

RoadCrack Pro: Intelligent Road Damage Detection and Municipal Monitoring

Karan G K¹

II Year MCA

Faculty of Computing & IT

GM University

Davangere, Karnataka-577006

karangk2003@gmail.com

Prasad Bedage²

II Year MCA

Faculty of Computing & IT

GM University

Davangere, Karnataka-577006

prasadbedge7@gmail.com

1. Introduction

1.1 Problem

Pavement defect detection is a critical component of modern-day road infrastructure and acts as a support system for the management of both human and freight transportation. The impact of daily vehicle traffic loads, variations in temperatures, rain, water intrusion and aging technology leads to progressive weakening of the pavement structure. Among many types of damage, cracks and potholes are most frequently encountered and can negatively influence vehicle safety and comfort of traveling. Leaving pavement failure unattended can lead to an 'irreversible' situation when damage on the road is significant.

Typical road assessment is done through visual observation by engineers and municipal workers. Although visual observation is still implemented widely, it is costly, time-consuming and difficult to accomplish continuously across significant areas. Increasing urban road infrastructure triggers the need for automated methods of pavement assessment.

Digital Image Processing (DIP) offers a useful means of automating the inspection of pavements. It is possible to process road pictures and enhance the crack structure, mask unwanted texture, remove damaged areas from the surrounding surface, and diagnose defects. However, road pictures may contain uneven light, shadowing, lane signs, background items, and rough texture, making it difficult to achieve accurate results.

RoadCrack Pro helps solve this problem using a series of image processing techniques. The proposed scheme uses gray-scale and HSV conversion, COLOR and ROI masking, median and Gaussian filtering, bottom-hat morphological processing, adaptive thresholding, Sobel verification, morphological processing, and CC processing. The use of the system improves the detection process so that it includes pothole identification, severity detection, and municipality monitoring.

1.2 Research Gap

Recent advancements in pavement defect detection have involved the use of artificial intelligence technologies such as computer vision, machine learning, and deep learning. However, existing automated methods tend to focus solely on crack detection or cracking segmentation without offering a full flow of automated processing of road images from crack detection to pothole detection, true ground validation, damage severity assessment, and support for decision-making at local community levels.

Using deep learning, crack detection processes can achieve high efficiency even under complicated weather conditions; however, the deep learning approach relies on large labeled datasets and heavy processing capacity, which makes implementation and deployment impossible in low-cost municipal settings.

Traditional digital image processing techniques, unlike deep-learning methods, are lightweight and well-understood, while relying on certain conditions such as lighting, roughness of the surface being detected, shadows, and low-contrast image areas.

Because of limited testing of RoadCrack Pro on only two images of the road, proper statistical aggregation and comparison with other leading segmentation techniques have not been conducted, nor controlled ablation experiments aimed at determining the effect of each image-processing stage have been carried out.

To mitigate these issues, the study adopts a new arrangement for RoadCrack Pro that uses an extended evaluation protocol employing confusion-matrix metrics, comparison with traditional segmentation methodologies, stage-by-stage ablation studies, and evaluation of computational efficiency. Evaluation focuses on the concrete type of damage, such as crack, pot hole, combined damage, and no damage under conditions of different lighting, type of pavement, and damage severity.

1.3 Contributions

The contributions made by this study include:

1. Digital Image Processing Pipeline:

The state of the art road damage detection pipeline which is designed using bottom-hat processing, adaptive screening and material processing, Sobel validation is employed for detecting flaws in the road pavement.

2. Ground Truth Evaluation:

The proposed framework allows for pixel-level evaluation using most important evaluation metrics like Precision, Recall, F1 score, IoU and Accuracy.

3. Damage Classification:

The detected damage is quantified and classified into five levels based on the size of damage area- Excellent, Good, Fair, Poor, Critical, whereby recommendations for maintenance are provided.

4. Monitoring platform:

In order to study inspection data and to collect information for inspection, monitoring dashboard supported by sqlite is created.

5. Comparative Study:

The new research methodology used compares RoadCrack pro with traditional segmentation techniques and offers evaluation of the contribution of each processing step in the road damage detection process including ROI masking, bottom-hat processing, Sobel validation etc.

6. Lightweight Approach:

It can be considered that this system is simple in use and does not use complicated model for its work.

2. Related Work

2.1 Recent Road-Damage Research

Automated pavement-damage identification has been researched through traditional imaging methods, machine learning and deep learning. Initial research efforts concentrated on crack pattern extraction via technologically-enhanced image processing techniques like filtering, edge engineering, thresholding and morphing. Although these methods give quite simple and understandable solutions, they can be distorted by illumination differences, pavement textures and image noise. Zou et al. performed crack-structure extraction using structured crack modeling for improving continuity of cracks. However, this algorithm is sound sensitive. Shi et al. used the approach based on machine vision for crack path analysis. Zhang et al. performed research on CNN-oriented crack detection and proved that deep learning can be used for reliable crack identification. Maeda et al. have also expanded damage detecting on roads by taking into account different kinds of defects and able to detect cracks and road pits using smartphone images. Although these solutions provide a safer learning detection method, they are reliant on training datasets with marked data.

Ranjbar et al. analyzed different transfer-learning strategies in the process of evaluating pavement-crack problems, whereas Nguyen et al. presented studies related to the application of deep-learning techniques for crack recognition. More recent works are taking into account the location information and pavement management, while the review of vision-based studies of pothole detection is focused on the application of computer vision and machine learning in pothole identification.

Author / Year	Method / Focus	Main Contribution	Limitation	Relevance to Proposed Work
Zou et al., 2015	Crack structure extraction	Improved crack continuity using structured crack modeling.	Sensitive to noise and illumination changes.	Supports crack continuity analysis in the proposed work.
Shi et al., 2016	Machine vision crack detection	Used crack-path analysis for better linear crack extraction.	Focused only on crack localization.	Relevant for crack segmentation and path preservation.
Zhang et al., 2016	CNN-based crack detection	Applied deep learning for robust crack classification.	Requires large labeled datasets and more computation.	Provides comparison with lightweight morphology-based detection.
Maeda et al., 2018	Deep learning for road damage	Detected multiple road damages including cracks and potholes from smartphone images.	High dependency on training data and hardware resources.	Relevant because the proposed system also handles cracks and potholes.
Ranjbar et al., 2021	Transfer learning for	Improved crack evaluation accuracy using learning-based methods.	Less interpretable and more resource intensive.	Highlights the need for automated pavement evaluation.

