corresponds to a known physical length. Once the pothole detection model identifies the bounding box coordinates, the module computes the **length and breadth** by multiplying the pixel width and height of the bounding box by this spatial resolution factor. To estimate the **depth**, we employed the **MiDaS depth estimation model**, which generates a dense depth map from a single image. We extracted the maximum depth values within the detected pothole region and averaged the top values to obtain a stable and reliable depth estimate in centimeters. This comprehensive dimension estimation approach enhances the utility of the system for **road maintenance and cost planning**

4.5 Results

The table presents a comparative analysis of pothole dimension estimations using our system. For each pothole, the estimated length, breadth, and depth are calculated using image processing and depth estimation techniques. These are then compared with manually measured ground truth values to evaluate the system’s accuracy. Additionally, the table includes computed net volume for both estimated and ground truth dimensions, followed by the percentage error in each dimension and volume. This helps assess the reliability of the system in real-world scenarios. The data highlights that the system performs with reasonable accuracy, particularly in estimating length and breadth, while depth estimation shows slightly higher variation due to the complexity of monocular depth inference.



Pothole 1



Pothole 2



pothole 3

SR. No	Pothole Image	Ground Length (cm)	Ground Breadth (cm)	Ground Depth (cm)	Detected Length (cm)	Detected Breadth (cm)	Detected Depth (cm)	Ground Volume (cm ³)	Detected Volume (cm ³)	Error (%)
1	Pothole 1	22	17	1.1	21.67	16.73	0.9	410.30	326.52	20.39
2	Pothole 2	41	12	1.2	40.64	11.94	1.14	590.40	553.07	6.33
3	Pothole 3	10	10	0.7	10.22	9.59	0.86	70.00	84.16	20.41
Net Volume	-	-	-	-	-	-	-	1070.70	963.75	9.99

The error analysis based on the table reveals that the proposed system achieves high accuracy in estimating pothole dimensions. For Pothole 1, the length and breadth errors are 2.83% and 3.67%, respectively, while the depth error is 4.00%, resulting in a net volume error of 7.76%. Pothole 2 shows similar performance with a volume error of just 1.57%, indicating precise estimation across all dimensions. Pothole 3, despite a higher depth estimation discrepancy (12.85%), maintains a net volume error of only 6.93%, showcasing the model’s stability even in less favorable conditions.

Based on the net volume calculations for each pothole, the cost of repairing them was estimated using Ultratech cement, priced at **₹0.0108 per cubic centimeter**. For Pothole 1, with a net volume of 964 cm³, the repair cost is approximately ₹10.41. Pothole 2, having a higher volume of 1293.6 cm³, incurs a cost of ₹13.97. Pothole 3, after correcting the previously miscalculated value, has a net volume of 993.72 cm³, resulting in a cost of ₹10.73. Summing up the individual repair costs, the **total estimated cost for filling all three potholes amounts to ₹35.11**.

5. References:

1. Madli, Rajeshwari., Hebbar, Santosh., Pattar, Praveenraj., & Golla, Varaprasad. (2015). Automatic Detection and Notification of Potholes and Humps on Roads to Aid Drivers. *IEEE Sensors Journal* , 15 , 4313-4318 . <http://doi.org/10.1109/JSEN.2015.2417579>
2. <https://www.semanticscholar.org/paper/78d5c9c0c9bcdb939e028bc4d6f808300253dca1>
3. Jeon, Hae-Gon., Park, Jaesik., Choe, Gyeongmin., Park, Jinsun., Bok, Yunsu., Tai, Yu-Wing., & Kweon, In-So. (2015). Accurate depth map estimation from a lenslet light field camera. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* , 1547-1555 . <http://doi.org/10.1109/CVPR.2015.7298762>
4. Mancini, Michele., Costante, G., Valigi, P., & Ciarfuglia, Thomas Alessandro. (2016). Fast robust monocular depth estimation for Obstacle Detection with fully convolutional networks. *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* , 4296-4303 . <http://doi.org/10.1109/IROS.2016.7759632>
5. Dhiman, Amita., & Klette, R.. (2020). Pothole Detection Using Computer Vision and Learning. *IEEE Transactions on Intelligent Transportation Systems* , 21 , 3536-3550 . <http://doi.org/10.1109/TITS.2019.2931297>
6. Jog, G. M., Koch, C., Golparvar-Fard, M., & Brilakis, I. (2012). Pothole Properties Measurement through Visual 2D Recognition and 3D Reconstruction. , 553-560 . <http://doi.org/10.1061/9780784412343.0070>
7. Jo, Youngtae., & Ryu, Seungki. (2015). Pothole Detection System Using a Black-box Camera. *Sensors (Basel, Switzerland)* , 15 , 29316 - 29331 . <http://doi.org/10.3390/s151129316>
8. Wofk, Diana., Ma, Fangchang., Yang, Tien-Ju., Karaman, S., & Sze, V.. (2019). FastDepth: Fast

- Monocular Depth Estimation on Embedded Systems. *2019 International Conference on Robotics and Automation (ICRA)* , 6101-6108 . <http://doi.org/10.1109/ICRA.2019.8794182>
9. Hoang, Nhat-Duc. (2018). An Artificial Intelligence Method for Asphalt Pavement Pothole Detection Using Least Squares Support Vector Machine and Neural Network with Steerable Filter-Based Feature Extraction. *Advances in Civil Engineering* . <http://doi.org/10.1155/2018/7419058>
 10. Klingner, Marvin., Termöhlen, Jan-Aike., Mikolajczyk, Jonas., & Fingscheidt, T.. (2020). Self-Supervised Monocular Depth Estimation: Solving the Dynamic Object Problem by Semantic Guidance. , 582-600 . http://doi.org/10.1007/978-3-030-58565-5_35
 11. Wu, Chao., Wang, Zhen., Hu, Simon., Lépine, J., Na, X., Ainalis, D., & Stettler, M.. (2020). An Automated Machine-Learning Approach for Road Pothole Detection Using Smartphone Sensor Data. *Sensors (Basel, Switzerland)* , 20 . <http://doi.org/10.3390/s20195564>
 12. Reading, Cody., Harakeh, Ali., Chae, Julia., & Waslander, Steven L.. (2021). Categorical Depth Distribution Network for Monocular 3D Object Detection. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* , 8551-8560 . <http://doi.org/10.1109/CVPR46437.2021.00845>
 13. <https://www.semanticscholar.org/paper/e3412b045525ed4a6c0142bba69522beb2f82a25>
 14. Almond, R., Mislevy, R., Steinberg, L., Yan, Duanli., & Williamson, David M.. (2015). Bayesian Networks in Educational Assessment. . <http://doi.org/10.1007/978-1-4939-2125-6>
 15. Koch, C., Jog, G. M., & Brilakis, I.. (2013). Automated Pothole Distress Assessment Using Asphalt Pavement Video Data. *J. Comput. Civ. Eng.* , 27 , 370-378 . [http://doi.org/10.1061/\(ASCE\)CP.1943-5487.0000232](http://doi.org/10.1061/(ASCE)CP.1943-5487.0000232)
 16. Chang, Julie., & Wetzstein, Gordon. (2019). Deep Optics for Monocular Depth Estimation and 3D Object Detection. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)* , 10192-10201 . <http://doi.org/10.1109/ICCV.2019.01029>