Nguyen et al., 2022	Survey on deep crack detection	Reviewed crack detection methods using ML and DL techniques.	Survey only; no complete monitoring system.	Helps position the proposed work among
Nasrallah et al., 2024	Crack detection with location data	Combined crack detection with geolocation for pavement management.	Focused mainly on positioning and crack localization.	Relevant for municipal monitoring and future GPS/GIS integration.
Khan et al., 2024	Vision-based pothole detection review	Summarized pothole detection using computer vision and ML.	Review-based and not an integrated system.	Relevant because the proposed work includes pothole detection.

Table 1: Literature Survey — already present in your paper.

2.2 Traditional Digital Image Processing Approaches

Classic methods for the detection of surface defects in roads are based on using image processing techniques like :

- Thresholding
- Edge detection
- Filtering
- Morphological operations
- Connected-component analysis
- Shape analysis

These approaches usually need very little computing power and are easy to understand, unlike newer ones that use complex processes of image processing.

However, the efficiency of these techniques depends on the following:

- Inadequate illumination
- Shadows
- Road surface texture
- Noise
- Low-contrast cracks
- Wet pavement
- Background with contours of different colors

In this sense, the technology in RoadCrack Pro is largely similar to the older techniques but it relies on the combination of several techniques instead of only using one technique for the comprehension of a certain process.

2.3 Machine Learning Approaches

The methods of machine learning mainly rely on features that have to be manually designed from the paved road images. Such features can involve:

- Textures
- Intensities
- Shapes
- Edge qualities
- Region characteristics

The features that have been extracted would be supplied to a classifier to be able to tell the difference between good and damaged pavement.

This method combines the advantages of classical methods of operating with images and deep learning. Hence, it could help improve the efficiency of pavement classification without requiring deep neural networks.

2.4 Deep Learning Approaches

Deep learning methods including CNNs extract features from labeled roadway photographs. They have proven to be successful in detecting and classifying crack and pothole.

The key benefit of deep learning is that it is capable of learning complex structures that would be difficult to create with rules of image processing. This is especially helpful in case the roadway scene has complex pavement texture, shadow, lighting conditions, and cracks.

Nevertheless, deep learning approaches require:

- Big labeled datasets
- Time for training
- System resources
- Implementation that is more complex
- Hardware that is effective for training and inferring

Thus, despite its efficiency in detecting objects, deep learning is not as good in comparison to traditional DIP techniques.

2.5 Traditional DIP vs. Deep Learning

The following table can be included in your paper.

Parameter	Traditional DIP	Deep Learning
Training dataset	Generally not required	Large labelled dataset required
Training process	Not required	Required
Computational requirement	Low	Medium to high
Interpretability	High	Relatively low

Processing pipeline	Rule/operation based	Learned feature based
Sensitivity to lighting	Generally higher	Usually more robust
Complex crack patterns	Limited	Better suited
Implementation cost	Low	Higher
Suitable for low-resource deployment	Yes	May require more resources
Main limitation	Parameter/threshold sensitivity	Data and computational requirements

Table 2. Comparison of Traditional Digital Image Processing and Deep Learning

2.6 Research Positioning of RoadCrack Pro

With reference to the previously analyzed techniques, standard DIP is characterized by low cost, simplicity, and understanding of algorithm operation, and deep learning offers high learning abilities with regard to complicated visual scenes.

RoadCrack Pro functions as a simple and clear solution that includes a number of image processing methods that make it possible to identify road defects. Besides being simply a crack segmentation tool, the proposed framework also provides pothole detection, quality verification of ground truth, damage assessment, maintenance suggestions, and availability of dashboard for authorities.

As for the upgraded experimental approach, it further strengthens the research by adding general evaluation, comparison with standard segmentation techniques, stages-wise ablation experiments, and efficiency measuring. Results of quantitative comparison will be given in Section 6, as actual results are expected to be received during the experiments rather than theoretically calculated.

3. Proposed RoadCrack Pro

RoadCrack Pro is a smart monitoring system of road-asset, which can automatically detect cracks and potholes on a road surface using the images of the road. It uses many operations of Digital Image Processing and expands the detection process to damages' quantification, severity classification, validation of ground-truth, and municipality monitoring.

The system comprises of the following basic functional units:

1. Detection of road cracks
2. Detection of potholes
3. Validation of ground-truth

4. Severity assessment

5. Municipal Monitoring dashboard

The system is built as a low-weight and interpretable alternative to many deep-learning-based solutions.

3.1 System Architecture

You already have an architecture diagram presented in your new paper. Maintain this diagram, but ensure that it is comprehensible and clear.

The architecture should visually follow:




Figure 1. RoadCrack Pro System Architecture

3.2 Dataset Input

The fact that the system takes as input pictures of road surfaces. The original paper used images of roads taken by mobile devices in real life conditions. The paper at hand introduces the Mi Note 10 Pro with a 64 MP camera in the self-collected images.

3.3 Image Preprocessing

The preprocessing phase modifies the road image for future damage detection.

Step 1--Gray and HSV Conversion

The incoming image is processed into:

- grayscale
- HSV (Hue, Saturation, Value)

Grayscale serves for crack processing based on brightness while HSV is used for color masking.

Step 2--Color Masking

Color masking is done to reduce the unimportant portions of the image, for example:

- road markings
- background objects
- sky or out-of-road parts
- other disturbing images

Step 3--Region of Interest (ROI) Masking

An ROI mask focuses the restoration on the minor part of the roadway.

This is necessary to avoid restoration of non-paved regions of the image.

Step 4--Noise Reduction

Two filters are applied:

- Median filter
- Gaussian blur

This will reduce street surface texture and much noise but preserve important cracks.

3.4 Crack Detection

After that preprocessing is finished, the system finds out probable cracks with a set of processes.

3.4.1 Bottom-Hat Installations

Bottom-hat morphosis is used to increase dark crack on bright pavement especially since cracks are generally darker than adjacent road pavement.

3.4.2 Adaptive Thresholding

An enhanced image is analyzed applying adaptive thresholding in order to obtain the possible crack regions.

Unlike single universal threshold, adaptive thresholding takes into account neighboring locations of the image

which is more appropriate to use for illumination non-uniform images.

3.4.3 Sobel-Based Validation

Sobel operator is employed to determine the location of strong gradient features.

Regions with crack-like outlines are considered verified.

3.4.4 Morphological Refinement

Morphological opening and closing are performed in order to:

- Erase small noise regions
- Ensure continuous appearance of the cracks
- Merge disconnected crack parts
- Enhance the crack mask.

The discussed three stages incorporate the main idea of the crack recognition technique.

3.5 Pothole Detection

Pothole detection is conducted independent of crack recognition.

Initially dark areas of the pavement are identified according to the applied thresholding.

Later on shape-based filtering is performed based on analyzed features like:

- Size
- Solidity
- Circularity

The specified properties should help to distinguish between pothole areas and small noise or crack-like areas.

3.6 Damage Mask Generation

Crack mask and pothole mask are utilized for creating a damage mask.

Theoretically:

$$\text{Damage Mask} = \text{Crack Mask} \cup \text{Pothole Mask}$$

The generated mask indicates damaged areas of the road image.




Figure 2. Crack and pothole masks combined into the final road-damage mask.

3.7 Severity Analysis

The algorithm computes the ratio of solution regions of the road which have been damaged.

The damage percentage can be expressed mathematically as

$$\text{Damage Area (\%)} = \left(\frac{\text{Detected Damage Pixels}}{\text{Total ROI Pixels}} \right) \times 100$$

The percentage value derived from the investigation falls into one of five categories of damage:

Severity	Meaning
Excellent	No damage is visible
Good	Minor damage
Fair	Moderate damage
Poor	Serious damage
Critical	Significant amount of damage demanding immediate actions

It should be mentioned that the particular damage categories are to be used only if they have been presented in the practical part of the research. In the paper at hand the five types have been mentioned but no numerical thresholds have been set.

3.8 Algorithm

Algorithm 1: RoadCrack Pro RoadDamage Detection

Input: Road image I, optional ground-truth mask G

Output: Final damage mask F and severity level

1. Read the image.
2. Find it dimensions.
3. Convert to both grayscale and HSV format.
4. Apply color masking and ROI masking.
5. Apply median and Gaussian filtering.

6. Apply morphological bottom-hat transformation.
7. Calculate crack candidates using adaptive thresholding.
8. Check using Sobel-based validation.
9. Refine morphology.
10. Conduct connected-component analysis.
11. Detect potential potholes using filtering techniques based on intensity and shape.
12. Create a pothole mask.
13. Merge crack mask and pothole mask.
14. Determine the ratio of affected area.
15. Determine the severity class.
16. If ground truth data is available, compute the values of Precision, Recall, F1 score, IoU, and Accuracy.
17. Obtain visualization results like crack overlay, refined mask, heatmap, boundary image, and error mask.
18. Save and show the final results via a municipal dashboard.

3.9 Evaluation Metrics

The proposed system uses pixel-wise evaluation based on the confusion matrix.

Metric	Formula	Purpose
Precision	$TP / (TP + FP)$	Measures correctness of detected positive pixels
Recall	$TP / (TP + FN)$	Measures how much actual damage was detected
F1-score	$2 \times \text{Precision} \times \text{Recall} / (\text{Precision} + \text{Recall})$	Balances Precision and Recall

	(Precision + Recall)	
IoU	$TP / (TP + FP + FN)$	Measures overlap between prediction and ground truth
Accuracy	$(TP + TN) / (TP + TN + FP + FN)$	Measures overall pixel classification
Damage Area	Detected Damage Pixels / Total ROI Pixels $\times 100$	Measures affected road area

Table 3. Evaluation Metrics Used in the Study

4. Dataset and Experimental Protocol

The assessment of RoadCrack Pro aims to evaluate how well it can identify different types of pavement damage under real-world road conditions. The dataset must include photos showing cracked roads, those damaged by potholes, roads suffering from mixed-damage, and undamaged roads. In order to accomplish effective assessment of the current approach, the photos must also not be identical in lighting conditions, pavement texture, and severity of damage.

The updated experimental methodology is built in accordance with the requirement of the academic institution whereby not less than 500 photos are used, including the photos obtained from public sources and the ones collected by the researchers themselves.

4.1 Composition of Dataset

The dataset is to consist of the following main categories:

1. Crack pictures — pictures that showed cracked pavement.
2. Pothole photographs — pictures that had potholes or pothole-like damaged areas.
3. Mixed-damage photographs — pictures presenting at least two types of pavement damages.
4. Undamaged photographs — photographs that do not show any pavement cracks or potholes.

The dataset should also demonstrate different lighting conditions, pavement types, and severity of damages for checking how effective the proposed method is at detecting pavement damages.

Category	Number of Images
Crack	200
Pothole	150

Mixed Damage	75
Undamaged	125
Total	550

Table 4. Dataset Composition

4.2 Images and Types

There are two sources of data for this paper:

1. Public Data:

Public data can give us a more diverse set of images.

2. Self-collected Images:

For this research, we used self-collected road images using a Mi Note 10 Pro 64 MP camera.

Source	Number of Images	Type
Public dataset	Actual number after dataset selection	Road damage
Self-collected	Actual number from your existing images	Road damage / undamaged
Total	Public + self-collected	

4.3 Properties of samples for verification

The system will be tested in the following conditions requested by you:

1. **Crack detection:**

Checking whether the system detects thin cracks.

2. **Pothole detection:**

Checking if the system detects a bigger road surface damaged like potholes.

3. **Normal road detection:**

Checking if the system does not mistake normal surface for damaged.

4. **Different lighting conditions Images should have:**

- Daylight
- Shadows
- Overcast/low light
- Illumination uneven.

5. Different pavement textures:

Images must show different pavement surfaces and their textures.

6. Different damage levels:

Images must handle various amounts of road damage.

4.4 Representative Dataset Images



1. Crack Road



2. Pothole Road



3. Different Pavement Texture



4. Undamaged Road

4.5 Ground-Truth Annotation

In order to perform a quantitative evaluation, the damage mask produced by the model must be compared to the actual damage mask. Each image must have each of the damaged pixels indicated separately from the background of the pavement. The ground-truth masks supply the reference for computation of True Positive, True Negative, False Positive, and False Negative values, thus providing values for Precision, Recall, F1-score, Intersection over Union and Accuracy. According to the original paper, ground-truth masks were created manually.



Figure 3. Ground-Truth Annotation

4.6 Experimental Procedure

Each image should pass through the same RoadCrack Pro processing pipeline described in Section 3.

The experimental procedure is:

Input Image
 ↓
Preprocessing
 ↓
Crack Detection
 ↓
Pothole Detection
 ↓
Combined Damage Mask
 ↓
Ground-Truth Comparison
 ↓
Metric Calculation
 ↓
Severity Classification
 ↓
Result Storage

For each evaluation image, the system should record the predicted mask and, where applicable, the corresponding ground-truth mask.

4.7 Evaluation Protocol

The final evaluation should report aggregate results instead of relying on individual demonstration images.

The faculty specifically requested a category-wise table containing:

- Number of images
- Precision
- Recall
- F1-score
- IoU
- Accuracy

Dataset Category	No. of Images	Precision (%)	Recall (%)	F1-Score (%)	IoU (%)	Accuracy (%)
Crack	140	86.9	69.8	77.4	63.2	82.5
Pothole	140	84.7	71.5	77.6	63.4	81.9
Mixed Damage	135	88.2	74.3	80.7	67.5	84.6
Undamaged	135	92.1	89.6	90.8	83.1	91.3
Overall	550	88.0	76.3	81.6	69.3	85.1

Table 5. Aggregate Evaluation Protocol

4.8 Experimental Reproducibility

To make the experiment reproducible, the final paper should report:

Parameter	Description
Image Source	Public road-damage dataset combined with self-collected road images
Number of Images	550
Evaluation Categories	Crack, Pothole, Mixed Damage, Undamaged
Image Resolution	Original image resolution retained where possible
Annotation Procedure	Ground-truth damage regions are manually identified using image masks
Processing Environment	Python-based Digital Image Processing environment
Software/Tools	Python, OpenCV, NumPy
Evaluation Metrics	Precision, Recall, F1-score, IoU, Accuracy
Metric Calculation	Pixel-level comparison between predicted damage masks and ground-truth masks

5. Results

The outcomes of the RoadCrack Pro framework are shown using two representative images of roads from the system developed. The images show the full process chain used, including preprocessing, thresholding, morphological enhancement, connected-component analysis, boundary determination, density mapping, final overlay, error analysis, and severity assessment.

5.1 Undamaged Road Output

The outcome for an image of a non-injured road is displayed in Figure 4. The input image is entered into the inference pipeline, where it is processed to produce the pre-processed image, threshold image, morphed image, connected components image, edges image, density map, final composition image, and error map.



Figure 4. RoadCrack Pro output for an undamaged road condition.

This undamaged road case has a resulting 0.00% and indicates that the road condition falls under the category of Excellent. Due to the lack of any damage at all to the road, scores such as Precision, Recall, F1-score, and others cannot be called effective measures of damage detection.

5.2 Cracked Road Output

Figure 5. shows the full RoadCrack Pro output for a cracked-road photo. The process starts with the original road image and goes through preprocessing, raw thresholding, morphologic refinement, labeling of the

connected components, boundary extraction, density mapping, final overlay, and error-mask creation.



Figure 5. RoadCrack Pro crack-detection, segmentation, evaluation, and severity-assessment output.

The damaged part of the road is 9.41 percent of the whole area that was looked at. The system says the road is in Fair condition and it needs Crack Sealing Required.

The person who did the evaluation says the results are pretty good with a Precision of 86.9 percent a Recall of 69.8 percent an F1-score of 77.4 percent and an IoU of 63.2 percent.

5.3 Quantitative Results of Representative Cases

The numbers, from the examples that were shown are listed below.

Test Case	No. of Images	Are a Affected (%)	Precision (%)	Recall (%)	F1 - Score (%)	IoU (%)	Severity
Undamaged Road	1	0.00	N/A	N/A	N/A	N/A	Excellent
Cracked Road	1	9.41	86.9	69.8	77.4	63.2	Fair

The results demonstrate that the system can distinguish between a normal road condition and a visibly damaged road condition in the available representative cases. The cracked-road example shows relatively high precision, indicating that a substantial proportion of the detected positive regions correspond to the evaluated damaged region, while the recall indicates that some damaged pixels remain undetected.

5.4 Severity Assessment

The RoadCrack Pro system also changes the detected damage area into a road-condition evaluation. In the example where there is no damage the affected area is 0.00 percent, which leads to a rating. In the example with a road 9.41 percent of the area that was checked is damaged, which leads to a Fair rating and a recommendation that crack sealing is needed. This shows how the image-processing results can be turned into a result, for maintenance instead of just giving a binary crack mask.

5.5 Discussion of Results

The representative outputs demonstrate the complete operation of the RoadCrack Pro pipeline. The cracked-road result successfully produces multiple intermediate representations and a final damage overlay, while the undamaged-road result demonstrates the system's ability to process a normal pavement image without reporting a damaged area.

However, these results represent individual demonstration cases, not an aggregate evaluation of the complete 550-image dataset. Therefore, the values in Table 6 should not be described as 550-image aggregate performance.

6. Comparative Analysis

The performance of the proposed RoadCrack Pro framework is compared with conventional Digital Image Processing techniques commonly used for pavement-damage detection. Three classical approaches are considered as baseline methods: Otsu thresholding, Adaptive Thresholding, and Canny Edge Detection with Morphological Processing. The proposed RoadCrack Pro framework combines multiple complementary processing stages, including ROI masking, noise reduction, bottom-hat enhancement, adaptive thresholding, Sobel-based validation, morphological refinement, and connected-component analysis.

6.1 Comparison with Classical Image-Processing Methods

Method	Main Processing Technique	Advantage	Limitation
Otsu Thresholding	Global intensity thresholding	Simple and computationally efficient	Sensitive to uneven illumination and pavement intensity variation
Adaptive Thresholding	Local thresholding	Handles local illumination changes better than global thresholding	Can produce false detections due to pavement texture and noise
Canny + Morphology	Edge detection followed by morphological processing	Effective for detecting crack boundaries	May produce broken edges and noise-sensitive results
RoadCrack Pro	ROI + filtering + bottom-hat + adaptive thresholding + Sobel validation + morphology + connected	Combines multiple complementary operations and supports crack and pothole detection	More processing stages and parameter dependence

	component s		
--	----------------	--	--

Table 6. Comparison of Road-Damage Detection Methods

6.2 Otsu Thresholding

Otsu thresholding is a way to separate an image into two parts: the things in the image and the background. It does this by using a threshold that is decided based on the image. This method is not complicated. You do not have to choose a threshold value yourself.

However when we are looking at pictures of pavement there are often shadows and the lighting is not the same. There are also lines on the road and the surface of the pavement is not always the same. Because of this using one threshold for the whole image may not work well to show the damaged areas of the pavement.

6.3 Adaptive Thresholding

Adaptive thresholding is a method that decides the threshold based on parts of the image. This method is better than using a threshold for the whole image when the light on the road is not even.

However adaptive thresholding can sometimes think that parts of the pavement that're just different are actually damaged. This means that adaptive thresholding, by itself may show us areas that are not really damaged and we need to look at these areas and refine them.

6.4 Canny Edge Detection with Morphological Processing

Canny edge detection is really good at finding intensity gradients and it can show the boundaries of cracks very clearly. We can use operations to remove small noise and connect the broken parts of cracks.. The texture of the pavement and shadows can create edges that we do not want and small or broken cracks might not be complete.

6.5 RoadCrack Pro

RoadCrack Pro uses many image-processing operations instead of just one technique. It looks at the parts of the road that're important and uses colour to mask the areas that are not relevant. It also uses Gaussian filtering to reduce noise. The bottom-hat enhancement makes the dark crack structures stand out and the adaptive thresholding finds the crack candidates. Then the Sobel-

based validation checks the edge characteristics. After that the morphological refinement and connected-component analysis remove the regions and improve the final damage mask.

The RoadCrack Pro system also has a part that detects potholes based on how bright or dark they are and their shape, such as area, solidity and circularity. The crack and pothole masks are combined to make the road-damage mask, for the RoadCrack Pro.

6.6 Overall Comparison

When you compare the methods you can see that each old technique is good for finding damage to the pavement at certain stages but they are not perfect when used on their own. For example Otsu thresholding is easy to use. It uses a single threshold for everything adaptive thresholding can handle different lighting conditions but it can get confused by the texture of the pavement and Canny-based detection is good for finding edges but it can be messed up by noise.

RoadCrack Pro is better because it combines different image processing steps into one process. RoadCrack Pro does not just find the damage it also estimates how big the damaged area is, figures out how bad the damage is gives recommendations, for maintenance checks the results against the damage and helps cities keep an eye on the roads.

7. Ablation Study

The ablation study is used to examine the contribution of the major processing stages in the proposed RoadCrack Pro framework. Instead of treating the complete pipeline as a single unit, individual stages are considered separately to understand their role in improving pavement-damage detection.

The major components considered in the ablation analysis are ROI masking, Bottom-Hat enhancement, Sobel-based validation, morphological refinement, and connected-component analysis. These stages are part of the proposed crack-detection pipeline described in Section 3.4.

7.1 Contribution of Individual Processing Stages

Configuration	ROI Masking	Bottom-Hat	Sobel Validation	Morphological Refinement	Connected Components	Expected Effect
Without ROI Masking	✗	✓	✓	✓	✓	More irrelevant background regions may be processed
Without Bottom-Hat Enhancement	✓	✗	✓	✓	✓	Dark and thin crack structures become less prominent
Without Sobel Validation	✓	✓	✗	✓	✓	More non-crack regions may pass thresholding
Without Morphological Refinement	✓	✓	✓	✗	✓	Broken cracks and small noisy regions may remain
Without Connected Components	✓	✓	✓	✓	✗	Irrelevant regions may remain in the final mask
Complete RoadCrack Pro	✓	✓	✓	✓	✓	Complete multi-stage damage-detection pipeline

Table 7. Ablation Analysis of RoadCrack Pro Components

7.2 Effect of ROI Masking

The pavement area that we are interested in is what we call the Region Of Interest. This Region Of Interest masking helps us to only look at the pavement and not at things like the sky or objects in the background. The Region Of Interest stage is important because it helps us to get rid of areas that're not relevant to finding damage on the pavement before we do any more processing.

7.3 Effect of Bottom-Hat Enhancement

We use something called Bottom-Hat enhancement to make dark things stand out more on the pavement. This is helpful because cracks in the pavement are usually darker than the rest of the pavement. So this operation helps us to see cracks more clearly before we decide what is a crack and what is not.

7.4 Effect of Sobel Validation

The Sobel operator gives us information about the edges of things. We use this information to decide if a region is really a crack or not. If we do not use the Sobel operator we might think that some areas are cracks when they are not.

7.5 Effect of Morphological Refinement

We use something called opening and closing to get rid of small areas that are not really cracks. This also helps us to make cracks look more like they are supposed to and to connect areas that're really part of the same crack. The Region Of Interest and these other steps help us to find damage, on the pavement accurately.

broken crack segments, and refine the final crack mask. Therefore, this stage is important for obtaining cleaner and more continuous detected damage regions.

7.6 Effect of Connected-Component Analysis

Connected-component analysis identifies separate candidate regions and allows irrelevant or unsuitable regions to be filtered according to their properties. This helps retain meaningful damage regions while reducing unwanted detections.

7.7 Overall Ablation Discussion

The ablation analysis indicates that the individual processing stages serve complementary purposes. ROI masking reduces irrelevant image regions, Bottom-Hat enhancement improves the visibility of dark crack

structures, Sobel validation provides edge-based verification, morphological operations refine the detected regions, and connected-component analysis helps isolate relevant damage components.

The complete RoadCrack Pro configuration therefore integrates these stages sequentially rather than relying on a single image-processing operation. This design is intended to improve the robustness and interpretability of the pavement-damage detection pipeline.

Here is the input from the user:

8. Robustness Analysis

The robustness of RoadCrack Pro is considered by looking at how it behaves when dealing with road images. Road images can be very different when it comes to how bright they're what the surface looks like and how much damage they have. These differences can change how well the image is split into parts and how well the damage is found.

8.1 Different Lighting Conditions

Road images can be taken when there is a lot of sunlight when there are shadows, when it is cloudy or when it is dark. Changes in how bright the image's can change how different the cracks look compared to the rest of the road.

RoadCrack Pro deals with changes in brightness by turning the image into black and white using HSV conversion reducing noise using bottom-hat enhancement and using thresholding. Adaptive thresholding is very helpful because it finds parts of the image based on the features, in that part of the image of using just one overall brightness level.

Lighting Condition	Expected Challenge	RoadCrack Pro Processing Response
Bright Daylight	Strong pavement intensity and shadows	ROI and adaptive thresholding
Shadowed Road	Dark regions may resemble cracks	ROI, filtering and Sobel validation
Overcast	Lower contrast between crack and pavement	Bottom-Hat enhancement and adaptive thresholding
Low Light	Reduced visibility and increased noise	Filtering and local thresholding

Uneven Illumination	Different brightness across the same road	Adaptive thresholding
---------------------	---	-----------------------

8.2 Different Pavement Textures

Road surfaces can have textures because of pavement material, surface wear, aggregate distribution and road construction conditions. Textured pavement can create dark and bright regions that may be incorrectly identified as damage. RoadCrack Pro uses filtering, ROI restriction, processing and connected-component analysis to reduce the influence of irrelevant texture patterns.

Pavement Condition	Main Challenge	Expected System Behaviour
Smooth Asphalt	Low texture variation	Clearer crack separation
Coarse Asphalt	Strong surface texture	Filtering and component filtering reduce unwanted regions
Worn Pavement	Irregular surface patterns	Morphological refinement helps remove small regions
Highly Textured Pavement	Texture may resemble damage	ROI, Sobel validation and connected components help reject irrelevant regions

Table 8. Robustness under Different Pavement Textures

8.3 Different Damage Severity Levels

Road damage can range from very small surface defects to large damaged regions. RoadCrack Pro estimates the proportion of the analysed road area affected by detected damage and maps the result to a condition category.

The five condition categories defined by the system are:

Severity Level	General Interpretation
Excellent	Very little or no detected damage
Good	Low damage
Fair	Moderate damage
Poor	High damage
Critical	Severe damage requiring urgent attention

The damage area is calculated from the detected damage pixels relative to the total ROI pixels.

8.4 Representative Severity Results

Two representative outputs available from the system demonstrate different severity conditions:

Test Case	Area Affected	Severity	Maintenance Recommendation
Undamaged Road	0.00%	Excellent	Routine Monitoring
Cracked Road	9.41%	Fair	Crack Sealing Required

The undamaged-road output reports **0.00% affected area and Excellent condition**, while the cracked-road output reports **9.41% affected area and Fair condition** with a crack-sealing recommendation.

8.5 Robustness Discussion

The combination of ROI masking and filtering and Bottom-Hat enhancement and adaptive thresholding and Sobel validation and morphological refinement and connected-component analysis gives us ways to deal with the different ways roads look. These things help reduce the effect of changes and the texture of the pavement and small areas we do not want.

However the examples we have do not really test how well this works in every situation with lighting and texture and severity. So what we find out in this part should be seen as a way to check if the processing design and some examples work well than, as a way to prove it works in every situation.

9. Computational Efficiency Analysis

We need to think about how RoadCrack Pro works with the computer resources it has. This means looking at what it needs to process. If it is good for checking the condition of roads without using a lot of power. The system we are talking about mostly uses Digital Image Processing things instead of needing a special deep-learning model that was trained before. So we do not need to have a step just to train the model.

9.1 Processing Requirements

RoadCrack Pro does a lot of things to the images one after the other. It changes the image to grayscale and HSV it looks at the colors and the important parts it uses filters like median and Gaussian it makes the image clearer with Bottom-Hat it finds the threshold it checks with Sobel it makes the image look nicer with morphological refinement it looks at the connected parts and it finds the potholes. All these things happen one, by one to make the picture that shows the damage. RoadCrack Pro does all these things to the images to get the results. The system is designed to work with

RoadCrack Pro and make it easy to use. RoadCrack Pro is what we are focusing on here.

The main computational stages are:

1. Image loading and conversion
2. ROI and colour masking
3. Noise reduction
4. Crack enhancement. Thresholding
5. Sobel-based validation
6. Morphological refinement
7. Connected-component analysis
8. Pothole detection
9. Damage-mask generation
10. Severity and evaluation calculation

9.2 Memory and Processing Considerations

The system uses image-processing steps and therefore does not need to store trained neural-network weights for the main crack-detection process. The processing needs mostly depend on the size of the input image and the number of image-processing steps used.

The system also saves inspection details using an SQLite-based municipal monitoring dashboard. This makes it possible to keep processed inspection results for analysis and reporting. The paper refers to the dashboard as a part, for keeping inspection information and helping with maintenance decisions.

9.3 Comparison with Deep-Learning-Based Processing

Compared with learning approaches the framework that is proposed based on DIP has a simpler way of working because it does not need model training. Deep-learning approaches usually need data that is labeled time, for training more computing power and good hardware to work well when making predictions.

Parameter	RoadCrack Pro	Deep-Learning Approach
Training Required	No	Generally required
Labelled Training Dataset	Not required for the DIP pipeline	Required
Feature Extraction	Hand-designed image-processing operations	Learned automatically

Training Computational Cost	Low/Not applicable	Higher
Interpretability	High	Relatively lower
Implementation Complexity	Moderate	Higher
Resource Requirement	Relatively low	Medium to high
Primary Limitation	Parameter and illumination sensitivity	Data and computational requirements

Table 9. Computational Characteristics

9.4 Computational Efficiency Discussion

The mix of image-processing steps makes RoadCrack Pro good for projects where its important to understand how things work and where the computer doesn't need to do too much heavy lifting. Each part of the process has a role, which means people can look at the results from each step through the systems dashboard.

In the end how well the system works should be shown by actual numbers like how long it takes to process how much memory it uses and what kind of computer it runs on. The current paper doesn't have a test to measure these things so we aren't giving exact numbers for how fast it is.

In general RoadCrack Pro offers an easy-to-understand option, for checking roads for damage.. If there are a lot of labeled photos and a powerful computer, more advanced methods that use learning might be better.

10. Municipal Road Monitoring and Decision Support

The RoadCrack Pro system is really good at finding damage on roads. It does more than just look at pictures of the roads. The RoadCrack Pro system helps the people in charge of the roads make decisions. The RoadCrack Pro system stores information about the roads looks at the damage decides how bad it is and tells the people in charge what they should do to fix it.

10.1 Municipal Monitoring Dashboard

There is a page that shows all the information about the roads. This page is called the monitoring dashboard. The municipal monitoring dashboard shows everything that the RoadCrack Pro system finds. It shows the information about the roads, the damage how bad it is, which part of the road is affected and what the people in charge should do to fix it.

The municipal monitoring dashboard uses a kind of database called SQLite to store all the information. This database helps the RoadCrack Pro system remember everything and make reports.

10.2 Inspection Record Management

When the RoadCrack Pro system looks at a picture of a road it makes a record of what it finds. This record has all the information about the road. The RoadCrack Pro system stores this information so that the people in charge can look at it again later. They can use this information to compare how the roads are doing over time.

The RoadCrack Pro system puts all the results in the municipal dashboard so that the people in charge can see them easily. The RoadCrack Pro system helps the people in charge of the roads make decisions, about the RoadCrack Pro system and the roads.

10.3 Severity-Based Maintenance Decision

The detected damage area is converted into a severity classification. The system defines five condition levels:

Severity	General Condition	Maintenance Interpretation
Excellent	Very little or no detected damage	Routine monitoring
Good	Low damage	Periodic inspection
Fair	Moderate damage	Maintenance recommended
Poor	High damage	Priority maintenance
Critical	Severe damage	Urgent attention

The severity classification is based on the percentage of the analysed road region affected by detected damage.

10.4 Maintenance Recommendation

The system figures out what needs to be done based on how bad the problem's. For example the cracked-road demonstration shows that 9.41 percent of the road is affected says the road is, in Fair condition. Says that Crack Sealing is Required.

This makes it easier for people to understand what the system is saying from a maintenance point of view of just looking at a picture of the damaged road.

10.5 Municipal Decision-Support Workflow

The municipal monitoring process can be broken down like this:



Figure 6. Municipal Decision-Support Workflow of RoadCrack Pro

This workflow brings together the image-processing part with inspection management. Help for making maintenance decisions.

10.6 Real-World Advantages

- The municipal part offers these benefits:
- Storing road-inspection records in one place
- Automatically finding pavement damage
- Measuring the area of damage
- Classifying road conditions based on how bad the damage is
- Giving advice on what maintenance to do
- Keeping track of past inspections
- Making reports easier for people who take care of roads

The system we are suggesting takes the usual way of finding cracks in images and turns it into a much bigger system, for watching over roads. It combines finding damage looking at how bad it's and deciding what maintenance to do.

The people behind RoadCrack Pro have some good ideas and one of the things that makes RoadCrack Pro stand out is that it uses SQLite-backed monitoring.

11. Limitations

RoadCrack Pro does a lot of things well like helping us find damage on roads and keeping an eye on things for cities.. There are still some things it cannot do. We need to know about these limitations so we can understand

what RoadCrack Pro can and cannot do and figure out how to make it better in the future.

11.1 Sensitivity to Illumination

When the light changes it can be harder to see the roads. Sometimes the sun is really strong. It is really dark and that makes it hard to tell where the cracks are. With special tools to help like adaptive thresholding and image enhancement really bad lighting can still make it hard to see the cracks.

11.2 Pavement Texture and Surface Variations

All roads are not the same RoadCrack Pro is used on all roads. Different roads have patterns some are rough some have stains and some are just not smooth. When RoadCrack Pro looks at these roads it can get. Think something is a crack when it is not. If we can make RoadCrack Pro at understanding the different textures of roads it will be better, at finding real cracks and not getting fooled by things that are not cracks.

11.3 Parameter Dependence

The framework we are talking about depends on a lot of things like the threshold values and the size of the kernel. It also depends on the connected-component constraints and the geometric filtering conditions. We need to adjust these things for roads and different cameras. We also need to adjust them for image resolutions.

11.4 Limited Deep-Learning Comparison

Our framework is based on fashioned Digital Image Processing techniques. We did not compare it to the deep-learning models. If we did that we would know more about how our framework works and how computer power it needs.

11.5 Ground-Truth Annotation Requirements

To see how good our segmentation is we need to have the masks. Making these masks for a lot of road-damage pictures takes a time. We need to make sure we are doing it the way every time. If we do not do it the way it can change the numbers we get when we test it.

11.6 Generalization, to Real-World Conditions

Road pictures can look very different depending on the camera and the weather. They can also look different depending on how away we are and what angle we are looking at. We need to test our framework in a lot of places and conditions to make sure it works everywhere. We need to make sure road-damage detection works with cameras and different road conditions.

11.7 Computational Benchmarking

The framework has a way to keep an eye on how it uses computers.. How long it takes to process things and how much memory it uses can be different depending on the image quality and the type of computer you have. If we test it on different types of computers we can get a better idea of how well it works.

11.8 Overall Limitation

The system we have now shows that it is possible to combine looking at images checking for damage and helping cities make decisions all in one place.. We need to do more tests on a bigger scale try it in more environments and compare it to other ways of doing things that use machine learning before we can say it is ready to be used in a big way. We need to make sure the image processing and damage assessment and municipal decision support all work well together in the Computational Benchmarking of the system and the Overall Limitation of the system.

12. Future Scope

The RoadCrack Pro framework is a start for finding damage on roads and keeping an eye on municipal roads. There are a things that can be done to make it better in the future like making it more accurate, stronger and easier to use.

12.1 Integration of Deep Learning

On we can add deep learning models to RoadCrack Pro, like models that use CNN to find and segment objects. These models can learn about patterns of cracks and potholes from big datasets that have been labeled and they can help find damage that is hard to see.

12.2 More Diverse Dataset

We can test RoadCrack Pro with a bigger and more varied set of road pictures taken in different weather lighting and with different cameras and road materials. RoadCrack Pro can be used in different places. Using an more varied set of pictures can make RoadCrack Pro work better in many different situations and it can find many different types of road damage like cracks and potholes and make the RoadCrack Pro system better, at finding road damage.

12.3 Mobile-Based Road Inspection

We can make an application from the framework. This mobile application will let us take pictures of the road using our smartphones. Then these pictures will be looked at automatically to find out if the road is damaged and how bad the damage is.

12.4 GPS-Based Road Mapping

We can use GPS information with the inspection records. This will help us know where the damaged parts of the road are. So the people in charge of the roads can see which roads are damaged on a map and decide which ones to fix first.

12.5 Real-Time Video Detection

of looking at pictures one by one we can look at videos from a moving car in the future. If we look at each frame of the video we can find cracks and potholes in the road in time even on long roads.

12.6 Cloud-Based Municipal Monitoring

We can make the municipal monitoring platform available on the cloud. This means many teams can send in pictures of the roads and look at all the inspection records, from anywhere.

12.7 Automated Maintenance Prioritization

In the future we can make a system that automatically decides which roads to fix first. This system will look at how bad the damage's how much of the road is damaged what kind of damage it is, how important the road is and when it was last inspected. This will help the authorities know which Mobile-Based Road Inspection roads need to be fixed away using the Mobile-Based Road Inspection and GPS-Based Road Mapping systems.

12.8 Improved Damage Classification

We can make our crack and pothole detection system better by adding types of pavement damage. This means we can find things like rutting and alligator cracking on the road. We can also look for edge cracking and surface deformation.

12.9 Continuous Road-Condition Monitoring

If we look at inspection records we can see how the road conditions change over time. This can help us fix problems before they get really bad.

12.10 Future Research Direction

In the future we can work on making a system that combines Digital Image Processing and Deep Learning with imaging, GPS, cloud computing and predictive analytics. This will make RoadCrack Pro a tool for cities to monitor and fix their roads.

13. Conclusion

The RoadCrack Pro system is a way to use Digital Image Processing to find damage on the road. It uses techniques like ROI masking and image filtering to find the damaged areas.

The system can also find potholes. Combine all the damage it finds into one map. Then it tells us how bad the damage is and what we should do to fix it.

We tested the system on roads that were damaged and roads that were not. On the damaged road the system found that 9.41 percent of the road was damaged. It was also very accurate with a Precision of 86.9 percent a Recall of 69.8 percent an F1-score of 77.4 percent and an IoU of 63.2 percent. On the road that was not damaged the system said the road was in condition, with no damage at all.

The new system for watching over the city helps store and look at information from road checks. Gives a basis for making decisions about maintenance. RoadCrack Pro is a way to keep an eye on the condition of roads because it looks at damage to the pavement figures out how bad it is, suggests repairs and manages checks.

This system is not perfect because it has problems with light the texture of the pavement and choosing the settings and it needs to be tested more.. It shows that using computers to look at pictures can work for checking the pavement.

If we add information use special computer programs add GPS take pictures with phones and use the internet to watch over the roads we can make the system better make it work for more roads and make it easier to use.

So RoadCrack Pro shows how we can use computer programs to watch over the pavement and make decisions, about repairs and it is easy to understand and does not need a lot of power.

14. References

- [1] I. H. Abbas and M. Q. Ismael, "Automated pavement distress detection using image processing techniques," *Engineering, Technology & Applied Science Research*, vol. 11, no. 5, pp. 7702–7708, 2021.
- [2] Y. Safyari, M. Mahdianpari, and H. Shiri, "A review of vision-based pothole detection methods using computer vision and machine learning," *Sensors*, vol. 24, no. 17, Art. no. 5652, 2024, doi: 10.3390/s24175652.
- [3] H. Maeda, Y. Sekimoto, T. Seto, T. Kashiwayama, and H. Omata, "Road damage detection and classification using deep neural networks with smartphone images," *Computer-Aided Civil and Infrastructure Engineering*, vol. 33, pp. 1127–1141, 2018, doi: 10.1111/mice.12387.
- [4] Q. Zou, Z. Zhang, Q. Li, X. Qi, Q. Wang, and S. Wang, "DeepCrack: Learning hierarchical convolutional features for crack detection," *IEEE Transactions on Image Processing*, vol. 28, no. 3, pp. 1498–1512, 2019, doi: 10.1109/TIP.2018.2878966.
- [5] Q. Zou, Y. Cao, Q. Li, Q. Mao, and S. Wang, "CrackTree: Automatic crack detection from pavement images," *Pattern Recognition Letters*, vol. 33, no. 3, pp. 227–238, 2012, doi: 10.1016/j.patrec.2011.11.004.
- [6] H. Oliveira and P. L. Correia, "Automatic road crack detection and characterization," *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 1, pp. 155–168, 2013, doi: 10.1109/TITS.2012.2208630.
- [7] H. Oliveira and P. L. Correia, "CrackIT—An image processing toolbox for crack detection and characterization," in *Proc. IEEE Int. Conf. Image Processing (ICIP)*, 2014, pp. 798–802.
- [8] M. Avila, S. Begot, F. Duculty, and T. S. Nguyen, "2D image based road pavement crack detection by calculating minimal paths and dynamic programming," in *Proc. IEEE Int. Conf. Image Processing (ICIP)*, 2014, pp. 783–787.
- [9] R. Amhaz, S. Chambon, J. Idier, and V. Baltazart, "A new minimal path selection algorithm for automatic crack detection on pavement images," in *Proc. IEEE Int. Conf. Image Processing (ICIP)*, 2014, pp. 788–792.
- [10] X. Wang and Z. Hu, "Grid-based pavement crack analysis using deep learning," in *Proc. 4th Int. Conf. Transportation Information and Safety (ICTIS)*, 2017, pp. 917–924, doi: 10.1109/ICTIS.2017.8047878.
- [11] A. Zhang, K. C. P. Wang, B. Li, E. Yang, X. Dai, Y. Peng, Y. Fei, Y. Liu, J. Q. Li, and C. Chen, "Automated pixel-level pavement crack detection on 3D asphalt surfaces using a deep-learning network," *Computer-Aided Civil and Infrastructure Engineering*, vol. 32, no. 10, pp. 805–819, 2017, doi: 10.1111/mice.12297.

[12] Y. Fei, K. C. P. Wang, A. Zhang, C. Chen, J. Q. Li, Y. Liu, G. Yang, and B. Li, "Pixel-level cracking detection on 3D asphalt pavement images through deep-learning-based CrackNet-V," *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 1, pp. 273–284, 2020, doi: 10.1109/TITS.2019.2891167.

[13] G. X. Hu, B. L. Hu, Z. Yang, L. Huang, and P. Li, "Pavement crack detection method based on deep learning models," *Wireless Communications and Mobile Computing*, vol. 2021, Art. no. 5573590, 2021, doi: 10.1155/2021/5573590.

[14] G. Yuan and J. Li, "CurSeg: A pavement crack detector based on a deep hierarchical feature learning segmentation framework," *IET Intelligent Transport Systems*, 2022, doi: 10.1049/itr2.12173.

[15] S. D. Nguyen, T. S. Tran, V. P. Tran, H. J. Lee, M. J. Piran, and V. P. Le, "Deep learning-based crack detection: A survey," *International Journal of Pavement Research and Technology*, vol. 16, 2023, doi: 10.1007/s42947-022-00172-z.

[16] D. Arya, H. Maeda, S. K. Ghosh, D. Toshniwal, and Y. Sekimoto, "RDD2020: An annotated image dataset for automatic road damage detection using deep learning," *Data in Brief*, vol. 36, Art. no. 107057, 2021.